RADIO FREQUENCY SENSING MEASUREMENTS AND METHODS FOR LOCATION CLASSIFICATION IN WIRELESS NETWORKS

by

Dustin C. Maas

A dissertation submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Electrical and Computer Engineering

The University of Utah

May 2014

The University of Utah Graduate School

STATEMENT OF DISSERTATION APPROVAL

The dissertation of **Dustin C. Maas**

has been approved by the following supervisory committee members:

Neal Patwari	, Chair	**3/5/2014** Date Approved
Sneha K. Kasera	, Member	**3/5/2014** Date Approved
Behrouz Farhang	, Member	**3/5/2014** Date Approved
Cynthia Furse	, Member	**3/5/2014** Date Approved
Rong-Rong Chen	, Member	**3/5/2014** Date Approved

and by **Gianluca Lazzi**, Chair/Dean of

the Department/College/School of **Electrical and Computer Engineering**

and by David B. Kieda, Dean of The Graduate School.

ABSTRACT

The wireless radio channel is typically thought of as a means to move information from transmitter to receiver, but the radio channel can also be used to detect changes in the environment of the radio link. This dissertation is focused on the measurements we can make at the physical layer of wireless networks, and how we can use those measurements to obtain information about the locations of transceivers and people.

The first contribution of this work is the development and testing of an open source, 802.11b sounder and receiver, which is capable of decoding packets and using them to estimate the channel impulse response (CIR) of a radio link at a fraction of the cost of traditional channel sounders. This receiver improves on previous implementations by performing optimized matched filtering on the field-programmable gate array (FPGA) of the Universal Software Radio Peripheral (USRP), allowing it to operate at full bandwidth.

The second contribution of this work is an extensive experimental evaluation of a technology called location distinction, i.e., the ability to identify changes in radio transceiver position, via CIR measurements. Previous location distinction work has focused on single-input single-output (SISO) radio links. We extend this work to the context of multiple-input multiple-output (MIMO) radio links, and study system design trade-offs which affect the performance of MIMO location distinction.

The third contribution of this work introduces the "exploiting radio windows" (ERW) attack, in which an attacker outside of a building surreptitiously uses the transmissions of an otherwise secure wireless network inside of the building to infer location information about people inside the building. This is possible because of the relative transparency of external walls to radio transmissions.

The final contribution of this dissertation is a feasibility study for building a rapidly deployable radio tomographic (RTI) imaging system for special operations

forces (SOF). We show that it is possible to obtain valuable tracking information using as few as 10 radios over a single floor of a typical suburban home, even without precise radio location measurements.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to first thank my advisor, Neal Patwari, for his encouragement to consider graduate school in the first place. If he had not given me the initial opportunity for research that he did, I do not think I would have continued my studies further than my Bachelor's degree, and I know that I would have regretted not going to graduate school. I also want to acknowledge and thank him for his continued support and counsel throughout my graduate career. His advice about the trajectory of my work and his great attitude about research and life in general have been invaluable to me.

I would also like to thank my fellow graduate students in the SPAN lab at the University of Utah. Brad, Jessica, Joey, Maurizio, Merrick, Peter, Piyush, and Yang have made life in the SPAN lab a lot of fun, and their feedback on my presentations and papers has been vital to my success.

Finally, I would like to thank my family. Thayne and Karen have offered nothing but encouragement from the outset. I would not have been able to accomplish this without them. Jamie, Chad, Cory, and Marilyn helped inspire the confidence necessary to complete this goal.

CHAPTER 1

INTRODUCTION

1.1 Radio Frequency Sensing in Wireless Networks

Wireless networks permeate the world around us. The past two decades have seen extensive expansion of cellular networks to cover most urban areas around the world and many rural areas as well. We use local area networking (LAN) technologies like WiFi to connect our personal computers and handhelds to the internet. We use personal area networking (PAN) technologies like Bluetooth to connect peripheral devices to our computers or stream audio from our cell phones. We also leverage wireless sensor networks (WSNs) for tasks as diverse as detecting forest fires [1], tracking the migration and mating patterns of animals [2], and monitoring the structural health of buildings and bridges [3].

All of these wireless networks transmit information at radio frequency (RF), i.e., between 3 kHz and 300 GHz, typically less than 6 GHz. In most cases, to the dismay of wireless communications engineers, these RF transmissions are distorted by their environment. Objects in the environment, including human beings, reflect, diffract, and attenuate these transmissions. Moving objects produce Doppler distortion. Typically, it is the job of the wireless communications engineer to design modems that mitigate the effects these distortions have on the packets of data being push around the network, via equalization for example. However, it is also possible to leverage the distortions in RF transmissions caused by the environment to "sense" the environment.

A commonly used model that captures the distortions of wireless signals caused by the environment is the channel impulse response (CIR). The CIR of a wireless link is linear representation of the "echoing" that an RF transmission experiences as

it follows multiple paths from the transmitter to the receiver, known as multipath propagation [4]. Its complex baseband representation can be written as

$$h(t,\tau) = \sum_{i=1}^{N(t)} \alpha_i(t)\delta(t - \tau_i(t)) \tag{1.1}$$

where, at time t, $N(t)$ represents the number of paths, $\alpha_i(t)$ the complex gain of the ith path, $\tau_i(t)$ represents the corresponding time delay, and δ is the Dirac delta function.

The effects of objects in the environment are apparent in measurements that estimate the CIR. For example, when a person moves in the environment of a wireless link, she will effect some subset of the multipath for the radio link, thereby changing their contribution to (1.1). Figure 1.1 illustrates multipath propagation in a single room.

Ultra wide-band (UWB) radar devices create very accurate estimates of the CIR by transmitting very short (in time) RF pulses, in order to approximate an impulse, and listening for reflections [5]. The high bandwidth signals employed by these devices allow them to very accurately localize reflectors in the environment.

However, lower bandwidth signals are also useful for detecting and localizing objects. WiFi signals with bandwidths of 20-40 MHz have been a used to measure the channel and detect changes in transceiver position [6,7], as well as the motion and position of people that are not carrying radios [8]. The lower bandwidth signals used in these papers translates to the use of a lower-order approximation of the impulse function in (1.1), since a true impulse requires infinite bandwidth. The channel sounder developed as a part of this dissertation replaces the impulse function with the waveform shown in Figure 1.2.

Another common metric used to detect changes in wireless transmissions is received signal strength (RSS). RSS is an estimate of the power of the transmission measured at the receiver. The measured power results from multiplying the transmitted power by the squared magnitude of the phasor sum of all of the multipath contributions, written as

$$P_R = P_T \left| \sum_{i=1}^{N(t)} \alpha_i(t) e^{-j2\pi f \tau_i(t)} \right|^2, \tag{1.2}$$

where P_R, P_T, and f represent the received power, transmitted power, and center frequency, respectively.

Changes in the mean and variance of RSS have been be used to track the locations of people who are not carrying radios [9,10]. This kind of localization has been called "device-free localization" (DFL), as it does not require the targets of the localization to carry a device. This attribute makes DFL highly valuable in the realms of law enforcement, security, home automation, retail analytics, and in-home monitoring of the elderly.

1.2 Leveraging RF Measurements for Location Classification

Several methods have been introduced in the literature that leverage RF measurements to infer location information. The following examples employ RSS measurements for location classification.

- In [11], WiFi clients are localized indoors via RSS measurements from multiple access points (APs).
- In [12], RSS measurements made at multiple APs are used to detect the motion of people who are not carrying radio devices and localize them to spatial regions within a small office building.
- In [10], a wireless mesh network surrounding a building is used to localize people moving within the building.

Other methods have been introduced which utilize the CIR or its Fourier pair to infer location information:

- In [13], the authors use frequency response measurements between the WiFi clients and nearby APs to localize the clients.
- In [6] and [14], the authors use CIR estimates to perform location distinction, i.e., determining if the position of a transceiver has changed.
- In [15], the authors use channel frequency response measurements to defend against Sybil attacks.

1.3 Contributions of this Dissertation

This dissertation focuses on the physical layer measurements we can make on the RF links that exist in wireless networks, as well as some of the ways that we can leverage these measurements in order to obtain information about the physical locations of people in the vicinity of the network and the transceivers that comprise the network. The publications, accepted and under submission, that have resulted from this work, as well as the specific contributions made by this author, are summarized below.

1. D. Maas, M.H. Firooz, J. Zhang, N. Patwari, and S.K. Kasera, "Channel Sounding for the Masses: Low Complexity GNU 802.11b Channel Impulse Response Estimation," *IEEE Transactions on Wireless Communications*, 2012. [16]

   RSS measurements are often used in the literature because they are made available at the application layer of the network stack by most wireless devices, requiring no special hardware to measure. The same is not true of CIR measurements. Although CIR measurements are very often made by RF devices for the purpose of equalization, these measurements are not made available at the application layer. It has therefore been necessary for researchers to build their own "channel sounders" in order to build prototypical systems that utilize CIR measurements.

   In this work, I helped build and test an open source 802.11b receiver and channel sounder, capable of estimating channel impulse responses from standard transmissions, including debugging the FPGA design and signal processing through experimental validation. This work is included as Chapter 2 of this dissertation.

2. D. Maas, N. Patwari, J. Zhang, S. Kasera, and M. Jensen, "Location Distinction in a MIMO Channel," in *Proc. 2009 Virginia Tech Wireless Symp.* Student Poster. [17]

   D. Maas, N. Patwari, S.K. Kasera, D. Wasden, and M. Jensen, "Experimental Performance Evaluation of Location Distinction for MIMO Links," in *Proc.*

4th IEEE International Conference on Communication Systems and Networks (COMSNETS), 2012. [18]

It is possible to use CIR measurements to perform location distinction because these measurements are directly related to the physical arrangement of the transceivers and interfering objects in the environment and are also spatially unique in an environment that contains many interfering objects. If two CIR measurements for the same radio link differ significantly, it is likely that the transmitter or receiver has moved during the interval between measurements. It is also possible to discriminate between changes in transceiver position and changes in the environment because changes in transceiver location affect all of the reflective paths traversed, while changes in the environment only affect some subset of the paths.

In this work, I explored the design space for location distinction using two different MIMO testbeds, leading to an understanding of how location distinction performance scales with transmission bandwidth and number of antennas, as well as other insights. This work is included as Chapter 3 of this dissertation.

3. A. Banerjee, D. Maas, M. Bocca, N. Patwari, S.K. Kasera, "Exploiting Radio Windows for Through-wall Location Information," (submitted Sep. 2013) *IEEE Transactions on Networking.*

   The RF transmissions from many of our wireless networks are able to propagate through exterior walls, they create so-called "radio windows," exposing some information about the locations of people inside the buildings where the networks are deployed. This location information can be inferred from passive RF measurements made with receivers placed on the outside of the building.

   In this multiple-author collaborative research project, I developed an algorithm capable of detecting the motion of people across through-wall WiFi links, including the ability to accurately count the number of times a person crosses the line-of-site of the link and determine the person's direction of motion relative to a single through-wall MIMO link. I led experiments and performed

analysis to validate the methodology. This work is included as Chapter 4 of this dissertation.

4. D. Maas, J. Wilson, N. Patwari, "Toward a Rapidly Deployable RTI System for Tactical Operations," in *Proc. 8th IEEE International Workshop on Practical Issues in Building Sensor Network Applications* (SenseApp), 2013. [19]

   Previous work on radio tomographic imaging has been focused on tracking results obtained by complicated and sensitive research-driven systems, in which radio locations must be measured precisely and the system must be properly calibrated before use.

   In this work, I conducted an extensive feasibility study for making a rapidly deployable and robust RTI system for special operations forces (SOF) like SWAT. I performed a variety of experiments to determine whether or not the radios can self-localize and calibrate on-the-fly and still yield the tracking results necessary to make the system useful to SOF. This work is included as Chapter 5 of this dissertation.

   An addendum to the published paper, which includes additional tracking results and link budget information, is included in Appendix A.

In support fot the research projects listed above, I have developed, implemented, and/or tested several channel measurement systems. These include:

- USRP / GNU Radio Channel Sounder: FPGA code for the USRP (v1) to estimate the CIR from received 802.11b signals, code that is publicly available at `http://span.ece.utah.edu/download/form7.html`.

- MIMO Measurements for Location Distinction: A system that measures the channel response on each antenna pair of a 2x2 MIMO link, using LabView and National Instruments VSA/VSG hardware.

- MIMO / OFDM Measurements using CSI Tool: A laptop-based system that measures the channel state information for up to a 3x3 MIMO link, with 30 complex-valued channel state values from each antenna pair, using an Intel 5300 WiFi card and drivers available at [20].

- High TX Power WSN: A Texas Instruments CC2530/CC2590-based system with directional antennas for through-building RTI.

- Multitransceiver SISO Channel Sounder: A pseudo noise (PN)-based multiple-transceiver 20 MHz channel sounder using a network of National Instruments VSAs/VSGs.

1.4 NSF Acknowledgment

This material is based upon work supported by the National Science Foundation under Grant Nos. 0748206, 1035565, and 1315685. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

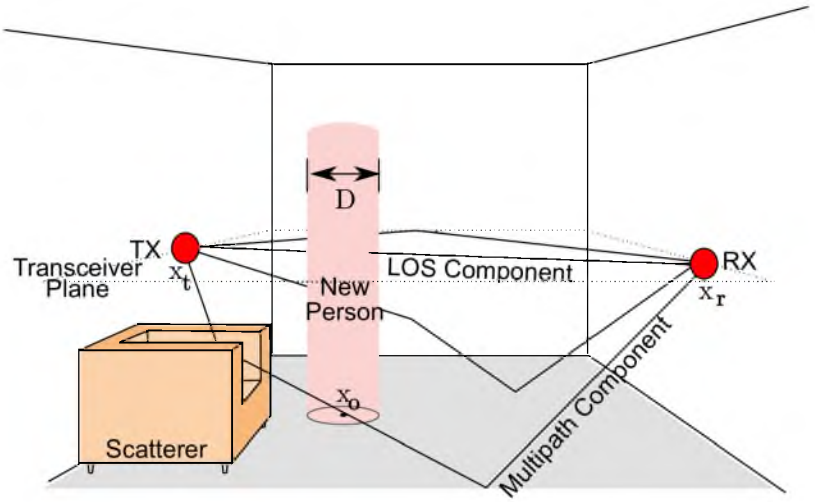

Figure 1.1. A diagram illustrating the multipath propagation of RF transmissions. The transmission follows many paths from the transmitter (TX) to the receiver (RX), each of which is determined by the environment. A person, modeled here as a cylinder with diameter D, attenuates the line-of-site (LOS) component.

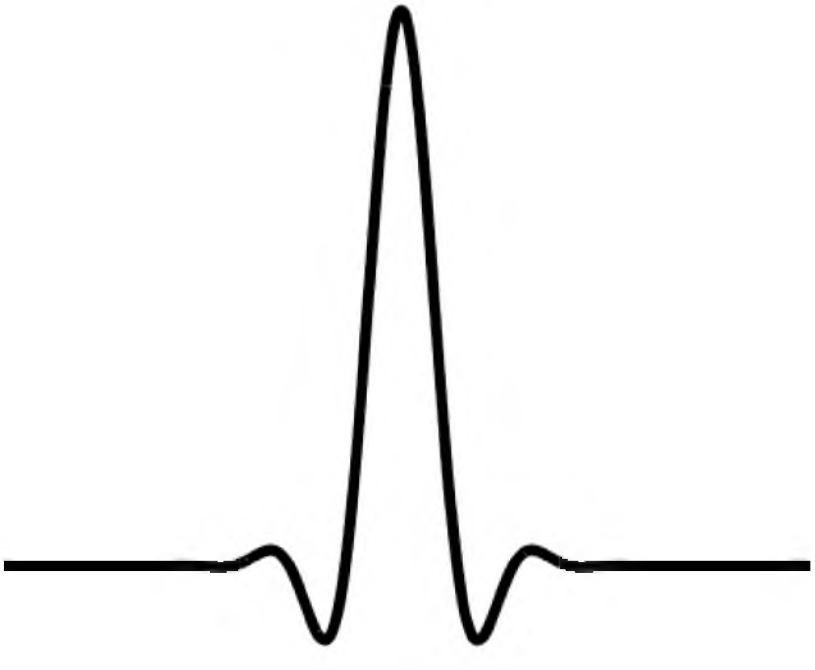

Figure 1.2. The finite bandwidth pulse used in an 802.11b channel sounder.

CHAPTER 2

CHANNEL SOUNDING FOR THE MASSES: LOW COMPLEXITY GNU 802.11B CHANNEL IMPULSE RESPONSE ESTIMATION [1]

2.1 Abstract

New techniques in cross-layer wireless networks are building demand for *ubiquitous channel sounding*, that is, the capability to measure channel impulse response (CIR) with any standard wireless network and node. Towards that goal, we present a software-defined IEEE 802.11b receiver and CIR measurement system with little additional computational complexity compared to 802.11b reception alone. The system implementation, using the universal software radio peripheral (USRP) and GNU[2] Radio, is described and compared to previous work. We validate the CIR measurement system and present the results of a measurement campaign which measures millions of CIRs between WiFi access points and a mobile receiver in urban and suburban areas.

2.2 Introduction

Channel impulse response (CIR) measurements have long held importance for communication system design [21–23]. The CIR describes the spreading, or echoing, that occurs when an impulse is sent through a channel. This spreading in time can

[1] ©[2012]. Reprinted, with permission, from D. Maas, M.H. Firooz, J. Zhang, N. Patwari, and S.K. Kasera, "Channel Sounding for the Masses: Low complexity GNU 802.11b channel impulse response estimation," *IEEE Transactions on Wireless Communications*, 2012.

[2] GNU is a recursive acronym for GNU is not unix.

lead to intersymbol interference (ISI), and frequency-selective or narrow band fading, depending on the symbol bandwidth. A knowledge of the CIR characteristics enables system designers to ensure that ISI does not dominate and hence lead to an excessive irreducible bit error rate [24]. Multipath channels can also be used to increase the bit rate and reliability of multiple-input multiple-output (MIMO) communications systems. Accurate MIMO channel models can be built from CIR measurements [25], and can be used to improve MIMO system design [26]. In general, measurements of CIR in wireless networks have become increasingly important to determine the real-world performance of many new technologies.

In addition, new cross-layer wireless networking technologies use measurements of the multipath channel for purposes of environmental awareness and security, such as fingerprint-based localization [27], RF-based multistatic radar [28], location distinction [29], secret key establishment [30]. These applications require CIR measurements to be performed in real time using commercial wireless devices, as opposed to with specialized measurement equipment or in postprocessing. Typical commercial wireless devices use the received signal in a demodulator to estimate the transmitted bits, but then discard the received signal samples. Information about the channel (besides the received signal strength) is not forwarded to higher networking layers, nor can it be estimated from the demodulated bits. For the mentioned applications to be viable, future commercial wireless devices must be able to rapidly calculate CIR information.

In this paper, we present the design of an inexpensive CIR measurement system. It is built upon GNU Radio, an open source framework for software-defined radio [31, 32], and the Universal Software Radio Peripheral (USRP), an open-source transceiver platform [33]. Compared to signal analyzers and oscilloscopes (a 3-GHz vector network analyzer (VNA) can cost US $20,000), our system is low cost. The cost of the proposed system is US $975 [34], which enables large-scale deployment as might be seen in a typical WiFi deployment. Our system works seamlessly with standard physical (PHY) layer signals from commercial 802.11b wireless devices. Essentially, our system provides an 802.11b receiver with the additional capability of CIR estimation.

However, this paper is not limited in scope to the USRP – the implementation

presented enables practical CIR estimation in hardware with strict computational limitations, such as field-programmable gate arrays (FPGAs) or application-specific integrated circuits (ASICs). It will not be feasible to compute a CIR estimate with commercial 802.11 hardware unless the computational complexity is low. The open-source platform chosen is an advantage, we believe, because it is likely to lead to cooperative improvement in the system capabilities and large-scale adoption. Providing a system implementation that works within the limitations of the hardware platform is, in part, a demonstration of the feasibility of the approach in future commercial systems.

The CIR measurement system we present in this paper for 802.11b is similar to sliding correlator channel sounding method in which a known pseudo noise (PN) signal is generated and continuously transmitted from a transmitter to the receiver [21,35]. However, our work is different from the existing method in the following four significant ways:

1. The PN sequence in 802.11b is fixed and not designed for high dynamic range CIR estimation.

2. Devices transmitting in 802.11b send PN-coded symbols modulated with data; modulation is undesirable from the perspective of CIR estimation.

3. Unlike sliding correlator measurement systems, which calculate the full cross-correlation signal after thousands of PN signal periods, our system calculates the full cross-correlation signal during each PN signal period [35].

4. No specialized transmitter is required, as any standard 802.11b transmitter (e.g., laptop or access point) may be used.

Note that IEEE 802.11b devices must support two mandatory bit rates (1 Mbps and 2 Mbps) and may optionally support two higher rates (5.5 Mbps and 11 Mbps) as specified in [36]. In this work, for simplicity, we only consider the standard rates. We note that the start of any 802.11b packet and some 802.11g packets (the first 192 symbols, known as the PLCP[2] frame), are sent at either the 1 or 2 Mbps rate.

[2]Physical Layer Convergence Procedure

Thus many packet sources exist which our system implementation can utilize for CIR estimation.

In a similar, but independent project, a channel sounder for 802.11b applications is reported in [37], which begins by recording the samples of a 192-bit segment of the 802.11b signal onto a PC. Then, the signal is despread and demodulated on the PC. Next, the transmitted signal for the 192-bit segment is recreated using the demodulated bits. Finally, the recorded received signal and recreated transmitted signal are convolved. Since both have many samples, the cross-correlation consumes significant PC computation time, on the order of $NB\log(NB)$, where N is the number of samples per bit, and B is the number of bits used. In comparison, our system involves PC computation on the order of NB. The system proposed in [37] uses proprietary software and VHDL implementations (ComBlock products from Mobile Satellite Services Inc.), while our implementation uses open-source hardware and software with a wide user base that utilizes and contributes to the code library. As an open source platform, our code has been downloaded from our website 1140 times since its first posting.

Our specific contributions to 802.11b CIR estimation system research are summarized as follows:

1. We provide an implementation of an 802.11b FPGA matched filtering method, the first, to our knowledge, to be presented for the USRP-based GNU Radio framework.

2. We provide a method to estimate the CIR from a modulated 802.11b signal. In particular, we use the output of the receiver's matched filter, which allows a lower-complexity CIR estimate compared to [37].

3. We perform extensive measurements, in both lab-controlled and real-world multipath channels.

2.3 Analysis Methods

In this section, we present a detailed analytical framework for CIR estimation using received 802.11b signals. We describe an 802.11b signal, how it is impacted by a

multipath channel, and how the proposed system estimates both: (1) the transmitted data, and (2) the amplitudes and delays of the multipath in the channel. This signal framework is used throughout this paper.

2.3.1 Transmitted Signal

The 802.11b physical layer uses direct-sequence spread spectrum (DSSS) modulation with symbol duration of $T_s = 1\mu$s. This transmitted symbol stream is multiplied by a pseudo-noise (PN) code signal, which also has duration T_s. Denoting the PN code signal as $c(t)$ and the jth transmitted data symbol as b_j, the transmitted signal in baseband is given by

$$s(t) = \sum_j b_j c(t - jT_s). \tag{2.1}$$

Note that b_j generally takes complex values, because data symbols may be modulated either using differential binary phase-shift keying (DBPSK) or differential quadrature phase-shift keying (DQPSK). Although our work is developed and tested for DBPSK, it is readily extendible to DQPSK.

The PN code in 802.11b is called the Barker code. This code consists of eleven *chips*, each with duration $T_c = T_s/11$ μs, thus "spreading" the bandwidth of the transmitted signal to eleven times the original bandwidth. The Barker code signal is a modulated sequence of pulses,

$$c(t) = \sum_{i=0}^{10} c_i p(t - iT_c), \tag{2.2}$$

where $p(t)$ is the pulse shape, and $c_i \in \{+1, -1\}$ as given in [36]. The pulse shape is chosen to meet the bandwidth limitations imposed by the 802.11b standard, but the precise shape of $p(t)$ is left to the designer. In this paper, when it is necessary to use a particular pulse shape, we have chosen to use a square root raised cosine (SRRC) pulse with roll-off factor $\alpha = 0.5$, which meets the spectral mask requirements and represents a good trade-off between temporal and frequency domain characteristics [38].

2.3.2 Received Signal

Because of the multipath radio channel, many copies of the transmitted signal arrive at the receiver with different time delay, amplitude, and phase. The multipath channel filter is modeled as [39]:

$$h(t) = \sum_{l=0}^{L-1} \alpha_l \delta(t - \tau_l), \tag{2.3}$$

where L is the total number of multipath components, $\alpha_l = |\alpha_l| e^{j\angle\alpha_l}$ is the complex amplitude gain of the lth multipath, τ_l is the delay of the lth multipath, and $\delta(\cdot)$ is the Dirac delta function. Since we are only interested in the relative time delay of each multipath, we let $\tau_0 = 0$, and then τ_l is the additional delay compared to the first arriving multipath.

The received signal $r(t)$ is the convolution of the transmitted signal and the channel filter. Applying (2.3) and (2.1),

$$r(t) = s(t) \star h(t) = \sum_{l=0}^{L-1} \sum_{j} \alpha_l b_j c(t - \tau_l - jT_s). \tag{2.4}$$

An 802.11b receiver "de-spreads" the signal, *i.e.*, performs matched filtering with the PN code signal $c(t)$ from (2.2), which results in signal $q(t)$,

$$q(t) = r(t) \star c(-t) = \sum_{l=0}^{L-1} \alpha_l \sum_{j} b_j R_c(t - \tau_l - jT_s), \tag{2.5}$$

where $R_c(t) = c(t) \star c(-t)$ and $R_c(0)$ is the energy in the signal $c(t)$, which we denote $\mathcal{E}_c$.

Standard 802.11b receivers must perform despreading, *i.e.*, the calculation of $q(t)$, in order to perform demodulation. We propose that $q(t)$ can be used directly in CIR estimation as well. By using an output that existing 802.11b receivers compute, we make it more feasible for future 802.11b receivers to estimate CIR without significant additional computational complexity.

We note that it is possible to estimate CIR from all symbols comprising the PLCP preamble and header. If the PLCP is known *a priori*, the reception range can be significantly increased by correlating with the entire PLCP, rather than $c(t)$. In this case, the "energy per bit" is essentially increased by a factor of 48, a 17 dB increase.

In this work, we present a CIR measurement system that works with any 802.11b transmitter, thus we cannot know the PLCP ahead of time. Further, correlating with the entire PLCP adds computational complexity.

The above formulation has not included interference. Inevitably, some packets will be unable to be received due to low SINR, and thus the CIR will not be estimated. Further, the SINR can be estimated from a received packet, and CIR estimates can be dropped if the desired SINR is not achieved.

2.3.3 CIR Estimation

The estimation of CIR from a received 802.11b signal is complicated by the modulated data $\{b_j\}$. That is, the PN code signal is modulated with data, presumably unknown to the receiver until after demodulation. For example, for BPSK, $b_j \in \{-1, +1\}$. In this section we first present the (unrealistic) case of an unmodulated signal, *i.e.*, where $b_j = 1$ for all j. We then describe how we estimate the CIR from a modulated 802.11b signal.

First, for an ideal unmodulated signal, (2.5) would simplify to

$$q(t) = \sum_{l=0}^{L-1} \alpha_l R_{pn}(t - \tau_l), \quad \text{where } R_{pn}(t) = \sum_j R_c(t - jT_s). \tag{2.6}$$

Here, R_{pn} is the correlation of a PN code signal with a repeating PN code signal with period T_s. The Barker code has the property that this correlation function $R_{pn}(t)$ peaks at $t = 0$ and integer multiples of T_s and is almost constant in between those peaks [40]. Figure 2.1(a) shows the signal $q(t)$ when there is exactly $L = 1$ path with amplitude $\alpha_0 = 1$. As multipath components correspond to time-delayed versions of $q(t)$, the almost constant correlation in between peaks makes it possible to identify multipath contributions even when their magnitude $|\alpha_l|$ is small.

When dealing with modulated signals, the correlation $q(t)$ may not be nearly constant between peaks, making low-amplitude multipath components harder to identify. In Figure 2.1(b), we use the transmitted symbols $\mathbf{b} = [1, 1, -1, 1, 1]$, and plot the correlation output signal $q(t)$ from (2.5), for the case that $L = 1$ and $\alpha_0 = 1$. Note that the normalized amplitude of $q(t)$ between the 2nd and 3rd peaks, and between the 3rd and 4th peaks, rapidly change between $\pm 1/11$. These periods of

varying correlation correspond to the times in between changes in symbol values b_j. When $b_j \neq b_{j+1}$, the value of $q(t)$ for $jT_s < t < (j+1)T_s$ is not almost constant.

However, note that when $b_j = b_{j+1}$, there is a nearly constant $-1/11$ correlation value in between the two peaks at jT_s and $(j+1)T_s$. *When subsequent symbols are identical, the almost constant correlation value in $q(t)$ can be exploited for improved CIR estimation.* To avoid the negative impact of symbol modulation, we use the correlator output signal $q(t)$ whenever the symbol value b_j repeats.

To be explicit, define two correlation functions, $R_o(t)$ and $R_s(t)$ (shown in Figure 2.2), as:

$$\begin{aligned} R_o(t) &= (R_c(t) - R_c(t - T_s))I_{(0,T_s)} \\ R_s(t) &= (R_c(t) + R_c(t - T_s))I_{(0,T_s)} \end{aligned} \tag{2.7}$$

where $I_{(0,T_s)}(t)$ has value 1 at interval $(0, T_s)$ and zero elsewhere. We also define two subsets, $J_s = \{j : b_j = b_{j+1}\}$ for symbol integers j when the next symbol value repeats, and $J_o = \{j : b_j \neq b_{j+1}\}$. Then we can write (2.5) as

$$\begin{aligned} q(t) &= \sum_{j \in J_s} b_j \sum_{l=0}^{L-1} \alpha_l R_s(t - jT_s - \tau_l) \\ &\quad + \sum_{j \in J_o} b_j \sum_{l=0}^{L-1} \alpha_l R_o(t - jT_s - \tau_l). \end{aligned} \tag{2.8}$$

This version of $q(t)$ contains terms $R_s(\cdot)$ and $R_o(\cdot)$ that have support only over one symbol period. We estimate the CIR by averaging only the symbol periods of $q(t)$ that correspond to repeated symbol values:

$$\hat{h}(t) = \frac{1}{|J_s|} \sum_{j \in J_s} b_j q(t - jT_s) I_{(0,T_s)}(t) \approx \sum_{l=0}^{L-1} \alpha_l R_s(t - \tau_l) \tag{2.9}$$

Essentially, the channel estimator in (2.9) averages together only the impulse responses estimated during periods when the symbol value is not switching and thus the correlation function is nearly constant. Note that symbol values b_j do not affect $\hat{h}(t)$. In the ideal case, the channel estimate is a sum of time-delayed, attenuated, and phase-shifted versions of $R_s(t)$. However, in a given hardware implementation, $R_s(t)$ may be affected by other filters, known or unknown, in the receiver chain. If

the overall filter of the receiver chain is unknown, it may be beneficial to estimate $R_s(t)$ using a known channel, *i.e.*, an interference-free cabled connection between the transmitter and receiver. We employ this method to generate an estimate of $R_s(t)$ from a single packet, which we call $\hat{R}_s(t)$.

The CIR estimate $\hat{h}(t)$ in (2.9) is a convolution of the true CIR in (2.3) with $\hat{R}_s(t)$, which has a zero-to-zero pulse width of approximately 188 ns. Since multipath arrive more closely spaced than 188 ns, the complex-valued, time-delayed pulse shapes $\hat{R}_s(t - \tau_l)$ overlap in time, making it difficult to visually inspect $\hat{h}(t)$ to identify multipath arrival delays.

We apply a deconvolution procedure based on [41] to estimate multipath time delays. This procedure is described in detail in [42]. In short, we discretize the CIR and write the measurement as a linear combination of the CIR amplitudes. Then, we solve a quadratic optimization problem using the well known convex optimization software [43] to perform the inversion.

The sampled measurement is written as,

$$\hat{h}[n] \;=\; \sum_{l=0}^{L-1} \alpha_l \hat{R}_s(n\,T_s - \tau_l) + w[n] \tag{2.10}$$

where $w[n]$ is measurement noise, assumed to be i.i.d. Gaussian. Equation (2.10) can be written as $\hat{\mathbf{h}} = \hat{\mathbf{R}}_s\boldsymbol{\alpha} + \mathbf{w}$, where $[\hat{\mathbf{R}}_s]_{k,l} = \hat{R}_s(kT_s - \tau_l)$ is an $M \times L$ matrix, $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_L]^T$, and $\mathbf{w}$ is an M-dimensional noise vector. An estimate of $\boldsymbol{\alpha}$ is the solution to the following optimization problem [41, 42],

$$\hat{\boldsymbol{\alpha}} = \min_{\boldsymbol{\alpha}} \|\hat{\mathbf{h}} - \hat{\mathbf{R}}_s\boldsymbol{\alpha}\|_2^2 + \lambda\|\boldsymbol{\alpha}\|_1 \tag{2.11}$$

where λ is a fixed parameter, the inverse of a Lagrange multiplier [44], which is set as discussed in [41].

2.4 Implementation

In this section, we present the system implementation of the WiFi-based receiver and CIR estimator using the proposed FPGA matched filtering method on a USRP. The USRP receiver path consists of a 64 MS/s (million samples per second) 12-bit ADC, an Altera Cyclone FPGA, and a USB controller. The USB 2.0 bandwidth is not

sufficient to stream an 802.11b signal sampled at the Nyquist rate. However, the rate limitations of a USB 2.0 link do not limit transfer of 802.11b *symbol information*, since symbols are sent at 1 Msymbols/s [45]. The spreading via DSSS adds no information, but causes the RF bandwidth to increase by a factor of 11. To reduce the received signal to samples at 1 MS/s, we must first despread on the FPGA. After despreading, symbol decisions can be made using only one sample per symbol, and, as we will show, a subset of samples per symbol can be used for CIR estimation.

A broad overview of the signal processing steps is shown in Figure 2.3. We first reduce the sampled data $r(t)$ to 32 MS/s. Then, we despread using (2.5). The output $q(t)$ has a sample rate of 32 MS/s, however, not every sample is necessary, so we send only samples near the peaks in $q(t)$, as described in detail in this section.

The result is that the average data rate sent to the PC is within the rate limitations of a USB 2.0 link. The PC then performs the symbol detection and bit decoding operations as specified in the IEEE 802.11b standard. Our receiver implementation can consistently receive 802.11b packets sent at the 2 Mbps rate, and the reception range is up to 20 m.

We compare our implementation to previous work [46], which we call the *bandwidth reduction method*. In this method, $r(t)$ is filtered and downsampled to a 8 MHz RF bandwidth, smaller than the RF bandwidth of the DSSS signal. Then, the samples are at a rate low enough to be transfered over USB to be processed on a PC. The downsampling reduces the range of the receiver, as we show in Section 2.5.

The main computational challenge in the proposed method is the implementation of matched filtering on the FPGA. We propose a computationally-efficient method to implement the 802.11b matched filter, valid for the strict limitations of the given FPGA, or any computationally limited ASIC or FPGA implementation. We describe three ways in which the implementation reduces computational complexity and data rate yet still provides a high-capability system implementation: multiplication reduction; use of two memories; and peak selection.

2.4.1 Multiplication Reduction

A direct implementation of the matched filter in (2.5) would require 32 multiplies and additions per sample. We reduce the complexity as follows. Figure 2.4 shows

$c(t)$ and its sampled version, $c(iT_s)$. We quantize each sample of the PN code signal, $c(iT_s)$, and denote the quantized values as $c_q(iT_s)$. In our implementation, we chose quantization to five bits, in a trade-off between resolution and multiplier space complexity.

Some values of $|c(iT_s)|$ are similar enough, that when quantized to five bits, $|c_q(iT_s)| = |c_q(jT_s)|$ for some $j \neq i$. Since summation is simpler than multiplication in an FPGA, it saves both time and complexity to first add (or subtract) samples with identical $|c_q|$ value, and then multiply the sum by its $|c_q|$ value.

$$q(nT_s) = \sum_{i=0}^{31} c_q(iT_s) r((n-i)T_s) = \sum_{g=1}^{15} c_g \left[\sum_{i \in S_g} \operatorname{sgn}\{c_q(iT_s)\} r((n-i)T_s) \right] \tag{2.12}$$

where $q(nT_s)$ is the nth sample of the match filter output $q(t)$, S_g is the set of all indexes in the gth group, $\operatorname{sgn}\{\cdot\}$ is the signum function, and c_g is the multiplicative factor $c_q(iT_s)$ for all $i \in S_g$. The S_g and c_g for each group g are listed in Table 2.1. Using this rearrangement, we require 15 multiplications, instead of the 32 that would be required in a direct implementation.

2.4.2 Two Memories

An FPGA requires parallelization in order to complete the several additions and multiplications required at each new sampling time. Our implementation allows two clock cycles (clock rate of 64 MHz) per sampling time (sampling rate of 32 MS/s). During these two clock cycles, we must perform addition and multiplication as described above, and shift samples to allow space for the new incoming signal sample.

For this purpose, we use two 32-length arrays, which we refer to as mem and bmem. When a new sample is received, it is located at mem[0] while mem[1] to mem[31] are filled by bmem[0] to bmem[30]. In the next cycle, mem[0] to mem[31] are put in bmem[0] to bmem[31]. This process is depicted in Figure 2.5. As explained in the previous paragraph, we first add the data in bmem, by group g, which is completed in one cycle. Then, multiplication by group multiplier c_g is performed, and the results summed.

2.4.3 Peak Selection

The output of the FIR filter, $q(nT_s)$, has a 32 MS/s rate. With a sampling period of 31.25 ns, we capture 344 ns (much more than the typical excess delay for short-range channels [47]) of the signal within 11 samples. These 11 samples per symbol can be reliably transfered via USB.

The peak selection algorithm selects 11 out of each 32 samples per symbol as follows. First, the FPGA computes the power values $|q(nT_s)|^2$, $n = 1, \ldots, 32$. The index of samples with maximum power is denoted $n_{max} = \text{argmax}_n |q(nT_s)|^2$. The FPGA sends through the USB the samples from three samples before to seven samples after the peak power sample, *i.e.*, $\{q((n_{max} + i)T_s)\}_{i=-3}^{7}$.

Note that the proposed CIR measurement system finds samples near the maximum-power peak, not necessarily the line-of-sight (LOS) path. In a non-LOS dominant channel, if the LOS path arrives within three samples prior to the maximum-power peak, the proposed system records the full CIR.

2.5 Experimental Results

In all cases, we load an Ettus Research USRP (rev 4.5) with the code described in Section 2.3.3. The RF front end is a RFX2400 daughterboard (rev 30), also from Ettus Research. The antenna is a 2400-2480 MHz sleeve dipole antenna with omnidirectional pattern in the horizontal plane and a 3 dBi gain. The USRP is connected to a Dell Inspiron laptop running Python and Matlab. The Python (GNU Radio) code collects data from the USB, demodulates the packet data, and writes to a file. The Matlab code performs the averaging required in (2.9) and then displays and stores the impulse response estimate $\hat{h}(t)$. From the stored $\hat{h}(t)$, the deconvolution described in Section 2.3.3 is performed in postprocessing.

2.5.1 Demodulator

We do not proceed with CIR estimation when packet data do not pass the CRC test. Equally important, the MAC address of a transmitter is included in the packet header, and this is necessary to distinguish packets originating from different transmitters. In this section, we measure the packet reception rate (PRR) of the implemented 802.11b CIR measurement system in an interference-free environment.

We configure a test transmitter, a D-Link 802.11b wireless router (model DI-614+), to broadcast a beacon packet at a basic rate (1 or 2 Mbps) every 200 ms (5 packets/sec). The router and receiver are placed in a shielded anechoic chamber and separated by 6.0 m. The packet reception rate is recorded for three minutes, and experiment repeated four times. The implementation presented in this paper receives an average of 724 packets, while the bandwidth reduction method receives an average of 454. The results show that the FPGA matched filtering method outperforms the bandwidth reduction method by successfully demodulating 1.6 times more packets.

2.5.2 Channel Measurement

In this section we first perform two experimental validations on our implementation using known channels between the transmitter and the receiver. Then, we perform an experimental measurement campaign to measure a large number of CIRs in outdoor areas in and around Salt Lake City, Utah. We provide measurement results and summarize the measured delay characteristics.

2.5.2.1 Validation

To validate the CIR estimation system, we create two channels with known impulse response out of RF hardware and cable, with which we connect the wireless router (transmitter) and receiver.

In the single-path experiment, the transmitter is connected to an attenuator, whose output is connected via cable to the receiver. We record several measured CIR estimates $\hat{h}_1(t)$. Figure 2.6-(a) shows three measurements $\hat{h}_1(t)$ and the estimated CIR for a single-path channel, $\hat{R}_s(t)$. Since $\hat{h}_1(t)$ is nearly identical to $\hat{R}_s(t)$, it is apparent that the channel has only one path, i.e., $L = 1$. Figure 2.6-(b) presents the deconvolved CIR estimate from (2.11).

In the double-path experiment, the transmitter cable is connected to a RF splitter with two outputs, one connected to a short cable (1.5m), and another to a long cable (25.9m). We first measure the CIR using a vector network analyzer, from which we find that the difference in delay between the two paths is 122 ns. The amplitude difference between the two paths is measured to be 9.5 dB by using a LadyBug power sensor (LB479A). Figure 2.6-(a) shows the CIR measurements for $\hat{h}_2(t)$. As can

be seen, $\hat{h}_2(t)$ is consistently higher in amplitude than $\hat{R}_s(t)$ between the samples 6 through 10, indicative of later-arriving multipath power. The deconvolution algorithm of (2.11) is applied and the resulting estimate shown in Figure 2.6-(c).

The results clearly show two paths, the later paths with a 125 ns relative delay and between 10 and 12 dB less received power. In the results where the estimated power of the second path is above -20 dB relative to the path with maximum power (the same noise level we use for the calculation of the dispersion statistics), we find we are able to achieve a standard deviation of 0.69 radians for the difference in phase between multipath components corresponding to the first and second paths. Additionally, the standard deviation of the relative power of the two multipath components is 3.5 dB. These statistics are good considering the hardware synchronization issues, phase noise, and the coarse sampling period for the CIR estimates.

Observation of Figures 2.6(b) and 2.6(c), as well as many other deconvolution results lead us to the conclusion that the dynamic range for the CIR measurement system is at least 20 dB, which is expected since the PN coding gain of the Barker code is $20 \log_{10} 11 \approx 20.8$ dB.

2.5.3 Drive-Test CIR Measurement Campaign

We use our system to measure CIRs in three residential, two commercial, and one downtown area in Salt Lake City. The residential areas are comprised of one to three story single-family homes and apartment buildings. The commercial areas include streets near strip malls, low-rise office buildings, and heavy vehicle traffic. The downtown area is an urban canyon of high-rise office buildings on both sides of the streets. In each area, the receiver antenna is on the outside of a vehicle that drives at typical speeds on city streets. In the course of six five-minute drive-test measurements, a total of three million CIR measurements are recorded. Figure 2.7 presents a typical deconvolved CIR estimates $\hat{\boldsymbol{\alpha}}$ from each area.

In order to compare different multipath channels and to develop some general design guidelines for wireless systems, parameters that grossly quantify the multipath channel are used. The time dispersive properties of wide band multipath channels are most commonly quantified by their mean excess delay $\bar{\tau}$ and RMS delay spread σ_τ, as defined in [47]. Table 2.2 presents the average mean excess delay, average RMS

delay spread, and maximum RMS delay spread of the measured channel responses for each area.

Delay spreads depend strongly on path length, antenna height, frequency, and environment. However, previous work has shown that, in general, rural and suburban delay spreads are smaller than in urban or dense urban areas [48–51]. Our results in Table 2.2 show a similar trend, since the residential and commercial areas can be considered suburban and have lower average RMS delay spreads than the urban downtown area. One of the few studies of RMS delay spreads for indoor-to-outdoor channels near 2.4 GHz reported average RMS delay spreads of 27-44 ns [52], but the studied path lengths were about 330 m, significantly longer than one would expect from 802.11b path lengths.

2.6 Conclusion

Future wireless networks are envisioned that rely on the real-time estimation of CIR from received WiFi packets for the purposes of cross-layer security, localization, and environmental imaging. We present a CIR estimation system using an inexpensive and open source hardware and software platform to enable these emerging areas of research. We show how accurate CIR estimation can be performed using a resource-constrained FPGA, which provides a proof-of-concept for future commercial devices.

Future work should address MIMO (e.g., 802.11n) CIR measurement using a bank of synchronized software radios. Low complexity MIMO CIR measurement will likely benefit the development of future cross-layer techniques in multiple-antenna wireless networks.

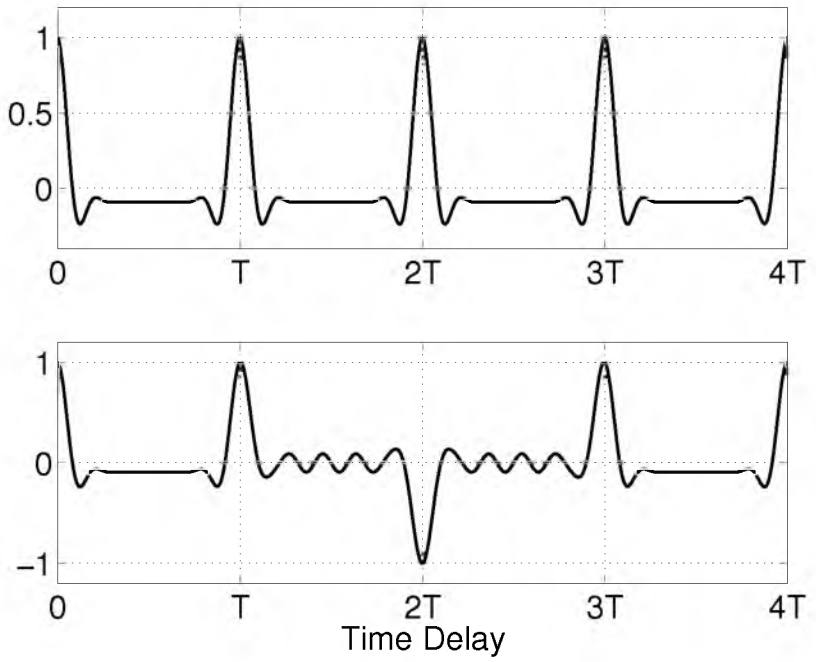

Figure 2.1. Correlation output signal $q(t)$ in one-path channel ($L = 1$ and $\alpha_0 = 1$) when (top) receiving an unmodulated signal (*i.e.*, $\mathbf{b} = [1, 1, 1, 1, 1]$); (bottom) receiving a signal modulated with $\mathbf{b} = [1, 1, -1, 1, 1]$.

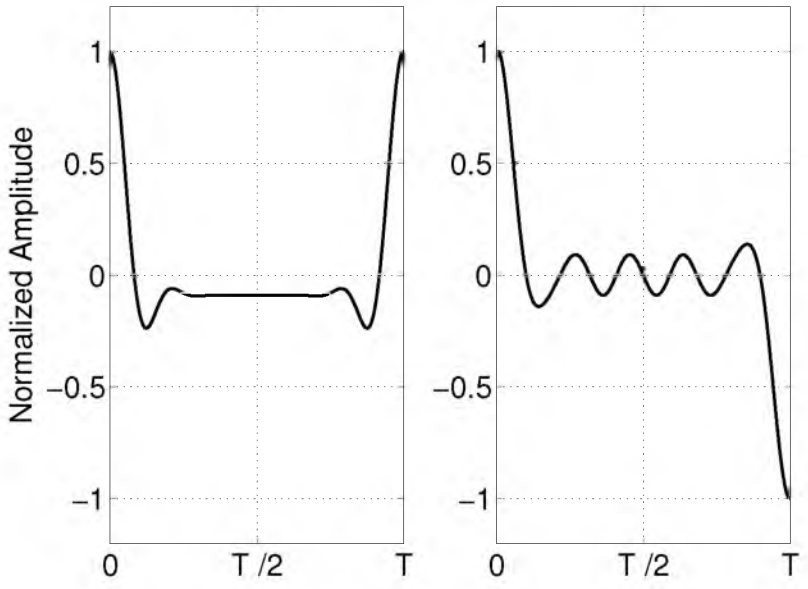

Figure 2.2. Normalized symbol-period length correlation functions (left) $R_s(t)$ and (right) $R_o(t)$, both given in (2.8).

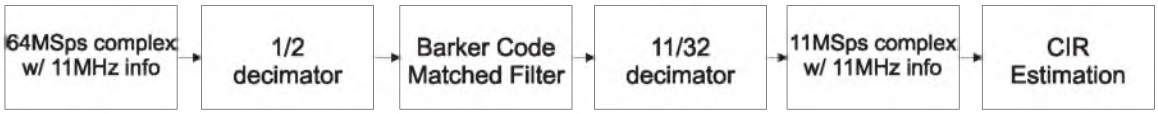

Figure 2.3. Block diagram of FPGA matched filtering method.

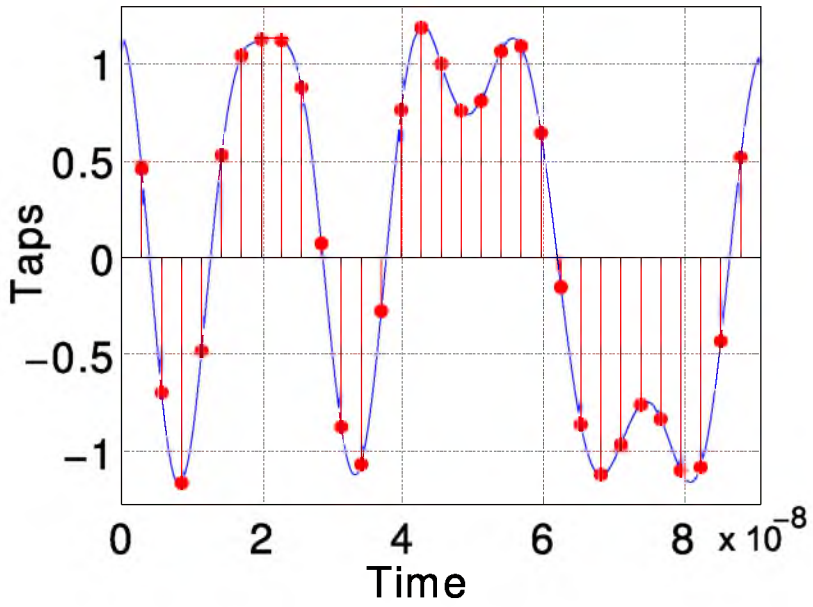

Figure 2.4. Samples (•) of PN code signal $c(t)$, or equivalently, taps of the matched filter.

Table 2.1. Indices by group g and the group's multiplier value c_g.

g	Multiplier c_g	Index Set S_g
1	19	$\{16, 28\}$
2	18	$\{3, 7, 23, 24, 31\}$
3	17	$\{11, 12, 19, 22, 15\}$
4	16	$\{25\}$
5	15	$\{6\}$
6	14	$\{8, 20, 22\}$
7	13	$\{4, 13\}$
8	12	$\{5, 14, 17\}$
9	11	$\{29\}$
10	10	$\{10\}$
11	8	$\{0, 26, 27\}$
12	7	$\{1, 30\}$
13	4	$\{18\}$
14	2	$\{9\}$
15	1	$\{21\}$

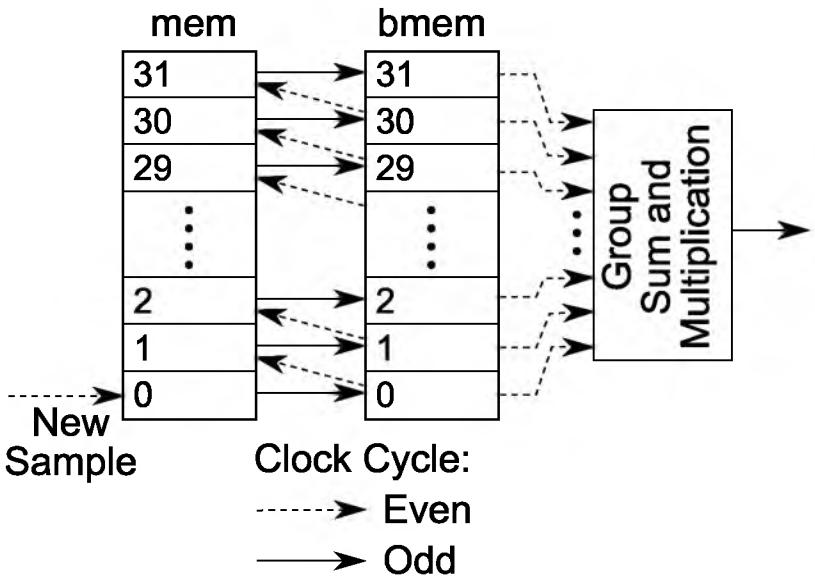

Figure 2.5. Memories mem and bmem are used to accept a new sample, and shift data, in two cycles, to allow for summation and multiplication.

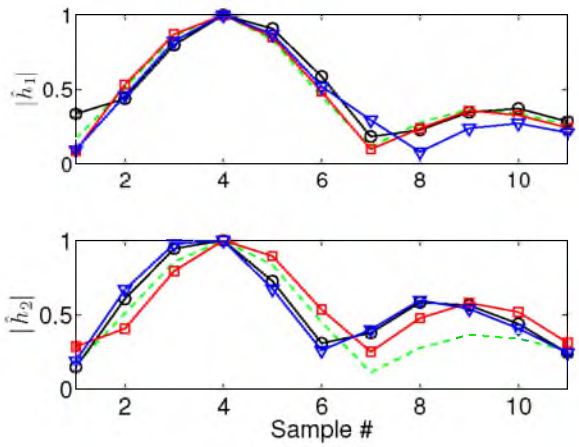

(a)

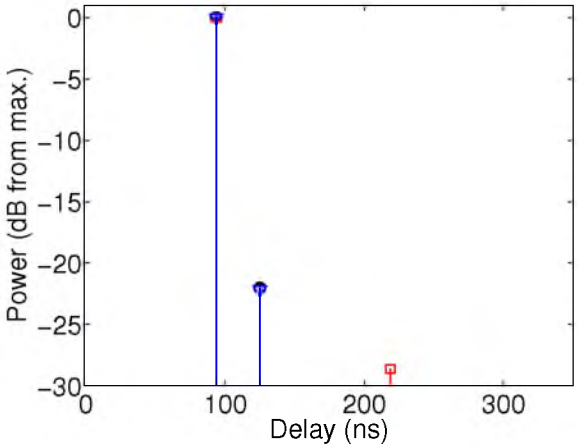

(b)

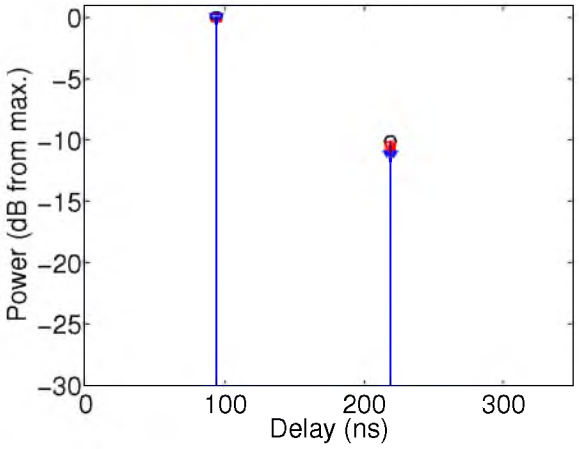

(c)

Figure 2.6. Single- and double-path results: (a) $\hat{h}$ for single-path (upper figure) and double-path (lower figure), both showing ideal CIR $\hat{R}_s[n]$ (- - - -); Deconvolved $\hat{\boldsymbol{\alpha}}$ for (b) single-path; and (c) double-path.

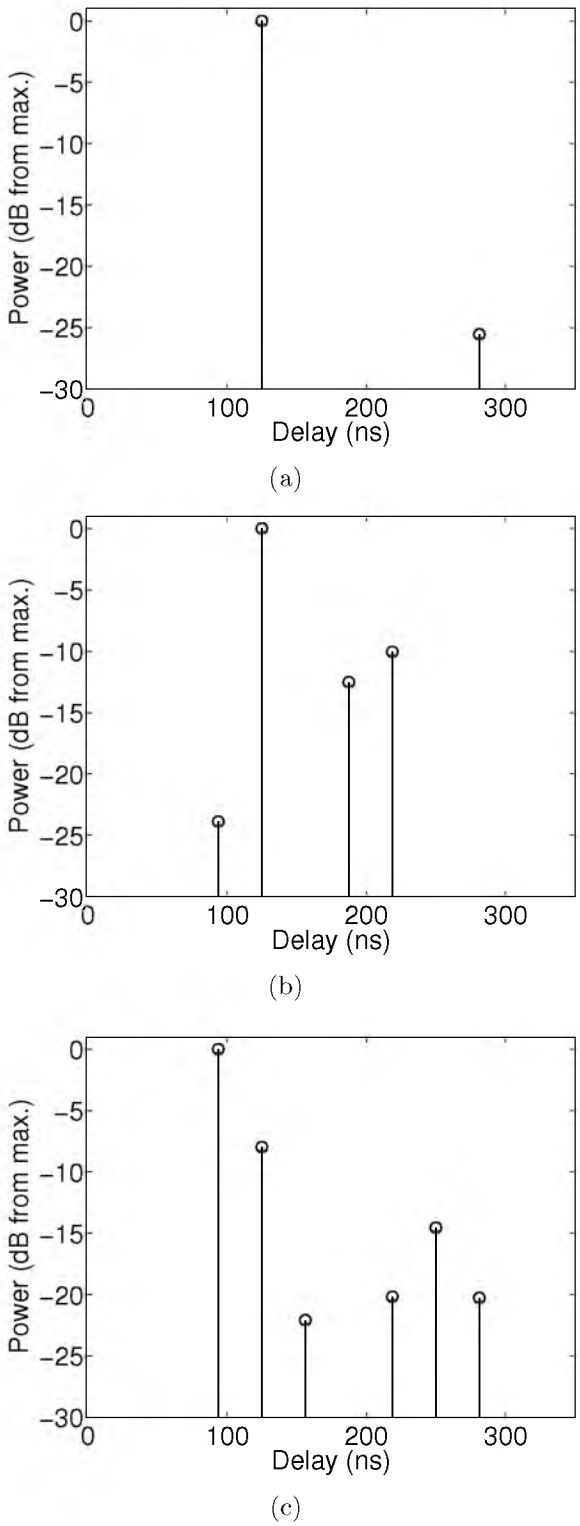

Figure 2.7. Typical CIR, $\hat{\boldsymbol{\alpha}}$, measured in (a) residential, (b) commercial, and (c) downtown areas.

Table 2.2. RMS delay spread and mean excess delay statistics for residential (Res.), commercial (Com.), and Downtown (DT) areas.

	Res. 1	Res. 2	Res. 3	Com. 1	Com. 2	DT
Average $\bar{\tau}$ (ns)	7.1	36.7	7.4	6.4	17.7	48.2
Average σ_τ (ns)	7.0	23.7	7.4	6.5	16.9	30.6
Max. σ_τ (ns)	47.4	86.8	35.2	22.8	80.7	88.6

CHAPTER 3

EXPERIMENTAL PERFORMANCE EVALUATION OF LOCATION DISTINCTION FOR MIMO LINKS [1]

3.1 Abstract

A radio channel-based location distinction system monitors physical layer measurements of received signals to detect if a transmitter has changed position since its previous transmission. This paper explores the design space for MIMO-based location distinction systems. Using extensive channel measurements collected with two different MIMO testbeds, we make several observations about the tradeoffs inherent in MIMO location distinction, and the scaling of performance with respect to bandwidth, history size and insertion delay, and number of antenna elements. We show that MIMO location distinction is very reliable. For example, a 2x2 MIMO channel with a bandwidth of 80 MHz allows a 64-fold reduction in miss rate over the single-input single-output (SISO) channel for a fixed false alarm rate, achieving false alarm rates as low as 4×10^{-4} for a 2.4×10^{-4} probability of missed detection.

3.2 Introduction

Location distinction is determining whether or not the position of a wireless device has changed. Detecting a change in position is fundamentally different from

[1] ©[2012]. Reprinted, with permission, from D. Maas, N. Patwari, S.K. Kasera, D. Wasden, and M. Jensen, "Experimental Performance Evaluation of Location Distinction for MIMO Links," in *Proc. 4th IEEE International Conference on Communication Systems and Networks* (COMSNETS), 2012.

estimating position (localization). In fact, location distinction can be performed without performing the more costly task of localizing wireless devices.

The ability to perform location distinction provides several benefits. These include an improved capability to monitor the positions of radio-tagged objects, better energy conservation in radio localization systems (by localizing only when devices have changed position), and a means to perform position-based authentication in wireless networks [14, 15, 53, 54]. Existing work has shown that characteristics of the physical layer of wireless networks, including received signal strength (RSS), channel impulse response (CIR), or channel frequency response can be exploited to detect changes in transmitter or receiver positions [6, 7, 55, 56].

Multiple-input multiple-output (MIMO)-capable devices represent the state-of-the-art in wireless networking and have enabled significantly improved spectral efficiencies in wireless networks. Many new wireless standards, such as 802.11n, WiMax, and 4G cellular, take advantage of MIMO technology. Therefore, it is very important to evaluate location distinction for MIMO networks. However, to the authors' knowledge, no extensive implementation or experimental evaluation of MIMO-based location distinction has been performed[2].

Intuitively, we expect that location distinction performance should increase with the transition from SISO to MIMO, because the higher number of channels leads to a richer link measurement. However, the rate at which location distinction performance scales with system parameters including the number of antennas, channel bandwidth, and others, remains to be seen. In this paper, we evaluate location distinction for MIMO links under varying system parameters in order to understand the benefits and limits of MIMO location distinction performance, as well as how system design choices contribute to this performance. Specifically, we explore the use of band-limited estimates of the CIR, called link signatures [6, 14], for location distinction.

We limit our exploration to methods that do not rely on calibration or supervised training. A training-based approach would require prohibitively extensive measurements because link signatures vary rapidly with position. Furthermore, changes in

[2]We had presented a preliminary experimental study of location distinction in MIMO networks in a poster [17]. We significantly expand on our preliminary study in this paper.

the environment, such as the rearrangement of scatterers, would render the training data useless, and require periodic retraining. Without training or calibration, a location distinction algorithm should be able to reliably detect changes in location, regardless of whether or not the transmitter or receiver was moving, by comparing the current link signature with the ones in a recent history of signature measurements. We present the following work in order to characterize the performance of temporal link-signature-based location distinction for MIMO links:

1. We perform extensive measurement experiments with two different experimental testbeds.

2. We evaluate spatially and temporally dense channel measurements in order to study the spatial evolution of link signatures.

3. We show how to design a robust location distinction system and evaluate several trade-offs between system design choices and performance, including: link signature history size and insertion delay, bandwidth, complex vs. magnitude-only signatures, and number of antenna elements.

Our experimental evaluation leads to a better understanding of the benefits and limits of location distinction performance, as well as a general guide for system design. We show that:

1. Measured link signatures should not immediately be inserted into the history. Instead, for robust detection, the insertion of measurements should be delayed.

2. The number of link signatures to store in the history depends on the amount of temporal variation in the link signatures when the wireless device is stationary. Channels with less temporal variation require smaller histories, while those with more temporal variation require larger histories.

3. We observe, based on empirical data, that the miss rate for a constant false alarm rate follows an inverse power law with the number of antennas, i.e., the miss rate decreases slowly for large numbers of antennas. However, we see very dramatic improvements when comparing SISO to 2x2 or 1x4 MIMO/SIMO

location distinction systems, which is advantageous for MIMO systems with few antennas.

4. When random phase shifts due to imperfect synchronization are removed, link signatures with phase information lead to better performance than magnitude-only link signatures.

5. Increasing the bandwidth of the link signatures offers diminishing returns after about 20 MHz. In fact, higher bandwidth measurements are more susceptible to synchronization errors.

Our empirical data also show that MIMO location distinction performs well in a variety of experimental conditions. For example, we achieve a 4×10^{-4} probability of false alarm for a 2.4×10^{-4} probability of missed detection using a 2x2 MIMO channel with a bandwidth of 80 MHz, and a 3×10^{-4} probability of missed detection for a false alarm rate of 0.01 using a 1x2 SIMO channel with a bandwidth of 20 MHz.

The rest of this paper is organized as follows. In Section 3.3, we describe the link signatures, metrics, and MIMO location distinction algorithm. In Section 3.4, we discuss two measurement experiments, which we will refer to as Experiment I and Experiment II. In Section 3.5, we present testing results and analysis of the MIMO location distinction algorithm. We discuss related work in Section 3.6. Conclusions and future work are presented in Section 3.7.

3.3 Methods

In this section, we first describe the wireless measurements, a.k.a. link signatures, we use for location distinction and the difference metrics we use to quantify changes in them. Next, we present a real-time location distinction algorithm. Please note that our definitions and methods described below for MIMO are similar to those we have used in our past work on SISO links [6, 14].

3.3.1 Link Signatures

We define the *complex temporal link signature* (CTLS) calculated for the cth transmitter/receiver antenna pair as

$$\mathbf{f}_c = [h_c(0), h_c(1T_s), \ldots, h_c((M-1)T_s)] \tag{3.1}$$

where $h_c(\tau)$ is the band-limited channel impulse response as a function of delay τ, M is the number of samples, T_s is the sampling period, and $c \in S$, where

$$S = \{1, ..., k_1\} \times \{1, ..., k_2\}. \tag{3.2}$$

The number of transmitter and receiver antennas are represented by k_1 and k_2, respectively. We also define the *temporal link signature* (TLS) calculated for the cth transmitter/receiver antenna pair as

$$\mathbf{g}_c = [|h_c(0)|, |h_c(1T_s)| \ldots, |h_c((M-1)T_s)|]. \tag{3.3}$$

The MIMO channel measurements used in this paper are gathered using either a multitone probe or preamble-based channel estimation, both of which are described in Section 3.4. In both cases, time-domain representations of the channel response are used for link signatures.

We let the *MIMO complex temporal link signature* (MIMO CTLS) be the concatenation of the set of complex temporal link signatures measured between the first $k_1 \times k_2$ transmitter and receiver antennas:

$$\mathbf{F} = [\mathbf{f}_{c_1}, \ldots, \mathbf{f}_{c_k}], \tag{3.4}$$

where $c_1, ..., c_k$ is a list of the elements of S.

Finally, we let the *MIMO temporal link signature* (MIMO TLS) be the concatenation of the set of temporal link signatures measured between the first $k_1 \times k_2$ transmitter and receiver antennas:

$$\mathbf{G} = [\mathbf{g}_{c_1}, \ldots, \mathbf{g}_{c_k}]. \tag{3.5}$$

3.3.2 Difference Metric

In this section, we define the metric for measuring the difference between the current MIMO link signature and the history of previous MIMO link signatures. The history $\mathcal{H}$ is a first-in first-out (FIFO) buffer that stores a set of N previous MIMO link signatures.

The difference metric we explore in this paper is

$$\Delta(\tilde{\mathbf{F}}, \mathcal{H}) = \frac{1}{\sigma} \min_{\mathbf{F} \in \mathcal{H}} \|\mathbf{F} - \tilde{\mathbf{F}}\| \tag{3.6}$$

where $\tilde{\mathbf{F}}$ is the current MIMO link signature and σ is the average distance between link signatures in the history, defined as

$$\sigma = \frac{1}{(N-1)(N-2)} \sum_{\mathbf{F}_1, \mathbf{F}_2 \in \mathcal{H}} \|\mathbf{F}_1 - \mathbf{F}_2\| \tag{3.7}$$

The magnitude-only TLS uses the ℓ_2 norm in (3.6) and (3.7); for the CTLS, these norms are the ϕ_2 norm, defined as

$$||\mathbf{X} - \mathbf{Y}||_{\phi_2} = \min_{\phi} ||\mathbf{X} - \mathbf{Y}e^{j\phi}||_{\ell_2} = ||\mathbf{X}||^2 + \|\mathbf{Y}\|^2 - 2\|\mathbf{X}^*\mathbf{Y}\|. \tag{3.8}$$

The ϕ_2 norm removes the effect of random phase shifts that occur between subsequent CTLS measurements [6].

3.3.3 Real-time Location Distinction

We evaluate a location distinction algorithm that operates in real-time without training. In such algorithms, recently recorded link signatures are stored in a "history" and compared to the most recent link signature. Specifically, we evaluate the following algorithm:

1. Measure the current link signature.

2. Calculate the difference metric Δ from (3.6) between the current link signature and the link signatures in the FIFO history $\mathcal{H}$.

3. Compare the difference Δ to a threshold γ. If $\Delta > \gamma$, raise an alarm to indicate that the receiver has moved since the last link signature was measured. If $\Delta < \gamma$, do not raise an alarm, thereby indicating that the receiver has not moved since the last link signature was measured.

4. Add the current link signature to a delay buffer and add any link signature in the delay buffer older than D to the FIFO history $\mathcal{H}$, where $D \in \mathbb{R}$ is a time delay.

5. Return to step 1.

We evaluate the performance of this algorithm for various thresholds γ in order to identify system tradeoffs and characterize achievable system performance.

3.4 Measurements

We describe two MIMO measurement experiments. One is performed at Brigham Young University [57], and another is performed at the University of Utah. These experiments provide an opportunity to examine the following two use cases for location distinction:

1. A wireless device sends packets while in motion so that each new packet is sent from a distinct location. In this case, the location distinction algorithm should detect the change with every new packet. Our Experiment I provides MIMO data to test the performance of location distinction in this use case.

2. A wireless device sends packets while stationary for a long period of time. Then, a new packet is sent from a distinct location, either because the wireless device has moved, or because a second wireless device is attempting to impersonate the first from a different location. In either case, the location distinction algorithm should detect the change with the final transmission. Our Experiment II provides MIMO data to test the performance of location distinction for this use case.

Under both use cases, in order to simulate MIMO antenna arrays of different sizes and examine the associated performance of temporal signature-based location distinction, we compile the MIMO link signatures, as in (3.4) and (3.5), from the subsets of the SISO link signatures, CTLS and TLS, measured with $1 \times k$ and $k \times k$ antenna arrays, where $k \in \{1, \ldots, 8\}$. In both of these experiments, the receivers change position and the transmitters are stationary, but the reciprocity of the radio channel allows us to operate as if the opposite were true.

3.4.1 Experiment I

In the first experiment, conducted at Brigham Young University by Wallace et al. [57], MIMO channel data are collected using an 8x8 MIMO channel sounder in

which a multitone baseband signal is mixed with a carrier frequency of 2.55 GHz and transmitted to stationary and moving receivers. The transmitter is stationary for these measurements. The multitone signal is 80 MHz wide.

Channel measurements are collected at eight different receiver locations on a single floor of an office building. Figure 3.1 is a diagram showing the first three receiver locations. The circled numbers represent each location.

In the cases where the receiver is moving, it moves with a speed of 31.75 cm/sec. Note that this speed is about 0.7 miles per hour or 1.1 km per hour, which corresponds to a relatively slow walking speed. At each receiver location, between 390 and 585 measurements are made, depending on the space available for receiver motion. In the measurements made with a moving receiver, the multitone probe is sent every 3.2 ms, or given the receiver speed of 31.75 cm/sec, every 1.016 mm. These spatially dense measurements are the reason we delay (D) inserting the most recently measured link signature into the history $\mathcal{H}$. As we show in Section 3.5, the performance of location distinction improves when this delay is increased, or equivalently, when the current location of the receiver is further from the location it occupied during the measurements in the history $\mathcal{H}$.

3.4.2 Experiment II

The second experiment is performed at the University of Utah. Channel measurements are made at a center frequency of 2.42 GHz using a MIMO-OFDM transceiver implemented with a National Instruments vector signal generator (VSG) and vector signal analyzer (VSA) and Labview software.

The transmitted signal is designed to emulate the IEEE 802.11n standard [58]. It is an OFDM signal and has 64 subcarriers contained in a total bandwidth of 20 MHz (312.5 kHz per subcarrier). The frame (timing) synchronization, carrier offset recovery, and channel estimation are using the greenfield preamble described in the physical layer specification of the IEEE 802.11n standard, but we omit the high throughput signal field. This field is normally used to convey MAC information regarding the coding, modulation scheme, etc., and is not necessary for the channel estimation required by this experiment. Moose's method is used for frame synchronization and carrier recovery [38, 59].

The MIMO channel state is estimated using mutually orthogonal sequences. A minimum mean-squared-error (MMSE) channel estimation algorithm with a structure derived from the MMSE estimator in [60] is employed, but we increase the number of transmit symbols used for estimating the channel from two symbols (for a 2x2 system) to four symbols.

In order to ensure their accuracy, channel measurements are only recorded for packets with low bit error rate. Similarly, future location distinction implementations may ensure accurate channel measurements and improve system performance by only recording measurements made on correctly decoded packets, thereby avoiding the effects of interference from other transmitters.

The data are collected in the Wireless Communication Lab at the University of Utah, an open plan office lab containing desks, bookcases, chairs, and measurement equipment. We take measurements at eighteen different receiver locations and four different transmitter locations, as shown in Figure 3.2, resulting in a total of 3600 measurements of 72 distinct radio links.

3.5 Results and Discussion

In order to evaluate the performance of our location distinction algorithm, we define the outputs of the difference metric (3.6) for the MIMO CTLS and MIMO TLS as

$$E_f = \Delta(\tilde{\mathbf{F}}, \mathcal{H})$$

and

$$E_g = \Delta(\tilde{\mathbf{G}}, \mathcal{H}),$$

respectively. These values are recorded under experimental conditions corresponding to the following null and alternate hypotheses:

$\mathbb{H}_0$: The receiver has not changed position.

$\mathbb{H}_1$: The receiver has changed position.

This allows us to frame location distinction as a standard threshold-based detection problem, as discussed in [61], and produce receiver operating characteristic (ROC) curves which quantify the tradeoff between false alarm P_{FA} and detection P_D rates

under various thresholds γ. The P_{FA} and P_D as a function of γ allow us to evaluate how well location distinction would have worked if a threshold of γ was used in the real-time algorithm. Thus the set of possible P_{FA}/P_D combinations provide a curve of feasible real-time detection performance.

In this section, we present and discuss these results in the context of four link signature characteristics: spatial distance between link signature measurements, the size of the history $\mathcal{H}$, the number of antennas in the MIMO array, and link signature bandwidth.

3.5.1 Spatial Distance / Delay

The results of both experiments show that differences in spatial location between link signatures are more significant than the temporal variations in link signatures measured for stationary receivers. In other words, changing the position of the transmitter/receiver has a more significant effect on the measured link signatures than moving scatterers. Figure 3.3 shows the magnitudes of the 1x1 TLS measured at a stationary or moving receiver in Experiment I. The variation of the signatures for the moving receiver is more significant. In the case of the MIMO TLS, the same effect can be seen in the empirical distributions of the difference metric (3.6). These distributions are shown in Figure 3.4(a). The mean difference metric is much higher in the case of a moving receiver. The same result can be seen in the empirical distributions of the difference metrics calculated for Experiment II. These distributions are shown in Figure 3.4(b).

Figure 3.5(a) shows the average ℓ_2 and ϕ_2 distances between 8x8 MIMO CTLSs as a function of receiver separation where the ϕ_2 distance is defined in (3.8). The average ℓ_2-distance reaches a maximum at a separation of approximately $\lambda/2$ ($\approx$ 12.5 cm for our testbeds), and then oscillates with a period of λ. This result agrees with a result of the Clarke fading model, which assumes incoming multipath are uniformly distributed about the receiver [47]. The average ϕ_2-distance peaks at a receiver separation of about λ and the oscillation is mitigated by the phase rotation inherent in the ϕ_2-distance. Figure 3.5(b) shows the average difference metrics E, calculated according to (3.7), as a function of receiver separation. These results indicate that the difference metrics perform best in the case where the receiver has moved about a half-wavelength

between measurements. We note that for the detection of impersonation attacks it is very reasonable to assume that the attacker's antennas are more than $\lambda/2$ away from the antennas of the device being impersonated.

Under use case #1, the spatial distance between the signatures in the history and the most recent signature is determined by the delay D. Choosing D to be larger than the coherence time of the channel ensures that the signatures will offer sufficient decorrelation. If D is less than the coherence time of the channel, missed detections will increase. A simple approach to decide on an appropriate value for D is to use the tighter estimate suggested in [47] for estimating 50% coherence time T_c:

$$T_c = \frac{9}{16\pi f_d} \tag{3.9}$$

where f_d is the maximum Doppler shift, which is proportional to the velocity of the moving transceiver. This Doppler shift can be estimated using one of the methods reviewed in [62], or it can be computed using the lowest transceiver velocity to be detected. We note that the delay improves performance under use case #1, but has no effect under use case #2.

The average maximum Doppler for a moving receiver in Experiment I is approximately 5 Hz [57]. Using (3.9), this yields a coherence time of approximately 85 ms. In our analysis we examined delays of 32, 64, 96, and 128 ms. While performance improves with delay, it stabilizes for $D > T_c$. An example of this can be seen in Figure 3.6, which presents the false alarm rates P_{FA} vs. D for $P_M = 1 \times 10^{-4}$ using the 8x8 MIMO TLS, where P_M is the probability of a missed detection. We note that this corresponds to approximately a quarter of a wavelength, suggesting that (3.9) leads to a smaller than ideal delay, but the performance gain associated with larger D is minimal.

3.5.2 History Size

The optimal number of signatures to include in the history depends on the distributions of the differences measured under $\mathbb{H}_0$ and $\mathbb{H}_1$. We examine a range of history sizes in both experiments in order to understand how this parameter affects location distinction performance. Because of the minimum operator in (3.6), increasing the history size can only lower the average difference metric, E, under both hypotheses.

This has the effect of decreasing false alarms at the expense of an increase in missed detections.

Figure 3.7(a) shows the ROC curve of the location distinction algorithm for the 8x8 MIMO CTLS of Experiment I and various history sizes. In this case, the best performance corresponds to a history containing fifteen previous link signatures. Figure 3.7(b) shows the ROC curve of the location distinction algorithm for the 2x2 CTLS of Experiment II and various history sizes. In this case, a history size of five offers the best performance. The difference in optimal history size can be understood in terms of the marginal distributions from Figure 3.4.

The difference metrics E_f measured under $\mathbb{H}_0$ in Experiment I have a significantly higher mean and variance than those measured under the same hypothesis in Experiment II, indicating that the temporal variations of the link signatures measured for a stationary receiver in Experiment I are more prominent than those in Experiment II. Therefore, a larger history size is necessary in Experiment I in order to capture the temporal variations of the stationary receiver. In general, the history size should increase with the temporal variations in the channel and/or system noise. Future work should investigate adaptively setting the history size based on current channel conditions.

3.5.3 Number of Antennas

The results show that as the size of the MIMO antenna array is increased, the performance of the location distinction algorithm improves. This is consistent with the simulation results of [7], which use a ray-tracing simulation to show that the average miss rate in a location distinction system decreases with the number of antenna elements.

Figure 3.8 shows the location distinction ROC curves for the data from Experiment I and MIMO antenna arrays with k_1 transmit antennas and k_2 receive antennas for various combinations of k_1 and k_2. Figure 3.9 shows the ROC curves for the same experiment, but using SIMO arrangements. The trend in these figures is toward better location distinction performance with the increase in size of the MIMO antenna array. Figure 3.10 shows the achievable miss rates for a false alarm rate of 2×10^{-3} for various SISO, SIMO, and MIMO arrays. The miss rates appear to follow the inverse power

law

$$P_M = \frac{b}{(k_1 k_2)^m} \tag{3.10}$$

where b and m are parameters that define the rate that the probability of missed detection approaches zero with the number of channels. A least-squares approximation yields $b \approx 10^{-1.44}$ and $m \approx 0.93$ for the data in Figure 3.10. As a rule of thumb, the achievable miss rate for a constant false alarm rate is approximately inversely proportional to $k_1 k_2$, the number of channels. In general, the power law relationship for P_m is not as conducive to rapid decrease in P_M as an exponential decrease would be, for example. If the relationship holds for $k_1 k_2 > 64$, then it would require significant increases in the number of antennas to further reduce P_M.

However, we note that the miss rate shows dramatic improvement for $k_1 k_2 = 4$ (2x2 or 1x4), compared to 1x1, MIMO systems. Table 3.1 shows the improvement of the location distinction algorithm in a 2x2 MIMO channel over the SISO channel in Experiment I. There is as much as a 108-fold reduction in the miss rate for a constant false alarm rate when changing from SISO to 2x2 MIMO.

3.5.4 MIMO CTLS and TLS

In comparing Figures 3.8(a) and 3.8(b), it is also apparent that the MIMO CTLS and its associated difference metric leads to better performance than the MIMO TLS in Experiment I. Table 3.1 shows the improvement of the location distinction algorithm when using the MIMO CTLS instead of the MIMO TLS. Using the MIMO CTLS results in as much as a 133-fold reduction in miss rate for a constant false alarm rate.

This result is also confirmed in Experiment II, as shown in Table 3.1. In Experiment II, the 1x1 CTLS results in a 3.5-fold improvement in miss rate over the 1x1 TLS. The 2x2 TLS and 2x2 CTLS both reach the lowest measurable miss rate in Experiment II.

3.5.5 Link Signature Bandwidth

Another crucial parameter in both experiments, and typically a limiting factor in radio design, is system bandwidth. We examine the performance of the location

distinction algorithm over a range of bandwidths by varying the number of tones included in the frequency-domain measurements from Experiment I.

Figure 3.11 shows that performance typically improves with bandwidth, but it does so with diminishing returns. This is consistent with the simulation results of [7], which show that the miss rate of a location distinction system decreases with system bandwidth, but that the performance gain of MIMO over SISO also decreases, because at high bandwidths the SISO link signatures offer sufficient decorrelation.

However, at high bandwidths the algorithm is more sensitive to timing-synchronization errors that might be hidden by lower bandwidth signatures. Figure 3.12 shows an example of two consecutively measured link signatures that exhibit this effect. These errors cause small drops in performance. The higher bandwidth of the link signatures measured in Experiment I (80 MHz) allows for better location distinction performance, but the results for the 2x2 MIMO link signatures of Experiment II (20 MHz) still offer a 3×10^{-4} probability of missed detection for a 7×10^{-3} probability of false alarm.

3.6 Related Work

The papers discussed in this section have contributed to this work in different aspects. The most closely related work is presented in [14] and [6]. In these two papers, a temporal link signature is defined to be used in the context of multiple transmitters/receivers and then refined to include phase information. We compliment that work by showing that a single MIMO transmitter/receiver pair can be used to perform reliable location distinction, and that lower false alarm rates are possible using a single receiver, when the communication system is a 1x2 or 2x2 MIMO system. In [6], the authors report a 9×10^{-3} miss rate for a 0.01 false alarm rate using three receivers. For the same false alarm rate, we are able to achieve a 3×10^{-4} miss rate using a single receiver and the 2x2 MIMO CTLS with less bandwidth. This net reduction in system complexity may enable location distinction in future wireless networking systems.

In [6], a *complex temporal link signature* is defined which allows for the exploitation of the phase information in the CIR. However, not all of the phase information

represented by the link signature is due to the channel. Some phase shifts occur due to a lack of time and/or frequency synchronization between the transmitter and receiver. The distance between two link signatures which minimizes the contribution of random phase shifts is shown to be (3.8); [6] calls this the ϕ_2-distance.

In [7], ray-tracing simulation results for MIMO location distinction in defense of impersonation attacks in an office building are presented. The authors assume that channel measurements made in the frequency domain are distributed as complex Gaussian random variables and derive ideal change metrics based on this assumption. We extend this work by offering an experimental validation of MIMO location distinction using two MIMO testbeds.

In [56], the authors propose some of the underlying ideas of this work, namely, that characteristics of the radio channel (rapid decorrelation in space, time, and frequency) can be exploited to secure wireless networks. They offer methods of probing the channel in order to determine, based on the channel gains between transmitters and receivers, whether or not communications are coming from an authentic user or a would-be attacker. Using the USRP/GNU Radio and a simple change-point detector, they show that they are able to detect a change in the wireless link via channel gains and thereby detect a possible spoofing attack.

In [55], the authors utilize similar principles in designing a method for identifying a transmitter by its *signalprint*, which consists of a vector of RSS values. These RSS values are gathered using wireless access points as sensors and a central authentication server for cataloging and comparing signalprints. Their results show that a stationary transmitter will produce a consistent signalprint and thereby allow for discrimination between authentic users and attackers whose signalprints will vary significantly because they are located in a different position in the multipath fading channel. The signalprint is limited in that it may be unable to detect attackers located near authentic transmitters, because they may have similar signalprints.

Much of the discussed work has suggested using channel measurements gathered at multiple receivers in order to perform location distinction [6,14,54,55,63]. However, in typical WiFi networks, adjacent access points are set to operate on different channels in order to reduce interference and clients operate on a single channel. This

makes collecting channel data at multiple access points difficult. Extending location distinction to MIMO allows robust location distinction to be performed with a single receiver.

3.7 Conclusion and Future Work

In order to design a reliable MIMO location distinction system with limited resources, a system designer should understand in what aspects increasing system complexity will lead to better system performance. We present an extensive experimental evaluation of MIMO location distinction using two experimental test beds. The results show that there are diminishing returns for certain aspects of system design. Increasing the number of antennas offers diminishing returns after 1x4 SIMO or 2x2 MIMO, which favors MIMO systems with a small number of antennas. Increasing system bandwidths beyond 20 MHz offers diminishing returns as well, partially because lower bandwidths tend to mask the effects of timing synchronization errors on link signatures. However, the experiments show that MIMO location distinction performs very well with just a single receiver. We detail performance tradeoffs regarding the size of the history for optimal performance. This combined knowledge will benefit anyone seeking to implement a location distinction algorithm.

In the future, it will be beneficial to further characterize the link signatures used for location distinction and explore other difference metrics. For instance, our current difference metric uses the minimum Euclidean or ϕ_2-distance between the most recent link signature and those in the history $\mathcal{H}$. This tends to increase the miss rate in the context of noisy measurements. A weighted average of distances, such as the Mahalanobis distance, may offer better performance. A broader experimental analysis of link signatures and their temporal and spatial variations will facilitate the design of better difference metrics.

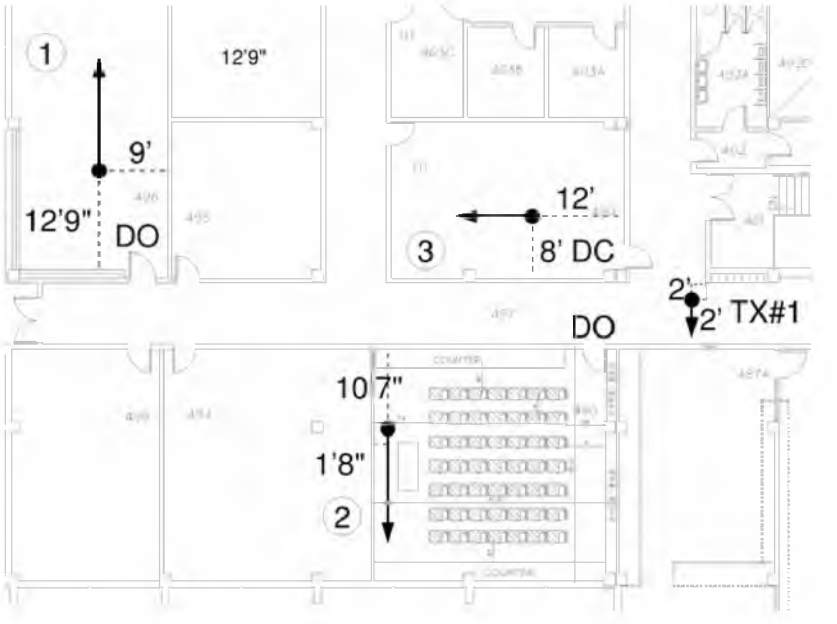

Figure 3.1. Diagram of a subset of receiver locations from Experiment I. Circled numbers represent the receiver locations for individual measurement sets. DO or DC indicate door open or door closed, respectively.

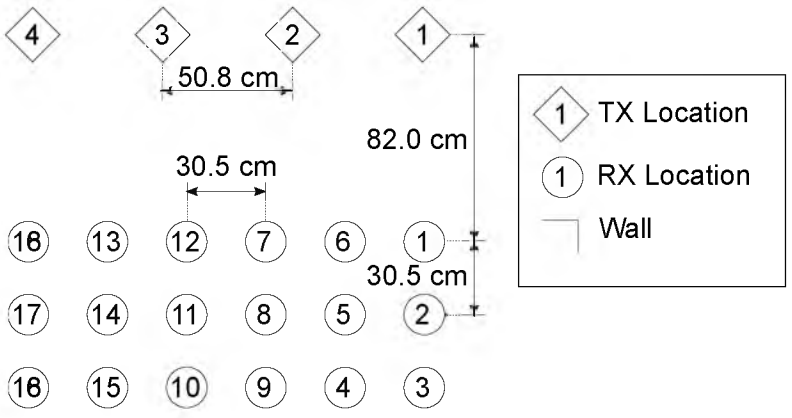

Figure 3.2. Diagram of Experiment II. Circles represent receiver locations, diamonds represent transmitter locations. The outer line represents the wall of the room. Channel measurements are made at each transmitter/receiver location. Desks, equipment, and other scatterers are present, but not depicted in this diagram.

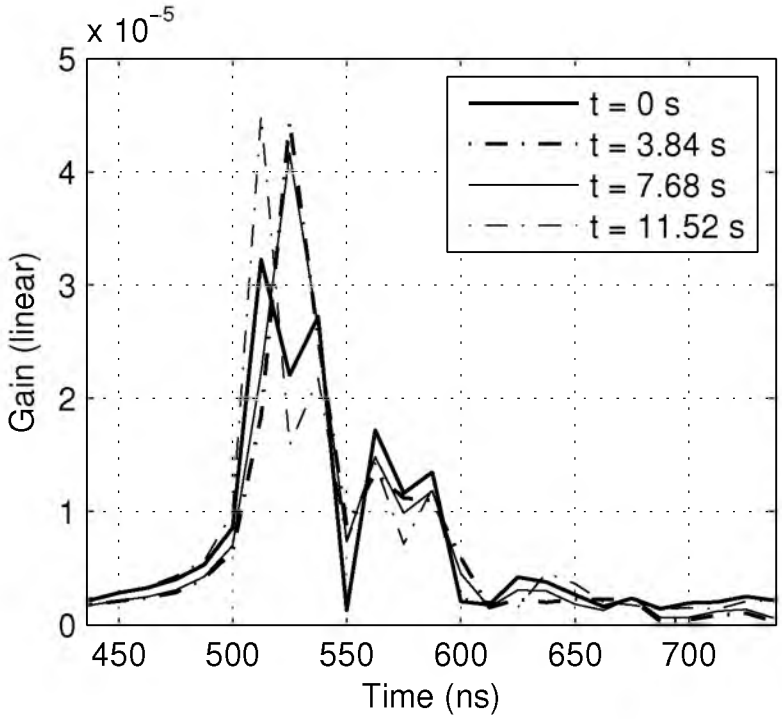

(a)

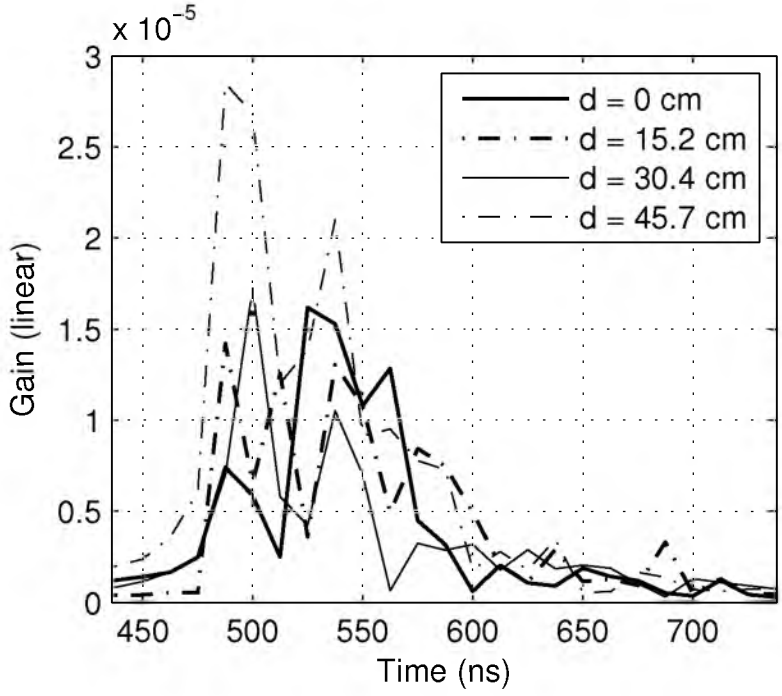

(b)

Figure 3.3. Link signatures measured (a) over time at a stationary receiver and (b) at a moving receiver. The signatures measured at a moving receiver fluctuate more than those measured at the stationary receiver.

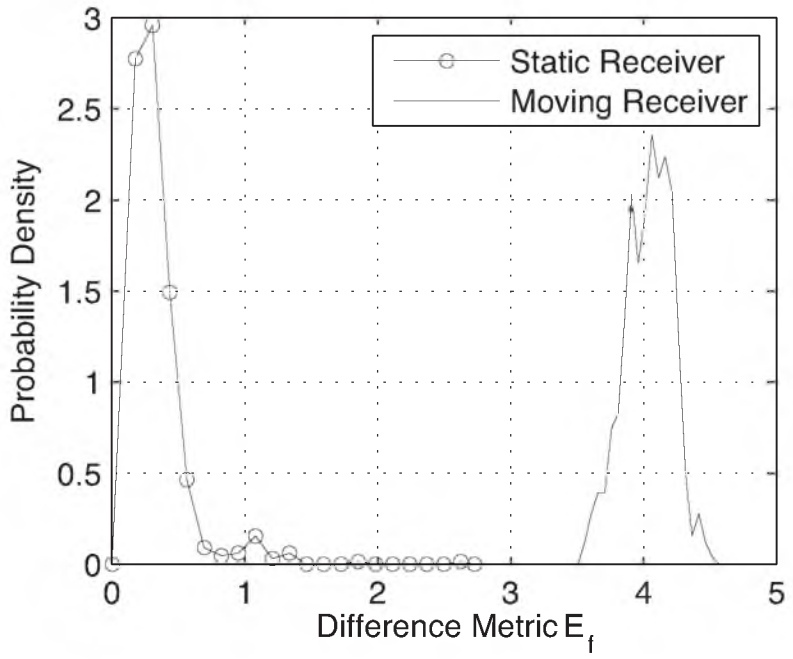

(a)

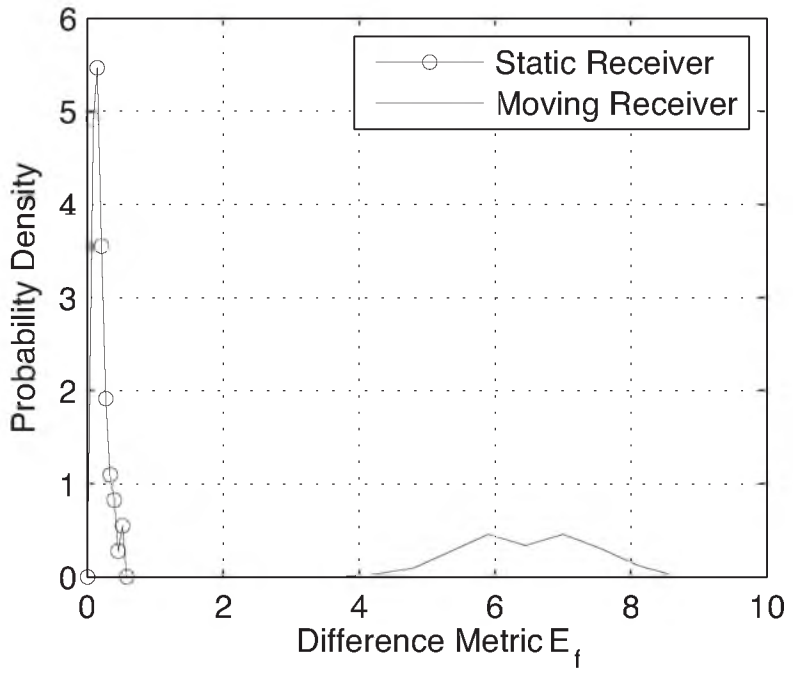

(b)

Figure 3.4. Empirical distributions of E_f for stationary and moving receiver from (a) Experiment I with 8x8 CTLS, and (b) Experiment II with the 2x2 CTLS. In both cases the mean difference metric for a moving receiver is significantly higher than for a stationary receiver.

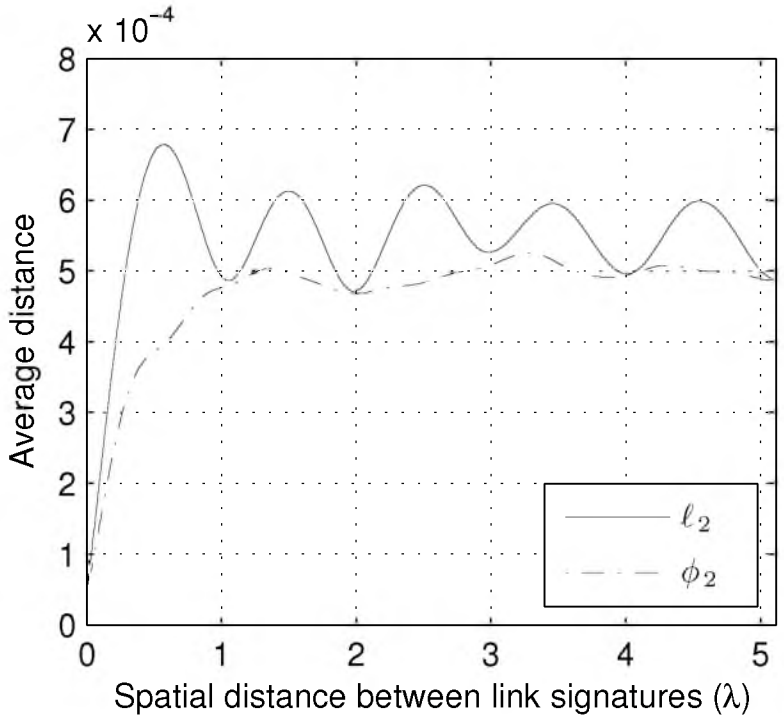

(a)

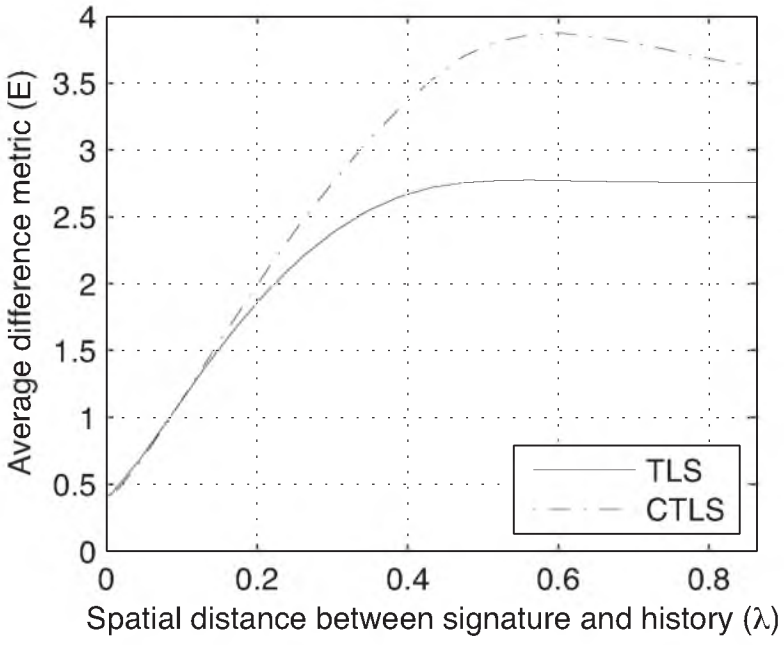

(b)

Figure 3.5. (a) Average ℓ_2 and ϕ_2-distances between 8x8 MIMO CTLS as a function of spacial separation. The average ℓ_2-distance peaks at a receiver separation of roughly $\lambda/2$. (b) Average difference metrics E for 8x8 CTLS/TLS as a function of spatial separation.

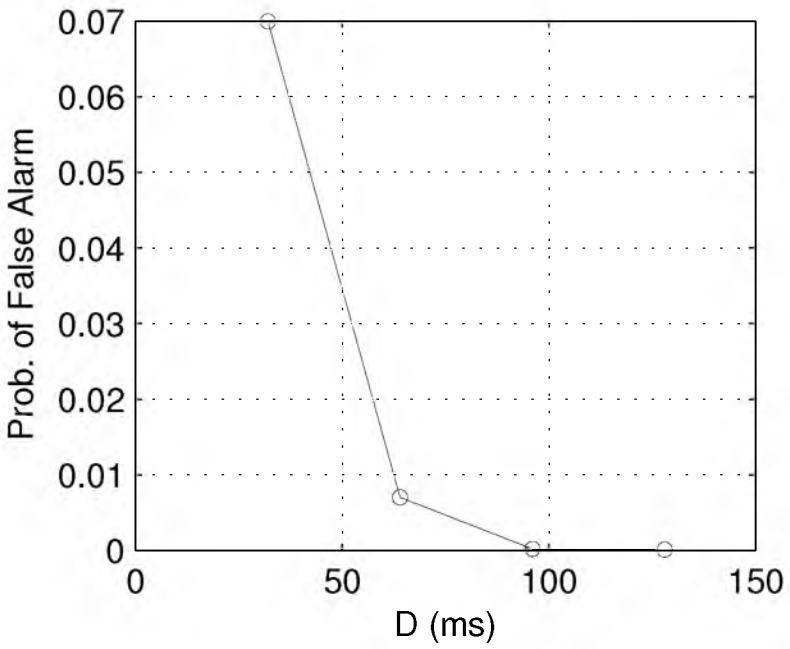

Figure 3.6. Probability of false alarm vs. delay D for a miss rate of 1×10^{-4} for the 8x8 MIMO TLS. Performance gain stabilizes for delays larger than 85 ms, the coherence time of the channel.

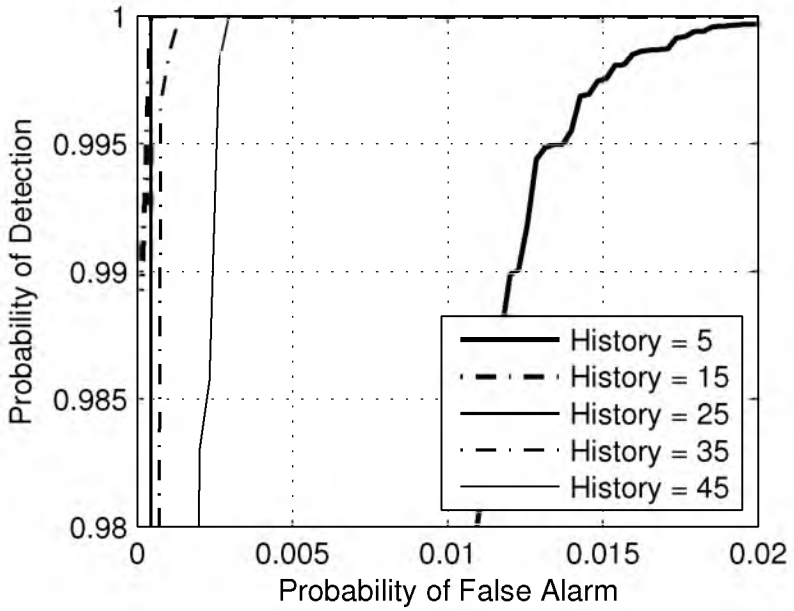

(a)

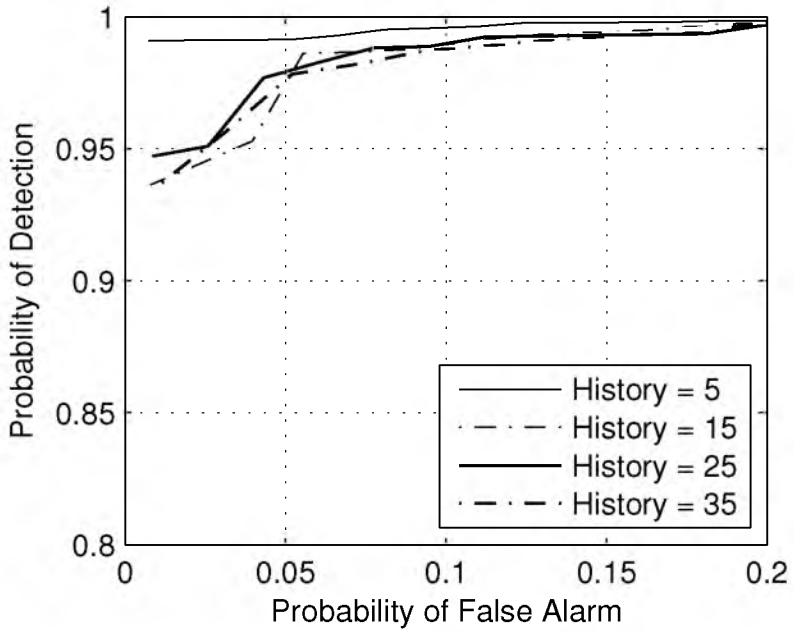

(b)

Figure 3.7. ROC curves for (a) Experiment I: 8x8 MIMO CTLS and (b) Experiment II: 1x1 CTLS for various history sizes. In Experiment I, a history size of fifteen link signatures yields the best performance. In Experiment II, a history size of five link signatures yields the best performance.

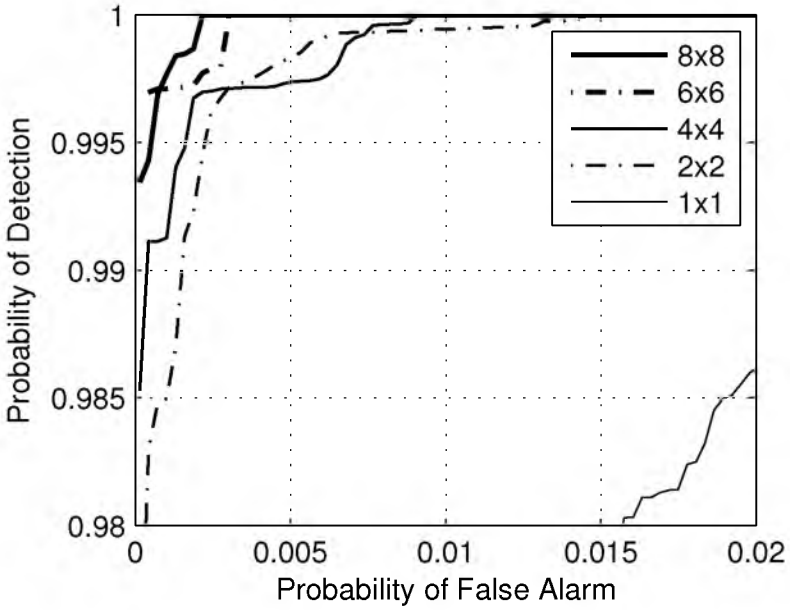

(a)

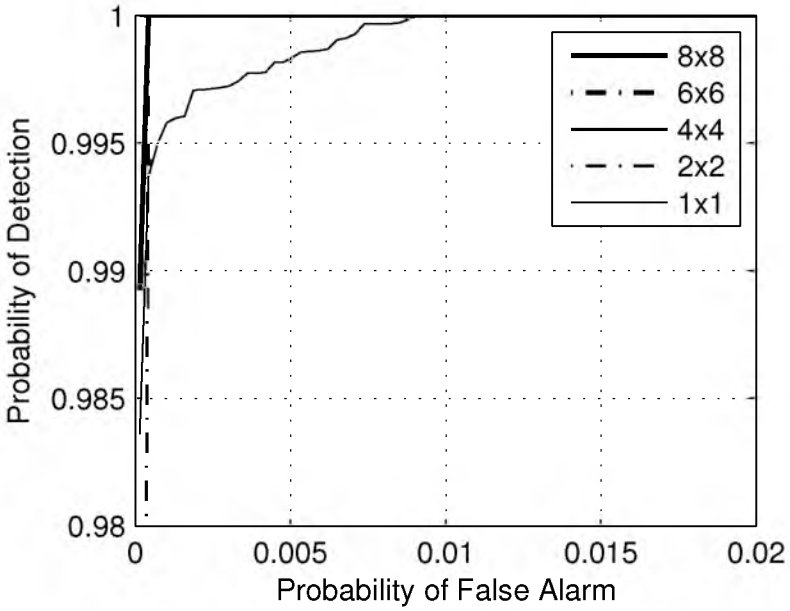

(b)

Figure 3.8. ROC curves for (a) MIMO TLS and (b) MIMO CTLS for various antenna array sizes. Location distinction performance improves with the number of antennas and the MIMO CTLS performs better than the MIMO TLS.

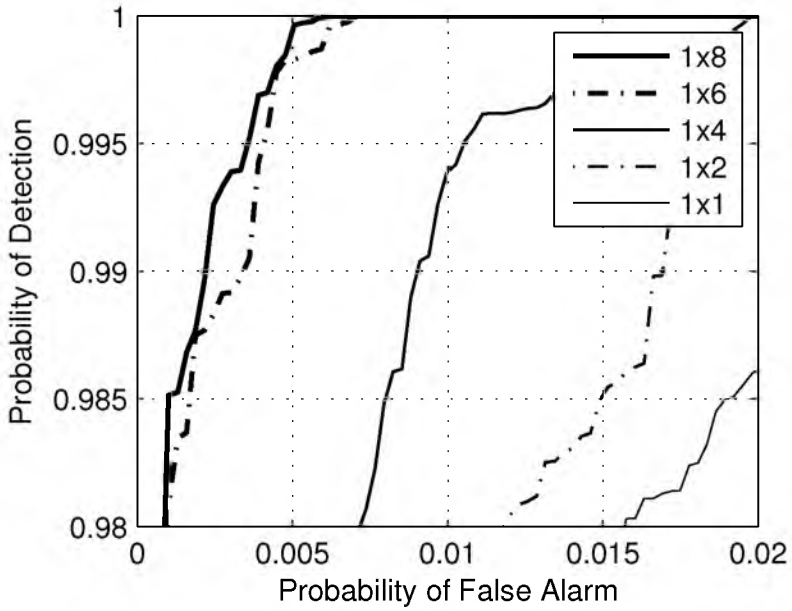

(a)

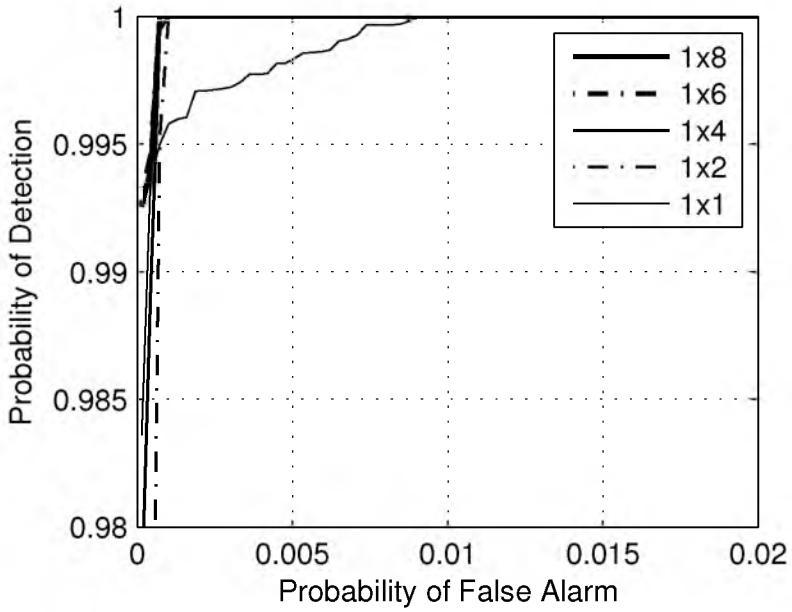

(b)

Figure 3.9. ROC curves for (a) SIMO TLS and (b) SIMO CTLS for various antenna array sizes. Location distinction performance improves with the number of antennas and the SIMO CTLS performs better than the SIMO TLS. The SIMO signatures nearly match the performance of the MIMO signatures.

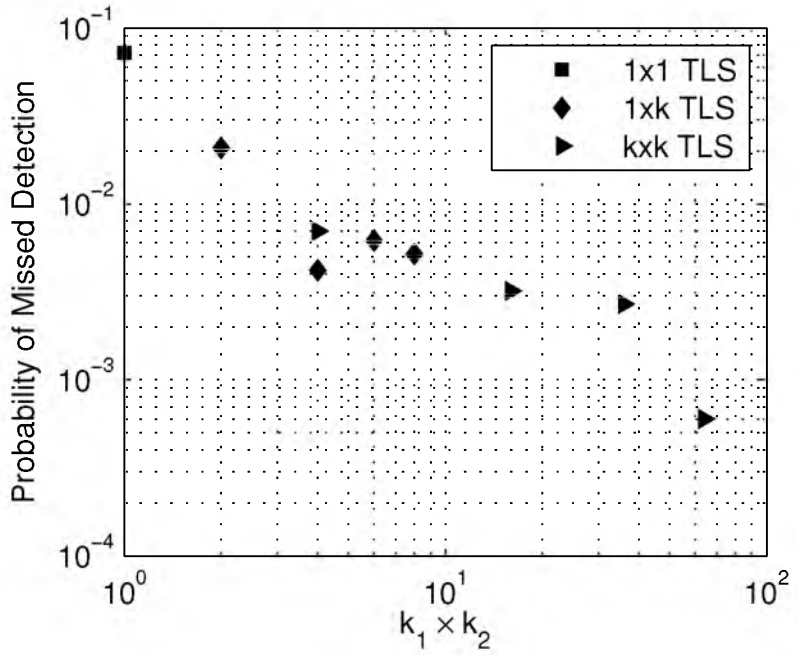

Figure 3.10. Experiment I: Probability of missed detection for a 2×10^{-3} probability of false alarm vs. k_1k_2 for different SISO, MIMO, and SIMO arrays.

Table 3.1. P_M for $P_{FA} = 10^{-2}$ for Experiments I and II.

	Experiment I			Experiment II		
(k_1, k_2)	MIMO TLS P_M	MIMO CTLS P_M	CTLS/TLS Improvement	MIMO TLS P_M	MIMO CTLS P_M	CTLS/TLS Improvement
(1,1)	0.032	$\leq$ 0.00024	$\geq$ 133x	0.0323	0.0092	$\approx$ 3.5x
(2,2)	0.0005	$\leq$ 0.00024	$\geq$ 2x	$\leq$ 0.0003	$\leq$ 0.0003	N/A
MIMO Improvement	$\approx$ 64x	N/A		$\geq$ 108x	$\geq$ 31x	

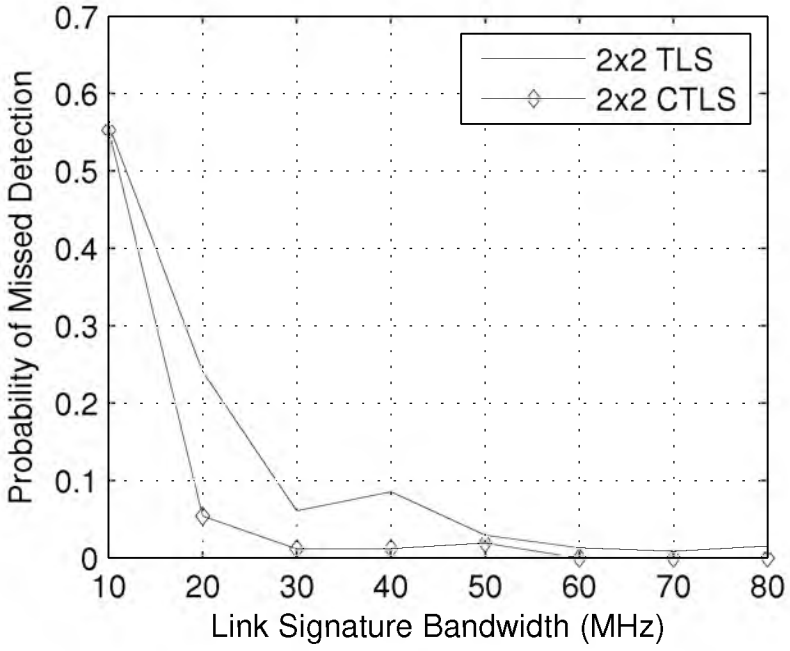

Figure 3.11. Location distinction miss rate vs. link signature bandwidth for a 7×10^{-4} false alarm rate in Experiment I. Increasing bandwidth offers diminishing returns.

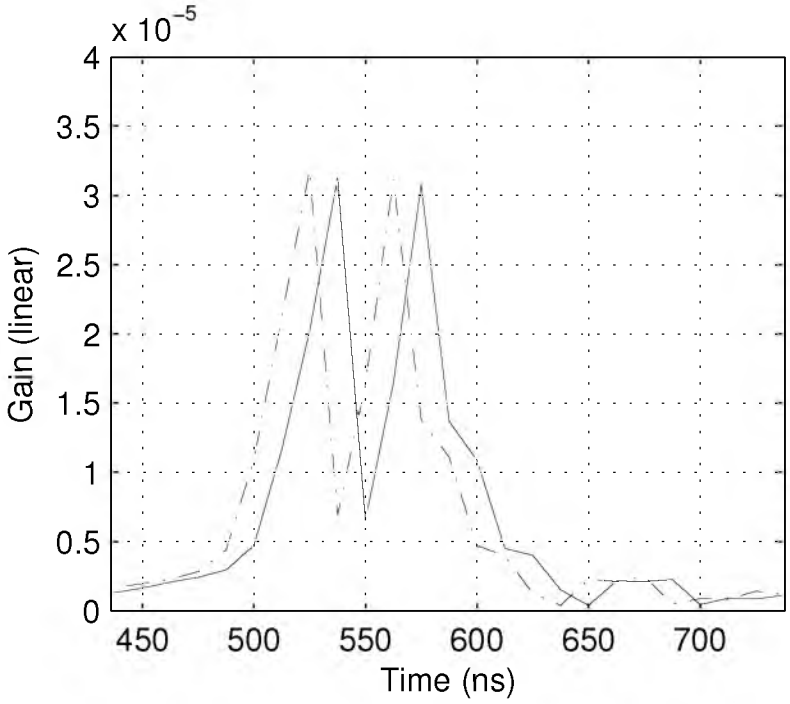

Figure 3.12. Two consecutive link signatures with 80 MHz bandwidth showing the results of a timing-synchronization error. The time-resolution of high-bandwidth link signatures cause an increased impact on location distinction performance.

CHAPTER 4

EXPLOITING RADIO WINDOWS FOR THROUGH-WALL LOCATION INFORMATION [1]

4.1 Abstract

We introduce and investigate the ability of an attacker to surreptitiously use an otherwise secure wireless network to detect moving people through walls, in an area in which people expect their location to be private. We call this attack on location privacy of people an "exploiting radio windows" (ERW) attack. We design and implement the ERW attack methodology for obtaining through wall people location information that relies on reliably detecting when people cross the link lines by using physical layer measurements between the legitimate transmitters and the attack receivers. We also develop a method to estimate the direction of movement of a person from the sequence of link lines crossed during a short time interval. Additionally, we describe how an attacker may estimate any artificial changes in transmit power (used as a countermeasure), compensate for these power changes using measurements from sufficient number of links, and still detect line crossings. We implement our methodology on WiFi and ZigBee nodes and experimentally evaluate the ERW attack by monitoring people movements through walls in two real-world settings. We find that our methods achieve very high accuracy in detecting line crossings and determining direction of motion.

---

[1] Submitted to *IEEE Transactions on Networking* as A. Banerjee, D. Maas, M. Bocca, N. Patwari, S.K. Kasera, "Exploiting Radio Windows for Through-wall Location Information," September, 2013.

4.2 Introduction

We investigate an attack on the privacy of the location of people moving in an area covered by a wireless network. People moving in an area covered by one or more wireless networks affect the way radio signals propagate. We demonstrate that the presence, location and direction of movement of people not carrying any wireless device can be "eavesdropped" by using the channel information of wireless links artificially created by an attacker by deploying sensing devices or *receivers* that can "hear" *transmitters* such as WiFi access points (APs), composing the legitimate wireless network. Signals from the transmitters passing through nonmetal external walls that allow radio waves to go through, are analogous to light from light bulbs passing through glass windows which an adversary can use to "see" where people are in a building. Hence, we call this attack on location privacy of people an "exploiting radio windows" (ERW) attack.

Consider a building where security is important, e.g., an embassy, with a concrete exterior wall. One or more wireless networks may have been set up in this building to transfer different types of data, including voice and video. We can expect these networks to implement advanced data security protocols to prevent eavesdropping of data. However, an attacker can still deploy receivers outside the wall of the building to measure different parameters of the received radio signals. By measuring the channel state information (CSI) or received signal strength (RSS), for example, of the links from the transmitters inside the building to the receivers deployed, the attacker can monitor the movements of people and objects inside the building in the area behind the wall in Figure 4.1. The information about people's movements can be put to malicious use including planning a physical attack on the personnel inside the building. On the contrary, law enforcement personnel can apply similar techniques in the case of a hostage situation to track activity inside a large building and plan their operation accordingly.

In this paper, we design and implement the ERW attack methodology for through wall people localization. Our methodology relies on reliably detecting when people cross the link lines between the legitimate transmitters and the attack receivers. We first develop a majority-vote based detection algorithm that reliably detects line of

sight (LOS) crossing between the legitimate transmitter and the attack receivers by comparing short-term variances in link channel information with its long-term counterpart. We also develop a method to estimate the direction of movement of a person from the sequence of link lines crossed during a short time interval. Next, we implement our methodology on WiFi and ZigBee nodes and experimentally evaluate the ERW attack by monitoring people movements through walls in two real-world settings — a hallway of a university building separated from the outside by a 1 ft thick concrete wall, and a residential house. When we use two WiFi 802.11n nodes with normal antenna separation, or two groups of ZigBee nodes as attack receivers, we find that our methods achieve close to 100% accuracy in detecting line crossings and the direction of movement. We also find that our methods achieve $90 - 100\%$ accuracy when we use a single 802.11n attack receiver.

To protect the privacy of the location information from the ERW attack, the owner of the legitimate network may choose to implement a countermeasure in which the transmitters vary their transmit power during successive transmissions. The artificial transmit power changes can be either random or follow a predefined profile replicating the typical channel variations introduced when a person crosses a link line. This countermeasure is expected to introduce additional variability in the received signal measured by the attack receivers and can be wrongly interpreted by the attacker as caused by moving people or objects crossing the link lines. In this paper, we demonstrate that an attacker who can measure a sufficient number of links can accurately estimate the artificial transmit power change, compensate for it, and ultimately locate people and monitor their movements. We base our compensation strategy on the following intuition: an artificial transmit power change at a transmitter will impact all the links between the transmitter and the attack receivers, whereas genuine power changes due to human movement are likely to impact only some of the links.

The ERW attack described in this paper is significantly different than device-free localization[2] (DFL) in that the ERW attack is practical for large buildings, is stealthy

[2]In which people who are not carrying any radio transmitters are located by a static deployed network.

because no transmitters are deployed by the attacker, and is immune from jamming. DFL systems such as the ones in [9, 64–70] require dozens of radio transceivers deployed throughout or on many sides of the target area. Further, through-building DFL systems such as [70, 71] assume the transmitted signal penetrates through two external walls and any internal walls in between, and as such have been tested only in buildings of small (18 - 42 m^2) size. In this paper, we show access to one side is sufficient for an ERW attack, and it requires a signal from inside a building to penetrate only one external wall. Other fingerprint-based DFL systems [12, 72–74] require collection of training data with a person in each possible location in the environment. In our ERW attack, we do not assume that an attacker has prior access to the inside of the building to be able to perform such data collection. Further, to perform DFL, an attacker must deploy some nodes which transmit, exposing them to being detected and located by RF source localization, while an ERW attack is stealthier in that purely passive receivers are deployed by an attacker. Finally, DFL systems' signals could be interfered with by a powerful jammer. In the method in this paper, any transmitter in the building, including a jammer, could be used as a source for ERW.

The remainder of the paper is organized as follows. In Section 4.3, we describe the adversary model. In Section 4.4, we formulate the methods used to detect link line crossings and estimate changes in transmit power. We also describe the method used to determine the direction of motion of the person. The experimental setup is presented in Section 4.5. In Section 4.6, we present the results of our experiments. Section 4.7 discusses the previous research in the area of location privacy attacks in wireless networks. Conclusions and directions for future work are given in Section 4.8.

4.3 Adversary Model

We make the following assumptions about the attacker[3]:

- The attacker is able to deploy multiple wireless sensing devices within the trans-

[3]In this paper, we use the term attacker for anyone, whether malicious or genuine, who is trying to localize humans.

mission range of the legitimate transmitter(s) outside the area being monitored. The attacker is able to measure the physical layer information (RSS and/or CSI) of the links between the transmitter(s) and the attack receivers.

- The attacker does not have access to the content of the packets transmitted by the legitimate network nodes.
- The attacker does not deploy any transmitters, nor does he have any control over the legitimate transmitters. However, he requires the legitimate transmitters to transmit packets frequently enough to perform line crossing detections.
- The attacker does not make any assumption regarding the transmit power profile of the transmitters.
- The attacker nodes do not associate or interfere in any manner with the transmissions of the legitimate transmitter(s).
- The attacker may not know the precise location of the transmitters or the arrangement of their antennas. However, we do assume that a transmitter is located well inside the perimeters of buildings for network coverage reasons ensuring that they do not lie between the people (being localized) and the attack receivers.

4.4 Methodology

In this section, we first develop a methodology to detect line crossings based on a majority vote for WiFi 802.11n receivers. We also develop a method that uses a sequence of line crossings to determine the direction of human movement. Next, we present our approaches for estimating transmit power change and its compensation, when the transmit power is artificially changed by the owner of the wireless transmitters, inside a secure building, with the hope of preserving location privacy. Last, we show how we adapt our methodology for IEEE 802.15.4 ZigBee attack receivers.

4.4.1 Line Crossing Detection

Many modern WiFi networks use the 802.11n standard, in which transceivers are equipped with multiple antennas in order to leverage the spatial diversity of

the wireless channel. While these multiple-input multiple-output (MIMO) systems provide high data rates, they also provide a rich source of channel information to an adversary interested in localizing people inside a building.

The 802.11n wireless standard uses the well-known orthogonal frequency-division multiplexing (OFDM) modulation scheme, which encodes and transmits data across multiple subcarriers for each transmitter-receiver antenna pair. When an 802.11n receiver receives a packet, it estimates the effect of the wireless channel on each MIMO OFDM subcarrier for the purpose of channel equalization. Since this channel state information (CSI), represented as a complex gain for each subcarrier, is measured during the unencrypted preamble of each WiFi packet, an adversary without legitimate access to data on the network can still measure the CSI for every packet.

We apply a windowed variance method for detecting abrupt changes in the CSI for a WiFi link. Let $H_{j,k}(n)$ be the magnitude of the signal strength for the jth transmitter-receiver antenna pair and the kth OFDM subcarrier for the nth packet. We define the windowed variance measurement at packet n as follows. Let

$$\bar{H}^w_{j,k}(n) = \frac{1}{w} \sum_{i=n-w+1}^{n} H_{j,k}(i), \tag{4.1}$$

$$v^w_{j,k}(n) = \frac{1}{w-1} \sum_{i=n-w+1}^{n} (H_{j,k}(i) - \bar{H}^w_{j,k})^2, \tag{4.2}$$

and

$$s^w_{j,k}(n) = \sqrt{v^w_{j,k}(n)}, \tag{4.3}$$

where, w is the number of previous CSI samples in the window. We define the subcarrier-average variance for packet n for a given antenna pair j as

$$V^w_j(n) = \frac{1}{N} \sum_k v^w_{j,k}(n), \tag{4.4}$$

where N is the number of subcarriers. We define the subcarrier-average standard deviation for packet n as

$$S^w_j(n) = \frac{1}{N} \sum_k s^w_{j,k}(n). \tag{4.5}$$

The quantities (4.4) and (4.5) represent the average CSI variance and standard deviation across all subcarriers for antenna pair j at packet n for a time window that

includes the past w packets. We track both (4.4) and (4.5) over a short-term time window w_s, and a long-term time window w_l, allowing us to compare the short-term and long-term statistics of the WiFi link and detect line crossings.

A line crossing is detected for antenna pair j when

$$\sum_{n \in D} V_j^{w_s}(n) - V_j^{w_l}(n) > \gamma(n), \tag{4.6}$$

where D is the most recent contiguous set of packets for which $V_j^{w_s}(n) - V_j^{w_l}(n) > 0$ and the threshold $\gamma(n)$ is defined as

$$\gamma(n) = V_j^{w_l}(n) + CS_j^{w_l}(n). \tag{4.7}$$

The constant C is included to allow the user to adjust the trade-off between false alarms and missed detections.

In the case where there are more than two antenna pairs, we take the majority vote between antenna pairs over the short-term window to decide if a line crossing has occurred. More specifically, when a receiver antenna detects a line crossing, we count the line crossing detections for all the receiver antennas over the short-term window, w_s. For a 3×3 MIMO transmitter and receiver, this would mean computing a majority vote over nine measurements. When the majority of the receiver antennas detect a line crossing within w_s, we infer that a person has crossed the link line between the transmitter and the receiver. We will show that this majority vote method improves the performance of our detector by decreasing false alarms and missed detections. We decrease the false alarm rate further by combining temporally close detections together. More specifically, if we detect a line crossing at time t_1 for a transmitter-receiver pair using the majority vote, we do not consider any other line crossing detected in the time interval $(t_1, t_1 + \Delta]$ for the same transmitter-receiver pair, i.e., all line crossings detected in the interval $[t_1, t_1 + \Delta]$ are considered as a single line crossing for a transmitter-receiver pair.

We note that our window-based variance method differs from the method presented in [72, 75]. In [72, 75], the authors compare recent window-based variance measurements of RSSI at multiple WiFi links to measurements made during a static calibration period when nobody is moving in the area of interest. If a certain number

of WiFi links within the area of interest detect motion within a certain time interval, a motion event is detected in the area of interest. Our attacker does not know if and/or when people are moving inside of the building, and therefore cannot create calibration measurements based on a static environment. Instead, we compare a short-term window variance to a long-term windowed variance. The long-term window allows us to capture the behavior of the wireless links when the majority of measurements are likely made while there is nobody crossing the link line. Additionally, in the case of 802.11n, we exploit the effect that line crossings have on each OFDM subcarrier and MIMO antenna pair.

4.4.2 Determining Direction of Motion

If the adversary measures the CSI at multiple receivers, or if a single receiver includes multiple antennas as is the case with 802.11n, it is also possible to infer the direction that a person is walking when line crossings are detected. The direction of motion is inferred from the time differences between the line crossing detections at each receiver, in the case of multiple receivers, or at each transmitter-receiver antenna pair, when the receivers include multiple antennas.

Consider the scenario where the attacker arranges the MIMO antenna array of an 802.11n receiver such that the antennas are roughly parallel to a hallway as shown in Figure 4.2(a). The spatial order of the antennas with reference to the hallway is known, and each transmitter-receiver antenna is given an index according to its spatial order. Based on the adversary model assumption that a transmitter is located well inside the perimeter, the attacker, even without knowing the precise location of the transmitter or the arrangement of its antennas, may treat the antennas of the wireless transmitter as if they are colocated and still achieve reliable results.

In the single WiFi receiver case, if a link crossing is detected by majority vote for a given short-term window, we find the line that best fits the set of points $\{(d_j, n_j) : j \in P\}$, where d_j is the spatial index of antenna pair j representing its location relative to the other links, n_j is the packet index indicating when a detection occurred at antenna pair j according to (4.6), and P is the set of antenna pairs ending at the WiFi receiver which detected a line crossing during the short-term window. The sign of the slope of this line indicates the direction of motion. Figure 4.2 shows an

example which uses CSI measurements from three antennas at the WiFi transmitter and three antennas at WiFi RX1 (nine antenna pairs). In the case of two single-input single-output (SISO) WiFi receivers, a similar method may be applied, but the two spatial and packet indexes directly determine the line and its slope.

4.4.3 Compensation of Transmit Power Change

In this subsection, we propose a methodology to detect artificial transmit power changes (if any) and compensate for the same. The signal strength for the jth transmitter-receiver antenna pair and the kth OFDM subcarrier for packet n is given by

$$H_{j,k}(n) = T_x(n) + G_t + G_r - L_{j,k}(n) + \Psi_{j,k}(n), \tag{4.8}$$

where $T_x(n)$ is the transmit power of the transmitter at time n, G_t and G_r are the transmitter and receiver antenna gains, respectively, $L_{j,k}(n)$ is the path loss, and $\Psi_{j,k}(i)$ is a noise term. The path loss includes all environmentally-dependent terms, including large-scale loss, shadowing, and small-scale fading. The noise term includes thermal noise, quantization noise, and other measurement noise at the attacker receiver.

The attacker cannot depend on knowing the transmit power or antenna gains. Instead, the attacker relies on the difference between the subcarrier signal strength for the packet n and the reference packet ($n = 0$) (the attacker may update the reference packet periodically to account for changes in the environment). This difference in subcarrier signal strength is given by

$$h_{j,k}(n) \triangleq H_{j,k}(n) - H_{j,k}(0). \tag{4.9}$$

From (4.8), we see that

$$h_{j,k}(n) = t_x(n) - l_{j,k}(n) + \psi_{j,k}(n), \tag{4.10}$$

where

$$\begin{aligned} t_x(n) &= T_x(n) - T_x(0), \\ l_{j,k}(n) &= L_{j,k}(n) - L_{j,k}(0), \\ \psi_{j,k}(n) &= \Psi_{j,k}(n) - \Psi_{j,k}(0). \end{aligned}$$

The subcarrier signal strength difference $h_{j,k}(n)$ above, contains transmit power changes and channel-induced changes between the nth packet and the reference packet, in addition to noise.

The ideal situation from the attacker's perspective would be that there is no artificial change in transmit powers, and that $t_x(n) = 0$ for all j and k. In this ideal situation, the subcarrier signal strength difference below is solely due to changes in the channel.

$$h_{j,k}(n) = -l_{j,k}(n) + \psi_{j,k}(n). \tag{4.11}$$

Furthermore, people crossing the line between the transmitter and receiver antennas typically cause a path loss change more significant than noise, and thus the $h_{j,k}$ signal allows direct inference of people's motion. However, when the transmitter artificially changes its transmit power, from (4.10), we cannot directly attribute a large magnitude of $h_{j,k}$ to environmental changes. In particular, if the magnitude of transmit power changes is high enough, the magnitude of $h_{j,k}(n)$ will be predominantly due to because of transmit power changes at the transmitter. A transmitter could thus presumably preserve location privacy by changing its transmit power frequently.

We now propose a method that a smart attacker can use to estimate and remove the artificial power changes and accurately detect line crossings. In our method, the attacker estimates the artificial transmit power change amplitude by correlating measurements across all antenna pairs and all subcarriers, and removes the effect of transmit power changes from the received signal strength measurements. We propose to use the median of $h_{k,j}(n)$ for all available transmitter-receiver antenna pairs and corresponding subcarriers, as an estimator of the artificial transmit power change, as shown in the equation below:

$$\hat{t}_x(n) = \text{median}\left\{h_{j,k}(n) \forall j, k\right\}. \tag{4.12}$$

Our choice of this estimator is based on the following observations. First, we observe that $t_x(n)$ appears in the equation for $h_{k,j}(n)$ for all j and k. This is because, any change in transmit power affects measurements across all transmitter-receiver antenna pairs and corresponding subcarriers simultaneously. Moreover, $t_x(n)$ is linearly related to $l_{j,k}$ and $h_{j,k}$. We also know that the change in the path loss $l_{j,k}$ is just as likely to

be positive as negative. Furthermore, any change due to human movement will not affect all the links simultaneously.

In the absence of an artificial transmit power change, $\hat{t}_x(n)$ is likely to be close to zero, i.e., our estimator does not require us to detect whether or not there is an artificial transmit power change for packet n.

The compensated signal strength for packet n, which we denote $\hat{H}_{j,k}(n)$, is given by

$$\hat{H}_{j,k}(n) = H_{j,k}(n) - \hat{t}_x(n). \tag{4.13}$$

Although the reference packet was sent with unknown transmit power $T_x(0)$, for $n > 0$, we consider $T_x(n)$ to be the relative dB shift in transmit power compared to $T_x(0)$. $\hat{H}_{j,k}(n)$ essentially, is an estimate of the subcarrier signal strength if there were no transmit power changes between the reference packet and packet n.

It is clear that, any error in the estimation of the transmit power changes amplitude will introduce additional noise in the measurements. However, the dynamics of the signal are still preserved and an attacker can use any variation in the signal over a short time period in order to notice motion of a person near the link line.

4.4.4 ZigBee Networks

The methodologies described above are also applicable for IEEE 802.15.4 ZigBee nodes. However, the ZigBee nodes are generally equipped with a single antenna, so the MIMO setup is not available. Moreover, ZigBee nodes do not use OFDM for communication, so we use channel information from a single frequency channel (instead of averaging across all subcarriers as in the case of OFDM) to evaluate our methodologies. Furthermore, there is no tool to get the complete CSI at the receiver. Instead, we rely on the RSS value obtained from the receiver hardware. Thus, in the case of ZigBee we set $H_{j,k}(n)$ to the RSS value measured in decibel units for the jth transmitter-receiver antenna pair for packet n, also $k = 1, \forall j$ as we have measurements from a single channel only.

In order to create spatial diversity we use three closely located ZigBee receivers together to form a group as described in Section 4.5. We detect line crossings by applying our majority vote approach on the three links formed between the

transmitter and the three receivers in the group. We detect direction of motion using two groups of receivers and observing sequence of groups crossed over a short time window. We estimate and compensate for artificial transmit power changes (if any) by applying the methods described in Section 4.4.3, and utilizing the fact that any change in transmit power affects all receivers simultaneously across all groups.

4.5 Experiments

In this section, we describe the experimental setup. Section 4.5.1 describes the tools we use to measure the wireless channel, Section 4.5.2 describes the transmit power changes we apply, and Section 4.5.3 describes two real-world experimental deployments.

4.5.1 Tool Description

We use the following tools to measure the wireless channel and detect line crossings.

4.5.1.1 WiFi

We use laptops with Intel 5300 NICs that have three-antenna MIMO 802.11n radios. We use the CSI Tool [20], that has been built for these radios, to get channel state information from the WiFi transmitter. The CSI tool extracts 802.11n channel state information for 30 subcarriers at each antenna pair. Since we use three antennas at each node for communication, for each transmitter-receiver pair, we have $3 \times 3 = 9$ links each with 30 subcarrier groups. We use two kinds of antenna separations — in the normal case (WiFi_NORM), we place the antennas 6 cm apart, in the other case (WiFi_SEP), we use a larger antenna separation of 30 cm. The increased separation is accomplished by connecting the antennas to the Intel 5300 NIC with standard RF cables that are long enough to provide up to 30 cm separation. We program the transmitter to transmit packets at a rate of 10 Hz which is similar to beacon frame rates of a standard wireless access point. The attack receivers use the CSI Tool to obtain channel state information from the received packets which in turn is used to detect line crossings as described in Section 4.4.1.

4.5.1.2 ZigBee

For the ZigBee experiments, we use Texas Instrument CC2531 USB dongles, which are equipped with low-power, IEEE 802.15.4-compliant radios operating in the 2.4 GHz ISM band. The transmission frequency in this case is 12 Hz. A laptop is used to process the measured data at the attack receivers. There is no tool to obtain the CSI information in the case of ZigBee nodes. Therefore, we use the RSS value (in dBm) measured by the receiver hardware for our analysis, as described in Section 4.4.4.

4.5.2 Transmit Power Variations

We consider three different settings of transmit power variations for our experiments: while simulating effects of transmit power change we rely on the fact that any change in the transmit power at a time instant is observed across all subcarriers for all transmitter-receiver antenna pairs in case of WiFi and across all receivers in case of ZigBee at the same instant and we change the received signal parameters accordingly. We also add a zero mean Gaussian random variable (with standard deviation 0.67) to each $H_{j,k}(n)$ measurement, in addition to the the transmit power change $t_x(n)$, to account for errors due to environmental noise.

4.5.3 Experimental Deployments

We evaluate our methodologies in two different real world settings.

4.5.3.1 University Hallway

We choose a hallway inside a university building as the area being monitored (Figure 4.4(a)). The hallway is adjacent to a 30 cm thick and 3.5 m tall rebar-reinforced concrete boundary wall (Figure 4.4(b)). We note that this type of a wall causes significant RF attenuation at WiFi frequencies and represents a worst-case scenario among typical exterior walls for our purposes [76]. We place the attack receivers outside the boundary wall parallel to the hallway approximately 1 m away from the wall.

For the WiFi experiment, we deploy one transmitter inside the building across the hallway, and two attack receivers separated by 3 m outside the concrete wall (Figure 4.3(a)). Similarly, for the ZigBee network, we deploy one transmitter across

the hallway and six receivers outside the boundary wall. The attack receivers are placed in two groups of three nodes each, with the distance between the groups being 3 m (Figure 4.3(a)). Nodes in the same group are almost 30 cm apart. We perform both TX_NORMAL and TX_RAND experiments with the same ZigBee setup. We also experiment with three different transmitter locations in the case of ZigBee.

During the experiment, a person walks back and forth along a predefined path (route in Figure 4.3(a)) along the corridor between the transmitter and the attack receivers. With the help of a metronome, the person walks at a constant speed of $0.5\frac{m}{s}$. We collect over $12,000$ data samples for WiFi and over $20,000$ data samples for ZigBee in this experiment. In our evaluation, we use $w_s = 4$ s (short time window), $w_l = 40$ s(long term window), and $\Delta = 4$ s (Section 4.4.1).

4.5.3.2 Residential House

In this experiment, we monitor two sides of a residential house (Figure 4.3(b)) to detect people movement. We perform two sets of experiment with the WiFi nodes. In the first experiment (House 1), we place the WiFi transmitter in a corridor centrally located inside the house and two WiFi receivers with normal antenna separation (WiFi_NORM) in the backyard of the house outside the external wall as shown in the Figure 4.3(b). The receivers are placed approximately 1 m away from each other. For the second experiment (House 2), we use two WiFi receivers with larger antenna separation (WiFi_SEP) and place one of them in the backyard and the other outside the front entrance. The transmitter is placed in the same position as in experiment House 1.

For the ZigBee network, we place two groups of receivers, each group with three nodes, on either side of the house outside the external walls. As shown in Figure 4.3(b), the ZigBee groups 1 and 2 are placed outside the front entrance, and groups 3 and 4 are placed in the backyard, approximately 1 m away from the walls. Nodes in the same group are almost 30 cm apart while the intergroup distance on either side is at least 1 m. The ZigBee transmitter is placed inside the house colocated with the WiFi transmitter. We perform two sets of experiments with the same network settings — in one experiment the ZigBee transmitter transmits with fixed transmit power of +4.5 dBm (TX_NORMAL), in the other experiment the transmitter is programmed

to vary its transmit power randomly with each transmission (TX_RANDOM).

During these residential experiments, a person walks inside the house at normal speed back and forth first near the front entrance of the house (route 1 in the Figure 4.3(b)), and then in the living room which is near the rear end of the house (shown as route 2 in the Figure 4.3(b)). Finally, the person makes a few rounds inside the house as shown in route 3 in the Figure 4.3(b). We collect over $10,000$ data samples for each set of ZigBee and WiFi experiments. We video record the line crossings to test the accuracy of our detection method against ground truth. For the residential experiments, we use $w_s = 2$ s (short-term window), $w_l = 20$ s (long-term window) and $\Delta = 4$ s (Section 4.4.1). We use smaller window sizes for detection of line crossings as the person walks at a faster speed as compared to the University Hallway experiments.

4.6 Results

We evaluate the performance of the ERW attack in terms of false alarm and missed detection rates. False alarm (FA) rates are calculated as the number of line crossings wrongly detected by the system over the number of sample points. Missed detection (MD) rates are calculated as the number of actual line crossings not detected by the system over the total number of actual line crossings.

4.6.1 Detection of Line Crossing

In this section, we present the accuracy of detection of line crossings using the methodology as described in Section 4.4.1.

4.6.1.1 University Hallway

Table 4.1 lists the results obtained in the University Hallway experiment using our majority vote detection. We achieve almost 100% detection rate with few false alarms and missed detections. Using a WiFi 802.11n receiver with normal antenna separation, we get zero false alarms and only 1.92% missed detections. We compare the detected crossing times with those in the recorded video footage of the experiment and find that we can detect the crossing times with an average error of 0.79 s, with minimum and maximum errors of 0.03 s and 2.73 s, respectively.

We obtain zero false alarms and missed detections when using a 802.11n WiFi receiver with a large spatial separation between antennas, the mean error in this case being 1.22 s. For ZigBee, using a group of three closely located receivers, we get a 2.66% false alarm rate and a 1.67% missed detection rate in line crossing detection with an average error of 1.22 s. We use two groups of receivers and experiment with three different transmitter locations in case of ZigBee. We obtain the above results by averaging over all transmitter location and receiver group pairs.

Note that while computing the errors as compared to the ground truth, we consider the line connecting the centroid of transmitter antenna locations (or the transmitter location in case of ZigBee) and the centroid of the receiver antenna locations (or the centroid of the receiver locations in the group in case of ZigBee) as the representative link line.

4.6.1.2 Residential House

We present the detection accuracy of the Residential House experiment in Table 4.2. We achieve greater than 94% detection accuracy with a 0.043% false alarm rate while using WiFi receivers with normal antenna separation (WiFi_NORM). With larger antenna separation (WiFi_SEP) the accuracy is above 95% with a 0.005% false alarm rate. The mean error in detection of line crossings is 1.06 s in case of WiFi_NORM, the same being 0.56 s for WiFi_SEP.

For ZigBee, we achieve above 99% accuracy in detection with a false alarm rate of 0.004% only. The mean time-of-crossing estimation in this case is 1.63 s. Note that during this experiment, we placed one group of ZigBee nodes (group 2) directly in front of the metal-plated entrance door. The packet reception rates for receivers in this group are much lower than the receivers in the other groups. Also, perhaps due to attenuation through the door, the RSS measurements made by this group are more noisy than those made by the other groups, leading to further degradation in performance. The missed detection rate for this group is almost 30%, about 60 times more than the average missed detection rate of other groups (results presented in Table 4.2 are averaged over the other three groups). Thus, we conclude that, although an ERW attack can penetrate concrete and brick walls, metallic structures in the line of sight path of the radio signals degrades the detection accuracy significantly.

4.6.2 Determining Direction of Motion

In the following sections, we present the accuracy we achieve in detecting the direction of motion of the person in each experiment.

4.6.2.1 University Hallway

In the University Hallway experiment, the corridor was crossed by a moving person an equal number of times in either direction. We achieve 100% accuracy in detecting direction of movement on either side of the corridor while using two WiFi receivers or two groups of ZigBee nodes using the method described in Section 4.4.2.

We also achieve an accuracy as high as 90.38% in detecting direction of motion with only a single WiFi 802.11n receiver by increasing the spatial separation of the MIMO antennas. The accuracy with a single WiFi receiver with standard antenna separation is 59.62%, which is slightly better than guessing the direction of motion.

4.6.2.2 Residential House

In the Residential House experiment, we achieve 100% accuracy in detection while using two WiFi receivers with standard antenna separation (experiment House 1) or two groups of ZigBee nodes on either side of the house. Individual detection accuracy of the two WiFi receivers (with standard antenna separation placed on the same side of the house as in experiment House 1) used are 100% (RX1) and 68% (RX2), respectively. Detection accuracy with spatially separated antennas for these receivers (when they are placed on opposite sides of the house as in experiment House 2) are 96% (RX1) and 52.6% (RX2), respectively. These results differ from the University Hallway experiment where we get better accuracy in detecting direction of movement while using large spatial separation between antennas as compared to using normal antenna separation. The degradation in accuracy with antenna separation in the Residential House experiment may be due to the fact that during the House 2 experiment, walking speed of the person was about 20% faster as compared to the House 1 experiment with normal antenna separation, hence crossing times for individual antennas overlapped with each other in some cases.

To summarize, our results indicate that an ERW adversary should use two WiFi receivers or two groups of ZigBee nodes at each side in order to detect direction of

motion accurately. It is possible to achieve very high accuracy even with a single WiFi receiver in some cases (e.g., RX1 in experiments House 1 and House 2), however, the results depend on the environment and need further investigation.

4.6.3 Advantages of Majority Vote

In this section, we show how our majority vote approach helps overcome inherent uncertainties in wireless links. All wireless links are not equally sensitive to motion occurring in their vicinity and the sensitivity varies with link fade level along with other factors. For example, Figure 4.5 shows the RSS for the three ZigBee receivers belonging to group 1 used in the Residential House experiment for a time interval during which the person crossed in front the group two times. For RX1 and RX3, the overall RSS variance is very small. When the person crosses the link line, she causes high short-term variation of the RSS, as can be seen during time intervals [113 s - 116 s] and [128 s - 131 s]. Thus, one can infer link crossing times monitoring for these high short-term variations in RSS for these links. However, the link to RX2 has very low mean RSS value with high variance overall. This link does not show clear short-term high variance region corresponding to actual link line crossings as compared to RX1 or RX3. Hence, a line crossing detection method that relies only on the link to RX2 will perform poorly.

Since it is not possible for an adversary to know beforehand whether a link is good or bad for detecting LOS crossings, he relies on correlation among multiple closely located links and infers a line crossing only when majority of these closely located links indicates a crossing. In our experiments, $3 \times 3 = 9$ links between the MIMO transmitter-receiver antenna pairs are considered for majority vote in the WiFi case, and groups of three single-antenna receivers in the ZigBee case. Figure 4.6 shows one scenario where our majority vote algorithm helps get rid of some false alarms and missed detections due to one bad WiFi link (for clarity we show three out of the nine links) from the University Hallway experiment. The link shown in Figure 4.6(b) fails to detect a line crossing that occurs around 100 s. However, the other two links (Figure 4.6(a) and Figure 4.6(c)) detect the crossing and a majority vote among these three links detects the crossing at that time (Figure 4.6(d)). Similarly, we see that

the link in Figure 4.6(b) flags a false alarm at 180 s but the other two links do not indicate any crossing. Again, the majority vote gets rid of the false alarm at time 180 s (Figure 4.6(d)), thereby improving the overall accuracy of the system.

We summarize our findings as follows — a single wireless link suffices in some cases in detection of line crossings between a transmitter and a receiver, however, the results are not always reliable due to inherent uncertainties in link sensitivity to object movements. We can improve accuracy and reliability by correlating detections across multiple colocated links using a majority vote approach. Our results confirm that we can get rid of most of the false alarms and missed detections caused by a bad link by applying the majority vote based detection method.

4.6.4 Compensation for Transmit Power Change

In this section, we show how transmit power changes (random or strategic) affect line crossing detection accuracy and how our compensation method nullifies the effect of such power changes.

Figure 4.7(a) shows the effect of random transmit power changes on line crossing detection for a WiFi link between a single transmitter-receiver antenna pair that is crossed three times by a moving person. The top figure corresponds to the case when there is no transmit power change. This figure clearly shows distinct short time periods of high variance in the CSI corresponding to the times when the person crosses the link. However, transmit power change masks these distinct short-term variance regions and renders line crossing detection ineffective as can be seen in the figure in the middle. The bottom figure plots the CSI for the same link after compensating for the transmit power changes as described in Section 4.4.3. Clearly, our compensation method almost nullifies the masking effect of transmit power changes and the attacker can detect three line crossings (high short-term variance region) from the compensated signal.

Similarly, Figure 4.7(b) shows how strategic power changes can be used to simulate link line crossings, and how our compensation method eliminates these artificial variations. The top figure plots the RSS in dBm for a ZigBee link that is crossed during the time interval 856-860 s. The figure in the middle shows one additional line crossing (high variance region) introduced in the link by strategic transmit power

changes during time interval 838-841 s. However, as seen from the bottom figure, our compensation method gets rid of the false alarm introduced by strategic power change and we can detect the original line crossing from the compensated signal.

In Figures 4.8 and 4.9, we show false alarms and missed detections induced by transmit power changes and the accuracy of our compensation method. In the figures, NORMAL corresponds to the case when the transmitter transmits with fixed transmit power, CRS is when strategic power changes are introduced in the data using TX_LINECROSS simulation, CRS_CMP corresponds to the results when we apply our compensation method on TX_CRS. Similarly, RND shows results when the transmitter is changing its transmit power randomly with each transmission, while RND_CMP is the corresponding compensation results. Note that the owner of the legitimate transmitter has full control over the transmitter node and can randomly select the periodicity with which to introduce transmit power changes in case of the TX_LINECROSS experiment. We present results for one such simulated scenario where the owner randomly selects a time period between $3 - 10$ s to change transmit power according to a profile that mimics typical channel variation introduced by a person crossing the link line.

We see that transmit power changes (for both TX_LINECROSS and TX_RANDOM experiments) introduce significant false alarms and missed detections while using either WiFi (with or without spatially separated antennas) or ZigBee nodes. As an example, in the University Hallway experiment, a strategic transmit power change at the WiFi transmitter increases the missed detections rate from 1.92% to 32.69% and the false alarms rate from 0% to 0.199% when using a WiFi receiver with normal antenna separation. However, our compensation method gets rid of all the additional false alarms and missed detections. Similarly, in the Residential House experiment, for random power changes at the ZigBee transmitter, the missed detections rate increases to 31.37% from 0.94% and the false alarms rate increases to 0.429% from 0.003% but our compensation method brings down the missed detection and false alarm rates to only 0.94% and 0.006%, respectively. Using Equation 4.12, we can estimate the transmit power change amplitude accurately in 98% cases if we allow an error margin of ± 2 dB.

To summarize our findings, transmit power changes (strategic or random) increase the false alarm and missed detection rates significantly. However, using our compensation method, an attacker can accurately estimate the transmit power change amplitude and compensate for the same to get rid of most the adverse effect caused by such changes and still sense people location and motion with high accuracy.

4.6.5 Detection with Varying Transmission Rate

ZigBee applications in modern facilities use different transmission rates for communication. In this section, we show how detection accuracy varies when the transmission rate for the ZigBee transmitter is lowered. We use the data from TX_NORMAL for both the University Hallway and Residential House experiment to simulate the effect of lower transmission rate. Note that the original transmission rate is approximately 12 Hz. We simulate three additional transmission rates — 6 Hz, 4 Hz, and 2 Hz, respectively, from the original data. Figure 4.10 shows the results of our simulation. We find that the overall detection rates decrease with lower transmission rates. For the transmission rate of six transmissions/second, accuracy of the detector is over 98% for the University Hallway experiment and over 96% for the Residential House experiment. These results are similar to what we observe for original transmission frequency of 12 Hz. The accuracy is worst for transmission frequency of 2 Hz with the detection rate being as low as 71% for the Residential House experiment. For the transmission rate of 4 Hz, the detection rate degrades to 87% in the University Hallway experiment, although it remains above 96% for the Residential House experiment. We do not see any noticeable change in the false alarm rates with varying transmission rate.

We summarize our findings as follows: detection accuracy with ZigBee nodes decreases as transmission rate is lowered. For an ERW attack to succeed with high accuracy, the transmission frequency must be at least 6 transmissions/second.

4.7 Related Work

Preserving the privacy of the location of mobile devices in wireless networks has been object of intense research [77, 78]. Location represents an important private information that can be used by malicious attackers for serious privacy violations and

potentially dangerous attacks. The work in [79] presents an evaluation of the privacy and security of wireless tire pressure monitoring systems. It shows that eavesdropping on these systems is possible through their static identifiers even at a distance of 40 m.

Other works have demonstrated that communicating wireless devices leak the current and past location of people carrying these devices. In [80], the authors show that distance bounding protocols [81] can leak distance and location information to an attacker overhearing the communication between the prover and the verifier to such an extent as to allow the attacker to estimate his own position relative to the two devices. They also introduce a location private distance bounding protocol that protects against malicious provers, passive eavesdroppers, and attackers trying to actively initiate a distance bounding session. In [82], the authors describe a system that can reveal the locations of WiFi-enabled mobile devices within the coverage area of a single high-gain antenna. By knowing the location and/or the maximum transmission range of the APs, an eavesdropper can set up a high-gain antenna to sniff the traffic between the *victim* mobile device and the APs on all the available wireless channels and estimate the position of the mobile device. The work in [83] proposes three countermeasures to improve the location privacy in wireless networks, i.e., anonymize the identity of the device by frequently changing its pseudonym during communications (as in [84]), unlink different pseudonyms of the same device with silent periods between different pseudonyms, reduce the transmission range of the devices through power control to minimize the number of APs that can collaborate to localize the devices' location (the precision to which a mobile device can be located depends on how many APs can hear from the device [85]).

The works focusing on location privacy typically assume that the *victims* of the attack are carrying a wireless device (e.g., a mobile phone, radio frequency identification (RFID) tag, low-power radio transceiver) that is actively communicating with the surrounding network infrastructure (e.g., WiFi APs, RFID readers, other radio transceivers). The work in [86] presents a through-walls passive WiFi radar system. In it, a receiver is situated outside the target building and a Wi-Fi AP placed inside the building and having a narrow-beamwidth directional antenna is used as

transmitter. The signal received by the passive radar detector is then used to create a range-Doppler surface and detect a moving target. Our work is complementary to [86] because through wall radar systems have limited range due to direct signal interference. Further, as they are based on transmission, they could be detected by source localization or counteracted by jamming. Other systems localize people by measuring the change in RSS of links traveling across an area where several WiFi APs or ZigBee radio transceivers are deployed. In the case of Wi-Fi based passive localization systems [72], a radio map of the environment is created by having a person standing at different locations while recording the RSS of all the links. This requires access to the target area for an initial calibration of the system. For radio tomographic systems [70], accurate localization of people requires a high density deployment of radio transceivers on all sides of the target area. In [87], the authors developed a method for through wall localization using WiFi signals. However, their method depends on active probing, i.e., a custom hardware sending WiFi signals through a barrier (e.g., a wall) and measuring the way it reflects back from objects on the other side. This active transmission of radio waves makes their work susceptible to detection and jamming. Our work relies on passive measurements using standard hardware and, hence, is immune to detection and jamming.

In this work, we demonstrate that the presence, location and movements of people not carrying any wireless device can still be eavesdropped by measuring the RSS of the links between the devices composing the legitimate network and few receivers positioned outside the target area. This can be achieved without requiring a complex network infrastructure or previous access to the target area for an initial calibration. In [88], the authors propose a method to detect an attack to a radio tomographic system in which some of the deployed radio transceivers are maliciously reprogrammed to change their transmit power. Our work is different in that we propose a method capable of correctly estimating the amplitude of the transmit power changes implemented by the legitimate devices as a countermeasure to the ERW attack. This enables reconstructing the true dynamics of the RSS signals and estimating people's locations. Moreover, in our work we do not make any assumption on the number of transmitters changing their transmit power and on the periodicity

and amplitude of such changes.

4.8 Conclusion and Future Work

We investigated the ability of an attacker to surreptitiously use an otherwise secure wireless network to detect moving people through walls. We designed and implemented an attack methodology for through wall people localization that relies on reliably detecting when people cross the link lines by using physical layer measurements between the legitimate transmitters and the attack receivers. We also developed a method to determine the direction of movement of a person from the sequence of link lines crossed during a short time interval. Additionally, we described how an attacker may estimate any artificial changes in transmit power (used as a countermeasure), compensate for these power changes using measurements from sufficient number of links, and still detect line crossings. We implemented our methodology on WiFi and ZigBee nodes and experimentally evaluated the ERW attack by monitoring people movements through walls in two real-world settings. We found that our methods achieve close to 100% accuracy in detecting line crossings and the direction of movement, when we use two WiFi 802.11n nodes with normal antenna separation, or two groups of ZigBee nodes as attack receivers. We also found that our methods achieve $90 - 100\%$ accuracy when we use a single 802.11n attack receiver.

Future work must develop more sophisticated protocols to prevent person location information leakage. Device hardware enhancements may be necessary for this purpose.

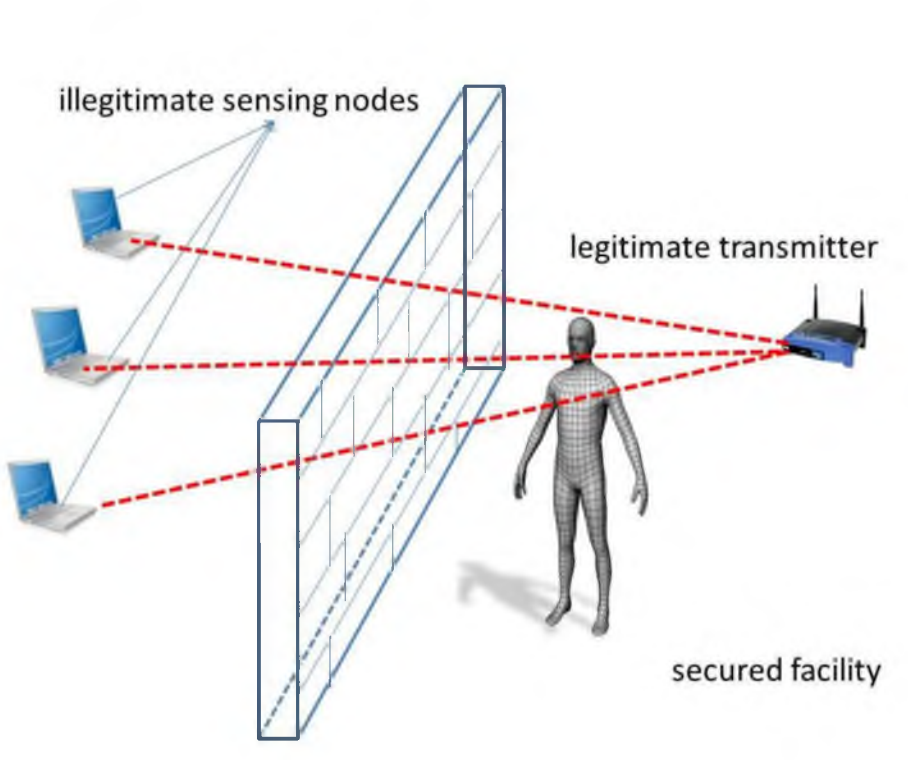

Figure 4.1. Exploiting radio windows (ERW) attack example.

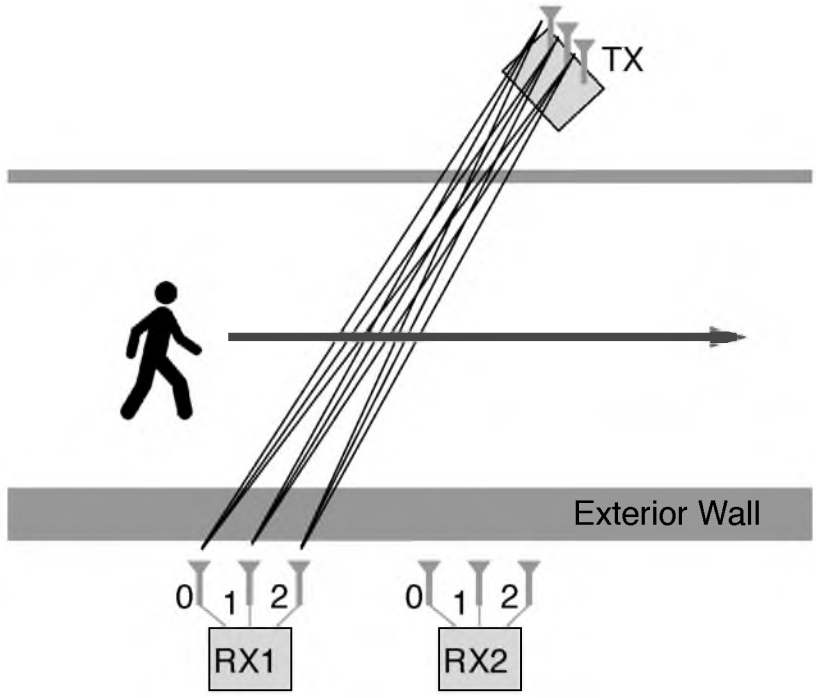

(a)

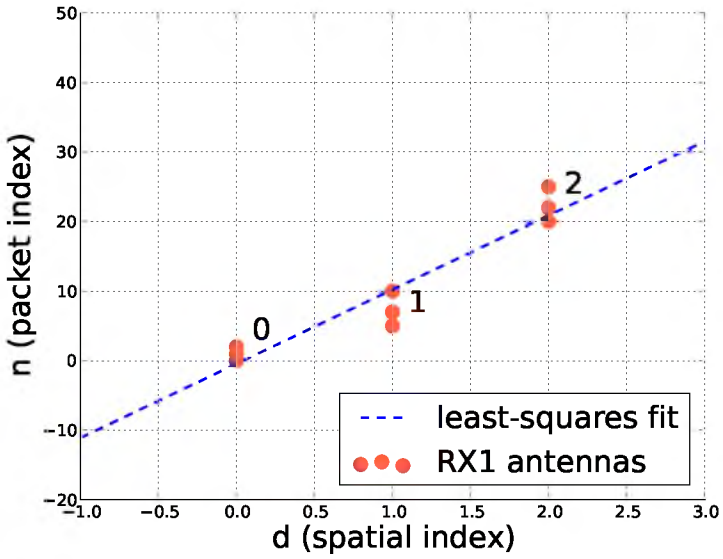

(b)

Figure 4.2. (a) Line crossing detection diagram. The attack receiver(s) measure channel state information from the legitimate transmitter. The MIMO antenna array at the receiver(s) allows the adversary to count line crossings and determine direction of motion. (b) Direction of motion is determined by fitting a line to the points created by the spatial indexes of the antennas which detect a line crossing and the corresponding packet indexes of the detections. The sign of the slope of the fitted line indicates the direction of motion.

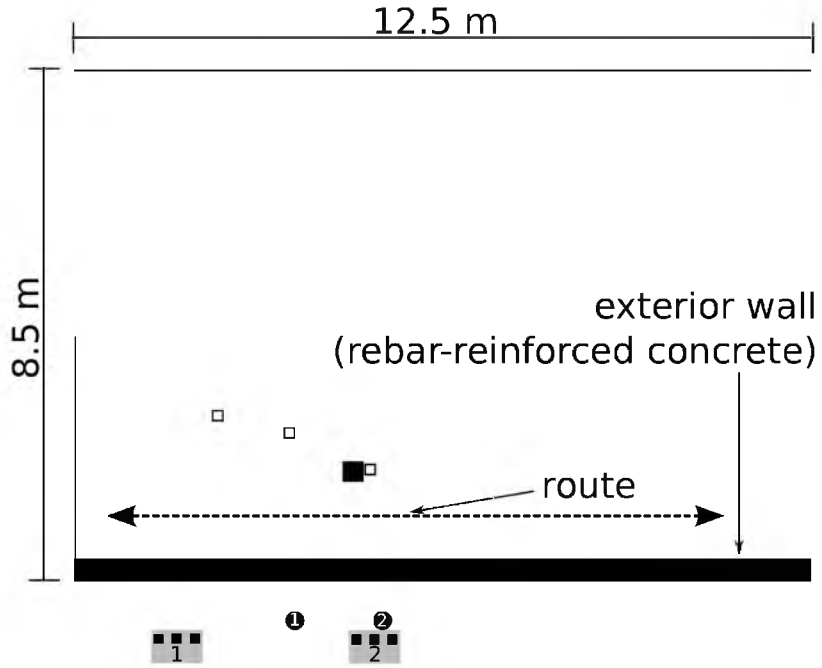

(a) University hallway.

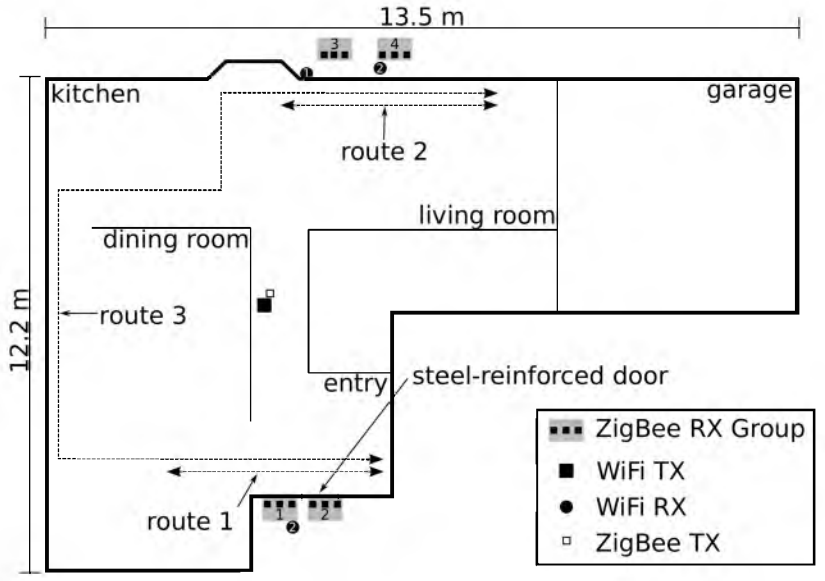

(b) Residential.

Figure 4.3. Experiment diagrams. We show maps of the (a) University Hallway and (b) the Residential House and mark the location of the legitimate transmitter(s) and the attack receivers. We also highlight the route(s) followed by the walking person.

Table 4.1. Detection Accuracy (Hallway experiment).

Hallway	Accuracy		Error (sec)		
Experiment:	FA%	MD%	Min	Max	Mean
WiFi_NORM	0	1.92	0.03	2.73	0.79
WiFi_SEP	0	0	0.27	2.37	1.22
ZigBee	0	1.02	0.27	2.37	1.22

Table 4.2. Detection Accuracy (House experiment).

House	Accuracy		Error (sec)		
Experiment:	FA%	MD%	Min	Max	Mean
WiFi_NORM	0.043	5.70	0.29	2.78	1.06
WiFi_SEP	0.005	4.35	0.03	1.82	0.56
ZigBee	0.004	0.49	0.10	3.55	1.63

(a) Transmitter deployment

(b) Attack receiver deployment

Figure 4.4. Experimental setup of ZigBee radios at the University Hallway experiment. The images (a) and (b) show the interior and exterior radios, respectively.

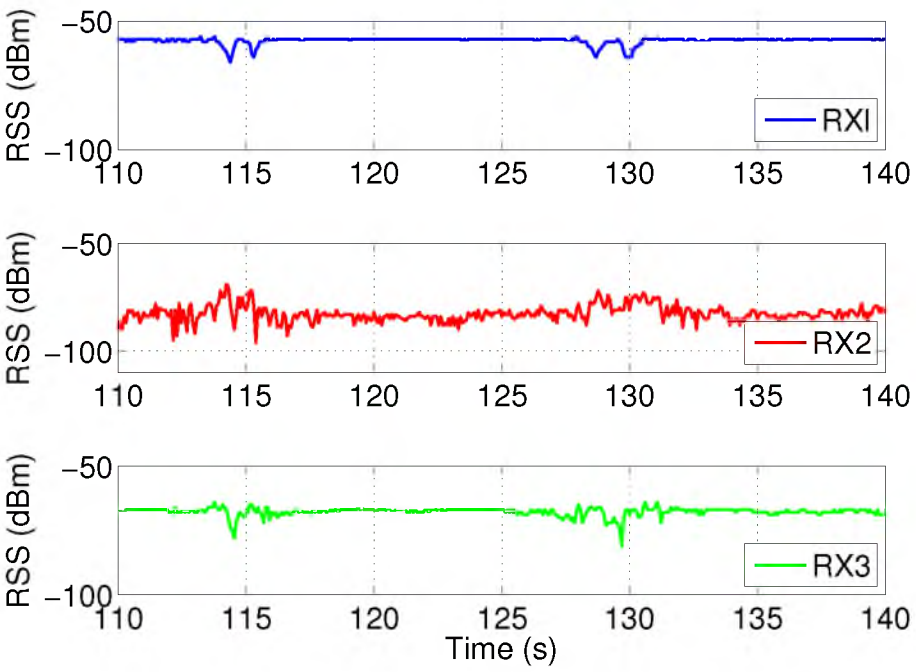

Figure 4.5. All links are not equally sensitive to movement — RX1 and RX3 measure high short-term variations in link RSS corresponding to person crossings (time intervals [113 s - 116 s] and [128 s - 131 s]). Such a distinct high variation region is not present in link to RX2.

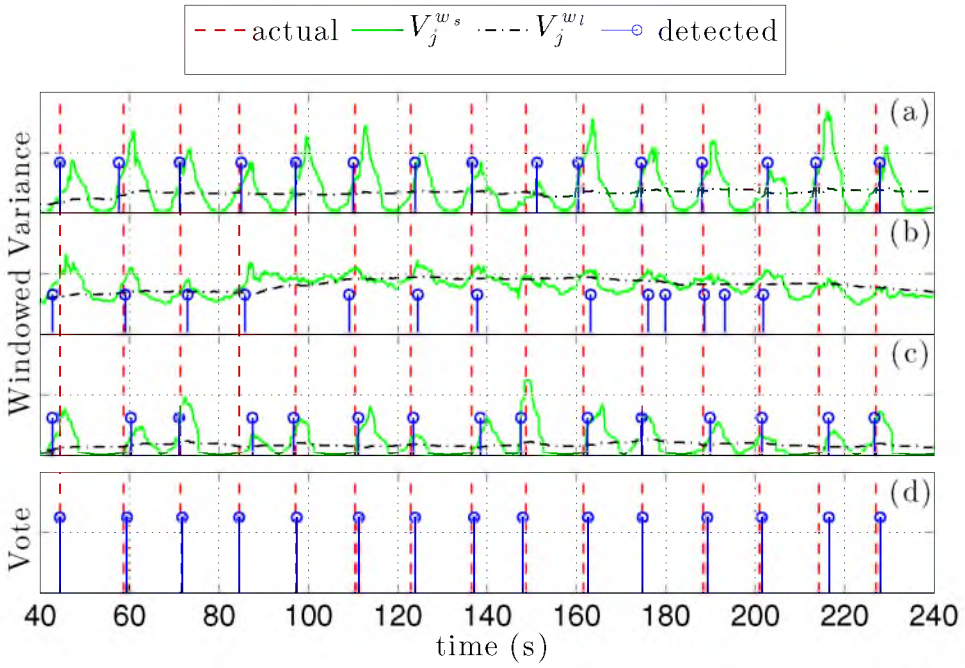

Figure 4.6. The majority vote over transmitter-receiver antenna pairs reduces false alarms and missed detections. (a),(b), and (c) show the results of the windowed variance based line crossing detection for a different antenna pair using WiFi. In (d), we see that the majority vote eliminates false alarms and missed detections.

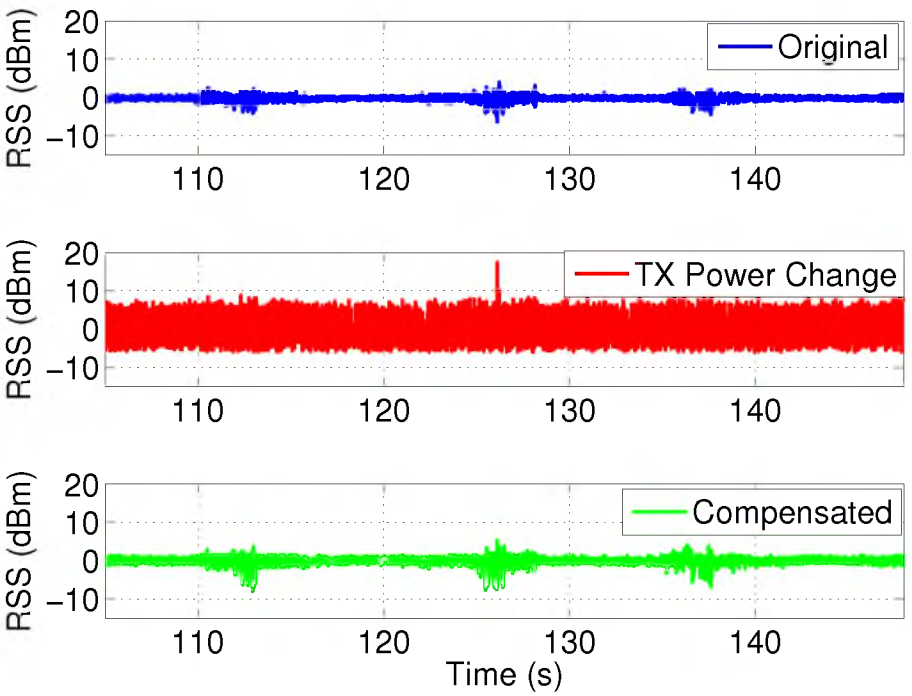

(a) TX_RANDOM (WiFi)

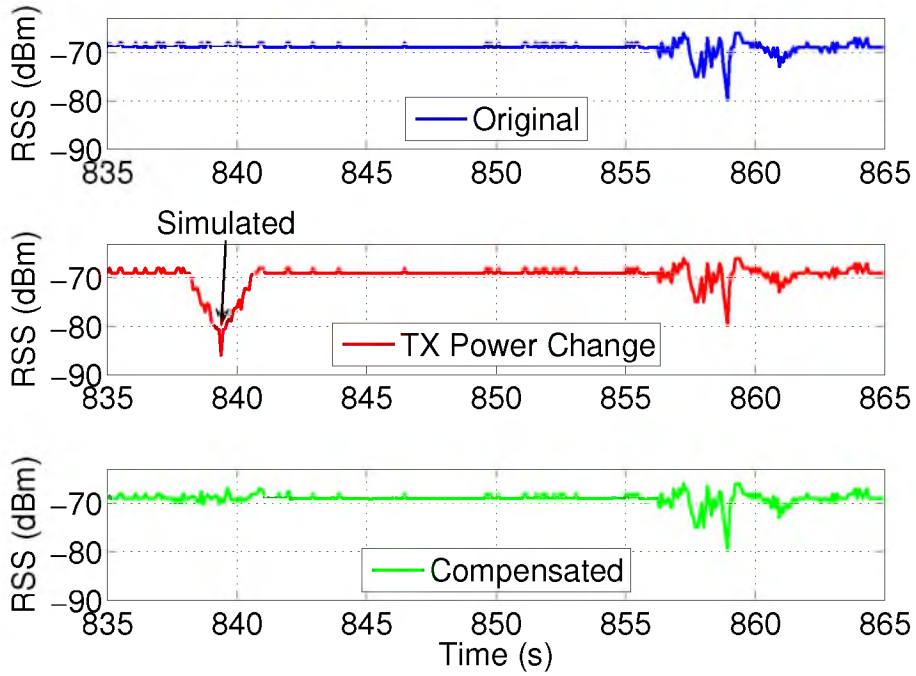

(b) TX_LINECROSS (ZigBee)

Figure 4.7. Measured CSI and RSS (top) without and (middle) with TX power change; and (bottom) after compensation, which nullifies the effect of TX power change. The changes are random in (a) and meant to spoof a line crossing in (b).

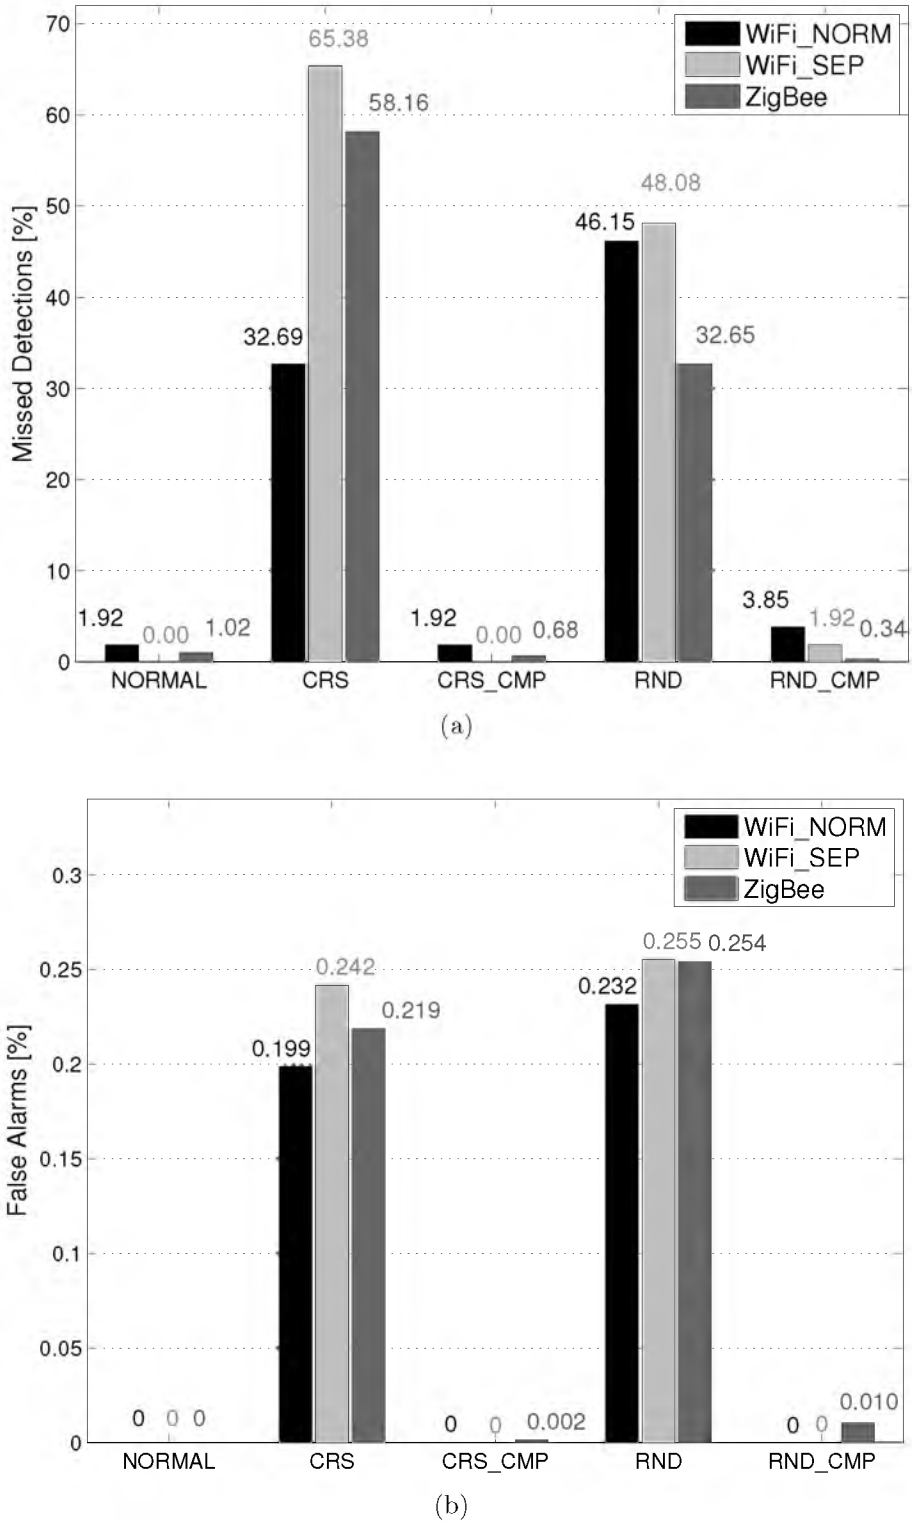

Figure 4.8. Compensation accuracy in the University Hallway experiment. Both strategic (CRS) and random (RND) transmit power variations increase (a) missed detections and (b) false alarms rate significantly. However, our compensation method eliminates most of these artificially induced missed detections and false alarms (see CRS_CMP and RND_CMP).

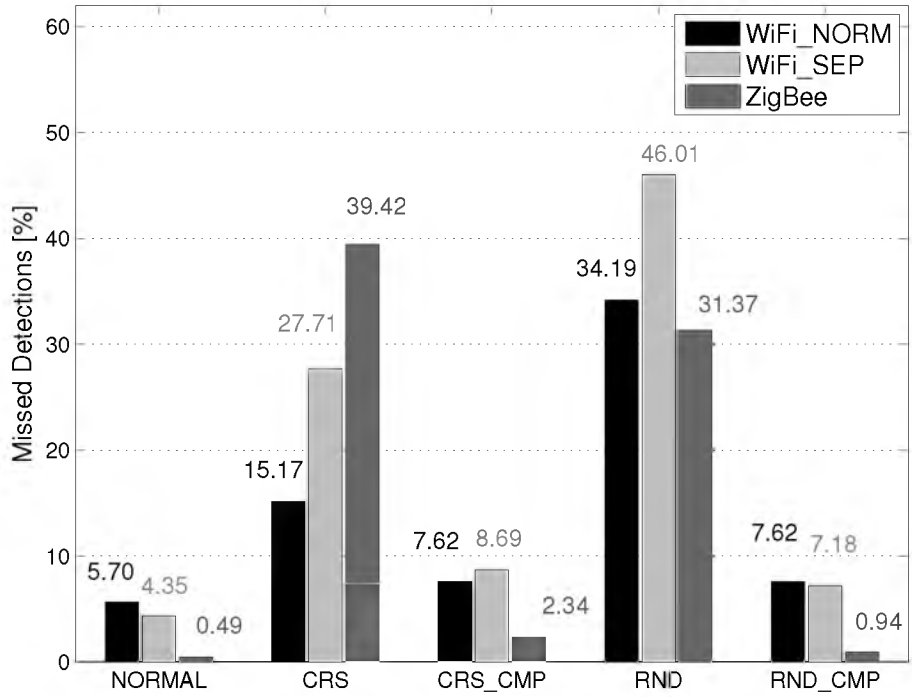

(a)

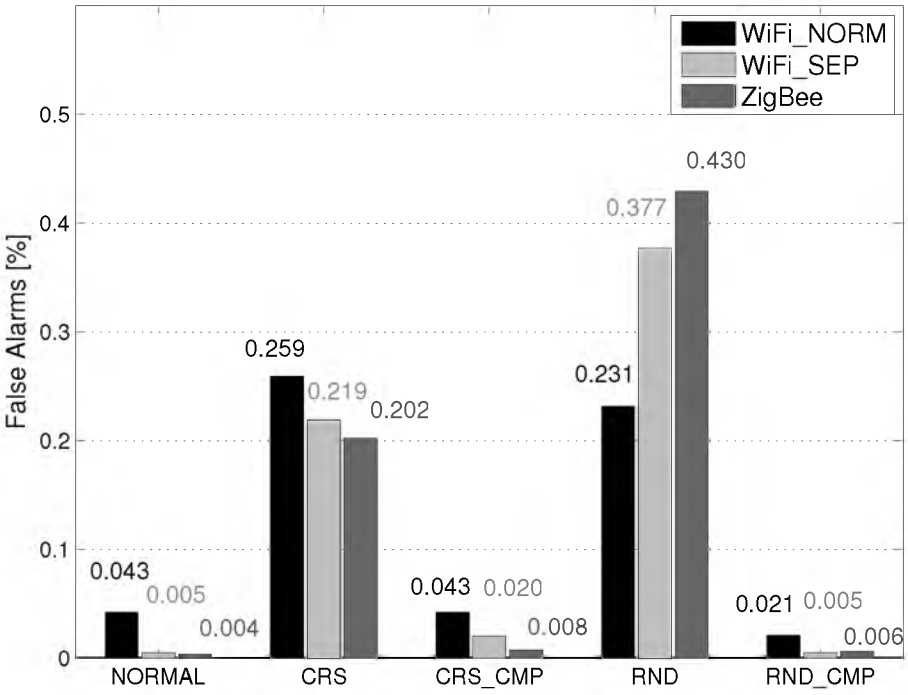

(b)

Figure 4.9. Compensation accuracy in the Residential House experiment. Both strategic (CRS) and random (RND) transmit power variations increase (a) missed detections and (b) false alarms rate significantly. However, our compensation method eliminates most of these artificially induced missed detections and false alarms (see CRS_CMP and RND_CMP).

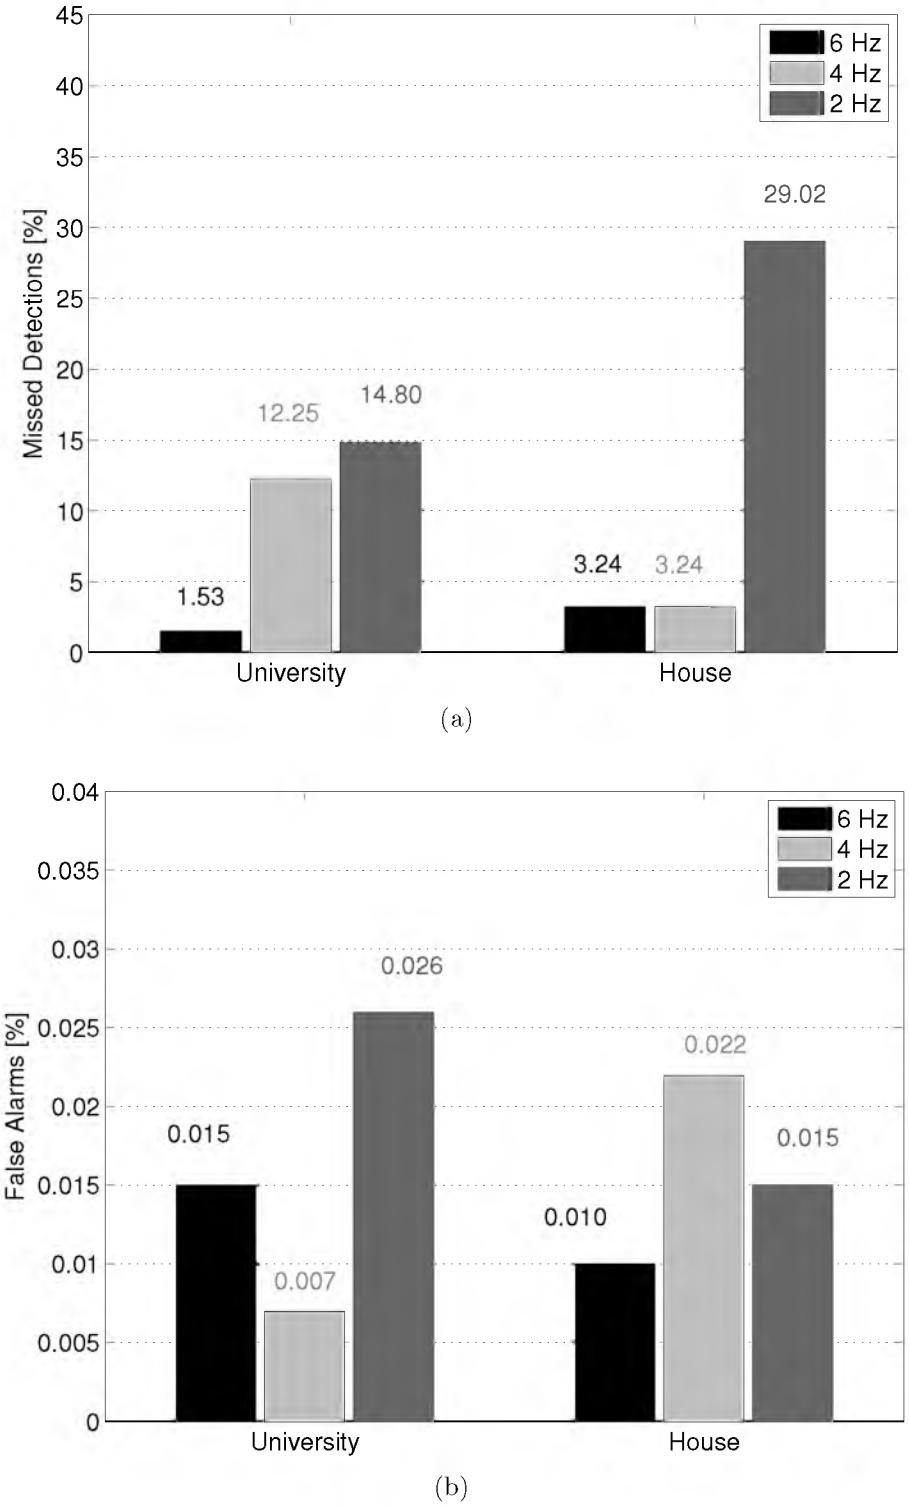

Figure 4.10. System performance in terms of (a) missed detection and (b) false alarm rates with varying ZigBee transmission rates during both experiments.

CHAPTER 5

TOWARD A RAPIDLY DEPLOYABLE RTI SYSTEM FOR TACTICAL OPERATIONS [1]

5.1 Abstract

The ability for special operations forces (SOF) to rapidly deploy a through-wall tracking system upon arrival at a tactical operation, e.g., a hostage scenario, and thereby estimate the approximate locations of the people within the building, has the potential to lower the risk of the operation and save lives. We study the feasibility of a rapidly deployed radio-tomographic imaging (RTI) system for use in tactical operations by SWAT and other SOF, in which several low-power radio devices are placed around a building and used to image and track the motion of humans inside the building. Specifically, we identify and study the constraints of this application, such as the need for the sensor network to self-localize and self-calibrate with minimal input from the SOF. We implement and test, in a wide variety of experimental deployments, a real-time RTI tracking system which adheres to these constraints and provides valuable situational intelligence. We work in concert with local law enforcement and SWAT in order to obtain valuable feedback from end users. We show that our system is capable of providing useful tracking information (average errors of less than 2 m) even when the self-localization results are inaccurate (up to 3 m average error).

5.2 Introduction

This report describes progress in determining the feasibility of a new radio frequency (RF)-based technology for through-building surveillance, specifically, deter-

[1] ©[2013]. Reprinted, with permission, from D. Maas, J. Wilson, N. Patwari, "Toward a Rapidly Deployable RTI System for Tactical Operations," in *Proc. 8th IEEE International Workshop on Practical Issues in Building Sensor Network Applications (SenseApp)*, 2013.

mining the positions of people inside a building using sensors placed only on the outside of the building. The enabling technology, RF tomographic imaging (RTI), uses a network of small, inexpensive wireless devices, placed around an area, to make measurements and estimate where people and objects are currently located in the area [9,65,67,69,89–93]. By using radio waves, the devices are able to image through walls, smoke, and other obstructions [10,94], a major advantage over light and infrared. We introduce the fundamentals of RTI in Section 5.3.1. The "see through walls" capability of RTI opens the door for many emergency response applications in which situational awareness is critical to save lives.

In this work we investigate the application of these technologies to a system for use in emergency response, specifically, for SWAT and military special operations forces (SOF). Consider the scenario of a SWAT team responding to a hostage situation. Upon arrival at the scene, golf-ball-sized RTI radios are placed, thrown, or tactically launched (from an M-32 or M203 launcher) around the building. Depending on the scenario, these might land on the ground, or be deployed so that they stick to the outside wall of the building. Once deployed, the radios communicate and form a mesh network. After the radios self-locate and form an accurate map of their own locations, they continuously measure received signal strength (RSS) on all of the pair-wise links in the network. The measurements are collected and processed in real-time to show the tracks and current locations of moving people and objects in the environment, as shown in Figure 5.1. These data from our system represent significant situational intelligence which may help save lives during the course of the SWAT operation. For example, SWAT commanders could decide which part of the building is furthest from people and thus may be their safest point of entry.

This paper details feasibility studies for a robust, rapidly deployable, commercial RTI system. In contrast to experimental research tests in which sensors are hand-placed, mapped, and manually calibrated, in a tactically-deployed system, sensors must self-localize, self-calibrate, and the network must automatically form and start to measuring RSS.

The sensors must self-localize because many tactical operations are time-critical, and SOF cannot take the time to map the locations of the nodes. Additionally,

precisely measuring the node locations may put SOF personnel at risk.

The network must self-calibrate regardless of the number of people already present within the building being monitored. Previous RTI methods [10, 94] have required empty-building calibration measurements in order to generate accurate tracking results, but empty-building calibration measurements may not be possible in many tactical operations. Additionally, since our system can measure RSS on multiple wireless channels, a part of the self-calibration process involves deciding which channels represent the best source for tracking measurements in real-time.

We show that these capabilities are feasible, that robust localization performance can be achieved, and that a complete system with these capabilities would be very compelling for end users. In summary, a through-building surveillance system with the capabilities we demonstrate would be useful to SWAT and other SOF, and would help save lives.

Specifically, our paper describes the following achievements towards a complete RTI tracking system that could be operated and used by SOF:

- We implement in real-time *kernel distance-based radio tomographic imaging* (KRTI), an RTI method that has improved performance compared to previously reported attenuation-based RTI [9] and variance-based RTI (VRTI) [10,94]. We describe KRTI and its performance in Sections 5.3.1 and 5.4.1.

- Since sensors must self-localize using a combination of GPS, received signal strength (RSS) measurements, and minimal user input, before they can estimate the positions of people in the environment, we study the effects of poor self-localization on tracking performance. Surprisingly, we find that the performance of KRTI degrades gracefully as the sensors' self-localization errors increase. This result is discussed further in Section 5.5.3.

- We implement and test a particular sensor self-localization method, called distributed weighted multidimensional scaling (dwMDS) [95,96] that combines GPS, RSS, and building layout information for sensor self-localization. We describe dwMDS in Section 5.3.2 and show in Section 5.5.1.1 that our experiments yield an average sensor self-localization accuracy of about 1.5 m.

- We implement two types of sensor self-calibration. First, the KRTI system must know the histogram of RSS values on each link. We show this can be calculated in real-time from RSS data, without requiring any "empty-building" calibration. Empty-building calibration is impossible in emergency response applications, but has been used in most previous research [9,89,97]. Second, we choose upon deployment the best frequency channel for each link according to its fade-level. These two self-calibration methods are discussed in Section 5.3.3.

- We examine the use of directional antennas for through-building KRTI. We find that equipping sensors with directional antennas, compared to omnidirectional antennas, reduces average tracking error further, by as much as 22%. This result is described in Section 5.5.2.

- Finally, we study the effects of using different sized networks for KRTI and find that the number of sensors can be dramatically reduced compared to the 30 or more used in previous research [9, 94]. With only ten sensors, accurate localization (less than 1 m RMSE) can be achieved. This development is described in Section 5.5.3.

In summary, we show that a tactically deployed RTI system with a small number of sensors can perform sensor self-localization with minimal input from end users, can self-calibrate, and still provide high accuracy localization and tracking of people in a variety of experimental deployments. In addition, we collaborate extensively with local SWAT in order to get feedback on system deployment and usability. End user observations of a testbed deployment are described in Section 5.5.4.

5.3 Methodology

In this section we introduce the RTI method we implement to produce the images used to track human motion. Next, we discuss the method we use to allow the nodes to self-localize. Finally, we discuss network self-calibration.

5.3.1 Radio Tomographic Imaging Implementation

Several methods for RTI-based location tracking have been introduced over the past few years [8,75,91,98]. In [9], the authors measure the average RSS on each link

while the tracking area is empty and then determine where people are in the network based on changes in the RSS values for each link. In [10], the authors monitor the variance of the RSS on each link in order to localize motion in the network. This method has the benefit that it does not require offline calibration, but it cannot detect stationary targets.

Recently a new RTI method, kernel distance radio-tomographic imaging (KRTI) [92,99], was introduced which detects stationary and moving targets without the need for offline calibration. KRTI uses a kernel distance metric to quantify the difference between two histograms of RSS measurements for each link in order to track people within the network. Using histograms of RSS measurements combines the benefits of the methods presented in [9] and [10], quantifying changes in both the mean and the variance of RSS measurements for each link. An example of the images generated with KRTI and used for tracking is shown in Figure 5.2. The hot point in the image represents the position of the person being tracked.

In KRTI, a long-term histogram is used as a baseline, while a short-term histogram is used to track recent changes in RSS on each link. When applied to these two histograms for a given link, the kernel distance metric is an indicator of motion on or near the link. The results we present in this work rely on KRTI in order to perform tracking because it is well-suited to hostage and barricade situations, in that it does not require empty-building calibration measurements and is capable of running in real-time. We note that a background subtraction method like the one presented in [100] is also capable of determining these distributions without empty-building measurements, but includes more computational complexity.

5.3.2 Sensor Network Self-Localization

The proposed tracking system requires knowledge about the relative locations of the radio transceivers deployed around the building within which the human targets are to be tracked. More precise node localization leads to more accurate tracking, which would be valuable to end users like SWAT. Since a SWAT team may not have the time or be willing to put their personnel at higher risk in order to precisely measure out the node locations, the nodes should self-localize and begin to track people within the network with little or no help from the team deploying the system.

There are several methods in the literature for localizing radios. They typically use the time of arrival (TOA) or received signal strength (RSS) of radio transmissions in order to estimate internode distances [101]. An ordination technique like multidimensional scaling (MDS) [96] can then be applied to find a map of sensors that best fits the measured internode distances.

In this work, we implement and augment a type of MDS called distributed weighted multidimensional scaling (dwMDS) [96], an ordination method which, given a noisy set of interpoint distances, attempts to find the most likely arrangement of these points. In our case, these points correspond to the locations of the network nodes, and the interpoint distances are estimated using the maximum-likelihood estimator (MLE) for the large-scale path loss model in [47] and the RSS measurements made by the nodes. In order to mitigate the effects of fading error on the RSS measurements for each link, we use the average RSS over five channels for each link.

The dwMDS cost function is

$$S = 2\sum_{i \neq j} w_{ij}(\delta_{ij} - d_{ij}(X))^2 + \sum_i r_i \|\mathbf{x}_i - \bar{\mathbf{x}}_i\|^2, \tag{5.1}$$

where δ_{ij} is the estimated distance between nodes i and j, $d_{ij}(X)$ is the distance between nodes i and j for the node location matrix X, w_{ij} is a weighting factor which represents the quality of the distance estimate, x_i is the ith node location, $\bar{x}_i$ represents an *a priori* estimate of the ith node location, and r_i is a weight that represents the quality of the *a priori* estimate. The $\bar{x}_i$ could come, for example, from GPS receivers attached to the nodes or from coarse location estimates contributed by the end user. The cost function is then minimized over the node location matrix X.

We envision that the user interface might include a method for the users to mark (for example, by tapping on a touch-screen) the approximate node locations ($\bar{x}_i$) on a map or aerial image, similar to those provided by Google Maps. Building shapes could be directly inferred from the satellite imagery using edge detection. In addition, end users like SOF have access to building plans, which could also function as input to the software interface. We leave the design of the user interface for future work, but note that the shape of the building around which the nodes are deployed further

constrains the locations of the nodes. We augment the dwMDS cost function in order to include the building shape constraint

$$S = 2\sum_{i\neq j} w_{ij}(\delta_{ij} - d_{ij}(X))^2 + \sum_i r_i\|\mathbf{x}_i - \bar{\mathbf{x}}_i\|^2 + a\sum_i \|x_i - p\|^2, \qquad (5.2)$$

where p is the nearest point on the perimeter to x_i and a is a weighting factor that represents the quality of the perimeter information.

Another possible way to improve network self-localization is to use nodes that include GPS capability. Current commercial GPS devices are capable of localizing to as little as 2 m within 30 s of deployment, and are inexpensive thanks to the rise of the mobile phone. The GPS-based node locations may be used as *a priori* information in (5.2) or in addition to it. In fact, if GPS can reliably localize to 2 m and the end user only requires coarse target tracking, RTI might be performed using the GPS measurements alone. It is important to note that GPS receivers require unobstructed views of the sky to accurately localize, so we may not be able to rely on GPS in situations where the nodes are not exposed to the sky.

5.3.3 Sensor Network Self-Calibration

The tracking system must also establish baseline RSS distributions for each link in order to quantify changes in RSS and localize motion. Since it is not possible for end users to remove the antagonists or hostages from a building in order to perform calibration measurements, these distributions must be estimated online. For multichannel KRTI, it is also necessary to decide which frequency channels to use. We describe our methods for online baseline RSS estimation and channel selection below.

5.3.3.1 Baseline RSS estimation

KRTI relies on keeping two histograms of RSS measurements for each link in the network and comparing those distributions in order to determine whether or not people are moving near each link. Self- calibration after deployment occurs in real-time by continuously calculating the long-term histogram and using it as a baseline for detecting human motion. The long-term histogram converges to what would be seen in a calibration to be useful for finding both moving and stationary people in the

building. The convergence speed is adjustable, but we find that good performance is achieved with parameters that require about 30 s for convergence. According to our end user contacts, most barricade scenarios last long enough (sometimes multiple days) to allow such a convergence time. We note that people must move periodically in order for our tracking system to locate them. If they remain stationary for a period of time beyond the memory of both histograms, they will disappear from the tracking image.

5.3.3.2 Channel Selection

We leverage frequency diversity in our test system and demonstrate through multiple experiments that it improves tracking performance. The fading experienced by each link in the network is frequency selective, i.e., the RSS is different due to the different constructive or destructive combinations of the multipath components as a function of frequency. Transmitting on multiple channels makes it more likely that a channel will be found on which each link can be used reliably for RTI.

The best channels for RTI are those in an antifade, because the RSS on these channels are typically strongly affected only when a human target is blocking the line between the two nodes of the link, and not when he is moving at other positions [97]. In other words, antifade links are the most spatially informative [69, 102]. For each link, we choose the channel with the highest average RSS, because these channels are most likely to be in an antifade. There are other options for combining information from multiple channels, e.g., using the best n channels, but we leave the exploration of these options for future work. For multichannel KRTI, we allow an additional 30 s for channel selection, leading to a total of 60 s for calibration and channel selection.

5.4 Experiments

In this work, we present results both from real-time experiments as well as experiments that were used in postprocessing for analysis of system design. We perform experiments at the following sites (the building layouts are presented in Figure 5.3):

- **Site A**: A 110 m^2 single floor of a modern home in a typical suburb, comprised of four rooms and a bathroom. (33 nodes deployed)

- **Site B**: A 50 m^2 building comprised of 2 rooms. (34 nodes deployed)
- **Site C**: A 55 m^2 living space comprised of a single room. (36 nodes deployed)

In each case, Texas Instruments CC253X-based nodes are deployed as uniformly as possible around the perimeter of the building and data are collected using 8 dBi directional and omnidirectional antennas while a human target follows planned routes throughout the building. The tracking data are analyzed in postprocessing to determine the accuracy of the system. We study the tracking performance when fewer nodes are used to surround each location by using RSS measurements made at a subset of the nodes from each deployment. Additionally, we study the effects of poor node self-localization by adding noise to the known locations of the nodes.

5.4.1 Tracking

Before pursuing our research objectives relating to self-localization and self-calibration, we first evaluate the tracking performance of a system when the node locations are known exactly for each of the three experiments. Knowing the performance with exact node locations is important as a baseline for evaluating the effect of automatic configuration on tracking accuracy during rapid deployment.

5.4.2 Node Self-Localization

In order to test the accuracy of node self-localization methods like dwMDS, we precisely record the positions of each node during each deployment. During the calibration phase immediately after each deployment, we apply dwMDS in order to estimate the relative locations of the nodes.

We are also interested in the performance of our tracking system in the presence of imperfect knowledge of the node locations. In order to understand the effects of poor node self-localization, we simulate the circumstance by adding Gaussian noise to the true node locations and comparing the corresponding tracking results to those we achieve with the correct node locations.

5.4.3 Antenna Type

The use of better radio hardware may improve the performance of an RTI tracking system. For example, we are interested in determining whether or not the use of

directional antennas results in better tracking performance. Previous research of RTI [9, 10] has relied primarily on omnidirectional antennas, which radiate more energy away from the tracking area than they do into it. We expect that the more focused gain pattern of the directional antennas should maximize the amount of power being radiated through the building, as opposed to away from it or around it, leading to a more connected network, a higher average fade level, and better tracking performance. Maximizing the power radiated into the building is especially important in through-building imaging, where the signal may need to propagate through multiple exterior and interior walls.

In order to examine the benefits of directional antennas for our application, we first perform each experiment with radios that include a PCB microstrip inverted-F antenna with an omnidirectional gain pattern. We then repeat the experiment using circularly polarized 8 dBi directional antennas. In each case, we set the transmitted power for our radios to the maximum power allowed by the hardware in order to increase network connectivity as much as possible.

5.4.4 Network Size

It is important to understand the trade-off between the number of nodes in the RTI network and the corresponding tracking accuracy, because the tracking system must offer a fast and simple deployment in order to be useful to the end users. In some barricade scenarios the hostile targets may be armed, making it dangerous for SOF to spend time setting up nodes around the perimeter of the building. In these cases, smaller networks may allow for safer deployments and still offer useful tracking data. For example, using 30-40 nodes may allow for tracking a person to within 0.3 m of their true location, but the end user may wish to sacrifice some accuracy in order deploy the system quickly in a dangerous situation, e.g., using 10 nodes and accepting a tracking error of 1 m.

We examine the tracking performance for networks which include 10 to 36 nodes. At each experimental deployment, the nodes are placed around the perimeter of the building in an approximately equally spaced pattern.

5.4.5 Collaboration with End Users

In order to understand the constraints of the potential end users of our system, we collaborate with one of Utah's largest SWAT operations, the Unified Police Department in Salt Lake City. The purpose of our collaboration is to obtain explicit feedback about our proposed system, whether it would actually help in tactical operations, and what physical constraints need to be addressed in order for such a system to become important and useful to end users.

We organize an extensive through-building tracking demonstration day for members of the SWAT team and other law enforcement agencies in order to deploy our tracking system around a home in Salt Lake City and simulate hostage and barricade scenarios while law enforcement officers offer valuable feedback about system deployment and performance.

5.5 Results

We present the major results from our experimental deployments below. We discuss general tracking results in Section 5.5.1, self-localization results and the corresponding effects on tracking performance in Section 5.5.1.1, a comparison of tracking performance for directional and omnidirectional antennas in Section 5.5.2, and the effects on tracking performance of using fewer nodes in Section 5.5.3. Finally, we discuss the feedback from local SOF after a real-time demonstration of the system in Section 5.5.4.

5.5.1 Tracking

At Site A, with exact locations of nodes known, an average tracking error of approximately 1.1 m was achieved with 33 nodes over 110 m^2. At Site B, an average tracking error of 0.46 m was achieved with 34 nodes over 50 m^2. At Site C, an average tracking error of 0.54 m was achieved with 36 nodes over 55 m^2. Some tracking results for Sites A and B are depicted in Figure 5.4.

We expect that with a higher density of nodes per unit area, we should see a lower average tracking error, and this can be seen in the results presented in Table 5.1. As seen in Figure 5.6, Sites B and C, which are approximately the same size and have similar node-to-area ratios, show similar average tracking results. Site A, which

represents a larger area and is covered with less nodes, shows slightly higher average tracking error.

While these tracking results would be beneficial according to our SOF contacts, they are achieved using near-perfect knowledge of the node locations, which SOF may not have access to in most scenarios. The sequel discusses node self-localization results.

5.5.1.1 Effect of Self-localization Error

In Figure 5.5(a), we show the dwMDS results for Site A without any *a priori* information about the node locations ($r_i = 0$ for all i), which yield an average error of 3.3 m. The reason for the high average error is the rich multipath environment of the building, which leads to small-scale fading, and the failing of the large scale path loss model. We will show later that we can still achieve acceptable human target tracking results with this level of error in the network self-localization, but we can improve the localization by including some information from the end user about the deployment, specifically, *a priori* estimates of the node locations and building perimeter shape.

Figure 5.5(b) shows the results of dwMDS for Site A with coarse (2 m average error) *a priori* node locations and an the augmented cost function (5.2). In this case we achieve an average error of 1.5 m.

We note that our work investigates the accuracy of target tracking vs. the accuracy of node locations regardless of the methods used to localize the nodes. As expected, the accuracy of tracking decreases as the error of node location increases. However, keeping the mean squared error (MSE) of the node location estimates below 4 m^2 allows for average tracking errors of less than 1.5 m. The results are presented in Figure 5.6.

5.5.2 Directional vs. Omnidirectional Antennas

Figure 5.6 shows the tracking performance at each experiment site vs. error variance in the network self-localization for both antenna types. Directional antennas offer better performance at Site A, but the two types of antennas result in similar performance at Sites B and C.

The difference in the performance may be due in part to network connectivity: Site A shows improved connectivity, in terms of packet reception rates, when using directional antennas instead of omnidirectional antennas, while Sites B and C exhibit similar network connectivity regardless of antenna type.

5.5.3 Number of Nodes

Figure 5.7 shows the tracking results from each site, for network sizes ranging from 10 to 30 nodes, and both antenna types. Tracking results for the maximum number of nodes at each site can be seen in Figure 5.6. Surprisingly, we find that with as few as ten nodes, we are able to achieve less than 1.3 m average tracking error in most cases. If we guess randomly and uniformly at the location of the target across the area of the deployments at Sites A, B, and C, we find average errors of 6.0 m, 3.6 m, and 3.8 m, respectively.

The tracking accuracy appears to improve with the number of nodes. Although our experiments used a maximum of 36 nodes, we would expect that further increasing the number of nodes will further decrease the tracking error.

5.5.4 End User Feedback

After demonstration of our through-building tracking system, we interviewed SWAT commander Lt. Jake Petersen to receive his feedback and advice regarding the system. The following are quotes from the interview with their respective times in the video. The interview in its entirety can be found at `http://www.youtube.com/watch?v=QnQKfz-AEi4`. Note that a portion of the tracking demo is contained in the video at time 1:00.

- "This is something that I would use on really any barricaded subject or any hostage situation. Pinpointing exactly where the individual is, or even the hostages, allows us to make a save for these victims much easier. Really it is going to save lives, that's the mission of SWAT." (1:45)

- "Making sure that the technology works is really important and you've given me a lot of confidence in that here today." (4:20)

- "I would not be here if I didn't think that this product could save lives, that's the honest truth." (16:16)
- "I want to save their lives, and I believe this kind of thing could help us do that." (17:00)

5.6 Conclusion

We have examined the feasibility of a rapidly deployable through-building RTI system for SWAT and other SOF. We have shown that our system can rapidly self-localize and self-calibrate after deployment. The self-localization process requires minimal input from the user, and the system produces useful tracking results even when the node self-localization contains errors. We have also seen that directional antennas help increase through-building tracking accuracies as more power is radiated through the area of interest. Future development may use higher-power transmitters that provide full connectivity for larger building sizes.

Finally, through our interviews with SOF end users, we have further validated the need for this technology in tactical operations. We have shown that a simple, rapidly deployable, and user-friendly through-building tracking system is technically feasible. Future work will include the development of a user interface for SOF that will allow them to input deployment information, e.g., the building shape and coarse node locations, into the system, and then coordinate operations on top of the tracking data it generates.

5.7 Acknowledgment

We would like to thank Matt Kankainen for helping with the many experiments necessary for this work.

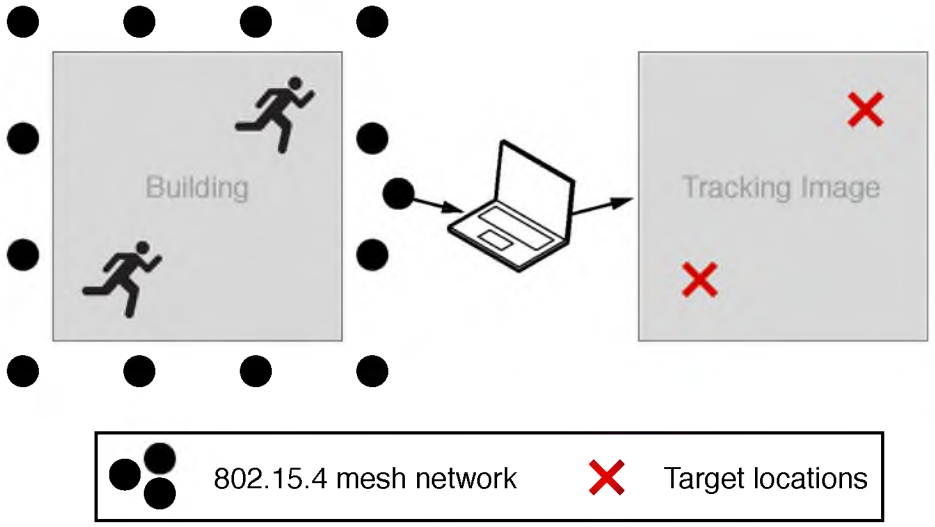

Figure 5.1. System overview. Special operations forces arrive at a building, deploy mesh network nodes around the perimeter of the building, and estimate the locations of people moving inside.

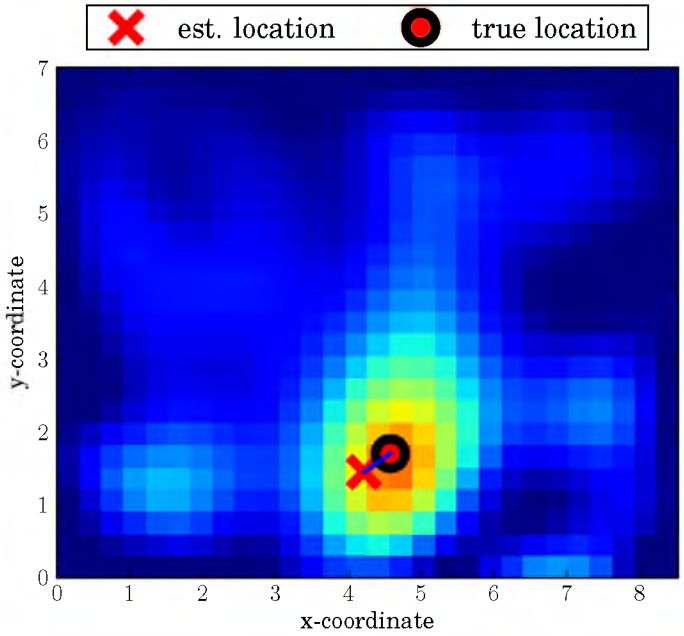

Figure 5.2. Example image for multichannel KRTI.

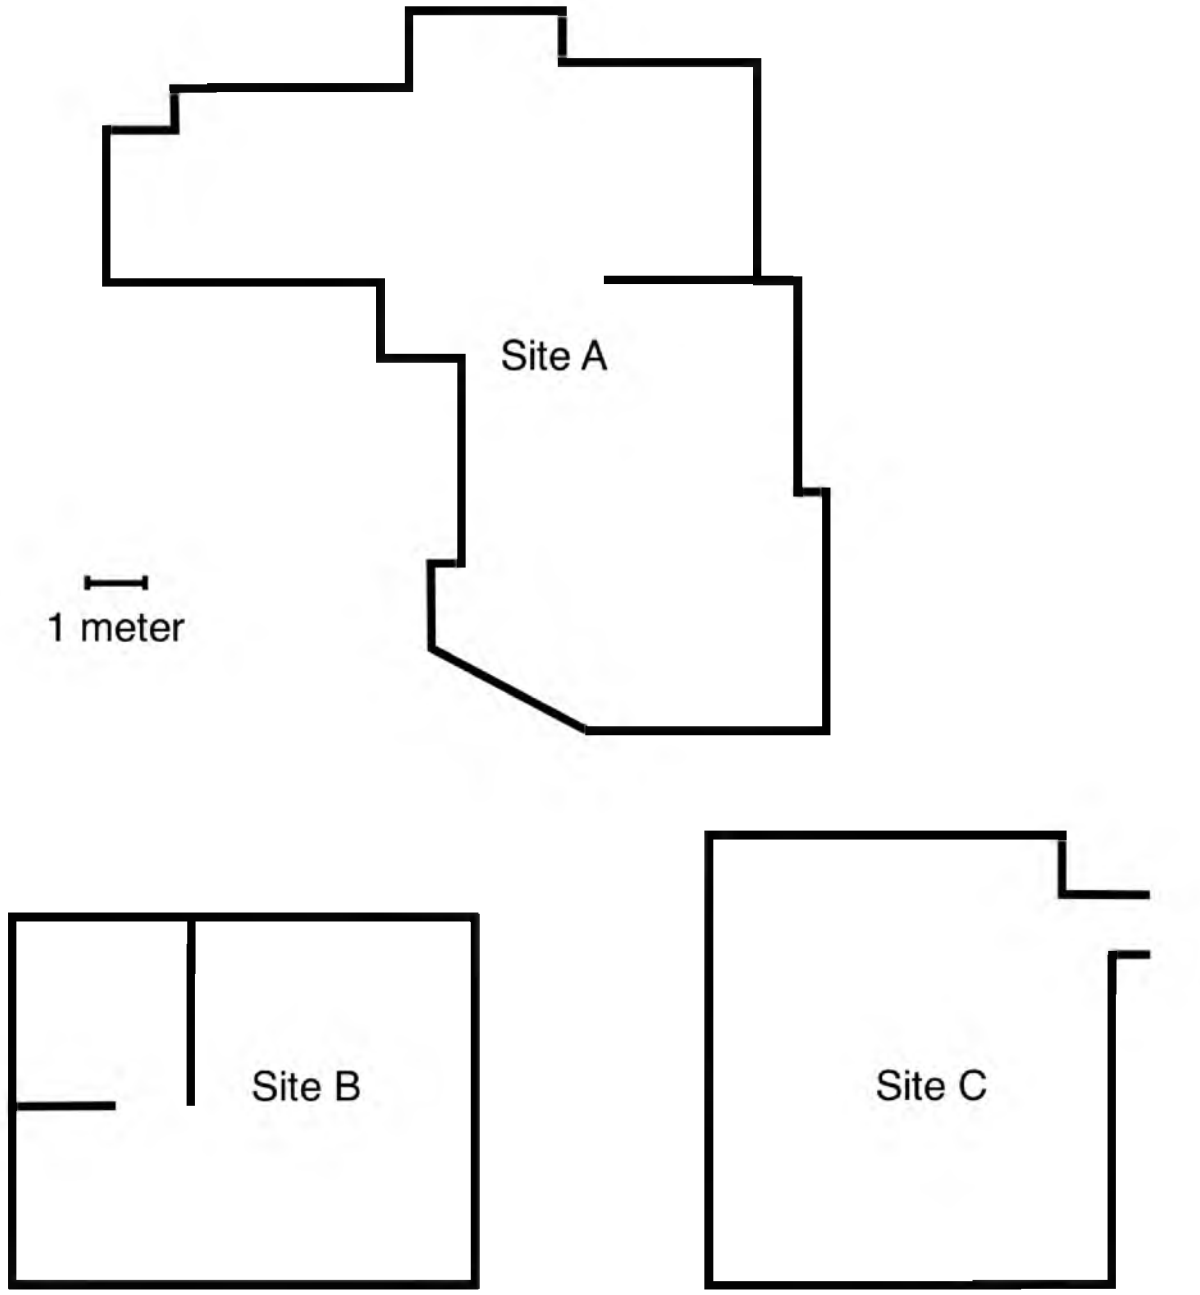

Figure 5.3. Experiments were conducted at three sites with different floorplans and building materials.

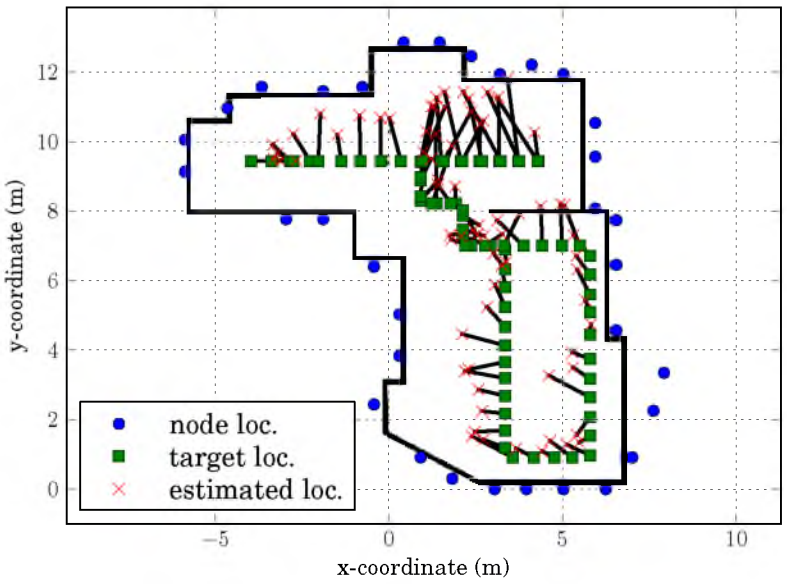

(a)

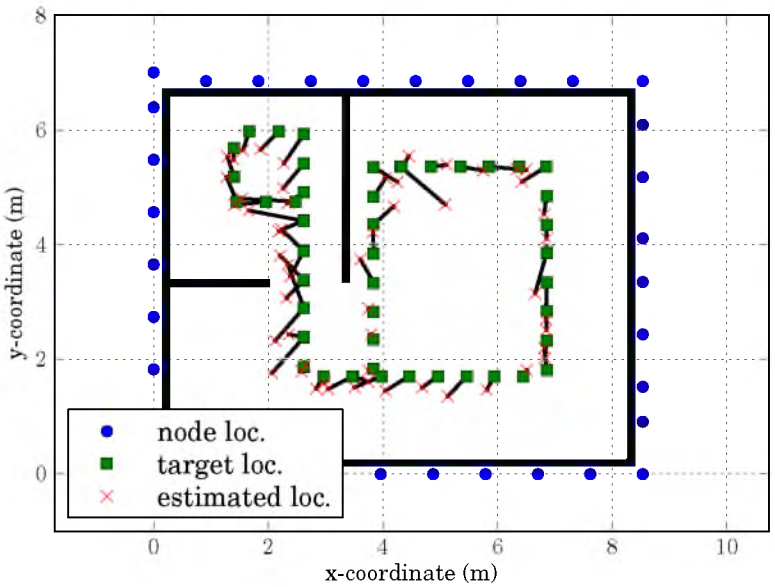

(b)

Figure 5.4. A subset of tracking results from: (a) Site A using directional antennas and multichannel KRTI resulting in an average error of 1.1 m; (b) Site B using directional antennas and multichannel KRTI resulting in an average error of 0.46 m.

Table 5.1. Average tracking error for best antenna type at each site compared to random estimator and number of nodes.

	Site A	Site B	Site C
10-node system	1.27 m	1.19 m	0.89 m
20-node system	1.22 m	0.70 m	0.68 m
30-node system	1.01 m	0.49 m	0.58 m
Random estimator	6.0 m	3.6 m	3.8 m

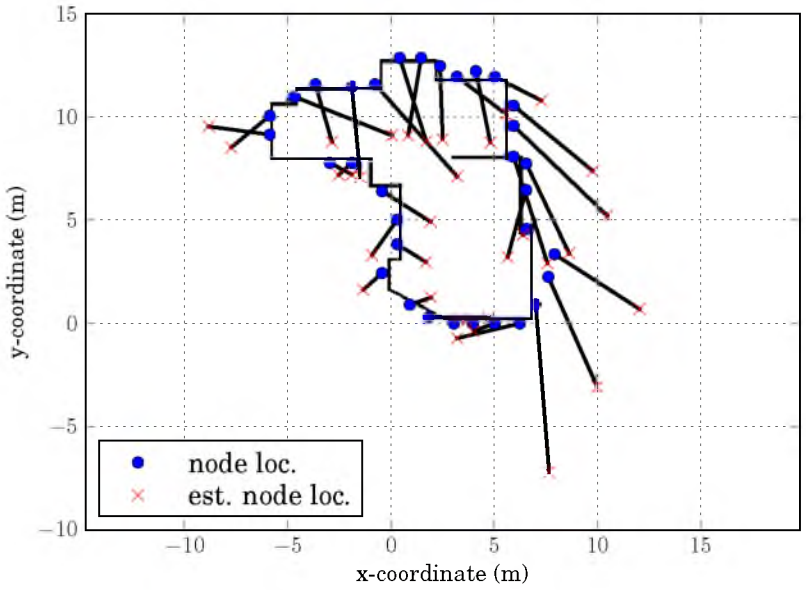

(a)

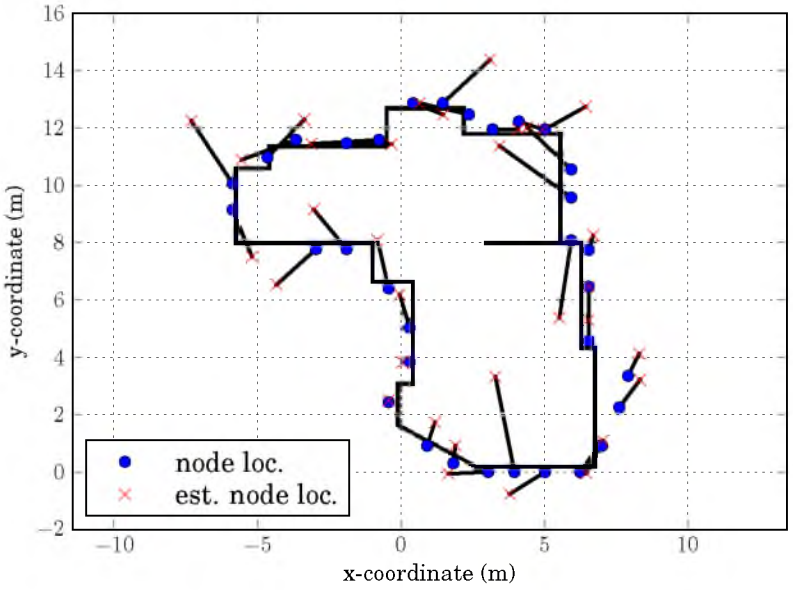

(b)

Figure 5.5. Multichannel dwMDS (a) without *a priori* information or augmented cost function and (b) with *a priori* information and augmented cost function.

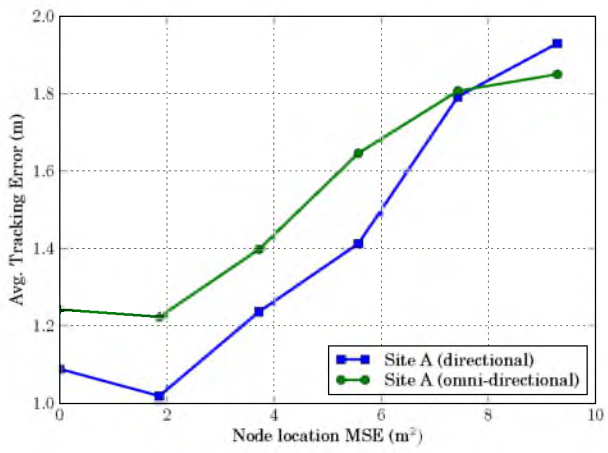

(a)

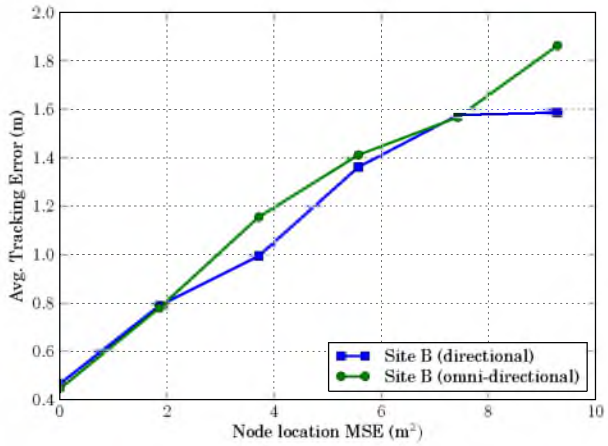

(b)

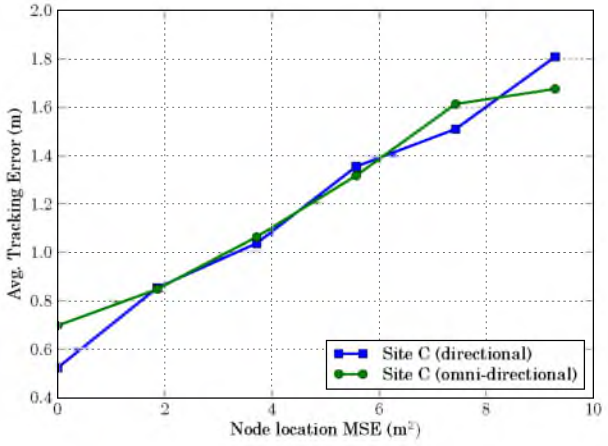

(c)

Figure 5.6. Average tracking error vs. mean squared error of node locations for directional and omnidirectional antennas at (a) Site A, (b) Site B, and (c) Site C.

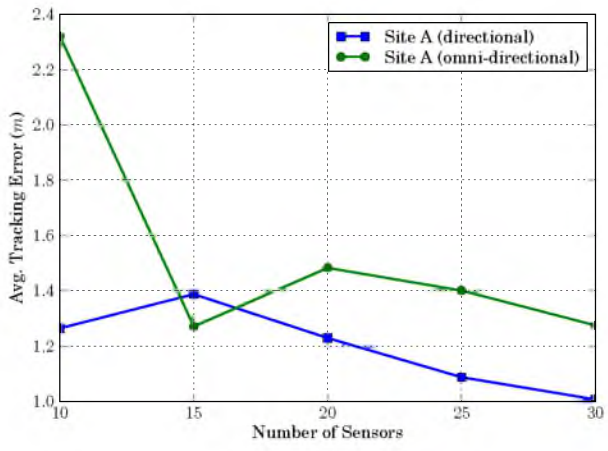

(a)

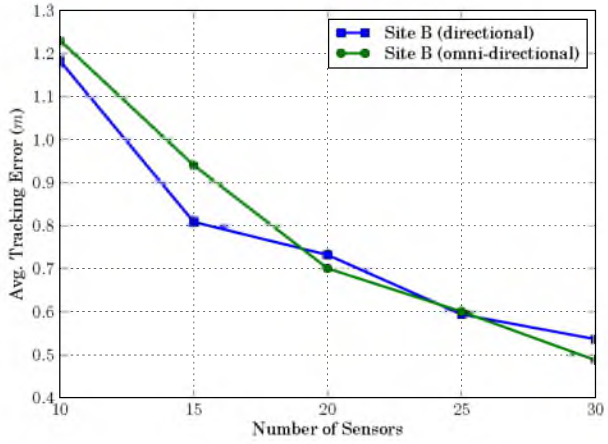

(b)

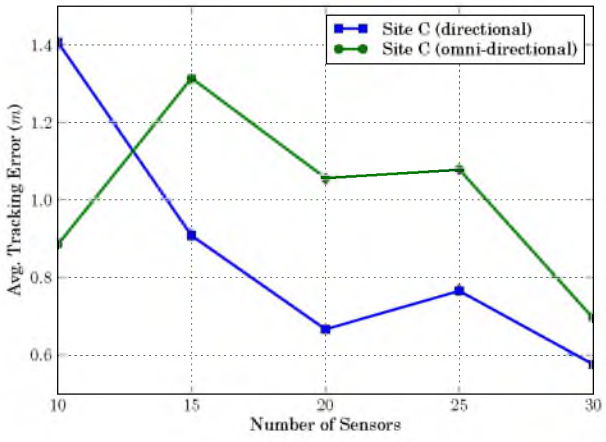

(c)

Figure 5.7. Average tracking error vs. number of nodes deployed for directional and omnidirectional antennas at (a) Site A, (b) Site B, and (c) Site C.

CHAPTER 6

CONCLUSION

This dissertation is concluded with a summary of the most important findings in in the work. Ongoing and future work are then discussed.

6.1 Research Summary

The ubiquity of wireless networks (Cellular, WSN, WiFi LAN, Bluetooth PAN) transmitting in the RF range is leading to new applications for measurements made at the physical layer of the network stack. Channel measurements, e.g., estimates of the CIR made with channel sounders or the RSS estimates commonly made available to the application layer of the network stack by commodity wireless devices, are now being leveraged to sense the environment of the network.

Part of this dissertation is an effort to improve the tool-set used for CIR impulse response measurement by wireless communication researchers. To that end, I have helped build an open source 802.11b channel sounder on top of the popular GNU Radio software defined radio platform. The sounder is capable of receiving and decoding standard 802.11b packets at 1 Mbps and 2 Mbps data rates at full bandwidth, improving on previous receivers [46] by match filtering the incoming packets on the FPGA of the USRP before sending data through USB to the host PC. This tool has been downloaded over 1000 times since release. The tool has been validated in controlled experiments in the lab as well as in a variety of real-world environments.

I have also conducted a study of the methods for using CIR measurements to perform location distinction, specifically, in the context of MIMO communications, leading to an understanding of the trade-offs between different aspects of system complexity and system performance. Key findings include the knowledge that a system with two antennas offers nearly as much improvement in location distinction

performance as a system with eight antennas, and that measurement bandwidths beyond 20 MHz offer diminishing returns because of the tendency of lower bandwidth measurements to mask the effects of timing synchronization errors. Finally, the MIMO location distinction system is capable of accurately detecting changes in receiver position using measurements for just one transmitter, while accurate detection in previous work using SISO devices has required multiple links.

Part of this dissertation is concerned with exploiting the so-called "radio windows" created by infrastructure wireless networks. This work shows that, using ambient secured radio signals emitted by a building's WiFi or ZigBee network, an attacker outside of the building can infer location information about the people within the building. This information includes their direction of motion and what side of the building they are on, and is based on the detection of people crossing the LOS radio links between a legitimate transmitter inside the building and an illegitimate receiver placed outside the building by the attacker. The work also examines possible defenses for this attack. A key finding is that we can identify line crossings and direction of motion through walls with greater than 90% accuracy using just a single 802.11n attack receiver.

It should be noted that we did not use GNU Radio / USRP channel sounder or the National Instruments 2x2 MIMO channel sounder for the "radio window" work because a cheaper and more portable measurement system, the CSI Tool [20], became available. This tool, which makes 3x3 MIMO OFDM channel measurements on WiFi packets, provides higher quality measurements across larger bandwidths than the other tools are capable of, and does so at a significantly lower cost.

The CSI Tool relies on a common Intel NIC with a modified firmware/driver. No other specialized hardware is necessary, showing that it is feasible to make information-rich channel measurements available to the application layer on commodity hardware. Location distinction performance using measurements made with this tool, or similar measurements made with other commodity NICs, would be as robust as the performance achieved using measurements made by the sounders used in Chapter 3. I hope that this and other useful applications, e.g., secret key sharing and fingerprint-based localization, encourage NIC designers to make channel measurements available higher

in the network stack.

The final part of this dissertation is concerned with bringing the benefits of through-wall imaging and tracking to SWAT teams and other special operations forces by making a rapidly deployable and robust RTI system. One key finding in this feasibility study are that the node locations in an RTI network do not need to be known very accurately to achieve useful tracking results, and can be estimated through an RSS-based method like dwMDS. Additionally, we find that as few as 10 nodes can be used in a deployment, while still offering useful tracking results. A method like KRTI can be used to alleviate the need for offline calibration, while maintaining the ability to image people that are stationary for short periods of time. Finally, directional antennas offer better connectivity and RTI performance, since less transmit power is being radiated away from the building.

6.2 Ongoing and Future Work

At the time of this writing, I am continuing research efforts to develop a rapidly deployable and robust RTI system. This includes studying the effects of system parameters on tracking performance and automating parameter selection, so the end user does not have to, and examining the possibility of online adaptation of the RTI weighting matrix in order to improve image quality. It also includes developing a 2.5D RTI algorithm that uses multiple radio channels and takes into consideration the level of multipath fading for each link when creating the weighting matrix.

Future work should address the radio window problem. Currently, countless homes and other buildings contain wireless networks that expose them, on some level, to anyone with a cellular phone, tablet, or laptop. A savvy thief, for example, can make a good guess as to whether or not a building is vacant based on measuring signal strength changes over time on its wireless network.

Perhaps the problem can be solved using algorithms that disguise the effects of humans on infrastructure wireless signals in subtle ways that are less susceptible to detection and removal than the method discussed in Chapter 4. The radio window is becoming even more of an issue with the introduction of less costly and more powerful consumer-grade software defined radios, which can go further than measuring signal

strengths on the network. For example, with a single software defined radio, it is possible to detect the respiration rate of a person. So, the most savvy of thieves could potentially find out whether or not you are home, even if you remain perfectly still.

This work has been focused on making radio channel measurements of different types and using them to (a) localize people moving in the area covered by a wireless network and (b) detect the motion of the radios that make up the wireless network. The former tends to be the harder of these two tasks, since a person moving in the environment of a radio link will only affect a subset of the multipath that contribute to the CIR, while the motion of one of the radios that make up the link will affect most of the multipath. Using RSS as a channel difference metric makes both problems harder, since multiple positions of a person or radio can lead to similar RSS measurements. Therefore, future research should apply RTI methods to CIR measurements instead of RSS measurements. It is possible that this would lead to RTI systems that require fewer radios and that are therefore less expensive and easier to deploy. Further, using CIR measurements, it should be possible to improve radio self-localization in the RTI network, also leading to better performance.

APPENDIX

RAPIDLY DEPLOYABLE RTI ADDENDUM

This addendum is included as a supplement to the reprinted paper [19]. In Section A.1, I briefly discuss a link budget analysis for through-wall RTI systems. In Section A.2, I discuss the effect of randomly selecting the subset of radios used on tracking performance for through-wall RTI.

A.1 Link Budget Considerations for Through-Wall RTI

The radio links in a through-wall RTI system must penetrate multiple interior and exterior walls, as well as the furniture, appliances, and people that are in the building. Since the radios in the system must operate with finite transmit power and receiver sensitivity, it is important to consider the link budgets of the radios that comprise the system. While it is impossible to perfectly model the effects of every possible deployment environment for RTI, a link budget allows us to determine whether or not a given deployment is likely to have the connectivity necessary to perform RTI given the power constraints of the system and the size of the deployment.

A log-normal path loss model, which includes the losses contributed by free-space propagation and shadowing, is augmented with a fade margin term in order to capture the effects of multipath fading, and an additional shadowing term which represents the losses caused by the external walls. This model represents the received power for any link in the network as

$$P_R(d)_{dB} = P_T + G_T + G_R - L_C - L_P(d_0) - 10n \log_{10} \frac{d}{d_0} - X_\sigma - L_f - L_w \quad \text{(A.1)}$$

where P_R is the received power, d is the distance from transmitter to receiver, P_T is

the transmitted power, G_T is the gain of the transmitter antenna, G_R is the gain of the receiver antenna, L_C is the loss caused by cables and connectors, $L_P(d_0)$ represents the path loss at reference distance d_0, n is the path loss exponent, X_σ is a zero mean Gaussian random variable with standard deviation σ, which accounts for random shadowing losses, L_f is the fade margin term, and L_w represents the wall losses.

The TI radios used for the experiments performed in [19] offer a typical transmit power of 4.5 dBm and a typical receiver sensitivity of -97 dBm. The antenna gain of the directional antennas used for the experiments (transmit and receive) is 8 dBi. In [103], a study of indoor propagation at 2.4 GHz, $L_P(d_0)$ is found to be 50 dB with $d_0 = 3\ m$, n is found to be 3.73, and σ is found to be 4.35 dB for NLOS links. The cable and connector losses are approximately 2 dB. A good rule-of-thumb fade margin for indoor environments is 25 dB [104].

The wall loss term L_w is specific to the material and thickness of the exterior walls. In [105], a study of propagation losses through common building materials, the authors find the transmission losses shown in Table A.1. If, for example, the exterior walls are made of two-layers of red brick, and all links must links pass through two exterior walls, $L_w \approx 17.7$ dB. Substituting the corresponding values into (A.1) and simplifying yields

$$P_R(d)_{dBm} = -74.2 - 37.3 \log_{10} \frac{d}{3} - X_{\sigma=4.35} \tag{A.2}$$

Figure A.1 shows (A.2) for distances up to 20 m, with the dashed lines representing the positive and negative standard deviations. In this example, a reliable network for RTI is probable as long as the radio links are less than 9 m in length. Longer links will require an increase in transmit power, antenna gain, or receiver sensitivity.

A.2 Further Tracking Results

In [19], I examined the effect of using fewer radios for through-wall RTI by only including measurements from a subset of the radios deployed in each experiment. A single subset was chosen such that the radios it contained surrounded the site with roughly even spacing between adjacent radios. However, it is possible to choose more than one subset of a given size. Since some radio links are better than others for RTI, e.g., due to fade level, different radio subsets may offer different tracking performance.

In this addendum, I have included Figure A.2, which shows the subset average tracking errors for each experiment site. The subset average tracking error is the result of averaging the mean tracking errors over multiple subsets for each network size. Subsets for a network of N radios are selected such that the N radios are evenly distributed around the experiment site. When N is less than half of the radios originally deployed for an experiment site, this is accomplished by selecting every kth radio, where k is as large as possible given the number of radios deployed. When N is greater than half of the radios originally deployed, this is accomplished by pruning every kth radio, where k is as large as possible given the number of radios deployed. Different subsets for a given network size are created by shifting the selection/pruning by a single radio. This leads to 4 subsets for $N = 10$, 3 subsets for $N = 15$, and 2 subsets for $N = 20, 25, 30$.

Although the results differ slightly from those presented in Figure 5.7, the general trends and conclusions about decreasing the number of radios in the deployment remain the same: (a) the directional antennas tend to offer better tracking performance than the omnidirectional antennas, (b) the tracking error goes gracefully with decreasing number of radios.

Table A.1. Transmission coefficients (T) at 2.3 GHz for common building materials.

Material	T (dB)
drywall	−0.49
fir lumber	−2.79
red brick	−4.43
plywood	−1.91
stucco	−14.9
cinder block	−4.43

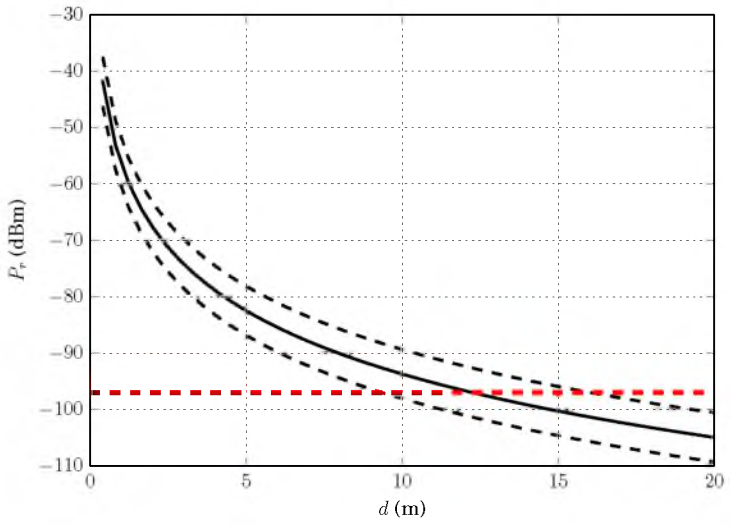

Figure A.1. Received power for a through-wall link according to (A.2). The black dashed lines represent a positive and negative standard deviation. The red line represents the receiver sensitivity, below which it is unlikely that the receiver will be able to detect and decode packets. In this case, to achieve a reliable network for RTI, it would be best to keep the link lengths below 9 m.

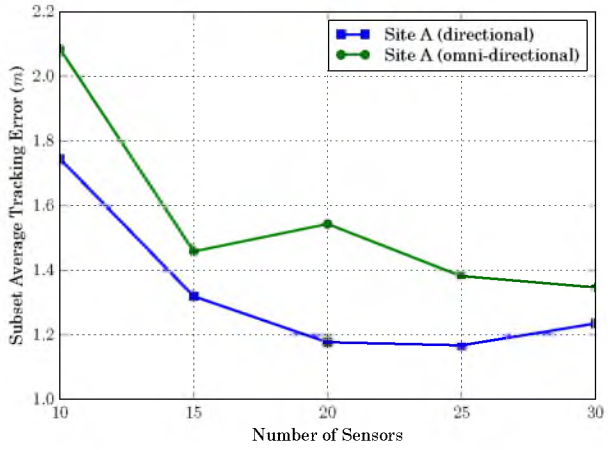

(a)

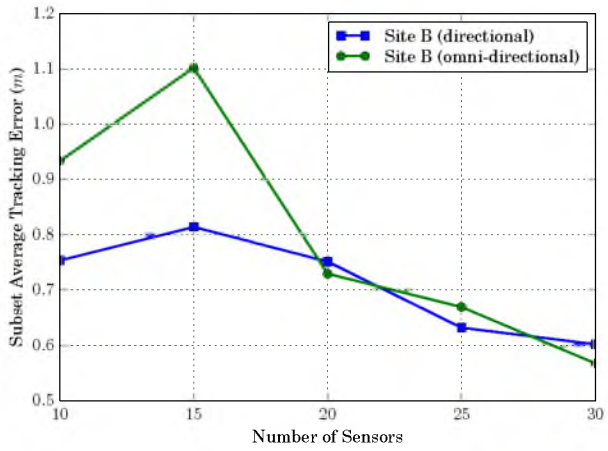

(b)

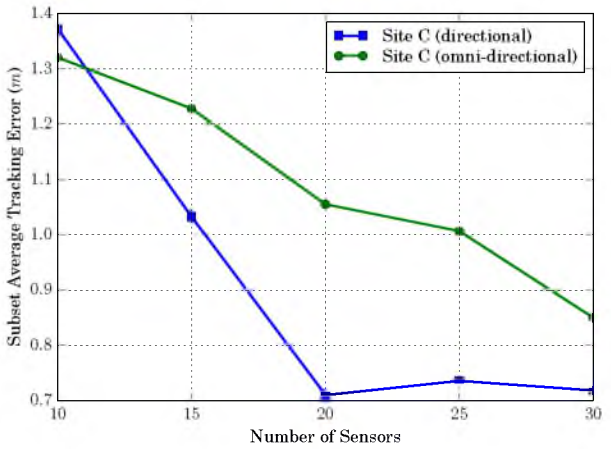

(c)

Figure A.2. Subset average tracking error vs. number of nodes deployed for directional and omnidirectional antennas at (a) Site A, (b) Site B, and (c) Site C. While the tracking results depicted in Figure 5.7 come from a single subset of radios for each network size, those shown here come from averaging the mean tracking error over multiple subsets of radios for each network size.

REFERENCES

[1] B. Son, Y.-s. Her, and J. Kim, "A design and implementation of forest-fires surveillance system based on wireless sensor networks for South Korea mountains," *International Journal of Computer Science and Network Security (IJCSNS)*, vol. 6, no. 9, pp. 124–130, 2006.

[2] R. N. Handcock, D. L. Swain, G. J. Bishop-Hurley, K. P. Patison, T. Wark, P. Valencia, P. Corke, and C. J. O'Neill, "Monitoring animal behavior and environmental interactions using wireless sensor networks, GPS collars and satellite remote sensing," *Sensors*, vol. 9, no. 5, pp. 3586–3603, 2009.

[3] J. P. Lynch and K. J. Loh, "A summary review of wireless sensors and sensor networks for structural health monitoring," *Shock and Vibration Digest*, vol. 38, no. 2, pp. 91–130, 2006.

[4] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication.* New York, NY: Cambridge University Press, 2005.

[5] B. Allen, M. Dohler, E. Okon, W. Malik, A. Brown, and D. Edwards, *Ultra Wideband Antennas and Propagation for Communications, Radar and Imaging.* New York, NY: Wiley, 2006.

[6] J. Zhang, M. H. Firooz, N. Patwari, and S. K. Kasera, "Advancing wireless link signatures for location distinction," in *14th ACM International Conference on Mobile Computing and Networking (MobiCom'08)*, pp. 26–37, 2008.

[7] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe, "MIMO-assisted channel-based authentication in wireless networks," in *42nd Annual Conference Information Sciences and Systems (CISS'08)*, pp. 642–646, 2008.

[8] M. Seifeldin, A. Saeed, A. Kosba, A. El-Keyi, and M. Youssef, "Nuzzer: A large-scale device-free passive localization system for wireless environments," *IEEE Transactions on Mobile Computing*, 2012.

[9] J. Wilson and N. Patwari, "Radio tomographic imaging with wireless networks," *IEEE Transactions on Mobile Computing*, vol. 9, no. 5, pp. 621–632, 2010.

[10] J. Wilson and N. Patwari, "See through walls: motion tracking using variance-based radio tomography networks," *IEEE Transactions on Mobile Computing*, vol. 10, pp. 612–621, May 2011. appeared online 23 September 2010.

[11] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan, "Indoor localization without the pain," in *16th Annual International Conference on Mobile Computing and Networking (MobiCom'10)*, (New York, NY), pp. 173–184, ACM, 2010.

[12] A. E. Kosba, A. Saeed, and M. Youssef, "Rasid: A robust WLAN device-free passive motion detection system," in *IEEE International Conference on Pervasive Computing and Communications (PerCom'12)*, pp. 180–189, 2012.

[13] S. Sen, R. R. Choudhury, B. Radunovic, and T. Minka, "Precise indoor localization using phy layer information," in *10th ACM Workshop on Hot Topics in Networks*, p. 18, ACM, 2011.

[14] N. Patwari and S. Kasera, "Robust location distinction using temporal link signatures," in *13th ACM International Conference on Mobile Computing and Networking (MobiCom'07)*, pp. 111–122, 2007.

[15] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe, "Channel-based detection of Sybil attacks in wireless networks," *IEEE Transactions on Information Forensics and Security*, vol. 4, pp. 492–503, Sept. 2009.

[16] D. Maas, M. Firooz, J. Zhang, N. Patwari, and S. Kasera, "Channel sounding for the masses: Low complexity GNU 802.11b channel impulse response estimation," *IEEE Transactions on Wireless Communications*, vol. 11, no. 1, pp. 1–8, 2012.

[17] D. Maas, N. Patwari, J. Zhang, S. Kasera, and M. Jensen, "Location distinction in a MIMO channel," in *Virginia Tech Wireless Symposium*, 2009. student poster.

[18] D. Maas, N. Patwari, S. Kasera, D. Wasden, and M. Jensen, "Experimental performance evaluation of location distinction for MIMO links," in *4th IEEE International Conference on Communication Systems and Networks (COMSNETS'12)*, pp. 1–10, 2012.

[19] D. Maas, J. Wilson, and N. Patwari, "Toward a rapidly deployable RTI system for tactical operations," in *8th IEEE International Workshop on Practical Issues in Building Sensor Network Applications (SenseApp'13)*, 2013.

[20] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Tool release: Gathering 802.11n traces with channel state information," *SIGCOMM Computer Communication Review*, vol. 41, pp. 53–53, Jan. 2011.

[21] D. C. Cox, "Delay Doppler characteristics of multipath propagation at 910 MHz in a suburban mobile radio environment," *IEEE Transactions on Antennas and Propagation*, vol. AP-20, pp. 625–635, Sept. 1972.

[22] G. L. Turin, F. D. Clapp, T. L. Johnston, S. B. Fine, and D. Lavry, "A statistical model of urban multipath propagation," *IEEE Transactions on Vehicular Technology*, vol. VT-21, pp. 1–9, Feb. 1972.

[23] D. M. Devasirvatham, "Time delay spread and signal level measurements of 850 MHz radio waves in building environments," *IEEE Transactions on Antenna and Propagation*, vol. 34, pp. 1300–1305, Nov. 1986.

[24] A. Kemp and S. Barton, "The impact of delay spread on irreducible errors for wideband channels on industrial sites," *Wireless Personal Communications*, vol. 34, no. 3, pp. 307–319, 2005.

[25] H. Xu, D. Chizhik, H. Huang, and R. Valenzuela, "A wave-based wideband MIMO channel modeling technique," in *13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, vol. 4, 2002.

[26] J.-M. Molina-Garcia-Pardo, J.-V. Rodriguez, and L. Juan-Llácer, "MIMO channel sounder based on two network analyzers," *IEEE Transactions on Instrumentation and Measurements*, 2008.

[27] D. Lee, K. Sowerby, and M. Neve, "Extracting fine multipath detail from measured data at 5.8 GHz," in *59th IEEE Vehicular Technology Conference*, vol. 1, pp. 74–78, 2004.

[28] M. C. Wicks, B. Himed, J. Bracken, H. Bascom, and J. Clancy, "Ultra narrow band adaptive tomographic radar," in *1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 36–39, Dec. 2005.

[29] J. Zhang, M. H. Firooz, N. Patwari, and S. K. Kasera, "Advancing wireless link signatures for location distinction," in *14th International Conference on Mobile Computing and Networking (MobiCom'08)*, pp. 26–37, Sept. 2008.

[30] B. Azimi-Sadjadi, A. Kiayias, A. Mercado, and B. Yener, "Robust key generation from signal envelopes in wireless networks," in *14th ACM Conference on Computer and Communications Security*, pp. 401–410, Nov. 2007.

[31] E. Blossom, "GNU radio: Tools for exploring the radio frequency spectrum," *Linux Journal*, p. 4, Jun. 2004.

[32] G. D. Troxel, E. Blossom, S. Boswell, A. Caro, I. Castineyra, A. Colvin, T. Dreier, J. Evans, N. Goffee, K. Haigh, and et al., "Adaptive dynamic radio open-source intelligent team (ADROIT): Cognitively-controlled collaboration among SDR nodes," in *IEEE Workshop on Networking Technologies for Software Defined Radio Networks*, pp. 8–17, Sept. 2006.

[33] Ettus Research L.L.C.. UHD Software. http://www.ettus.com.

[34] M. Firooz, D. Maas, and N. Patwari. http://span.ece.utah.edu/pmwiki/pmwiki.php?n=Main.80211bReceiver.

[35] R. J. Pirkl and G. D. Durgin, "Optimal sliding correlator channel sounder design," *IEEE Transactions on Wireless Communications*, vol. 7, pp. 3488–3497, September 2008.

[36] "Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: Higher-speed physical layer extension in the 2.4 GHz band," *IEEE Std. 802.11b, Part 11*, 1999.

[37] J. Jemai and T. Kurner, "Broadband WLAN channel sounder for IEEE 802.11b," *IEEE Transactions on Vehicular Technology*, 2008.

[38] B. Farhang-Boroujeny, *Signal Processing Techniques for Software Radios, 2nd ed.* Raleigh, NC: Lulu Press Inc., 2009.

[39] A. Goldsmith, *Wireless Communications*. Cambridge, UK: Cambridge University Press, 2005.

[40] N. Levanon and E. Mozeson, *Radar Signals*. New York, NY: Wiley, 2004.

[41] J.-J. Fuchs, "Multipath time-delay detection and estimation," *IEEE Transactions on Signal Processing*, vol. 47, pp. 237–243, January 1999.

[42] D. Maas, M. Firooz, J. Zhang, N. Patwari, and S. Kasera, "Channel sounding for the masses: Low complexity gnu 802.11b channel impulse response estimation," *arXiv:1007.3476v1 [cs.CR]*, July 2010.

[43] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21." http://cvxr.com/cvx, Apr. 2011.

[44] D. G. Luenberger, *Introduction to Linear and Nonlinear Programming*. Reading, MA: Addison-Wesley, 1973.

[45] IEEE wireless standards zone. http://standards.ieee.org/wireless/.

[46] D. Sumorok, "bbn_80211b_rx.py, version 1.7." http://acert.ir.bbn.com/.

[47] T. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ: Prentice Hall, 1996.

[48] J. Jorgensen, I. Kostanic, and W. Foose, "Application of channel sounding to CDMA PCS design at 1900 MHz," in *47th IEEE Vehicular Technology Conference*, vol. 3, pp. 1937–1941, May 1997.

[49] L. Greenstein, V. Erceg, Y. Yeh, and M. Clark, "A new path-gain/delay-spread propagation model for digital cellular channels," *IEEE Transactions on Vehicular Technology*, vol. 46, pp. 477–485, May 1997.

[50] X. Zhao, J. Kivinen, P. Vainikainen, and K. Skog, "Propagation characteristics for wideband outdoor mobile communications at 5.3 GHz," *IEEE Journal on Selected Areas in Communications*, vol. 20, pp. 507–514, Apr. 2002.

[51] E. Sousa, V. Jovanovic, and C. Daigneault, "Delay spread measurements for the digital cellular channel in Toronto," *IEEE Transactions on Vehicular Technology*, vol. 43, pp. 837–847, Nov. 1994.

[52] G. Durgin, V. Kukshya, and T. Rappaport, "Wideband measurements of angle and delay dispersion for outdoor and indoor peer-to-peer radio channels at 1920 MHz," *IEEE Transactions on Antennas and Propagation*, vol. 51, pp. 936–944, may 2003.

[53] Y. Sheng, K. Tan, G. Chen, D. Kotz, and A. Campbell, "Detecting 802.11 MAC layer spoofing using received signal strength," in *27th IEEE Conference on Computer Communications (INFOCOM'08)*, pp. 1768–1776, 2008.

[54] M. Demirbas and Y. Song, "An RSSI-based scheme for Sybil attack detection in wireless sensor networks," in *IEEE International Symposium World of Wireless, Mobile, and Multimedia Networks (WoWMoM'06)*, pp. 564–570, 2006.

[55] D. Faria and D. Cheriton, "Detecting identity-based attacks in wireless networks using signalprints," in *5th ACM Workshop Wireless Security (WiSe'06)*, pp. 43–52, 2006.

[56] Z. Li, W. Xu, R. Miller, and W. Trappe, "Securing wireless systems via lower layer enforcements," in *5th ACM Workshop Wireless Security (WiSe'06)*, pp. 33–42, 2006.

[57] J. Wallace and M. Jensen, "Time-varying MIMO channels: Measurement, analysis, and modeling," *IEEE Transactions on Antennas Propagation*, vol. 54, pp. 3265–3273, Nov. 2006.

[58] "Wireless LAN medium access control (MAC) and physical layer (PHY) specifications amendment 5: Enhancements for higher throughput," *IEEE Std. 802.11n-2009. Part 11*, Oct. 2009.

[59] P. Moose, "A technique for orthogonal frequency division multiplexing frequency offset correction," *IEEE Transactions on Communications*, vol. 42, pp. 2908–2914, Oct. 1994.

[60] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Upper Saddle River, NJ: Prentice Hall, 1993.

[61] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume II: Detection Theory*. Upper Saddle River, NJ: Prentice Hall, 1998.

[62] C. Tepedelenlioğlu, A. Abdi, G. B. Giannakis, and M. Kaveh, "Estimation of Doppler spread and signal strength in mobile communications with applications to handoff and adaptive transmission," *Wireless Communication and Mobile Computing*, vol. 1, no. 2, pp. 221–242, 2001.

[63] Y. Chen, W. Trappe, and R. Martin, "Detecting and localizing wireless spoofing attacks," in *4th Annual IEEE Conference Sensor, Mesh, and Ad Hoc Communications and Networks (SECON'07)*, pp. 193—202, 2007.

[64] D. Zhang, J. Ma, Q. Chen, and L. M. Ni, "An RF-based system for tracking transceiver-free objects," in *IEEE International Conference on Pervasive Computing and Communications (PerCom'07)*, pp. 135–144, 2007.

[65] M. A. Kanso and M. G. Rabbat, "Compressed RF tomography for wireless sensor networks: Centralized and decentralized approaches," in *5th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS'09)*, June 2009.

[66] X. Chen, A. Edelstein, Y. Li, M. Coates, M. Rabbat, and M. Aidong, "Sequential Monte Carlo for simultaneous passive device-free tracking and sensor localization using received signal strength measurements," in *10th International Conference on Information Processing in Sensor Networks (IPSN'10)*, April 2011.

[67] O. Kaltiokallio and M. Bocca, "Real-time intrusion detection and tracking in indoor environment through distributed RSSI processing," in *Embedded and Real-Time Computing Systems and Applications (RTCSA), 2011 IEEE 17th International Conference on*, vol. 1, pp. 61–70, IEEE, 2011.

[68] R. K. Martin, C. Anderson, R. W. Thomas, and A. S. King, "Modelling and analysis of radio tomography," in *4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pp. 377–380, 2011.

[69] O. Kaltiokallio, M. Bocca, and N. Patwari, "Follow@ grandma: Long-term device-free localization for residential monitoring," in *37th IEEE Conference on Local Computer Networks Workshop*, pp. 991–998, IEEE, 2012.

[70] J. Wilson and N. Patwari, "See through walls: Motion tracking using variance-based radio tomography networks," *IEEE Transactions on Mobile Computing*, vol. 10, no. 5, pp. 612–621, 2010.

[71] Y. Zheng and A. Men, "Through-wall tracking with radio tomography networks using foreground detection," in *IEEE Wireless Communications and Networking Conference (WCNC'12)*, pp. 3278–3283, 2012.

[72] M. Seifeldin, A. Saeed, A. Kosba, A. El-Keyi, and M. Youssef, "Nuzzer: A large-scale device-free passive localization system for wireless environments," *IEEE Transactions on Mobile Computing*, vol. PP, no. 99, p. 1, 2012.

[73] F. Viani, P. Rocca, M. Benedetti, G. Oliveri, and A. Massa, "Electromagnetic passive localization and tracking of moving targets in a WSN-infrastructured environment," *Inverse Problems*, vol. 26, pp. 1–15, March 2010.

[74] C. Xu, B. Firner, Y. Zhang, R. Howard, J. Li, and X. Lin, "Improving RF-based device-free passive localization in cluttered indoor environments through probabilistic classification methods," in *11th International Conference on Information Processing in Sensor Networks (IPSN'12)*, pp. 209–220, 2012.

[75] M. Youssef, M. Mah, and A. Agrawala, "Challenges: Device-free passive localization for wireless environments," in *13th International Conference on Mobile Computing and Networking (MobiCom'07)*, pp. 222–229, 2007.

[76] W. C. Stone, "NIST construction automation program report no. 3: Electromagnetic signal attenuation in construction materials," *Building Fire Res. Lab., Nat. Inst. Standards Technol., Gaithersburg, MD, Tech. Rep. NISTIR*, vol. 6055, 1997.

[77] B. Greenstein, R. Gummadi, J. Pang, M. Y. Chen, T. Kohno, S. Seshan, and D. Wetherall, "Can Ferris Bueller still have his day off? Protecting privacy in the wireless era," in *11th USENIX Workshop on Hot Topics in Operating Systems*, HOTOS'07, pp. 101–106, 2007.

[78] B. Danev, D. Zanetti, and S. Capkun, "On physical-layer identification of wireless devices," *ACM Computer Survey*, vol. 45, no. 1, pp. 6:1–6:29, 2012.

[79] 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 *19th USENIX Conference on Security*, USENIX Security'10, pp. 21–21, 2010.

[80] K. B. Rasmussen and S. Čapkun, "Location privacy of distance bounding protocols," in *15th ACM Conference on Computer and Communications Security (CCS'08)*, pp. 149–160, 2008.

[81] S. Brands and D. Chaum, "Distance-bounding protocols," in *Advances in Cryptology (EUROCRYPT'93)*, pp. 344–359, 1993.

[82] X. Fu, N. Zhang, A. Pingley, W. Yu, J. Wang, and W. Zhao, "The digital marauder's map: A new threat to location privacy," in *29th IEEE International Conference on Distributed Computing Systems*, 2009.

[83] T. Jiang, H. J. Wang, and Y. chun Hu, "Preserving location privacy in wireless LANs," in *5th International Conference on Mobile Systems, Applications, and Services (MobiSys 2007)*, pp. 246–257, 2007.

[84] M. Gruteser and D. Grunwald, "Enhancing location privacy in wireless lan through disposable interface identifiers: A quantitative analysis," *Mobile Networking Applications*, vol. 10, pp. 315–325, June 2005.

[85] P. Bahl, V. N. Padmanabhan, and A. Balachandran, "Enhancements to the radar user location and tracking system," *Microsoft Research Technical Report*, 2000. http://research.microsoft.com/bahl/papers/pdf/msr-tr-2000-12.pdf.

[86] K. Chetty, G. Smith, and K. Woodbridge, "Through-the-wall sensing of personnel using passive bistatic WiFi RADAR at standoff distances," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, pp. 1218 –1226, april 2012.

[87] F. Adib and D. Katabi, "See through walls with Wi-Fi!," in *ACM SIGCOMM*, pp. 75–86, 2013.

[88] M. Maheshwari, S. Ananthanarayanan, A. Banerjee, , S. K. Kasera, and N. Patwari, "Detecting malicious nodes in RSS-based localization," in *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pp. 1–6, 2011.

[89] N. Patwari and P. Agrawal, "Effects of correlated shadowing: Connectivity, localization, and RF tomography," in *IEEE/ACM International Conference on Information Processing in Sensor Networks (IPSN'08)*, pp. 82–93, April 2008.

[90] J. Wilson, N. Patwari, and F. G. Vasquez, "Regularization methods for radio tomographic imaging," in *2009 Virginia Tech Symposium on Wireless Personal Communications*, 2009.

[91] N. Patwari and J. Wilson, "RF sensor networks for device-free localization: Measurements, models, and algorithms," *Proceedings of the IEEE*, vol. 98, no. 11, pp. 1961–1973, 2010.

[92] Y. Zhao and N. Patwari, "Histogram distance-based radio tomographic localization," in *11th International Conference on Information Processing in Sensor Networks (IPSN'12)*, pp. 129–130, 2012.

[93] V. Koster, A. Lewandowski, and C. Wietfeld, "A segmentation-based radio tomographic imaging approach for interference reduction in hostile industrial environments," in *IEEE/ION Position, Location, and Navigation Symposium (PLANS)*, pp. 1074–1081, IEEE, 2012.

[94] Y. Zhao and N. Patwari, "Noise reduction for variance-based device-free localization and tracking," in *8th IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON'11)*, June 2011.

[95] J. A. Costa, N. Patwari, and A. O. Hero III, "Achieving high-accuracy distributed localization in sensor networks," in *IEEE International Conference on Acoustic, Speech, & Signal Processing (ICASSP'05)*, pp. 642–644, March 2005.

[96] J. A. Costa, N. Patwari, and A. O. Hero III, "Distributed multidimensional scaling with adaptive weighting for node localization in sensor networks," *IEEE/ACM Transactions on Sensor Networks*, vol. 2, pp. 39–64, Feb. 2006.

[97] J. Wilson and N. Patwari, "A fade level skew-Laplace signal strength model for device-free localization with wireless networks," *IEEE Transactions on Mobile Computing*, vol. 11, no. 6, pp. 947–958, 2012.

[98] M. Moussa and M. Youssef, "Smart devices for smart environments: Device-free passive detection in real environments," in *IEEE International Conference on Pervasive Computing and Communications (PerCom'09)*, pp. 1–6, 2009.

[99] J. M. Phillips and S. Venkatasubramanian, "A gentle introduction to the kernel distance," *preprint*, 2011. arXiv:1103.1625.

[100] A. Edelstein and M. Rabbat, "Background subtraction for online calibration of baseline RSS in RF sensing networks," *IEEE Transactions on Mobile Computing*, vol. 12, no. 12, pp. 2386–2398, 2013.

[101] N. Patwari, R. J. O'Dea, and Y. Wang, "Relative location in wireless networks," in *IEEE Vehicular Technology Conference (VTC)*, vol. 2, pp. 1149–1153, May 2001.

[102] O. Kaltiokallio, M. Bocca, and N. Patwari, "Enhancing the accuracy of radio tomographic imaging using channel diversity," in *9th IEEE International Conference on Mobile Ad Hoc and Sensor Systems*, 2012.

[103] D. Cheung and C. Prettie, "A path loss comparison between the 5 GHz UNII band (802.11a) and the 2.4 GHz ISM band (802.11b)," *Intel Labs Report*, 2002.

[104] J. Zyren and A. Petrick, "Tutorial on basic link budget analysis," *Application Note AN9804, Harris Semiconductor*, 1998.

[105] R. Wilson, "Propagation losses through common building materials 2.4 GHz vs 5 GHz," 2002. http://www.magisnetworks.com.