ASSURING ROBUSTNESS OF RADIO SPECTRUM TELEMETRY AGAINST VANDALISM
AND EXPLOITATION

BY

OMID FATEMIEH

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2011

Urbana, Illinois

Doctoral Committee:

Professor Carl A. Gunter, Chair, Director of Research
Professor Tarek Abdelzaher
Professor Nikita Borisov
Doctor Ranveer Chandra, Microsoft Research

Abstract

The emerging paradigm for using the wireless spectrum more efficiently is based on enabling secondary users to exploit white-space frequencies that are not occupied by primary users. An enabling technology for forming networks over white spaces is distributed spectrum measurement to identify and assess the quality of unused channels. This spectrum availability data is often aggregated at a central base station or database to govern the usage of spectrum. This process, also referred to as *radio spectrum telemetry*, is vulnerable to integrity violations if the devices are malicious and misreport spectrum sensing results. There may be nodes that seek to *exploit* a spectrum in a given region by falsely reporting that a primary signal is present, or, dually, seek to *vandalize* a primary by reporting that its signal is not present, thereby encouraging interference from secondaries.

This dissertation focuses on assuring robustness of radio spectrum telemetry against exploitation and vandalism attacks. This problem is particularly challenging when: (1) attackers are omniscient and coordinated, and constitute a large fraction of the nodes in an area, and (2) the quantity under measurement (signal power) faces natural spatial and temporal variations, as well as uncertainties due to noise, shadowing, and fading. These circumstances make it easier for sophisticated attackers to hide the likely abnormalities of their reports. As illustrative examples, this work investigates two new applications for white-space networks; the communications of the advanced meter infrastructure (AMI) and broadband Internet access for public school students. These applications are utilized to underline the importance of the considered security attacks in practical settings. In addition, the thesis offers communication architectures for these applications and shows the practical, economical, and societal benefits.

This work formulates the problem of robust radio spectrum telemetry using a grid-based model and offers a range of solutions. The solutions include (1) model-based techniques that probabilistically detect abnormalities using knowledge about signal propagation and shadowing formulas, (2) data-based techniques based on machine learning classifiers that do not assume prior knowledge about signal propagation models and only rely on direct training data, and (3) trust-based techniques that use a small subset of remotely-attestable nodes as a foundation for trust, and subsequently deter attacks using a combination of statistical sequential estimation and classification techniques. The proposed techniques are evaluated using a novel methodology that relies in part on predicted propagation data derived from real-world registered TV transmitter and terrain data (from the databases of the FCC and NASA) for areas in Illinois and Pennsylvania.

To Mother, Father, Tannaz, and Parya.

Acknowledgments

I am grateful to my advisor, Professsor Carl A. Gunter, for his patient academic guidance and support in the past 6 years. He was the main reason I became interested in the field of computer security. I significantly benefitted from the breadth and depth of his knowledge, insights, and vision, as well as his principled approach to student training. His contributions to my PhD research and this thesis cannot be overstated. In addition, his encouragements and moral support made a big difference in my academic progress, and life.

I would like to thank Dr. Ranveer Chandra for his contributions as a thesis committee member and research collaborator since 2008. His expertise and help in the area of white-space cognitive radio networks proved invaluable. His vision and ideas particularly influenced Chapters 1 and 5. I would also like to thank my other thesis committee members, Professor Tarek Abdelzaher and Professor Nikita Borisov, for their valuable feedback on this work during the thesis proposal exam.

Special thanks are due to Michael LeMay, Ali Farhadi, and Amir Nayyeri, for conversations and collaborations. Michael LeMay helped with remote attestation (Chapter 6), Ali Farhadi offered assistance on the subjects related to machine learning and classification (Chapter 5), and Amir Nayyeri was a great help throughout, particularly while writing Chapter 4 and preparing for the thesis proposal exam.

I would like to thank my uncle, Dr. Farid Kianifard, for being a source of motivation, encouragement, and academic guidance throughout my college education. In addition, he always offered valuable technical assistance on topics related to probability and statistics, for which Section 6.1.3 is a prime example.

I would also like to thank my group members and collaborators in the past and present: Hossein Ahmadi, Dr. Rakesh Bobba, Jodie Boyer, Dr. Michael Greenwald, Sonia Jahid, Hee-Dong Jung, Dr. Albert Harris, Prof. Susan Hinrichs, Fariba Khan, Arindam Khan, Prof. Sanjeev Khanna, Dr. Lars Olson, Ravinder Shankesi, Dr. Himanshu Khurana, Prof. Manoj Prabhakaran, Prof. Santosh Venkatesh, Reza Zamani Nasab, and Dr. Jianqing Zhang.

My fiance, Parya, has been a tremendous help. She offered invaluable emotional support and encouragement, as well as significant technical assistance through conversations and proofreading.

And above all, I am grateful to my parents and grandparents for their unconditional love and support; without them, I would have never made it this far.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The proliferation of smartphones, and a subsequent demand for wireless Internet services, has highlighted the scarcity of spectrum for data communications. The Cellular Telecommunications Industry Association (CTIA), which includes AT&T and Verizon, recently requested the Federal Communications Commission (FCC) to grant an additional 800 MHz of spectrum for data communications by 2015 [13]. However, nearly all the spectrum that is ideal for long-range data communications, *i.e.*, between 300 MHz and 3 GHz, has been allocated to various primary users.

The FCC's recent 'white space' ruling, which allows unlicensed devices to operate in unused TV spectrum is a significant step towards alleviating this spectrum crunch. White spaces refer to portions of spectrum that have been allocated to licensed users but are not in use at that time. Devices determine if a TV channel is not in use at their location before using it to send and receive data. This ruling has met with enthusiasm from industry, academia, and policy makers. The key reason for this enthusiasm is two-fold. White spaces not only provide additional spectrum, they also enable long-range communication since they are in the lower frequencies (below 700 MHz).

An important functionality when forming networks over white spaces is the *aggregation of spectrum availability data* from multiple white space devices. The need for aggregation arises in several contexts. First, nearly all existing standards or proposals for white space networks, *i.e.* CogNeA, IEEE 802.22, IEEE 802.11af and WhiteFi [1,4,18], require the white space base station to receive spectrum availability reports from clients and operate on TV channels that are available at all nodes in the network. The spectrum reports from clients can be very diverse, since white space networks are expected to span a radius of up to 100 km [70]. Second, it has been shown that aggregating spectrum sensing data from other devices (also called collaborative sensing) enables white space

1

devices to sense at a higher threshold than when sensing alone. This is very useful since sensing at low thresholds is very challenging [36, 74]. Finally, aggregation of spectrum sensing data from white space devices can be used to build a nationwide database of spectrum availability across locations [31]. This is similar to Wi-Fi wardriving data, and can be used for several purposes, for example to improve the accuracy of the white space geo-location database that is being mandated by the FCC [2, 3].

A threat to aggregating spectrum sensing reports, also referred to as *radio spectrum telemetry*, is that some nodes may maliciously report inaccurate data. There may be nodes that seek to *exploit* a spectrum in a given region by falsely reporting that a primary signal is present, or, dually, nodes that seek to *vandalize* a primary by reporting that its signal is *not* present, thereby encouraging interference from secondaries. The first attack denies the legitimate users' access to the spectrum and provides exclusive access to attackers, whereas the second attack creates chaos and interference for primary and secondary users. Detecting these attacks is particularly challenging when (1) attackers are coordinated and sophisticated and constitute a large fraction of nodes in an area, and (2) the quantity under measurement (signal power) faces natural spatial and temporal variations, as well as uncertainties due to noise, shadowing, and fading. These circumstances make it easier for attackers to hide the potential abnormalities of their reports.

Countermeasures to prevent mischief are a key enabling technology for white space networks. Existing strategies have focused on an instance of this problem – in the context of collaborative sensing – for the detection of malicious nodes by identifying them as abnormal or outlier nodes within a small 'cell' [26, 45, 59]. If one divides a service region into cells of sufficiently small size, then nodes within a given cell can be expected to give similar readings. If a preponderance of nodes in a given cell provide a reading in a common range, then other readings may be discarded as outliers. Ideally this will prevent malicious nodes from being effective. Unfortunately, this strategy suffers from a key drawback; there is the possibility that a given cell will have a preponderance of malicious nodes. In addition, the solutions are often based on unrealistic assumptions about the models and parameters of signal propagation, depend on detection threshold parameters which are

usually tuned manually, or are often too conservative and not able to detect nimble manipulations of data by sophisticated attackers.

In this thesis, we provide a range of mechanisms by which, despite the existence of coordinated malicious false reporting attacks, spectrum measurement data can be robustly aggregated. The following *thesis statement* summarizes what this work aims to prove:

Radio spectrum telemetry can be made quantifiably robust against coordinated malicious misreporting attacks through the use of security-aware protocols and feasible trust infrastructure.

In Chapter 2, we lay the foundation for addressing the problem. First, we provide background information on white space networking, spectrum sensing, and remote attestation. Second, we describe the general setting, assumptions, and problem statement for the rest of the thesis. In particular, we define the range of attacker models we consider throughout the thesis. This includes non-collaborating adversaries who act individually (also *uncoordinated*), collaborating adversaries who act as a group (also *coordinated*), and *omniscient* adversaries who act as a group and possess complete knowledge of the defense mechanism and sensor data, including that of the non-adversaries.

In Chapter 3 (also [30]), we elaborate on *Advanced Meter Infrastructure (AMI)* communications and broadband access for public school students as two applications for white-space networking. We provide communication architectures for using white spaces, and show the practical, economical, and societal benefits. We also investigate the associated security issues, and further motivate the importance of robust radio spectrum telemetry in this context.

In Chapter 4 (*model-based* protection [31]), we propose viewing the area of interest for detecting primary presence (or absence) as a grid of square cells and use it to identify false reports. The proposed mechanism starts by identifying outlier measurements inside each cell and 'punishing' them. The punishment is in the form of exclusion or a low weight assignment in the proposed weighted aggregation process. The mechanism proceeds by corroboration and merging of neighboring cells in a hierarchical structure to identify cells with outlier aggregates, as a sign of significant malicious

3

node presence in a cell. The solution uses a simple model based on exponential decay and log-normal distribution to account for the uncertainties in signal propagation. The chapter includes a novel framework for quantifying the expected legitimate variations in measurements, which systematically reduces the likelihood of inaccurate classification of valid measurements as outliers. We use simulations to evaluate the effectiveness of the proposed approach against attackers with varying degrees of sophistication. The results show that depending on the attacker type and the distance from primary to the region of interest, in the worst case we can nullify the effect of up to 41% of attackers nodes. This figure is as high as 100% for areas that are not near the border of primarys protection region.

In Chapter 5 (*data-based* protection [32]), we address an important limitation of the model-based solution in Chapter 4. The model-based solution requires fairly accurate knowledge about the signal propagation formula and parameters. In addition, one of the introduced detection thresholds requires manual tuning, a strategy that is error prone and not easily scalable. Instead, we offer an alternative called CUSP (for Classification Using Signal Propagation) using which a central aggregation server can protect against malicious reports of spectrum availability. The key idea is to *let the data speak for itself*. CUSP uses natural signal propagation data in a region to learn a *classifier* that effectively understands the patterns of signal propagation in the region. It can then use the learned classifier to efficiently filter out the malicious spectrum reports as they often represent unnatural propagation patterns. We evaluate the performance of CUSP in detail. We drive our evaluation on predicted propagation data derived from registered digital TV stations and terrain data from the FCC and National Aeronautics and Space Administration (NASA), as well as house density data from the US Census Bureau. We compute the signal strengths from the TV stations for two regions in the states of Illinois and Pennsylvania. We find that our techniques are quite effective with all three types of attacks, but regional variations have an impact that must be properly addressed to assure consistent quality of detection. In particular, areas with hilly terrain and urban activity must be treated in smaller cells. The resulting approach is practical and effective for application in all areas and avoids arbitrary assumptions about models, parameters,

and thresholds in favor of direct training data.

In Chapter 6 (*trust-based* protection [33]), we initiate a new direction in reliable radio spectrum telemetry by relying on a small subset of nodes that can perform *remote attestation*. These nodes can securely attest their operating state to a remote server, and will be excluded if they are detected as compromised. Otherwise, they will be used as a foundation for security and reliability. To that end, we propose a practical framework for using data from both attested and regular nodes to deter attacks, while achieving precise results in the absence of attacks. More specifically, we explore a strategy based on statistical sequential sampling and inference to obtain an estimate for signal power in each small region. The sampling method uses data from all of the attested nodes, as well as the minimum required data from the rest of the nodes to achieve results with a pre-specified margin of error. Next, the data contributed by non-attested nodes is verified against data from attested nodes in the neighboring areas. This step is performed using SVM classifiers with quadratic kernels that are trained with an initial set of trusted data in the region of interest. We evaluate this scheme using predicted signal power data obtained from applying empirical signal propagation data on real-world TV transmitter and terrain data from the FCC and NASA databases. We instantiate the evaluations to a hilly urban/suburban area in Pennsylvania and measure the performance of our approach in the absence and presence of omniscient coordinated attackers. We show the scheme is effective against such attacks even in cases where only a small subset of the sensors can be remotely attested. In addition, we systematically enumerate the costs associated with remote attestation and shed light on these costs for prototypes based on Trusted Platform Modules (TPMs) and AVR32 microcontrollers. The data shows attestation may introduce non-trivial costs, which motivates our approach to leveraging attestation efficiently to establish trust in spectrum sensing results.

In Chapter 7, we provide a comprehensive review of the related work in two categories, white-space networks, and sensor and ad-hoc networks. In each category, we enumerate the closest pieces of related work and contrast them with this thesis. Finally, in Chapter 8 we conclude the thesis and identify areas of future work.

To summarize the original contributions, this thesis:

- Identifies and formulates a key threat to distributed spectrum measurements in white space networks; attacks in which omniscient and coordinated attackers report false spectrum sensing results in order to obtain exclusive spectrum access (exploitation) or create chaos (vandalism).

- Proposes two novel applications for white spaces, namely advanced meter infrastructure communications and Internet access for public school students, and further motivates the significance of exploitation and vandalism attacks through the lens of these applications.

- Provides a comprehensive treatment of the considered attacks using a grid-based model and offers three general solutions: (1) model-based techniques that probabilistically detect abnormalities using knowledge about signal propagation and shadowing formulas; (2) data-based techniques based on machine-learning classifiers; and (3) trust-based techniques that rely on a small subset of trusted nodes.

- Introduces a novel method to build classifiers from location-tagged signal propagation data. This obviates the need for relying on closed-form formulas, models, parameters, and thresholds when analyzing signal propagation data. This approach detects misreporting attacks in the process of centrally aggregating spectrum sensing data by building SVM classifiers.

- Creates a new direction in secure radio spectrum telemetry against coordinated misreporting attacks that relies on a small subset of attestation-capable sensors, and offers a practical framework for using statistical sequential estimation coupled with classifiers to deter attacks and achieve quantifiably precise outcome.

- Presents a novel way to evaluate white-space applications using real-world transmitter and terrain data.

Chapter 2

Background and Problem Formulation

In this chapter we provide background information on white space networking, spectrum sensing, and remote attestation. Next, we describe the general setting, assumptions, problem statement, and attacker models considered in the rest of the thesis.

2.1   Background

2.1.1   White Space Networks

On November 4, 2008 (and subsequently on September 23, 2010) the FCC made historic rulings that allowed unlicensed devices to operate over the licensed TV bands [2,3]. Wireless communications in this spectrum (below 700 MHz) benefit from favorable signal propagation and penetration properties, which enable long transmission ranges. The opening of these bands for unlicensed use represents the first significant increase in unlicensed spectrum below 5 GHz in over 20 years, and is expected to promote more efficient spectrum use.

Access to this spectrum could enable more powerful public Internet connections ('super Wi-Fi' hot spots) with extended range, fewer dead spots, and improved individual speeds as a result of reduced congestion on existing networks. Many other applications are possible, such as broadband access to schools particularly in rural areas, campus networks that are better able to keep pace with user's increasing demands for bandwidth, home networks that are better able to support real time streaming video applications, remote sensing of water supplies by municipalities and support for the communications of the advanced meter infrastructure in the smart grid [30].

A number of TV band device applications are already operating on an experimental basis. The

city of Wilmington North Carolina is trialing 'smart city' applications, including public 'hot spots,' low-cost broadband to a low-income housing development, and water level and water purity sensors for compliance with Environmental Protection Agency requirements and flood controls. In addition, a demonstration project in Claudville Virginia is bringing broadband access to a rural elementary school, as well as to consumers in their homes, and newly established public hot spots in the community. Plumas County California has undertaken a 'smart grid' trial for electricity networks, which allows the electric cooperative to manage the electrical system, obtain data from substations, and manage power flow. The network in that trial also enables free energy monitoring tools that allow consumers to save energy and money, for example, by identifying appliances that are always on and using energy [3].

2.1.2 Spectrum Sensing and Aggregation

Sensing the spectrum to identify unused channels can be used to improve the performance of white space networks. This is despite the FCC's September 2010 ruling which exempts the devices that incorporate geo-location and can access a new TV band *database* from mandatory spectrum sensing: (1) The ruling still allows for operation of sensing-only devices that cannot or do not access the database. (2) The database is built from conservative propagation models, which results in declaring many unused channels as occupied in places far from the transmitters. Real-time spectrum sensing data can provide a more accurate view of spectrum availability, or be used to improve the database results. (3) In places where multiple channels are available, the spectrum sensing details can reveal the highest quality channels for communications.

Energy detection is the most popular approach for signal detection. This is often attributed to its simplicity and small sensing time (less than 1ms). An energy detector measures the signal power on a target frequency and compares it against a *detection threshold* λ to determine whether a primary is present. For example, in the case of primary digital TV (DTV) transmitters, FCC has mandated -114 dBm as the detection threshold [2]. If a specific signature of a signal such as pilot, field sync, or segment sync is known, the more sophisticated *feature detectors* may be employed

to detect primary signals. Feature detectors are often more accurate, but are more complex to implement, and require additional information and sensing time (up to 24ms) [36, 47].

There exist three scenarios for centrally aggregating spectrum sensing results from sensors in a large region.

- Spectrum sensing data from deployed spectrum sensors or volunteer (mobile) white-space devices can be used to build a regional or nationwide spectrum availability database. Such a database can be used to augment the white space geo-location database mandated by the FCC, or to learn spectrum usage as part of the recently passed Spectrum Inventory Bill [8].

- A white-space service provider or base station may collect spectrum sensing data to determine areas of primary presence from cognitive radios in its network. This centralized approach has been endorsed by the IEEE 802.22 WRAN standard draft [4], CogNeA [1] and recent research publications [18]. The spectrum sensing data collected by the service provider may be provided by not only in-network cognitive radios, but also deployed spectrum sensors, and additional volunteer (mobile) devices to determine areas of primary presence.

- Collaborative Sensing, which refers to the process of combining spectrum sensing results from cognitive radios for the purpose of primary detection. The main benefit of this approach is the mitigation of multi-path fading and shadowing effects, which improves the detection accuracy in highly shadowed environments [36]. In addition, it allows for relaxation of sensitivity requirements at individual CRs [74].

To capture the common nature of the above scenarios, we focus on the case of building a regional spectrum availability database by a service provider. The database may then be combined with databases from other regions to form a nationwide database of spectrum sensing. The spectrum sensing data used to populate the database is provided by one or more of the following sources.

- *Volunteer Radios:* a set of (mobile) devices with different owners. The data would be collected by a modern 'mobile app' built to perform spectrum sensing at its current location and

9

report the results to a central server. This form of participatory sensing is often referred to as *crowdsourcing* in this context.

- *In-Network Cognitive Radios:* cognitive radios that are part of the service provider's network.

- *Dedicated Sensors:* sensors (in the form of a wireless sensor network) deployed for the specific task of spectrum sensing alongside the main white-space network [27].

At a fine-grained level, there exist two broad classes of strategies for combining individual spectrum sensing reports (within a small region, or cell). *Soft-combining* techniques combine raw signal power measurements from CRs, whereas *hard-combining* techniques combine binary decisions from CRs. Note that directly combining individual results happens only within small cells where nodes are expected to provide similar readings.

One of the most popular methods for soft-combining is Equal Gain Combining (EGC). In EGC, each node N_i of the m nodes inside a small area periodically provides its signal power measurement p_i to the central server. Assuming a vector of received power observations $(p_1, p_2, \ldots, p_m)$, and a nominal Gaussian model for shadowing and multi-path distribution, EGC is the maximum likelihood detector. It simply averages the power measurements of individual nodes and compares it to a detection threshold λ. That is, the primary is present if $P_{\text{avg}} = \frac{1}{m} \sum_{i=1}^{m} p_i \geq \lambda$. The threshold λ is determined based on the power of the transmitter and the radius around it, r, that needs to be protected. This is done such that the probability of missed detection stays below a threshold (*e.g.* .95), while the probability of false alerts are minimized [72]. EGC is known to have near-optimal performance in a diverse set of fading channels with more realistic assumptions [71].

2.1.3   Remote Attestation

Remote attestation is a technique for a system to provide certified information about its operating state (*i.e.* software, firmware, or configuration) to a remote party. This process is typically initiated by a request from the remote party. Upon receipt of the request, the queried system creates

a (signed) record of the system's operating state and sends it to the initiator. To securely record and certify its current state, the system needs to contain a number of components. Trusted hardware components are often used to this end, although software can also be used in some cases. Regardless, remote attestation imposes additional computational, storage, energy, time, and potentially manufacturing costs on both parties. On desktop PCs, the Trusted Platform Module (TPM) is often used to provide remote attestation functionality. The Trusted Computing Group (TCG) is developing trusted computing standards specifically for mobile devices to minimize costs and support appropriate usage models, and have specified several primitives for a Mobile Trusted Module (MTM). MTMs are expected to be available for many new mobile applications in the near future [10]. Previous work has also shown that remote attestation can feasibly be implemented in software on-chip for embedded processors such as AVR32 micro-controllers [53].

2.2 Setting and Problem Statement

We consider building a spectrum availability database from received signal power data from a combination of volunteer radios, in-network cognitive radios, and deployed sensors. Unless specifically differentiated, we refer to all of them as nodes or sensors in the rest of this paper. Due to its widespread adoption, ease of implementation, and small sensing time, we consider sensors to be energy detectors [15, 70]. We also assume the primary signal faces path loss and shadow fading due to irregular terrain and obstacles such as trees, buildings, walls, and windows.

The spectrum availability database represents the region of interest as a grid of small cells (or *tiles*) on the map of the region. Each cell may be a 1km × 1km square and is the unit in which combining individual results, or collaborative sensing, occurs. Within a tile, we combine the raw signal power measurements from nodes to determine primary presence (as opposed to binary yes/no results). This allows for using signal power as a measure of quality among the available channels and enables us to detect misreporting attacks. Unless otherwise specified, we consider EGC to be the method for combining individual measurements in each cell. EGC simply averages the power

measurements of individual nodes in each frequency channel and compares it to a *detection threshold* λ. In the case of primary Digital TV (DTV) transmitters, we often use the FCC-mandated -114 dBm as the detection threshold.

2.2.1   Exploitation and Vandalism

We are specifically focused on addressing robust aggregation of spectrum measurements in presence of compromised nodes among the distributed set of nodes. An attacker may compromise a (large) subset of the nodes and make them act in cooperation in order to change the spectrum sensing outcome. For example, they may seek to change the primary signal power for a tile (cell) in the database from a value below threshold (*-120 dBm*) to a value above threshold (*-100 dBm*), or vice versa. The first attack is called *exploitation*, and the second is called *vandalism*.

In Chapters 4 and 5 we assume no prior knowledge about the legitimacy of nodes, and therefore we mostly focus on detecting attacker nodes, or attacker-dominated cells using (the irregularities) in their measurements. However, in Chapter 6 we assume an additional means to establishing trust; we assume that a small subset of nodes are able to perform remote attestation. For any such *attestation-capable* node, the aggregation server can detect whether it is compromised and thus running illegitimate code. In that chapter, we investigate ways to efficiently and effectively use this capability to obtain reliable spectrum sensing results. This question is particularly important when the attestation-capable nodes constitute a small fraction of the nodes. This may be due to the low penetration of the technology among the volunteer nodes, or cost considerations of deploying *and* using this capability by the service providers in the deployed sensor scenarios. Regardless of the cause, it is desirable to achieve highly reliable spectrum aggregation results using only a small set of attestation-capable nodes.

2.2.2 Attacker Sophistication

We consider the following general attacker models throughout the thesis. Note that the attackers' behavior should be considered through the lens of a particular cell that the attackers aim to dominate. The exact details of domination depend on the combining rule. For example, for the EGC rule it involves changing the average signal power from a status indicating primary absence to one indicating primary presence, or vice versa. In addition, note that due to differences in the proposed defense mechanisms (model-based in Chapter 4, data-based in Chapter 5, and trust-based in Chapter 6) the specific instantiations if these models may vary slightly. The details of such instantiations and differences will be explained in corresponding chapters.

1. **Uncoordinated** attackers do not have precise information on the number and power measurements of other legitimate or attacker nodes in the cell. Each attacker node aims to dominate the cell without cooperation with other attackers, if any. This may be due to lack of information, unavailability of communication channels, or to reduce the likelihood of being detected as a result of communicating with peers. In this case, a compromised node that senses a signal power below (above) the detection threshold may falsely report a value such that the average power in the cell changes to a value below (above) the detection threshold. The attacker may use rough estimates of the number and measurements of other nodes for this purpose (for example, for the latter it would be a close value to the attacker's true measurement).

2. **Coordinated** attackers do not know the number and power measurements of the legitimate nodes in the cell, but may roughly estimate them. They do, however, know their own number and measurements, and act according to a coordinated strategy; they collude and use the estimates to calculate the value that each of them should report so that they can dominate the cell and change the detection outcome to a value above (or below) threshold.

3. **Omniscient** attackers are coordinated attackers that know the exact number and measurements of other legitimate users. Therefore, they can simply calculate the exact power levels they should report to change the average power level to a value *slightly* above (or below) threshold, *e.g.* 1dB. In addition, when possible, we assume the attackers know the exact details of the deployed

defense mechanism(s) and can carefully craft their misreporting strategy to avoid or minimize their chances of being detected.

2.2.3 Additional Assumptions

While some of the nodes may be unreliable, malicious, or compromised insider attackers, we assume that each node maintains a secure link to the base station for sending spectrum sensing results, and that attackers are unable to fabricate nodes or identities arbitrarily ('Sybil' attacks [62]). The secure links can be realized using pre-shared keys or a PKI, which may also serve as a foundation for preventing Sybil attacks by being associated with the identity of each node. Alternatively, one can take the dual view that we aim to demonstrate a method that forces adversaries to discover and deploy a practical Sybil attack, which requires a higher level of sophistication on the attacker's side (*e.g.* faking multiple link layer addresses). We also assume that the locations of nodes are reliably known through GPS or other localization techniques and nodes do not misreport their locations. This assumption is easily achievable in two of the most popular proposed applications of white space networking that assume fixed nodes with known locations: (1) Internet access for consumer premises using IEEE 802.22 wireless regional area networks [70], and (2) advanced meter infrastructure communications [30]. In cases where the network contains untrustworthy mobile devices, secure localization and location verification techniques may be employed to ensure nodes' locations are not forged [23, 51, 54, 55, 64]. In addition, for exploitation, attackers do not gain any tangible benefit from misreporting their location. The above assumptions are common for the type of analysis we perform here [26, 45, 59]; if they are violated then additional protective measures are required.

Chapter 3

Illustrative Examples

In this chapter we propose two speculative applications for white-space networking: *Advanced Meter Infrastructure (AMI)* communications, and high-speed Internet for public school students. For each application, we provide a communication architecture using white-spaces and show the bandwidth, deployment, cost, and societal benefits. We also investigate security issues associated with the proposed architectures, and emphasize the importance of robust spectrum data aggregation[1].

3.1 Advanced Meter Infrastructure

Advanced Meter Infrastructure is an integral part of the recent smart grid initiatives. It refers to systems that measure, collect, and analyze energy usage and interact with smart (advanced) meters through some communication media. The reconfigurable nature and communication capabilities of smart meters allow for deploying a rich set of applications in the smart grid. Prime application instances are automated meter reading, outage management, demand response, electricity theft detection, and support for distributed power generation. The communication architecture for AMI must meet the needs of current and future applications in a cost-effective, scalable, reliable, and secure way. Of particular interest are two-way communications between the smart meters and service providers such as the utility companies. Figure 3.1 depicts a common approach to AMI communication in the existing deployments.

In this model, hundreds to thousands of meters form a mesh network using proprietary protocols in the public industrial scientific and medical (ISM) frequency bands. The mesh network is used to

---

[1]A substantial portion of the material in this chapter is adopted from Fatemieh, Chandra, and Gunter's recent publication [30].

Figure 3.1: Current AMI communication architecture.

route the data to an access point (often mounted on a telephone pole). The access point aggregates and relays data between the meters and the utility. This part is often performed using cellular data services such as GPRS or EVDO. This approach suffers from at least three shortcomings. First, the ISM bands are noisy and crowded in urban areas and not well suited to the distances needed in rural areas. Second, cellular links incur the extra expense associated with licensed bands. Moreover, there is considerable competition for this bandwidth in urban areas and limited availability in rural areas. Third, the use of proprietary mesh network technology reduces inter-operability and impedes meter diversity.

In this section we consider the idea of using white spaces as part of AMI. White space communications leverage licensed spectrum opportunistically when it is not being used by incumbent transmitters such as digital TV transmitters. We believe the high bandwidth and long transmission ranges offered by white spaces can provide substantial benefits to the AMI. To that end, we propose a two-layer architecture for AMI communication using a combination of standardized protocols and successful research prototypes. We show that the proposed architecture can address some limitations of the state of the art, particularly in terms of bandwidth, deployability, and cost. In addition, we investigate reliability and security issues associated with the proposed architecture.

Power Grids in Rural Regions   AMI provides a number of general benefits, some of which we enumerated earlier, but are there any benefits of white space AMI for rural regions beyond these

16

general benefits? We speculate here on at least one such possibility. Power grids in rural regions typically have the characteristic that loads are sparsely distributed along the power lines, often only at farm houses and machine sheds. Such loads must be metered and meters are expensive to monitor. Thus a 200 acre farm might have a half a mile of adjacent power line but have power only from a meter at the farm house on a corner of the property. If, for example, there is an electric fence on a remote part of the property then power must be supplied with a battery since it is impractical to run a power line from the meter. The value of power harvesting for military purposes is well recognized [6, 7] but civilian uses will require metering. If it were feasible to add meters and power links easily along the utility power lines then this problem could be significantly diminished. Generally power companies do not wish to add new meters for such low loads, but when the cost of collecting billing data is diminished by suitable wireless communications infrastructure (based on white space for instance), then costs can be contained and a valuable service becomes feasible.

3.1.1 Proposed Architecture

We believe the bandwidth, range, and cost improvements offered by white spaces can provide substantial benefits to AMI. Figure 3.2 depicts the proposed architecture for AMI communication. This architecture involves two types of wireless networks in a hierarchy. At the lower level, there are small-scale white space networks which are represented as small circles in the picture. A prime candidate for implementing such networks would be WhiteFi [18]. For simplicity we refer to the general class of such networks as WhiteFi in the rest of this chapter. Due to the favorable propagation characteristics of the TV spectrum, WhiteFi networks can easily expand in areas with radius of up to 2km, while using commodity Wi-Fi transmitters and conforming to FCC regulations. The WhiteFi networks are envisioned to be established and maintained by the utility companies.

At the upper level in the hierarchy there exist 802.22 networks that provide connectivity between WhiteFi access points and the utility company. As it will be shown below, this provides benefits in terms of cost and broadband penetration in rural areas, while the standardization improves interoperability. The 802.22 networks do not need to be operated by the utility companies. We envision
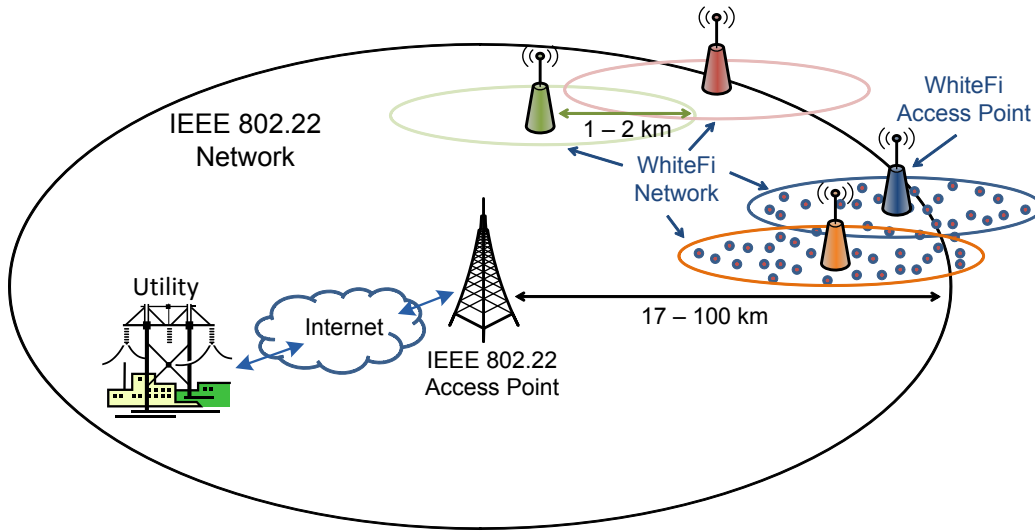
Figure 3.2: Proposed architecture for AMI communications over white spaces.

them to be operated by independent broadband service providers that offer service to a utility by admitting the utility's access points in their network. The 802.22 service provider may serve other clients such as residential households and mobile devices as well. This provides broadband access similar to ADSL and cable modems. The two-tier architectures provides a balance between independence and cost savings for the utility, while maintaining high data rate connections to the meters.

The large number and geographical separation of smart meters makes them a valuable resource for distributed spectrum sensing. The smart meters are owned by the utility, however, the 802.22 service provider can obtain the spectrum sensing data from the meters through the WhiteFi base stations. This reduces the number of spectrum sensing units the 802.22 service provider needs to deploy in order to build dynamic spectrum availability maps. This can be an important service that the utility can provide to the 802.22 service provider, and in exchange, receive low service rates and reimbursement for deploying meters with spectrum sensing capability.

In order to diversify its sources for spectrum sensing and avoid relying only on meters that are owned by another entity (the utility), the 802.22 service provider may need to collect spectrum sensing data from other means as well. Such sources could be other clients of the 802.22 service or sensors deployed specifically for this purpose [68]. In addition, in the case of availability of

transmitter geo-location databases, the list of available channels from both sources (*i.e.* spectrum sensing and transmitter databases) should be combined to derive the prioritized list of available channels. In all these scenarios, the WhiteFi and 802.22 base stations must coordinate their usage of the spectrum using co-existence techniques similar to those proposed in the 802.22 standard draft.

We argue that the proposed architecture provides the following benefits. First, compared to the state of the art, it allows for higher data rates at an economical cost for communication between the meters and the utility. Second, the penetration and long-range transmission properties in white spaces allow for direct communication between the meters and the (WhiteFi) access points. This obviates the need to form complex and unreliable mesh networks that consume considerable power for maintenance and routing. Third, it provides a valuable base of spectrum sensors (the smart meters) for the 802.22 service providers, which may lower their costs and improve their spectrum sensing. This will also provide a leverage to the utilities for discounts from the 802.22 service providers. In addition, this will result in better protection for primary transmitters, which has been the subject of substantial concern by FCC and spectrum license holders. Fourth, since the proposed solution provides cost savings and a revenue stream for 802.22 service providers, it contributes to the cause of providing affordable broadband service to rural communities. Fifth, since the approach insists on standardized protocols, it allows for inter-operability between products from different vendors.

One may consider the following limitations for the proposed approach. First, it requires a one-time cost of equipping smart meters with cognitive radios. The cost, however, may be small if the meters are produced at a large scale and could be covered by the spectrum sensing service they provide to the 802.22 provider. Second, there might be times or locations where no white space is available. In this case, the networks can temporarily operate in the ISM bands at lower bit rates. Therefore, in the worst-case scenario the performance would be similar to that of the existing architectures. Alternatively, a narrow band can be purchased at a small cost for emergency backup usage. Either of the above approaches guarantee that the network maintains minimum connectivity.

Third, there may exist various security concerns associated with the proposed architecture. These concerns are discussed in Section 3.3.

3.2 Public School Network

Providing free Wi-Fi broadband access in public areas or entire cities has been pursued by a number of municipalities around the world. Such efforts are often aimed at making wireless access to the Internet a universal service. While there exist a number of modestly successful instances (*e.g.* Luxembourg [14]), many attempts have failed due to assorted economical and technical challenges (*e.g.* Philadelphia [12]). One particularly important technical challenge is the limited range of service offered by commodity Wi-Fi base stations. This results in the need for deploying a large number of (routing-capable) base stations in order to form meshes, maintain connectivity, and provide extensive coverage.

In this section, we speculate on a sample application for achieving a similar goal to that of city-wide Wi-Fi. More specifically, we propose using white-spaces for providing city-wide (or neighborhood-wide) broadband Internet access to students enrolled in public schools; both at times when they are in school and out. In the case of small to mid-size cities, such a service may cover the entire city, while in the more urban metropolitan areas, the unit of coverage may be a neighborhood. This service increases broadband access for students, which can alleviate the 'digital divide' problem affecting impoverished communities, and potentially improve the quality of education.

3.2.1 Proposed Architecture

We believe the long range and favorable penetration properties offered by white spaces can play a key role in providing widespread mobile Internet access to public school students. To that end, we propose an architecture similar to that of AMI communications, which is illustrated in Figure 3.3. The large tower in the figure represents an 802.22-like long-range access point that is directly

connected to the Internet. These long-range access points extend Internet access to the WhiteFi base stations. The idea of using white spaces to provide the 'middle mile' connection between 802.22-like white space access points and short-range wireless hot-spots has been successfully deployed in rural Virgina [9]. In that instantiation the hot spots use Wi-Fi technology to reach to the end devices, which suffers from limited range and bandwidth. However, as discussed earlier in this chapter, each WhiteFi base station can provide coverage in areas up to 2km, and can potentially access large portions of unused spectrum to provide higher bit rates. Therefore, for this application we propose using white spaces for both the 'middle mile' and 'last mile' communications.
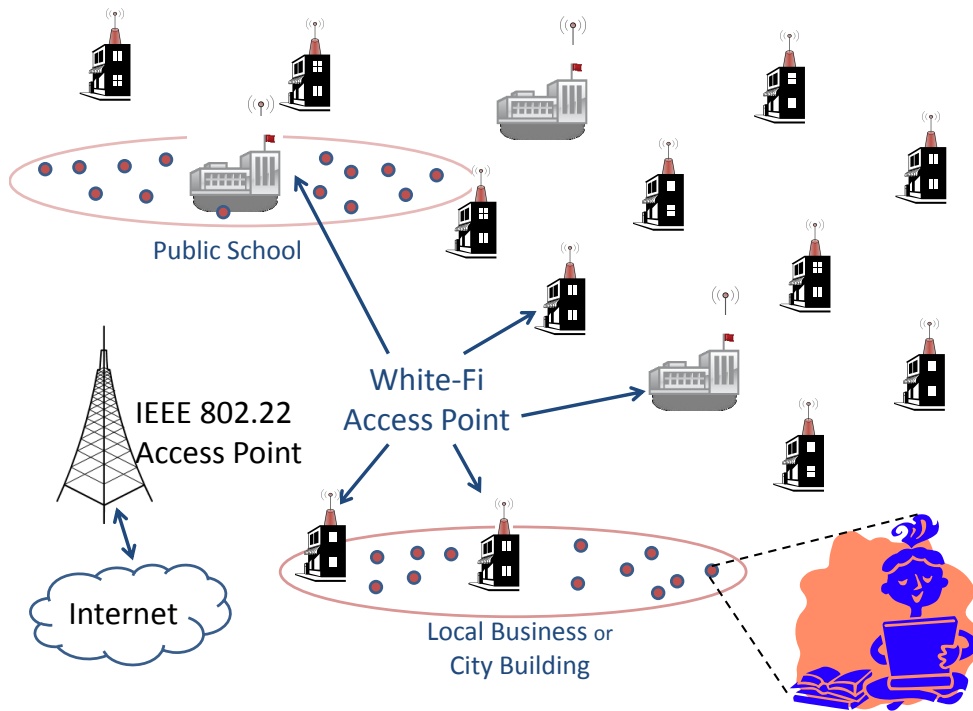
Figure 3.3: Proposed architecture for white space public school network.

As it can be seen in Figure 3.3, a small number of the WhiteFi base stations would be installed at the schools, while a larger portion of them would have to be installed throughout the area to provide comprehensive coverage. With proper prior configuration, the installation can be as simple as plugging an access point box to power. Plausible places for deploying access points would be city buildings and local businesses, which in return can be provided with window stickers showing their support for the local school system.

The end-users are students with (mobile) laptops or tablets that are equipped with cognitive radios; either built-in like WiFi, or through external PCMCIA cards or USB tokens. The students connect to their closest WhiteFi access point to obtain internet access. The authentication is performed through username and password provided by their school. Similar to AMI, on a periodic or on-demand basis, these devices perform spectrum sensing and report the results to the central (802.22) base station, or a regional/national spectrum availability database. Similarly, the WhiteFi base station can also perform spectrum sensing and report the results. The results are centrally aggregated to improve knowledge about spectrum availability, which is used to govern spectrum usage between the 802.22 access point and WhiteFi bases stations, as well as between WhiteFi base stations and end-devices. Therefore, the spectrum sensing task is effectively crowdsourced to the mobile student devices and a few WhiteFi base stations.

3.3   Security Threats

In both applications, primary emulation attacks can disrupt communications over white spaces. In a primary emulation attack, an attacker may modify the air interface of a CR to mimic a primary transmitter signal's characteristics, thereby causing legitimate secondary users to erroneously identify the attacker as a primary user, and abandon the channel. Of the body of existing work, LocDef utilizes both signal characteristics and location of the transmitter to verify a primary signal transmitters' location [27]. If it does not match the known locations for primary transmitters, the signal is from an attacker. This approach, however, requires knowledge of the location of the primary transmitter, and thus may not be practical in some circumstances.

An alternative is using cryptographic and wireless link signatures to authenticate primary users' signal in presence of attackers that may mimic the same signal [58]. This is achieved by using a helper node close to a primary user to enable a secondary user to verify cryptographic signatures carried by the helper node's signals and then obtain the helper node's authentic link signatures to verify the primary users signals. Liu *et al.* [57] also study the problem of detecting unauthorized

spectrum usage, where the authorized transmitter may be mobile. They propose two analytical methods and a solution based on machine learning to detect anomalous transmission by using the characteristics of radio propagation.

Another key threat, which is the focus of this thesis, is malicious false reporting of spectrum sensing results by end devices. The majority of end devices are cognitive-radio equipped advanced meters in the case of AMI, and mobile laptop or tablet PCs in the case of public school network. In both applications, spectrum measurement and reporting by end devices can significantly improve the operation of the network. The base stations collect spectrum sensing reports from the cognitive radios and store them in a database to govern the usage of the spectrum. In both cases, there exists the possibility that the software on end-devices be compromised by viruses or self-propagating worms. Therefore, the attacker may gain control of multiple such devices and use them as a platform to launch coordinated exploitation and vandalism attacks. For example, in the case of AMI an adversary may be interested in compromising the software on advanced meters in a neighborhood to make them falsely declare available spectrum as occupied. He can then exclusively use that spectrum for his/own benefit. Similarly, college students' computers may be manipulated by themselves, or their classmates, in order to perform vandalism attacks and create chaos and interference. Defense against these attacks is the subject of this thesis, and will be explored in the rest of this thesis.

Chapter 4

Model-Based Protection

In this chapter we propose viewing the area of interest for detecting primary presence (or absence) as a grid of square cells and use it to identify and disregard false reports. The proposed mechanism starts by identifying outlier measurements inside each cell and 'punishing' them. The punishment is in the form of exclusion or a low weight assignment in the proposed weighted aggregation process. The mechanism proceeds by corroboration and merging of neighboring cells in a hierarchical structure to identify cells with outlier aggregates, as a sign of significant malicious node presence in a cell. The solution uses a simple model based on exponential decay and log-normal distribution to account for the uncertainties in signal propagation. The chapter includes a novel framework for quantifying the expected legitimate variations in measurements, which systematically reduces the likelihood of inaccurate classification of valid measurements as outliers. We use simulations to evaluate the effectiveness of the proposed approach against attackers with varying degrees of sophistication. The results show that depending on the attacker-type and the distance from primary to the region of interest, in the worst case we can nullify the effect of up to 41% of attackers nodes. This figure is as high as 100% for areas that are not near the border of primarys protection region[1].

4.1 Grid-Based Hierarchical Approach

In this section we first introduce the soft-combining technique for collaborative sensing based on maximum likelihood estimation. Our approach is based on this method of collaboration among nodes. Next we provide a 'basic approach' which incorporates the basic ideas in the proposed

---

[1]The majority of the material in this chapter is adopted from Fatemieh, Chandra, and Gunter's recent publication [31].

scheme. We extend the basic approach to a 'weighted approach' which is the main protocol we use to evaluate our solution in Section 4.3.

4.1.1 Maximum Likelihood Detector

Consider a square grid consisting of $n \times n$ square cells. Each cell is the basic unit of collaborative sensing and we call it a *level 0 cell*, or simply *cell*. The dimensions of a level 0 cell C are denoted by $r_0 \times r_0$. Consider a level 0 cell containing m nodes. The outcome of sensing by node N_i is p_i, which represents an estimate of the received primary power at node N_i. In dB, this is written as $p_i = p_t - (10 \log_{10} r_i^\alpha + S_i + M_i)$ where p_t is the transmit power of the primary signal, r_i is the distance from N_i to the primary transmitter, $10 \log_{10} r_i^\alpha$ represents the signal attenuation with exponent α (typically $2 < \alpha < 4$), and S_i and M_i are losses due to shadowing and multipath fading. We adopt the log-normal shadowing model [63] and therefore consider S_i and M_i to follow a Gaussian distribution ($S_i + M_i \sim N(\mu_s, \sigma^2)$) on the dB scale. Therefore we have $p_i \sim N(\mu(r), \sigma^2)$, where $\mu(r) = p_t - (10 \log_{10} r_i^\alpha + \mu_s)$. For simplicity of analysis, unless otherwise noted, we consider μ_s to be 0 and σ to be independent of the distance to the transmitter [72].

Given a vector of received power observations $(p_1, p_2, \ldots, p_m)$ for this cell, a maximum likelihood (ML) detector would determine the primary presence by averaging the power measurements of individual nodes and comparing it to detection threshold λ [36, 72]:

$$\text{Primary is } \begin{cases} \text{Present,} & \text{if } P_{\text{avg}} = \frac{1}{\text{m}} \sum_{i=1}^{m} p_i \geq \lambda \\ \text{Absent,} & \text{otherwise.} \end{cases} \tag{4.1}$$

λ is determined based on the power of the transmitter and the radius around it, r, that needs to be protected. This is done such that the probability of missed detection stays below a threshold (*e.g.* .95), while the probability of false alerts are minimized. λ can be determined for a cell with m nodes as follows. If each measurement at distance r is distributed according to a normal distribution with mean $p_r = p_t - (10 \log_{10} r^\alpha)$ and standard deviation σ, we have $P_{\text{avg}} \sim N(p_r, \frac{\sigma^2}{m})$.

We can determine λ such that:

$$\mathbf{Pr}(P_{\text{avg}} \geq \lambda) = .95 \Rightarrow \lambda = \frac{\sigma}{\sqrt{m}}Q^{-1}(.95) + p_r \qquad (4.2)$$

where Q is the standard Gaussian distribution tail function and Q^{-1} is its inverse.

4.1.2 Basic Approach

It is easy to show that a few malicious nodes that report extremely high or low measurements can significantly skew the average in Equation 4.1, and thus alter the detection outcome. To that end, we propose a hierarchical structure for reducing or eliminating the effect of maliciously misreporting nodes. At the lowest level of the hierarchy (level 0) there exist level 0 cells. At higher levels of the hierarchy, each level l cell constitutes the area covered by b level $l-1$ cells that are adjacent. b is called the *branching factor* of the hierarchy and we assume $\sqrt{b}$ is an integer greater than one. Figure 4.1 provides an illustration for $b = 4$.

In simple words, our scheme first aims to detect outlier measurements inside a cell by peer comparisons. If the attackers compromise a large fraction of nodes in a cell, they effectively take over the call and may no longer be detectable as outliers in the cell. Therefore, we use the hierarchy to compare each cell's average with its neighbors to detect if it is unexpectedly high or low. This corroboration allows our protocol to identify 'outlier cells' with significant attacker presence.

Consider a level 0 cell C_j that contains m secondary nodes. We define a *dispute threshold* for level 0, d^0, as the maximum acceptable difference between the measurements of two nodes inside that cell. In Section 4.2 we provide a disciplined mechanism for deriving the dispute thresholds. As it will be shown, the dispute thresholds may vary for different cells. At level 0, pairwise comparisons between measurements of individual nodes are performed inside each cell. In each pairwise comparison between nodes N_i and N_j, if the difference is greater than d^0, N_i and N_j's dispute counts, c_i and c_j, are increased by one. After all pairwise comparisons, if $\frac{c_i}{m}$ is greater than or equal the *outlier threshold* for level 0, τ_0 ($0 < \tau_0 < 1$, *e.g.* $\tau_0 = .75$), the node is flagged as an

26

outlier and is excluded in the primary presence calculation in Equation 4.1. In other words, for a node *not* to be an outlier, at least a fraction $(1 - \tau_0)$ of the nodes in its cell should be within its d^0 distance. This method for outlier detection is an instance of the well-known *distance-based outlier detection* techniques in the literature. Formally, an object o in a data set D is a distance-based outlier with parameters *pct* and *dmin* if at least a fraction *pct* of the objects in D lie at a distance grater than *dmin* from o [40].

Figure 4.1: Cells of different levels in a hierarchy with branching factor $b = 4$. $C_i : L_j$ denotes cell C_i at level j.

The higher levels of the hierarchy are formed as follows. A collection of b adjacent level 0 cells form a $r_1 \times r_1$ level 1 cell, where $r_1 = \sqrt{b}r_0$. At this step, after discarding outliers, the average signal presence at each of the consisting level 0 cells is calculated. The b resulting averages are compared in a pairwise fashion, and at each comparison, the dispute count is increased for a level 0 cell that has a difference greater than the dispute threshold for level 1, d^1, with a neighboring cell. Again, after all comparisons, if a cell's dispute count divided by the number of cells (b) is greater than the outlier threshold ratio τ_1 ($0 < \tau_1 < 1$, *e.g.* $\tau_1 = .75$), the cell is flagged as an outlier and its result is considered unacceptable. The same procedure (averaging and neighbor comparison) is applied for up to l_{max} levels, and at each level if a cell j is flagged as outlier, all the cells it contains are flagged as 'indeterminate' for which primary presence cannot be accurately

determined. For example in Figure 4.1, if C_5 is an outlier, the primary presence at C_1, C_2, C_3, and C_4 is indeterminate. For indeterminate cells, we consider primary presence to be difficult to tell, in which case an alternative source of information or method should be used for decision-making. For example, if there exist out-of-band mechanisms for establishing high trust in a subset of nodes, we can rely only on the measurements of the (few but trusted) nodes in that region. We do not explore this particular method in this chapter, and leave it as an item of future work. Therefore, in our first set of simulations (see Section 4.3) we simply report these cells as indeterminate. However, in Section 4.4, once we provide other means (based on median) to identify indeterminate cells, we propose and evaluate a simple method based on the average of 8 surrounding cells for primary detection in indeterminate cells.

The following Propositions state the limits that the basic approach imposes on exploitation attacks. Similar results can be derived for vandalism attacks.

Proposition 1 *Consider a level 0 cell with dispute threshold d^0 and outlier threshold τ_0 under an exploitation attack. Let $\alpha < (1 - \tau_0)$ be the fraction of compromised nodes. If the average power of the un-compromised nodes and the average power including compromised nodes are denoted by m and m', under the basic approach we have: $m' \leq m + 2d^0\alpha$.*

Proposition 2 *Consider level i cells $C_l, \ldots, C_{l+b}$ with averages $m_l, \ldots, m_{l+b}$ that constitute the level $i+1$ cell C_t with dispute threshold d^{i+1}. Let the outlier threshold $\tau_j = (b-1)/b$ for all level $j > 0$ cells. Consider a level i cell $C_n \in \{C_l, \ldots C_{l+b}\}$ under an exploitation attack. In order for C_n with the attacker influenced average m'_n to stay undetected as an outlier under the basic approach, the following property should hold: $m'_n \leq \max_{k \in \{l,\ldots,l+b\}-\{n\}}(m_k) + d^{i+1}$.*

As an example for exploitation, consider a level 0 cell C_k (dispute threshold $= d^0_k$) with a fraction α of attackers. Assume the conditions of Propositions 1 and 2 hold, and $l_{max} = 2$. C_k is in level 1 cell C_l (dispute threshold $= d^1_l$). For ease of exposition, we represent all level 0 cells (excluding C_k) that are in C_l by $C_{k+1}, \ldots, C_{k+b-1}$. Also assume that C_l is in level 2 cell

C_m (dispute threshold $= d_m^2$). For ease of exposition, we represent all level 1 cells (excluding C_l) that are in C_m by $C_{l+1}, \ldots, C_{l+b-1}$. Propositions 1 and 2 provide the following constraints on the attacker influenced average for C_k, denoted by m_k':

(1) $m_k' \leq m_k + 2d_k^0 \alpha$,

(2) $m_k' \leq \max\limits_{k+1 \leq i \leq k+b-1} (m_i) + d_l^1$,

(3) $m_k' \leq b \max\limits_{l+1 \leq i \leq l+b-1} (m_i) + bd_m^2 - \sum\limits_{i=k+1}^{k+b-1} m_i$.

4.1.3 Weighted Approach

The basic approach may result in flagging a number of nodes, level 0 cells, and higher level cells as outliers. The outlier nodes are excluded in the averaging for their respective cells. Likewise, the outlier cells are excluded in the averaging at higher levels, and the primary presence status in them is considered indeterminate. We propose using the results of the basic approach to assign and update weights to individual nodes over time (at the end of each round). In a level 0 cell C_i, each node N_j is assigned a weight w_j such that $\sum_{N_j \in C_i} w_j = 1$. In a cell with m nodes, each node's weight is initialized to $\frac{1}{m}$. We do not assign weight to cells. At level 0, the weighted sum of node's measurements is compared to the detection threshold:

$$\text{Primary is} \begin{cases} \text{Present,} & \text{if } \sum_{i=1}^{m} w_i p_i \geq \lambda \\ \text{Absent,} & \text{otherwise.} \end{cases} \tag{4.3}$$

Outlier detection is performed similar to the basic scheme. The only difference is that cells (not nodes) that are flagged as outliers can be assigned a 'low' or 'high' label; if the average value at an outlier cell is considered too low compared to its peers, it is flagged as a *low-outlier*, otherwise it is a *high-outlier*. After all the outlier detection and averaging is performed (starting from level 0, up to level l_{max}), Algorithms 1 and 2 are used to update the weights of nodes for the next round. In these algorithms, functions LowOutlier(C) (HighOutlier(C)) are considered to return 'true' if C is

Algorithm 1 Determine Level 0 Cell Status
___
    Input: Level 0 cell C
    lowCount $\leftarrow 0$; *highCount* $\leftarrow 0$
    for each $C_i \in \big(Ancestors(C) \cup \{C\}\big)$ s.t. $\text{Level}(C_i) \leq l_{max}$
        if *LowOutlier*(C_i) **then**
            lowCount ++
        else if *HighOutlier*(C_i) **then**
            highCount ++
        end if
    if *highCount* + *lowCount* > 1 **then**
        UpdateWeights$(C,\text{'conflicted'})$
    else if *highCount* == 1 **then**
        UpdateWeights$(C,\text{'high'})$
    else if *lowCount* == 1 **then**
        UpdateWeights$(C,\text{'low'})$
    else   // Neither C nor any of its ancestors is an outlier
        UpdateWeights$(C,\text{'neutral'})$
    end if
___

a low-outlier (high-outlier) cell.

4.2   Dispute Threshold Calculation

The dispute thresholds introduced in Section 4.1 aim to define maximum 'reasonable' differences between the observed signal powers among nodes (or averaged measurements among cells), beyond which the differences are highly questionable. Deriving the thresholds entails identifying and analyzing the sources of such power differences. The observed signal strength p (in dBm) at a secondary node is determined by the power of the transmitted signal p_t minus losses in power due to (1) attenuation at a distance r from the transmitter $l(r)$, (2) shadowing S, and (3) multi-path fading M, that is $p = p_t - (l(r) + S + M)$ [72]. Therefore, in order to characterize the differences, we need to study the effects of these three factors. We study the problem of determining thresholds at two different levels: (1) *Intra-cell* dispute thresholds (d^0) that are used to compare individual power measurements between nodes in a level 0 cell (2) *Inter-cell* dispute thresholds ($d^i,\ \ 1 \leq i \leq l_{max}$) that are used to compare averaged measurements from each of the level $i-1$ cells contained in a level i cell.

Algorithm 2 UpdateWeights (C, *status*)

    Input: Level 0 cell C, and *status* $\in$ {'conflicted', 'high', 'low', 'neutral'}
    switch (*status*)
        case 'conflicted': **return**
        case 'high':
            sort the nodes in C based on power measurement
            cut the weights of the last 25% by half and equally
            distribute it to others in C
        case 'low':
            sort the nodes in C based on power measurement
            cut the weights of the first 25% by half and equally
            distribute it to others in C
        case 'neutral':
            cut the weights of the outlier nodes in C by half
            and equally distribute it to others in C
    end switch

4.2.1   Intra-Cell Dispute Thresholds

Consider honest nodes N_i and N_j in a level 0 cell at distances r_i and r_j from the primary transmitter.
Without loss of generality assume $r_j > r_i$. Therefore, $r_j = r_i + \Delta r_{i,j}$ ($0 < \Delta r_{i,j} \leq \sqrt{2}r_0$). Our
goal is to find a value d^0 such that with high probability (*e.g.* 0.9) we have: $p_i - p_j \leq d^0$. Assuming
independent, identically distributed (i.i.d.) Gaussian shadowing and fading at both nodes we have
$p_i \sim N(p_t - 10\log_{10}(r_i^\alpha) - \mu_s, \sigma^2)$ and $p_j \sim N(p_t - 10\log_{10}(r_j^\alpha) - \mu_s, \sigma^2)$. Therefore we obtain
the distribution of the difference as:

$$p_i - p_j = N(10\alpha \log_{10} \frac{r_i + \Delta r_{i,j}}{r_i}, 2\sigma^2)$$

For a fixed r_i, choosing $\Delta r_{i,j} = \sqrt{2}r_0$ maximizes the mean of the distribution. However, since
we do not know the exact location of the transmitter, we do not know r_i. In an ideal world
where α is accurately known, and there is no loss due to shadowing and fading, one can use
$p_i = p_t - 10\log_{10}(r_i^\alpha)$ to obtain r_i. We propose the following approach to estimate r_i in a more
realistic environment where α is not accurately known and the effect of shadowing and fading is
not negligible.

    In order to reduce the uncertainty due to α, we take a conservative approach by taking the value

of α that creates the largest attenuation from r_i to $r_i + \sqrt{2}r_0$. This is achieved by assuming a large α (*e.g.* $\alpha = 4$). In addition, the signal power, p_i, may have faced shadowing and fading. Therefore, p_i may not be the most valid choice for determining r_i. Since the size of a level 0 cell is relatively small compared to the distance to the transmitter, the average power reported by the nodes inside a cell may seem as an obvious candidate to estimate p_i. This average, however, is highly vulnerable to excessively high (or low) reports by malicious or deeply faded nodes. Therefore we opt for using the robust statistic of median [75] of the reported powers inside the cell for determining a conservative estimate of r_i. For a level 0 cell C, if p_{rep} is the representative power of this cell, and r_{rep} is the representative distance from this cell to the transmitter, we have:

$$p_{rep} = \text{median}(p_j), \text{ for all nodes } N_j \text{ in level 0 cell } C$$

$$p_{rep} = p_t - 10\log_{10}(r_i^\alpha) \Rightarrow r_i \sim r_{rep} = 10^{\frac{p_t - p_{rep}}{10\alpha}}$$

Therefore, if we aim to determine d^0 such that $\mathbf{Pr}\big(p_i - p_j < d^0\big) > .9$, we have:

$$p_i - p_j \sim N(10\alpha \log_{10} \frac{r_{rep} + \sqrt{2}r_0}{r_{rep}}, 2\sigma^2)$$

$$\mathbf{Pr}\big(p_i - p_j \geq d^0\big) \leq .1 \Rightarrow$$

$$Q\Big(\frac{d^0 - 10\alpha \log_{10} \frac{r_{rep}+\sqrt{2}r_0}{r_{rep}}}{\sqrt{2}\sigma}\Big) = .1$$

$$\boxed{d^0 = \sqrt{2}\sigma Q^{-1}(.1) + 10\alpha \log_{10} \frac{r_{rep} + \sqrt{2}r_0}{r_{rep}}}$$

where Q is the standard Gaussian tail probability function. Note that using this scheme, the dispute thresholds for different cells will likely be different. For future use, we denote $10\alpha \log_{10} \frac{r_{rep}+\sqrt{2}r_0}{r_{rep}}$ for a level 0 cell C_k by $\Delta\mu_k^{rep}$. We introduce the notation of d_k^0 to represent the dispute threshold for a level 0 cell C_k. We generalize this notation to represent the dispute threshold and average power for a level i cell C_k by d_k^i and p_k^i.

4.2.2   Inter-Cell Dispute Thresholds

For simplicity, we first provide details on how d_k^1, the dispute threshold for a level 1 cell C_k, is calculated. Then we generalize the obtained result to higher levels. Consider a hierarchy with branching factor b. After outlier nodes are detected, and (weighted) averages for level 0 cells are calculated, we advance to level 1. At level 1, we perform pairwise comparisons between averages provided by each of the b level 0 cells contained in C_k, identify and leave-out outliers, and average the values of the rest to be passed to level 2. Consider two neighboring level 0 cells C_i and C_j (in C_k), with corresponding computed average powers p_i^0, and p_j^0. Assume there are m nodes in each cell. We have:

$$p_i^0 \sim N(\mu_i, \frac{\sigma^2}{m}), \; p_j^0 \sim N(\mu_j, \frac{\sigma^2}{m})$$
$$p_i^0 - p_j^0 \sim N(\mu_i - \mu_j, \frac{2\sigma^2}{m})$$

Ideally, if we were absolutely sure about the integrity of the majority of the nodes in each of the cells C_i and C_j, we could have replaced μ_i and μ_j by the averages of the corresponding cells. However, either of the cells may be populated by a large number of malicious nodes in a way not detectable at level 0. Hence, either of the averages could be highly skewed. As a result, using the difference between the sample averages is not a safe way to determine the probability distribution of the difference. Otherwise, very high dispute thresholds may be created that allow attackers to hide their presence. Besides, for simplicity, we are interested in using only one dispute threshold for each level 1 cell (as opposed to one dispute threshold for each level 0 pair). We employ a similar strategy to the intra-cell case and estimate $\mu_i - \mu_j$ by: $\Delta\mu_k^{rep} = \sqrt{b} \times \text{median}(\Delta\mu_i^{rep})$, for all level 0 cells $C_i \in C_k$. We can generalize this method to any level greater than 0. Therefore, we have $p_i^0 - p_j^0 \sim N(\Delta\mu_k^{rep}, \frac{2\sigma^2}{m})$. If we aim to determine d_k^1 such that $\mathbf{Pr}(p_i^0 - p_j^0 < d_k^1) > .9$, we obtain:

33

$$\mathbf{Pr}\left(p_i^0 - p_j^0 \geq d_k^1\right) \leq .1 \Rightarrow$$

$$Q\Big(\frac{d_k^1 - \Delta\mu_k^{rep}}{\frac{\sqrt{2}\sigma}{\sqrt{m}}}\Big) = .1$$

$$\boxed{d_k^1 = \frac{\sqrt{2}\sigma}{\sqrt{m}}Q^{-1}(.1) + \Delta\mu_k^{rep}}$$

It can be easily shown that the same argument could be used for higher layers of hierarchy. Therefore, if we represent the dispute threshold for a level i cell C_k by d_k^i, we have:

$$\boxed{d_k^i = \frac{\sqrt{2}\sigma}{\sqrt{b^{i-1}m}}Q^{-1}(.1) + \Delta\mu_k^{rep}}$$

where $\Delta\mu_k^{rep} = \sqrt{b} \times \text{median}(\Delta\mu_j^{rep})$, for all level $i-1$ cells $C_j \in C_k$.

Note that the dispute thresholds do not depend on the detection threshold, λ. It is easy to verify that as we go up in the hierarchy, the mean of the distribution for determining the dispute threshold is increased, while its standard deviation is decreased. This stems from the fact that the mean of the distribution mainly represents variation due to signal power attenuation over distance, whereas the standard deviation represents variations due to shadowing, which (as expected) is reduced as a result of aggregating increasing number of individual measurements.

4.3   Simulation Study (Part 1)

In this section we first provide the simulation setup used for evaluating the proposed scheme. The attacker model, results, and a brief analysis of the results are followed.

4.3.1   Simulation Setup

The simulation environment is a 4096m $\times$ 4096m area in which secondary users are deployed uniformly at random with the density of 0.0008 per square meter. The branching factor, b, is 4.

34

The size of each level 0 cell is 128m $\times$ 128m, creating a total of 1024 level 0 cells. Therefore, the expected number of nodes per cell is about 13. A primary transmitter with transmission power of 50mW (17 dBm) is located at the center of the area to represent a wireless microphone [21]. We consider a circular area with radius 1000m around the primary as the area that needs to be protected. This represents the area in which the primary signal must be detected with high probability. In particular, we require that primary signal be detectable by collaborating nodes in a level 0 cell with probability greater than .95 (max false negative rate of 5%). Using the formulation in Equation 4.2, this translates to the detection threshold of $\lambda = -74.4$dBm. We set the attenuation exponent, α, to 3 [72], and the standard deviation for the fading and shadowing process, σ, to 3 (in dB scale) [36]. The dispute threshold for each cell is determined based on the framework proposed in Section 4.2. The outlier threshold for level 0 cells, τ_0, is 0.6, and for all $i > 0, \tau_i = \frac{b-1}{b} = .75$.

4.3.2 Attack Scenarios

We first study exploitation attacks. We pick two cells outside the protection radius of the primary transmitter. First cell is selected randomly in such a way that is located at a distance marginally greater than the protection radius. This choice helps us gauge the worst-case performance of our protocol. We call this cell the *borderline-outside* cell. Next we randomly select another cell with the constraint that it is located at about two times the protection radius of the transmitter. We call this cell the *well-outside* cell. In each scenario, the attacker has compromised a certain fraction of the nodes inside the cell. Compromised nodes work in cooperation to report values higher than their true measurements to change the detection outcome. For a given cell and attack type, we vary the fraction of compromised nodes and study the results. The compromised nodes' behavior is according to one of the following models.

- **Uncoordinated** attackers do not have any information about the number or measurements of others in their cell. They only know λ, and simply report measurements that are a fixed amount greater than λ.

- **Coordinated** attackers do not know the exact number or measurements of others. They know true measurements of themselves, the 'expected' number of nodes per cell, and λ. Assuming similar measurements by non-compromised nodes, they report measurements such that the estimated cell average is a few decibels (*e.g.* 4 dB) over λ. This to guarantee that if they underestimate the total number of nodes, they still succeed.

- **Omniscient** attackers know the number and measurements of all nodes in their cell, and λ. Using this information, they report measurements such that the final average for the cell is slightly over λ and not greater than $(p_{\mathrm{avg}} + d^0)$. Here, p_{avg} is the average power of the cell assuming honest reports, and d^0 is the dispute threshold for the cell. This helps attackers reduce the chances of being detected as outliers at level 0.

For vandalism attacks, similar to exploitation scenarios, we pick two cells inside the protection radius of the primary transmitter. One cell is randomly selected from cells at a distance marginally smaller than the protection radius of the primary transmitter. We call this cell the *borderline-inside* cell. We randomly select another cell located at about half the protection radius away from the transmitter. We call this cell a *well-inside* cell. Attacker strategies are defined similar to the exploitation case, except here they aim to lower the average power measurement below λ for their cell.

4.3.3 Results

Figures 4.2, 4.3, and 4.4 (pictures on left) depict the measured average and final detection outcome for exploitation attacks. Results are collected after running the simulations for enough number of runs so that the weights (and thus the final outcomes) are stabilized. Note that in all graphs the y-axis represents the weighted and outlier-excluded average of the power of the nodes in the cell. For indeterminate cells (low, high, or conflicted based on Algorithm 1; represented by a '$\times$'), we do not provide any disambiguation solution in this section. Later, in Section 4.4, we introduce and evaluate one such solution. The results from an unmodified ML detector are provided in the

captions for comparison.

It can be seen that for the well-outside cell, none of the attacker models can succeed. As a commonly observed pattern (except for omniscient attackers), when attackers constitute small fractions of the population in the cell, they are detected as outliers, and their weights are reduced. Therefore, they cannot move the cell average above the threshold. Once the attackers gain enough population to meaningfully increase the average without being individually detected, the entire cell is detected as an outlier at higher levels (level 1 here).

The picture is not as rosy in the case of the borderline-outside cell. It can be seen that once attackers obtain enough population (23% to 35% depending on the attacker model), they are able to successfully flip the detection outcome. This is not a surprise, and is in fact a direct consequence of our uncertainty model. In other words, once the mean of the distribution is close to λ, very few compromised nodes with reported measurements that *are* acceptable by the uncertainty model can move the average and flip the outcome without being detected. Note that such measurements could have come from a valid distribution, and thus been legitimate.

Figures 4.2, 4.3, and 4.4 (pictures on right) depict the measured average and final detection outcomes for vandalism attacks. The results from an unmodified ML detector are provided in the captions for comparison. Since the results and analysis are similar to the exploitation attacks we do not discuss them in detail.

4.4   Extensions & Simulation Study (Part 2)

The simulation results in Section 4.3 show that in areas where the average (mean) of signal power is close to the detection threshold, a modest fraction of compromised nodes in a cell can change the outcome of spectrum sensing without being detected. This is due to the difficulty of distinguishing between legitimate variations in signal power and slightly skewed false reports by attackers. Therefore, the attackers succeed by effectively 'hiding' under the 'acceptable' measurement variations. In this section we propose extending our solution by using median as a safeguard, in conjunction
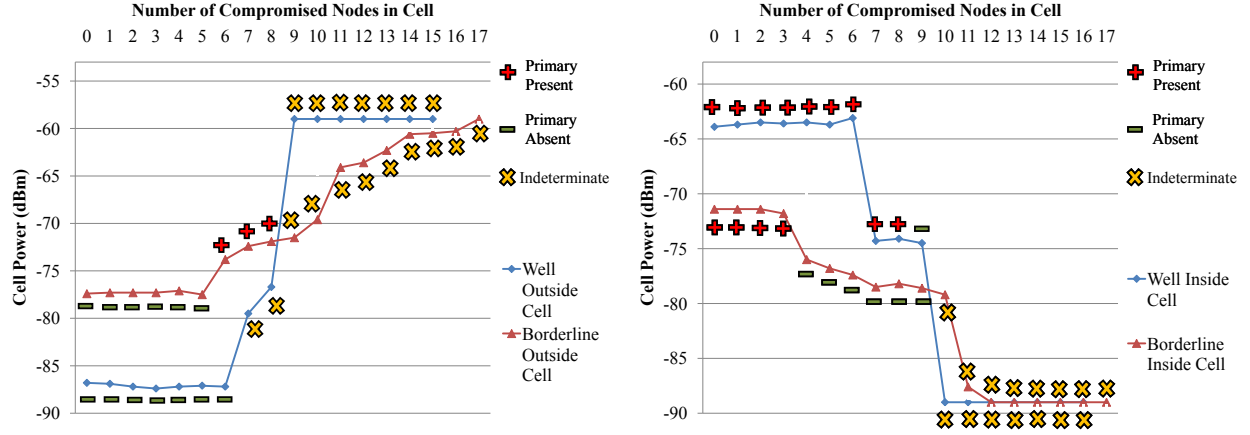
Figure 4.2: (left) Exploitation by +15dB uncoordinated attackers. ML detector is beat when 7 (4) nodes are compromised in the well-outside (borderline-outside) cell. (right) Vandalism by -15dB uncoordinated attackers. ML detector is beat when 7 (3) nodes are compromised in the well-inside (borderline-inside) cell.

with mean, for secure primary detection. We show that our solution achieves a desirable mix of accuracy (from mean), and robustness (from median).

4.4.1 Median: A Safeguard for Collaborative Sensing

An alternative estimator for signal power in a cell is the median of measurements. The median of a sample is known to be robust to outliers. The median, however, has the disadvantage that it does not use all the data available in the sample, and therefore is often not as accurate as the mean [75]. For a normal distribution, it is well known that the sample mean is the most 'efficient' estimator, that is no other unbiased statistic for estimating μ can have smaller variance. The efficiency of median, measured as the ratio of the variance of the mean to the variance of the median, depends on the sample size $m = 2n + 1$ as $\frac{4n}{\pi(2n+1)}$, which tends to the value $2/\pi \approx .63$ as m becomes large [46]. So, we can consider the following distribution for the median power in a cell: $P_{\text{med}} \sim N(\mu, \frac{\pi\sigma^2}{2m})$. Therefore, similar to Equation 4.2, in order to use median for primary detection we can derive the threshold λ' such that the probability of missed detection stays below a certain value (*e.g.* .95):

$$\lambda' = \frac{\sqrt{\pi}\sigma}{\sqrt{2m}}Q^{-1}(.95) + p_r \qquad (4.4)$$

38
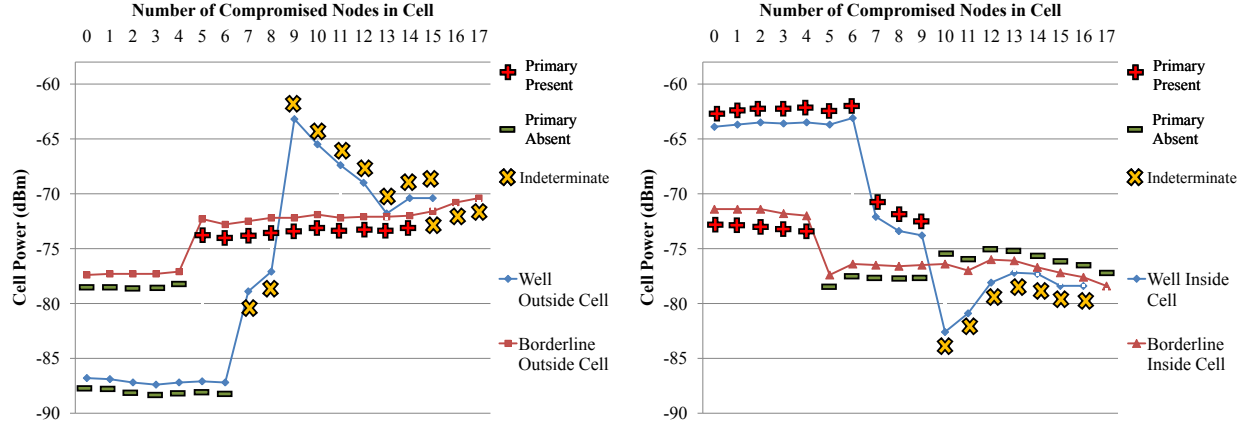
Figure 4.3: (left) Exploitation by coordinated attackers. ML detector is beat when one node is compromised in both the well-outside and borderline-outside cases. (right) Vandalism by coordinated attackers. ML detector is beat when one node is compromised in both the well-inside and borderline-inside cases.

where Q^{-1} is the inverse of standard Gaussian distribution tail function. Note that since $Q^{-1}(.95) < 0$ we have $\lambda' < \lambda$.

The next question that arises is how we can integrate median into our existing approach. To that end, we propose a framework based on the following principles: (1) safety (in terms of causing interference to primaries) is not compromised, and (2) a reasonable combination of efficiency (*i.e.* mean) and robustness (*i.e.* median) is achieved. In a given cell, we first perform the hierarchical grid-based scheme proposed in Section 4.1. If the status of the cell is 'neutral' (see Algorithm 1), then we perform the following *additional* operations. Consider P_{med} and P_{avg} to be the median and weighted mean of the power measurements. We have the following four cases:

1. $P_{\mathrm{med}} >= \lambda'$ and $P_{\mathrm{avg}} >= \lambda$: Since both estimators agree on the positive outcome, we consider primary signal to be **present**.

2. $P_{\mathrm{med}} < \lambda'$ and $P_{\mathrm{avg}} < \lambda$: Since both estimators agree on the negative outcome, we consider primary signal to be **absent**.

3. $P_{\mathrm{med}} >= \lambda'$ and $P_{\mathrm{avg}} < \lambda$: There exists a conflict; primary is present based on the median, but is absent based on the mean. Considering the importance of not causing interference to primary users, we disregard the potential optimality of the outcome from mean and opt

39
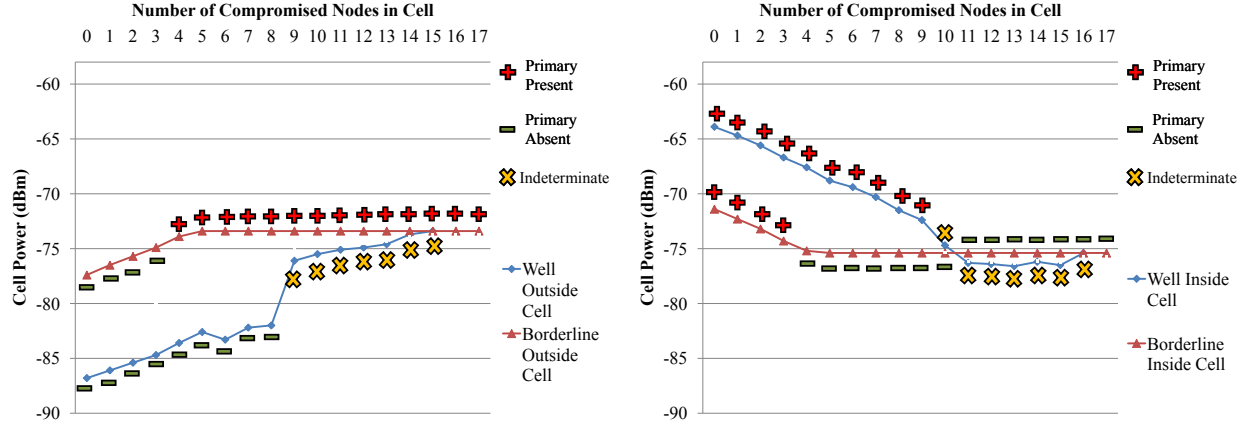
Figure 4.4: (left) Exploitation by omniscient attackers. ML detector is beat if one node is compromised in both the well-outside and borderline-outside cases. (right) Vandalism by omniscient attackers. ML detector is beat if one node is compromised in both the well-inside and borderline-inside cases.

for the conservative choice of declaring primary **present**. This choice is expected to reduce the chances of successful vandalism attacks, but may increase the chances of mistakenly declaring a borderline-outside cell as occupied (due to the relative inefficiency of median).

4. $P_{\text{med}} < \lambda'$ and $P_{\text{avg}} >= \lambda$: There exists a conflict; primary is present based on the mean, but is absent based on the median. The difference in opinions may be caused by an exploitation attack, or simply a legitimate inaccuracy by either of the two estimators. Given the previous choices, if we go with the mean's decision, we are effectively taking the decision to be the 'or' of the two. This choice has the drawback that would not make exploitation attacks any harder to launch. On the other hand, if we go with the median's decision, we are effectively ignoring mean in all the four cases, which is not desirable. Since we know that $\lambda' < \lambda$, the mean and median are at least separated by $\lambda - \lambda'$. This may be a sign of anomaly (*e.g.* an exploitation attack). We propose considering this cell as **indeterminate** and using the average power of the 8 neighboring cells (and compare it to λ) to determine the cell's status. We propose to use a similar disambiguation technique for **indeterminate** cells from Section 4.1 ('conflicted', 'low,' or 'high' in Algorithm 1). For the cells at the border of the area of interest (that do not have 8 neighbors), we consider the status to stay indeterminate.

4.4.2 Simulation Study (Part 2)

We first study the effect of using median in conjunction with mean in *absence* of attackers. This evaluation is done in terms of false positive and false negative rates.

Table 4.1: The number of false positives and false negatives.

Algorithm	False Positives	False Negatives
Hierarchical Average-Based (Section 4.1)	16	10
Extended Median-Based (Section 4.4)	49	0

Consider any cell that is *entirely* outside the no-talk radius of the primary transmitter. If either of our approaches mistakenly declare primary to be present in this cell, we count this as a false positive. Similarly, consider a cell that is (in part) in the no-talk radius of the primary. If either of our approaches mistakenly declare primary to be absent for this cell, we count this as a false negative. We measure false positive and false negative rates in two cases: (1) when only the average-based framework in Section 4.1 is used, and (2) when it is combined with the median-based framework introduced in this section. The results are summarized in Table 4.1. The table shows the number of false positives and false negatives for the final decision (after disambiguating indeterminate cells) for both approaches. The total number of cells is 1024. It can be seen that in the absence of attackers the extended approach provides an extra level of safety. This comes at the cost of higher false positive rates.

Next, we study the effectiveness of the extensions in this section against exploitation and vandalism attacks. For this purpose, we run the same experiments as in Section 4.3 with the added extensions in this section. Figures 4.5, 4.6, and 4.7 represent the results for uncoordinated, coordinated, and omniscient attackers respectively. The new changes are represented by arrows. In particular, arrows originating from a '+' or '−' represent scenarios for which cases (3) or (4) apply, that is when the mean and median do not agree. The symbol at the head of the arrow represents the final decision. Arrows originating from a × represent cases that are considered indeterminate based on the hierarchical scheme in Section 4.1, and the sign at the head of the arrow represents

the final outcome after neighbor averaging rule.
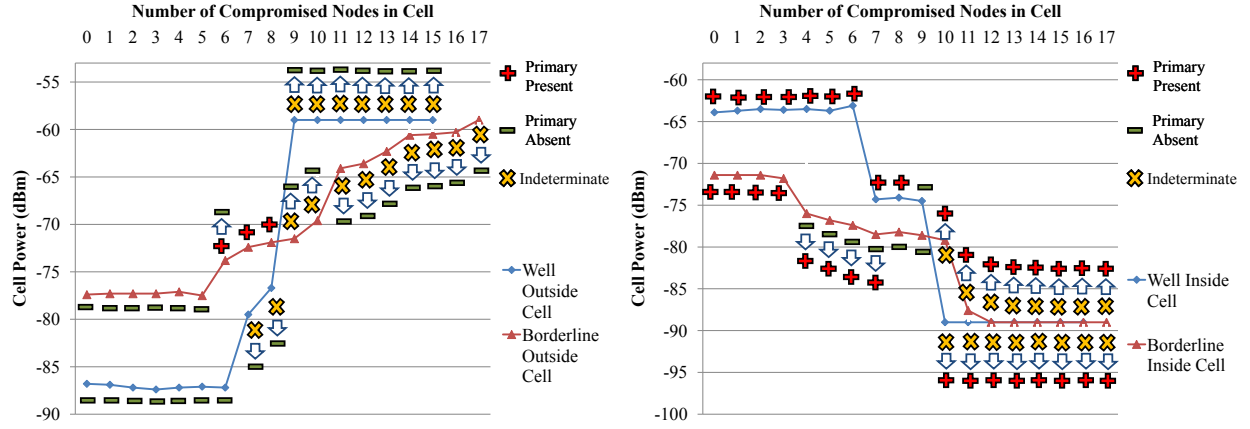
Figure 4.5: Exploitation (left) and Vandalism (right) by uncoordinated attackers. Arrows represent change of final detection outcome based on extensions in Section 4.4.
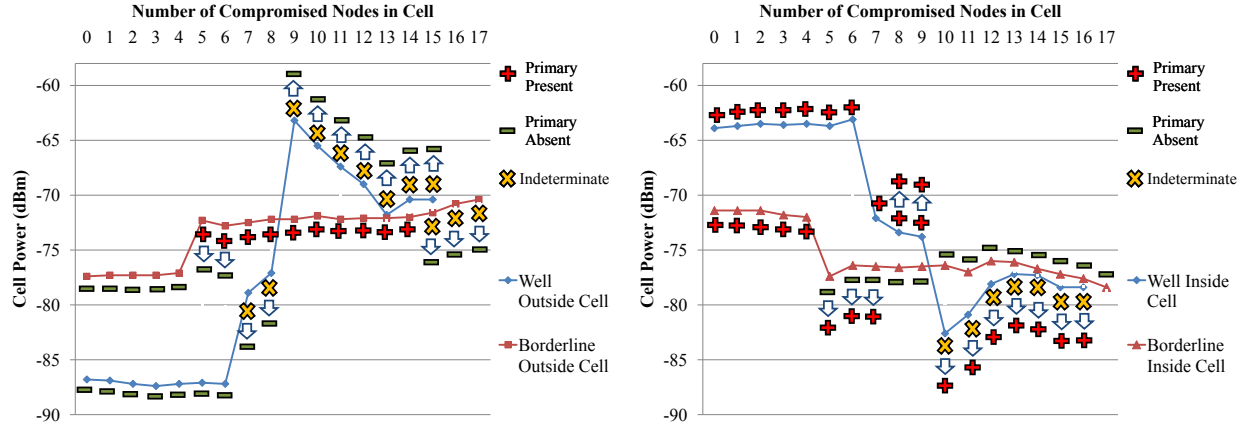
Figure 4.6: Exploitation (left) and Vandalism (right) by coordinated attackers. Arrows represent change of final detection outcome based on extensions in Section 4.4.

The results show that for the well-inside (well-outside) case, in almost all scenarios, our solution completely nullifies the effect of attackers. For borderline-inside (borderline-outside) case, the attackers need to compromise at least 47% (41%) of nodes to be able to succeed. Note that these ratios are higher for the cases of less sophisticated attackers. The difference between the results for exploitation and vandalism can be explained by our conservative approach that prioritizes safety (non-interference) over security. Overall, the results show a considerable improvement over the original grid-based hierarchical scheme.

Figure 4.7: Exploitation (left) and Vandalism (right) by omniscient attackers. Arrows represent change of final detection outcome based on extensions in Section 4.4.

4.5 Conclusions

In this chapter, we provided an approach to achieving robust radio spectrum telemetry that is applicable to large regions where no single ground truth is viable at all places. Our solution uses outlier detection at two levels: (1) intra-cell among individual CR measurements and (2) inter-cell by corroboration among cells in a hierarchical structure. The results are used in a weighted detection mechanism, in conjunction with a median-based framework, to eliminate or lower the effect of the attackers. We provided a novel framework for deriving the dispute thresholds for outlier detection based on the underlying propagation and uncertainty model of the signal power.

We provided analytical and simulation results to quantify the extent to which attackers can succeed. The attackers in the simulations ranged from ones with very little sophistication, to those with complete knowledge about their neighbors and the detection mechanism (who use it to avoid detection). Our results showed that in cases where attackers are not near the border of the primary's protection area, we can detect and fully eliminate the effect attackers in a particular region. For our worst-case scenarios, that is cells that are close to the border of primary's protection area in which coordinated omniscient attackers employ smart strategies, we can nullify the effect of up to 41% of attackers nodes.

Chapter 5

Data-Based Protection with Classifiers

In this chapter we propose CUSP, a technique based on machine learning that uses a trusted initial set of signal propagation data in a region as input to build a classifier using Support Vector Machines. The classifier is subsequently used to detect integrity violations. Using classification eliminates the need for arbitrary assumptions about signal propagation models and parameters or thresholds in favor of direct training data. Extensive evaluations using TV transmitter data from the FCC, terrain data from NASA, and house density data from the US Census Bureau for areas in Illinois and Pennsylvania show that our technique is effective against attackers of varying sophistication, while accommodating for regional terrain and shadowing diversity[1].

5.1   Motivation and Approach

The two problems of detecting individual maliciously false reporting nodes [45, 59] and that of detecting attacker-dominated cells [31] have been mainly formulated as abnormality or outlier detection problems. Despite moderate degrees of success, these approaches suffer from several technical and practical issues. First, they often involve unrealistic assumptions about the models and parameters of signal propagation. Second, the performance of almost all of these methods highly depend on detection threshold parameters which are usually tuned by hand, or depend on the parameters of the signal propagation model. This is impractical, because it requires too much 'conjecturing' and 'manual tuning' for each given region and frequency band of interest. In addition, outlier detection techniques are often very conservative and are not designed for detecting

---

[1]The majority of the material in this chapter is adopted from Fatemieh, Farhadi, Chandra, and Gunter's recent publication [32].

nimble manipulations of data by sophisticated attackers. This limitation is particularly important in the context of spectrum sensing in which there exist natural variations in signal power due to factors such as fading and noise [72].

As an illustrative example consider Figure 5.1 to be a subset of the area of interest. Each cell is the unit for averaging signal power measurements from sensors to determine primary presence. The average power from the nodes inside a cell are represented by a number (in dBm) in that cell, and the primary detection threshold is -114dBm. Cells A and B are normal, whereas C is dominated by attacker nodes. Therefore, the attacker nodes are able to decrease the average power to -115, which, if undetected, results in a successful vandalism attack. It is tempting to devise heuristics or simple outlier detection techniques based on approximate signal propagation formulas to catch cells like C. For example one may claim the difference between B's average power and its neighbors looks normal since its average is smaller than a threshold α, but this is not true for C, therefore C is compromised. But 'why is comparing the average distance to α is a good idea?' Why is C suspicious, but A is not? Many other questions may still linger; for example 'how do we know we chose the right threshold', 'how do we know we are not mistaking an attacker-dominated cell with one behind a hill', 'how do we make sure we have taken all the factors into account', or '*can we do better*'?

We believe that we should directly use signal propagation data for this purpose. Leveraging patterns latent in the data will lead to more practical, robust and accurate solutions. The key intuition is to learn the propagation *behavior* of the signal from the observed signal propagation data (we will discuss the practicality of obtaining data later). There are patterns in which the signal propagates. We can extract these 'patterns' and utilize them to predict how we expect the signal to behave in the (often large) region of interest. Naturally, the actual behavior of the signal should be similar to what we can predict from the observed propagation patterns. This is mainly because we learn to predict the patterns of propagation from the signal itself. We claim that if the propagation of signal in a given location within the region of interest is not similar to patterns of signal propagation extracted from the same or 'similar' signals in the region of interest, the location
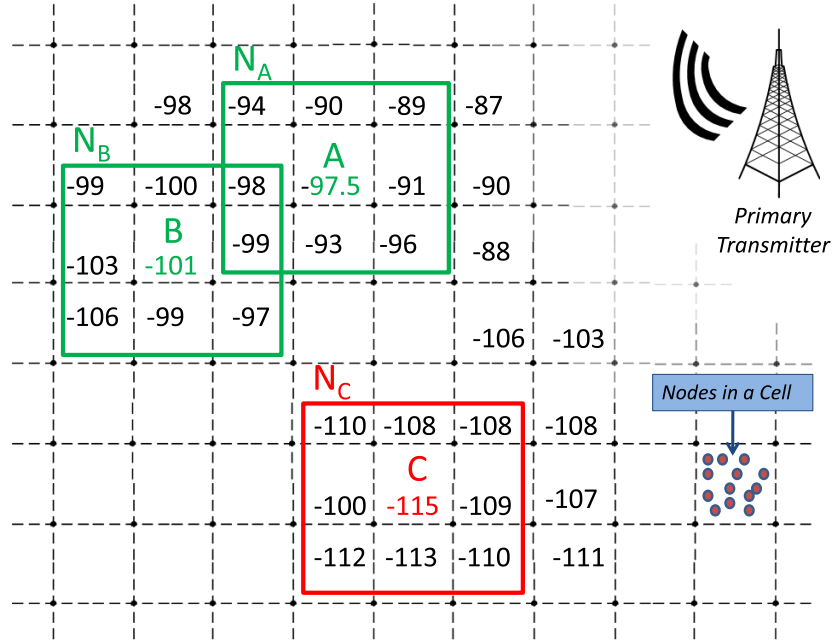
Figure 5.1: Sample grid with normal and attacker-dominated cells.

should be considered suspicious or un-natural. As a simplistic example, assume we somehow learn that in a particular flat desert, digital TV signals weaker than -70dBm attenuate by at most 5dB every 5 kilometers. Then, a 10dB decrease or an 8dB increase over a three kilometer distance may be considered suspicious, or at least unusual.

We believe that we can spot unnatural propagation of signal in local neighborhoods of adjacent cells by carefully analyzing samples from the actual signals in the same and several different neighborhoods (within the region of interest) in the past. For a given neighborhood, we are now concerned with a new type of question. *Is the propagation of the signal natural in this neighborhood?* Before answering this question, we must define and show how to represent the pattern of signal behavior in a neighborhood of cells. So, we first address the following question. *How to represent the pattern of signal propagation in neighborhoods?*

5.1.1 Representation of Signal Propagation

In order to better understand the patterns by which the signal propagates, we need to define a way to represent them. We start by a simple representation as follows. We consider the *local*

neighborhood of any cell A to contain A and its 8 neighboring cells. For example, in Figure 5.1 the local neighborhood for A, B, and C are shown and referred to as N_A, N_B, and N_C respectively. Using this definition for a local neighborhood, we represent a cell A by a 9-element tuple containing the power level in A, and the difference in power between A and the rest of the neighbors in a pre-specified order. For example the neighborhood for cell A, is represented as $\langle -97.5, -.5, 3.5, 7.5, 8.5, 6.5, 1.5, 4.5, -1.5 \rangle$. We call this the *neighborhood representation* of A. Note that the representation can be expanded to include, for example, the neighbors of neighbors of A as well to provide additional context for learning patterns. However, as we will show later, the 9-tuple representation is sufficient for our purposes. This representation provides us with a way to encode the pattern by which the data propagates in this neighborhood. Using this definition for the neighborhood of a cell, our original question can be re-phrased as: *For a given cell, is the propagation of the signal natural in its local neighborhood?*

5.1.2   Using Patterns of Signal Propagation

Let us assume that we have access to reliable power measurements in all of the region of interest. An example for a region would be a 50km by 50km area with a roughly uniform (flat, hilly, etc.) terrain. It is easy to see that the data can be used to create one neighborhood representation for each cell in the region. We refer to each of such representations as an 'example.' Therefore, we can assume access to a large number of such examples representing the 'natural' propagation of signal in local neighborhoods. Also, for now assume that we are magically provided with the neighborhood representation for a sufficiently large and diverse set of 'un-natural' (attacker-dominated) cells.

Having access to representations for patterns of signal propagation as natural and un-natural examples, we believe the best way of approaching our question is to learn the common characteristics in each group and use it to differentiate between natural and un-natural examples. This means that by discovering the key characteristics of signal propagation patterns, we can superimpose a boundary in our space of representations. This boundary works as the decision making module. For a

new example, we need to check which side of the boundary the example lies; the natural side or the un-natural side. This is a classic *classification* problem. We have now reduced our problem to a more specific question: *How to cast the problem of detecting attacker-dominated cells as a classification problem.* Before answering this question, we provide an analogy and the background on classification.

A useful analogy to this problem is that of spam detection in email systems: given a set of emails each marked as spam or normal, the goal is to learn the common characteristics among the normal emails, the common characteristics among the spam emails, and characteristics that differentiate between the two groups. Going back to our problem, we would like to discover a model that best describes the behavior of signal, and use it to make predictions about the normalcy of signal propagation in subsequent examples.

5.1.3   Background on Classification

Classification is one of the mainstreams of machine learning and has been widely adopted in many domains ranging from spam email detection [39] and unauthorized spectrum usage [57] to fraud detection [48], object detection [34], and speech recognition [69].

In a binary classification problem we are given a set of *training* examples with their corresponding *labels*, $(\overrightarrow{x_i}, y_i)$ where $\overrightarrow{x_i}$ is the representation of the i^{th} example in the *feature space* and $y_i \in \{1, -1\}$ (yes or no?) is the corresponding binary label. Each example is described by a vector of its attributes which is often called the *feature vector.* For example, in detecting if a person has a significant risk of heart attack, the features can be the blood pressure, cholesterol level, and body mass index. The goal is to predict a binary label for an example for which we do not know the label (*a.k.a.* a *test* example) using the training examples [19]. In the heart attack example, we want to predict whether a person is under a certain risk of heart attack, given her feature vector. We do this by learning the patterns in the features of several different persons with and without the risk of the heart attack.

Looking underneath the surface, a classifier tries to partition the input feature space into regions

where positive examples lie versus regions where negative examples lie. The boundary between regions for positive and negative examples is called the *decision boundary*. Training involves learning the decision boundary and classification involves determining on which side of the decision boundary a test example lies. In the simplest case, it is assumed that the decision boundary is a linear function of the input feature vector $\overrightarrow{x}$. Later, we relax this assumption and consider more complex decision boundaries. This linear function usually takes the form of

$$y(\overrightarrow{x}) = \overrightarrow{w}.\overrightarrow{x} + w_0 \tag{5.1}$$

where $\overrightarrow{w}$ is the *weight vector* and w_0 is the *bias* [19]. One might think about the decision boundary as a $(N-1)$-dimensional hyperplane in the N-dimensional feature space. The classification is done by determining the side of the hyperplane on which each point in the feature space lies. If $y(\overrightarrow{x}) \geq 0$ then $\overrightarrow{x}$ gets the label 1 and if $y(\overrightarrow{x}) < 0$ it gets the label -1.

5.1.4   Casting Attacker-Dominated Cell Detection as a Classification Problem

We need to learn a classifier to predict whether a cell seems natural or not. To that end, we represent signal propagation in a local neighborhood of a cell, by the power average in the cell, as well as the 8 numbers representing the difference between the power averages of the cell with its neighbors. We denote these features by $\overrightarrow{x}$. To automatically discover these patterns we search for parameters $\overrightarrow{w}$ and w_0 that best explain the training data and provide reliable generalization properties. To be more specific, we are optimizing for $\overrightarrow{w}$ and w_0 that, if used for classification, provide the best prediction accuracy over the training data set while not overfit to it. More formally, the prediction of train set label y, which takes the form of $\overrightarrow{w}.\overrightarrow{x} + w_0$ should be similar to the actual train set label y. At the same time, to avoid too much fine tuning to the train set examples, the size (norm) of the weight vector $\overrightarrow{w}$ should be controlled. One drawback of this model is the assumption of linear separability. Our predictions are linear in the feature space, thus form a linear decision boundary. To be able to model nonlinear decision boundaries, we project the data $\overrightarrow{x}$ to higher

dimensional spaces where the decision boundaries are linear on that higher dimensional space. Our new predictions take the form of $\overrightarrow{W}.\Phi(\overrightarrow{x}) + W_0$ where Φ is a mapping to the higher dimensional feature space. We postulate that the decision boundaries in the feature space can be modeled more reliably by quadratic functions, thus modeling Φ by a quadratic kernel. To be more specific, we are solving the following optimization problem:

$$\min \frac{1}{2}\|\overrightarrow{W}\|^2 + \gamma \sum_{i=1}^{N} \xi_i \tag{5.2}$$
$$\text{subject to } y_i(\overrightarrow{W}.\Phi(\overrightarrow{x}) + W_0) \geq 1 - \xi_i \quad \forall i$$

where N is the number of training examples, ξ_i is a collection of non-negative *slack variables* that account for possible misclassifications and γ is the *tradeoff factor* between the slack variables and the regularization on the norm of the weight vector $\overrightarrow{W}$. The constraint in this minimization implies that we want our predictions, $\overrightarrow{W}.\Phi(\overrightarrow{x}) + W_0$, to be similar to labels $y_i \in \{1, -1\}$. The objective function works as a regularizer to avoid overfitting to the training data set. We solve this optimization by quadratic programming in dual. This is an example of SVMs [29].

The only parameter that needs to be estimated is γ. We estimate the γ by cross validating it in the validation set, a part of train set which set aside for parameter estimation. This parameter is set using the data itself and there is no need of any assumption about data distribution.

Given a $\overrightarrow{W}^*$ and $\overrightarrow{W}_0^*$, which are the outputs of the Optimization 5.2, we can predict whether a cell is natural or not by looking at the sign of $\overrightarrow{W^*}.\Phi(\overrightarrow{x}) + W_0^*$.

Data Collection. The main remaining question is how to obtain the training examples needed to build the classifier. We argue that *normal* (negative) instances can be obtained in a practical *one-time* process based on a trusted sensor grid. By one-time we mean that in a particular region, we only need to collect signal propagation data once to build the classifier for that region. Once the classifier is built, it can be used forever (or until there is a significant environmental change

in the region). A typical strategy for collecting this data is war-driving where a sensor is moved though the region collecting training data as it goes. This data can also be extrapolated by signal propagation models such as Longley-Rice, but our approach does not require the use of any such model. War-driving for collecting spectrum data is similar to the current practice of taking images for street-view capabilities of online map applications in Google and Bing.

An alternative may be realized in the context of 802.22 internet service for residences, as well as the envisioned application of white-spaces for advanced meter communications [30]. In this case, the (one-time) measurements may be collected at the time of deploying radios (meters) at each house by the operator. They may also be collected by a temporary sensor network developed for this purpose alongside the main CR network [68].

Once negative instances are collected, we use a methodology to inject *attacker-dominated* (positive) training instances to incorporate attacker-dominated cells containing attackers of varying degrees of sophistication. For further details please refer to Section 5.2.1.

5.1.5 A Unified Classifier for each Region

At this point, provided with labeled training examples for a transmitter, we are able to build a classifier that can predict if a cell is attacker-dominated or not. This is still far from practical for the following reasons. First, it requires training and maintaining a classifier for each transmitter. Second, as it will be concretely shown in Section 5.2, each transmitter may only provide a particular distribution of power levels in the region of interest. This leads to insufficient or non-existent training examples for some power levels, which can lead to low classification accuracy. Given enough training examples for a frequency range (*e.g.* 620-698 MHz for DTVs), we argue that our classifiers are capable of discovering decision boundaries in the feature space which are independent of the transmitter. This is due to the fact that signal propagation is mainly a function of power, propagation environment, and the frequency of transmission. From a practical perspective, this means that we do not need to learn a separate classifier for each transmitter in the same frequency range. We show this property in Table 5.3 in the context of six DTV transmitters in Illinois.

We introduce the concept of *Unified Classifiers* that are trained by pooling data from multiple transmitters in such a way that there exist sufficient number of training examples at any power level in the power range of interest. For example for DTV transmitters this range will be between 90 dBm (maximum DTV transmission power) and -130 dBm (weakest signals considered). The new question we are facing is which transmitters to select so that we can ensure sufficient number of examples at any power level; the 'transmitter selection problem.' This problem can be reduced to the set covering problem, which is a well-known NP-Complete problem [73]. We divide the larger power spectrum of interest to a number of smaller power ranges and aim to enforce a lower bound on the number of examples per power range. Our goal is to select the minimum number of transmitters so that we are guaranteed to have at least a fixed number of examples per power range. We greedily select the transmitter that covers the largest number of uncovered power ranges at each stage. This is known to have an approximation ratio of $ln(n) + 1$ where n is the number of power ranges [73]. Having selected transmitters that cover the entire power range, we can now learn a classifier from the data from all selected transmitters. This is our unified classifier. We show that we can detect attacker-dominated cells for transmitters we never observe during training. This is of practical significance as one does not need to be concerned with providing information from all the transmitters in a frequency range, or those that may start transmission in future.

Another practical property of our unified classifier is its relative independence to the frequency. We later show that the unified classifier is not considerably sensitive to the frequency change in DTV transmitters in the UHF channels 14-51 ($470 - 698$MHz). This means that our unified classifier is capable of detecting attacker-dominated cells when trained with data from transmitters in different frequency ranges. Therefore, as we will show with evaluations for both Illinois and Pennsylvania, it is sufficient to build one classifier for the entire 470-698 MHz range.

Once the unified classifier detects a cell as compromised, the detection outcome in that cell should be reversed to cancel the attackers' misreporting effect. In cases where the actual power level is important, the power level should be replaced by the average powers reported by the majority of its neighboring cells. This strategy, which is motivated from image smoothing techniques

in vision applications, has been validated in the context of white space networks [31]. This strategy, when combined with a multi-resolution deployment of CUSP enables nullifying the affect of attackers at different granularity levels, as well as those that are able to dominate multiple adjacent cells. Alternatively, in the case of using white-spaces for AMI communications, or 802.22 Internet, the firmware for the suspicious devices may be (physically or remotely) examined by the utility or 802.22 service provider.

5.2 Instantiating CUSP

In this section we show how CUSP can be instantiated in a region to provide protection against attacker-dominated cells when aggregating spectrum sensing reports at a central server. To that end, we provide general guidelines as well as specific details for an illustrative environment, namely East-central Illinois. Since it was not practical for us to do wardriving through this region we instead rely on the FCC and NASA databases and the Longley-Rice empirical outdoor signal propagation model to generate sensor data (see [5, 61] for more details). Longley-Rice is endorsed by FCC for determining propagation contours in the TV spectrum and takes into account the effects of terrain as well as transmitter's location, height, and power. For the purpose of these experiments we treat these models as the ground truth provided by sensors and use this to test our method. Note, however, that our method does not rely on any specific choice of a model. Hence if these models have some inaccuracies then we believe that accurate training data and proper application of CUSP will achieve the necessary foundation for integrity protections. We defer the experiments in which we account for additional variations and uncertainties in signal propagation to Section 5.3.

5.2.1 Environment and Data Collection

We start by considering a 160km $\times$ 160km square area in the flat Midwest area in the US. The following points in (*latitude, longitude*) format define the boundaries of the region: $\langle$ (39.56, -89.4), (41, -89.4), (41, -87.5), (39.56, -87.5) $\rangle$. The area is located in East-central Illinois and mainly

consists of rural farmlands and a few small cities with populations under 100,000. Figure 5.2 depicts this area. We use registered DTV transmitter data from the FCC databases as well terrain data from the NASA database to build our grid-based crowdsourcing data. For any given location we can retrieve the list of nearby DTV transmitters as well as their properties such as channel (frequency), transmission power, and antenna height. We then combine this data with terrain data and use the Longley-Rice propagation model to estimate signal power from each of the DTV transmitters at that location.

Figure 5.2: Initial evaluation area and the first set of considered DTV transmitters in East-central Illinois.

Cell Size and Density. An important factor when using CUSP in any environment is the cell size and density of sensors (or wardriving samples). To make an informed decision about the cell size and sensor density, the following factors should be taken into consideration. First, the cell size must be large enough that about 10 to 20 sensors exist in each cell. Mishra *et al.* [60] show that this many independent sensors provide as much collaborative gain as many more correlated sensors whose collaborative gain is limited by geographical correlation in shadowing. Second, the variation of average signal power in a cell must not be significant (*e.g.* less than 5dB) in order for combining individual reports to be meaningful. Using a similar criteria, Kim *et al.* [47] proposed a maximum radius of 5.6 km for a circular cell for detecting the TV transmitters at the edges of their contours.

54

Third, collaborative sensing often works best when there exists independence in the (shadow) fading among different sensors. Using Godmunson's exponential shadow correlation decay model, it is shown that the maximum sensor density of 3.2 sensors/km^2 ensures independence between individual reports [47]. This factor, however, is more a recommendation than a requirement.

Considering an application such as advanced meter communications or 802.22 Internet, one may use the estimate of one sensor per house for spectrum sensing. To that end, we studied house density per square kilometer of the 102 counties in the state of Illinois from the US Census Bureau data [11]. The results show that the least dense county (Pope county) contains 2.5 houses/km^2. The 5th percentile of the data is 3.5 and median is 8.5 houses/km^2. In view of the discussion above, we opt for the following parameters. We consider base cells of size 2km×2km with the average density of 3.2 sensors per km^2. We consider nodes to be uniformly distributed at random. It is known that the actual distribution of the sensors (houses in this case) may not be uniform in real-world, however, for the following reasons we argue that this assumption is reasonable for the evaluations. First, since we take conservative estimates for sensor density, it is likely that in most areas there exist more than the assumed 3.2 sensors per km^2. In such cases, the central server can choose from the existing nodes in order to create a relatively uniform distribution. Second, in the rare cases (given the conservative choice of density) that some cells contain less number of sensors, or sensors are closely clustered, the service provider may deploy additional sensing units.

Attacker Model and Example Generation. Based on the assumptions in Chapter 2, the following four attacker models may be considered in this context. Note that the attackers' behavior should be considered through the lens of a particular cell that the attackers aim to dominate.

1. Uncoordinated attackers do not have precise information on the number and power measurements of other legitimate or attacker nodes in the cell. Each attacker node aims to dominate the cell without cooperation with other attackers, if any. This may be due to lack of information, unavailability of communication channels, or to reduce the likelihood of being detected as a result of communicating with peers. In this case, a compromised node that senses a signal power below

(above) the detection threshold may falsely report a value such that the average power in the cell changes to a value below (above) the detection threshold. The attacker may use rough estimates of the number and measurements of other nodes for this purpose (for example, for the latter it would be a close value to the attacker's true measurement).

2. Coordinated attackers do not know the number and power measurements of the legitimate nodes in the cell, but may roughly estimate them. They do, however, know their own number and measurements, and act according to a coordinated strategy; they collude and use the estimates to calculate the value that each of them should report so that they can dominate the cell and change the detection outcome to a value above (or below) threshold.

3. Omniscient attackers are coordinated attackers that know the exact number and measurements of other legitimate users. Therefore, they can simply calculate the exact power levels they should report to change the average power level to a value *slightly* above (or below) threshold, *e.g.* 1dB. This is to reduce the chances of being detected.

4. Mimicry-capable Omniscient attackers are omniscient attackers that have the (non-trivial) resources to build a classifier similar to that used in our detection technique. However, we can hide (or simply randomize) the schedule, frequencies, and locations in which we enable the detection scheme. Therefore, before any misreporting attempt the attackers can predict whether our classifier can catch them *if it is enabled* at that particular time, location, and frequency. In the small percentage of cases that they know it cannot detect them (even if enabled), they will misreport according to the omniscient strategy above. Otherwise, they may choose to misreport based on their risk appetite. In any case, if they choose not to misreport, we have achieved our goal of preventing attackers from manipulating the detection outcome. Otherwise, we will detect them as we would have detected omniscient attackers. Therefore, we do not report separate results for this class of attacks and rely on results for omniscient attackers.

For each selected cell, we include the value of the cell's average power (*e.g.* -65) as well the difference of this cell with its immediate neighbors as the features for a normal example. Therefore, a normal example takes the form $\langle -65, 5, -2.5, 0.6, -3, 3, 2, -3, -1.2 \rangle$. Generating attacker

instances is a non-trivial problem. The instances have to be general enough to train the classifier in such a way that it is able to detect attacks mounted using unknown strategies with varying fractions of attackers inside the cell. We opt for a randomized approach for generating attacker data in order to provide substantial variations in the training data. For uncoordinated attackers, we replace the actual power in the cell with $\text{Rand}(\lambda + 1, \lambda + 10)$ for exploitation attacks and with $\text{Rand}(\lambda - 1, \lambda - 10)$ for vandalism attacks, where $\text{Rand}(a, b)$ returns a random number between a and b and $\lambda = -114$ is the primary detection threshold. Similarly, we use $\text{Rand}(\lambda + 1, \lambda + 5)$, or $\text{Rand}(\lambda - 1, \lambda - 5)$ for coordinated attackers. For omniscient attackers, we simply replace the value with $\lambda + 1$ or $\lambda - 1$ for exploitation and vandalism attacks respectively. These attackers are knowledgeable and coordinated, and therefore they can only move the average exactly as much as needed to flip the detection outcome (1dB is the unit of measurements). This minimizes the attacker's chance of being detected.

5.2.2 Initial Evaluation

Of the tens of DTV transmitters in this area, we initially choose six DTV transmitters listed in Table 5.1 as a representative set. These transmitters are identified in Figure 5.2 as green antennas. This choice aims to serve two purposes; first, geographical diversity, and second, obtaining a wide range of received power levels across the area. Figure 5.3 represents the distribution of received signal powers from each of the six transmitters in the area. Later, we use the lessens learned in this section to perform a comprehensive analysis on other transmitters in the area of interest in Illinois, as well as all the transmitters that affect the area of interest in Pennsylvania.

Table 5.1: Initially-selected DTV transmitters.

Call Sign	Chan.	Fq. (MHz)	Tx Pow. (kW)
WAOE (MyN)	39	620-626	151
WCIA (CBS)	48	674-680	1000
WEIU-TV (PBS)	50	686-692	255
WICS (ABC)	42	638-644	954
WQAD-TV (ABC)	38	614-620	1000
KTVI (Fox)	43	644-650	1000

Figure 5.3: Distribution of received signal powers from six DTV transmitters in Illinois.

In our first set of experiments, we consider each transmitter separately. For the labeled data for each transmitter, we perform 'K-fold cross validation,' which is a commonly used technique to evaluate the performance of classifiers. We randomly partition the data into K subsamples. Of the K subsamples, a single subsample is retained as the test data for testing the model, and the remaining K - 1 subsamples are used as training data. The cross-validation process is then repeated K times (the folds), with each of the K subsamples used exactly once as the validation data. The K results from the folds then are averaged to produce a single estimation. The advantage of this method over repeated random sub-sampling is that all observations are used for both training and validation, and each observation is used for validation exactly once. In our experiments we set $K = 10$. The results are summarized in Table 5.2. Note that these results are obtained with an equal mix among the three attacker models. We will provide further breakdown based on the attack-type later in this section.

Table 5.2: Detection accuracy (D.A.) and false positive (F.P.) for six DTV transmitters in Illinois.

	WAOE		WCIA		WEIU-TV		WICS		WQAD-TV		KTVI	
	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.
$P > -65$	100	0	100	0	100	0	100	0	-	-	-	-
$-65 \geq P > -85$	100	0	100	0	100	0	100	0	99	0	-	-
$-85 \geq P > -105$	100	0	100	0	100	0	100	0	99.8	0	100	0
$-105 \geq P > -114$	99.1	2.2	-	-	99.8	4.8	99.7	2	99.8	1.5	98.7	2.9
$-114 \geq P$	95.3	8.7	-	-	87.8	15	87.1	8.6	95.5	11.9	99.2	2.3
Overall	**98.9**	**2**	**100**	**0**	**99**	**3**	**99.5**	**1**	**99.4**	**2.1**	**99**	**2.5**

58

5.2.3 Building a Unified Classifier

The results in Table 5.2 are obtained by considering each transmitter separately. From a practical perspective, it is ideal to use just one classifier. Such a classifier is trained by pooling data from multiple transmitters in a way that there exist sufficient number of training examples at any power level. According to CUSP's greedy method for transmitter selection (Section 5.1.5), we pick the data from WEIU-TV and KTVI for training the classifier. We test the classifier on the data from the other four transmitters. Table 5.3 summarizes the performance of the unified classifier. The important outcome is that the unified classifier trained with data from only two transmitters can perform very well on data from four other transmitters.

Table 5.3: Unified classifier's performance; detection accuracy (D.A.) and false positive (F.P.) for four DTV transmitter using the unified classifier trained with WEIU-TV and KTVI data.

	WAOE		WCIA		WICS		WQAD-TV	
	D.A. (%)	F.P. (%)	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.
$P > -65$	100	0	99.8	0	100	0	-	-
$-65 \geq P > -85$	100	0	100	0	99.7	0	100	0
$-85 \geq P > -105$	100	0	100	0	99.9	0	100	0
$-105 \geq P > -114$	99.1	.9	-	-	99.7	1.6	99.6	.8
$-114 \geq P$	97.3	3.2	-	-	97	2.4	95.1	7.6
Overall	**99.3**	**.8**	**99.9**	**0**	**99.7**	**.5**	**99.3**	**1.3**

It is well-known that signal path loss is directly proportional to the logarithm of frequency [63]. However, the approach of considering a unified classifier appears to ignore the difference in path loss between different frequency channels. We argue that in practice, for the limited frequency ranges of our interest, this factor can be ignored in favor of other dominating factors such as the environment and terrain. We show this here and later when we consider a hilly urban/suburban area in Pennsylvania. The success of the unified classifier in detecting attackers in frequencies that differ from its training data (Table 5.3) only validates this assumption for DTVs in the channels 38-50 (614 - 692 MHz). Ideally it is best to have a unified classifier for up to 100 MHz of spectrum. For example, for the current UHF DTV channels in the US (Channel 14-50; 470-698 MHz), one may consider building three classifiers; one for approximately each 75 MHz of spectrum. However, due to practical considerations such as insufficient data or increased complexity, *we argue in favor*

of building only one classifier for the entire 470-698 MHz range. To study this idea, we evaluate the effectiveness of our classifier, which is trained on data from the last third of the UHF DTV spectrum, for detecting attackers operating in frequencies near the first third of the spectrum. For this purpose, we consider the few DTV transmitters in this range in Table 5.4.

Table 5.4: Three DTV transmitters in the 400 MHz UHF channels.

Call Sign	Chan.	Fq. (MHz)	Tx Power (kW)
KNLC (IND)	14	470-476	891
WAND (NBC)	18	494 - 500	347
WYIN (PBS)	17	488-494	301

The performance of the unified classifier on this data is represented in Table 5.5. The results approve our statements about the unified classifier.

Table 5.5: Unified classifier's performance; detection accuracy (DA) and false positive (FP) for three DTV transmitters in the 400 MHz UHF channels.

	KNLC		WAND		WYIN	
	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.
$P > -65$	-	-	100	0	100	0
$-65 \geq P > -85$	-	-	100	0	100	0
$-85 \geq P > -105$	100	0	100	0	99.9	0
$-105 \geq P > -114$	98.8	3.4	100	1.2	98.3	9
$-114 \geq P$	98.6	3.1	-	-	99.3	2.5
Overall	**98.7**	**3.2**	**100**	**.1**	**99.2**	**3.3**

Effect of Attack-Type. In order to evaluate the effect of attack-type on the performance of our detection scheme, we create test datasets that only include normal examples and attacker examples of one type. We next evaluate these datasets using the unified classifier. We studied the four transmitters in Table 5.3. The results for WCIA were identical to the results reported earlier for all three attackers. This can be attributed to the fact that the data from this transmitter are mostly far away from λ and mostly in the first three power brackets, where detection is very accurate and robust. For the other 3 transmitters, we observed that the results in the first two brackets ($P > -65$ and $-65 \geq P > -85$), are identical to the results in the third bracket ($-85 \geq P > -105$). Therefore, we only report the results in the last three brackets for WAOE, WICS, and WQAD-

TV. Figures 5.4 and 5.5 report detection accuracy and false positive rates for these transmitters. The results show decreased detection accuracy and increased false positive rates as the attackers gain more sophistication. Overall, the results show that our scheme performs well even against omniscient attackers.

Figure 5.4: Detection accuracy classified by attacker-type for WAOE (left), WICS (center), and WQAD-TV (right).

Figure 5.5: False positive rates classified by attacker-type for WAOE (left), WICS (center), and WQAD-TV (right).

5.3 Stress Test and Comparison

In this section, we extend the initial evaluations in the relatively flat and detection-favorable Illinois environment, to a particularly unfavorable one, *i.e.* urban/suburban areas in hilly Southwest Pennsylvania. To account for additional shadow fading and signal variations in urban/suburban environments (not represented by Longley-Rice), we probabilistically add extra variations to the predicted signal powers. In a subset of our evaluations, where we simulate wireless microphones to compare our work to the state-of-the art, we use the log-distance path loss and log-normal shadow-fading [63] to model signal propagation.

5.3.1 Hilly Urban/Suburban Area: Southwest Pennsylvania
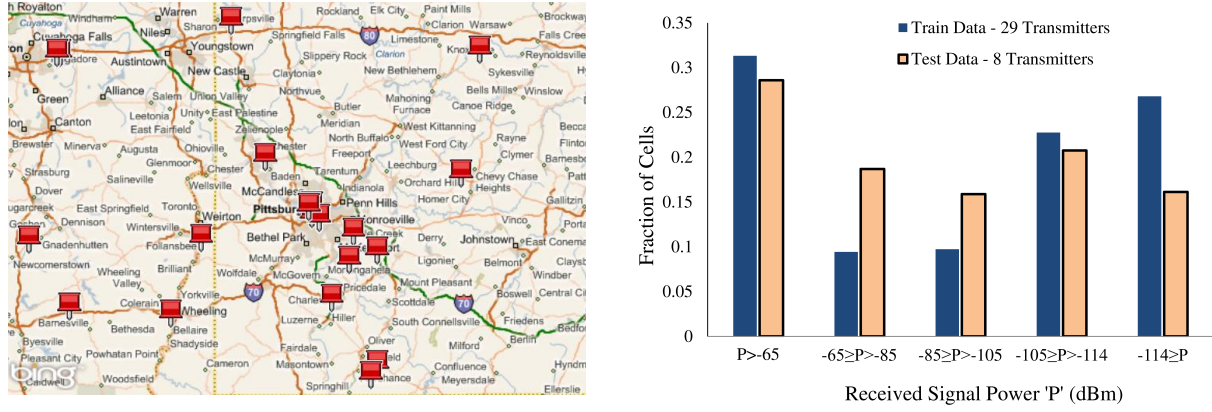
Figure 5.6: (a) Transmitters in parts of Southwest Pennsylvania / East Ohio. (b) Distribution of received signal for the training and testing data in Southwest Pennsylvania.

Table 5.6: Detection accuracy (D.A.) and false positive (F.P.) percentages when variations with dB-spread of σ is added to test data from 8 DTVs. The classifier is trained with data from a disjoint set of 29 DTVs with no added variations.

| | Standard Deviation of Added Variations in Test Data | | | | | | | |
| | $\sigma = 0$ | | $\sigma = 2$ | | $\sigma = 4$ | | $\sigma = 6$ | |
	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.	D.A.	F.P.
$P > -65$	100	0	100	0	100	0	100	0
$-65 \geq P > -85$	100	0	100	0	100	0	100	0
$-85 \geq P > -105$	99.8	.5	99.9	.5	99.8	.8	99.8	1.5
$-105 \geq P > -114$	92.7	6.8	92.2	8.3	91	12	89.2	17
$-114 \geq P$	92.1	9	92.5	9.8	92.4	15	91	21
Overall	97.2	2.9	97.1	3.4	96.5	5.2	96.3	7.3

In this section we evaluate the performance of CUSP when instantiated to a hilly urban/suburban area near Pittsburgh in Southwest Pennsylvania. We focus on signal from all DTV transmitters within 150 mile radius of this 20km by 20km area with estimated received powers higher than -130dBm. This results in a list of 37 DTV transmitters. As before, we use the Longley-Rice model to take into consideration the effect of terrain in signal propagation. In addition, in order to represent un-accounted fading and signal variations in urban/suburban environments, we supplement the data with Gaussian variations mean zero and standard deviation σ (dB-spread) of up to 6dB. This is in line with the log-normal distribution model commonly used in this context [72]. Figure 5.6(a)

depicts the majority of transmitters affecting this area.

We pool the data from different frequencies to obtain a sufficiently large set of training and testing examples across all power levels. To evaluate the performance of CUSP in cases that it is not practical to use the algorithm in Section 5.1.5 to carefully choose the training data, we randomly divide the set of transmitters to subsets of size 29 and 8, for training and testing respectively. We call these 29-DTV and 8-DTV data. We train a unified classifier from the 29-DTV data, and test it on the 8-DTV data. The distribution of the received signal powers for training and testing data are provided in Figure 5.6(b). The cell sizes are 500m by 500m, resulting in a 40×40 grid of cells. The area is assumed to be populated with sensors at the density of 20 per km^2, which is achievable in suburban/urban areas. In particular, this is well below the average house density the Pittsburgh area [11]. The results before adding any additional variations are illustrated in the first column of Table 5.6.

Training and Testing Under Different Conditions. To test the classifier in an extremely unfavorable setting, we add Gaussian variations with mean 0 and standard deviation σ to each power measurement in the test data. The classifier, however, remains trained with the data with no added signal variations. Table 5.6 summarizes the results. It can be seen that despite the significant amount of variation we added to signal propagation data, the classifier still performs reasonably well. As expected, the gradual degradation of performance is explained by the difference of examples that the classifier is trained with and those on which it is being tested. In particular, the relatively high false positive rates at high variation levels reflect the case that some of the variations seem 'too much' to the classifier, and therefore it mistakenly classifies them as malicious.

In general, the effectiveness of our approach can be reduced in environments with considerable natural variations in signal power within short distances. The reduced effectiveness presents itself as lower detection accuracy and higher false positive rates compared to environments in which signal propagation is 'smoother.' This is attributed to the descriptive power of our choice of features; there might be neighborhoods in which the classifier has difficulty differentiating between

63

significant natural variations and an unusual signal propagation pattern created by the false reports of attackers. At a high-level, a remedy would entail modifying the feature space to increase its descriptive power. As an item of future work, we consider adding elevation data to the feature space to improve the classifier's performance (see Section 5.4). In addition, the cell-size may be optimized for maximized classifier performance.

Effect of Attack-Type. Table 5.6 provides results for an equal mix of the three attack-types (note that we assume each cell is occupied by attackers of one type only). Here, we break the results by the type of attack. Since the false positive results are similar to those of Table 5.6, we only provide results for detection accuracy. The results are summarized in Table 5.7. It can be seen that the classifier provides respectable detection accuracies, even for the most difficult scenarios, that is defending against omniscient attacks in a hilly area with added variations of up to 6dB.

Table 5.7: Breakdown by attacker type; detection accuracy (D.A.) when variations with dB-spread of σ is added to the test data from 8 DTVs. The classifier is trained with data from a disjoint set of 29 DTVs with no added variations. Uncoordinated, coordinated, and omniscient attacks are represented by UC, CO, and OM.

	Standard Deviation of Added Variations in Test Data											
	$\sigma = 0$			$\sigma = 2$			$\sigma = 4$			$\sigma = 6$		
	Type of Attacker											
	UC	CO	OM	UC	CO	OM	UC	CO	OM	UC	CO	OM
$P > -65$	100	100	100	100	100	100	100	100	100	100	100	100
$-65 \geq P > -85$	100	100	100	100	100	100	100	100	100	100	100	100
$-85 \geq P > -105$	100	100	100	100	100	100	100	100	99	100	100	99
$-105 \geq P > -114$	97	93	88	97	93	88	95	91	88	93	89	87
$-114 \geq P$	92	87	84	92	87	84	91	85	84	89	85	84
Overall	98	96	95	98	96	95	97	96	95	97	95	94

5.3.2 Comparison to Model-Based Scheme

The solution in Chapter 4 requires knowledge of the parameters of the log-normal shadowing model in order to detect compromised cells. We are not able to evaluate that approach in the evaluation environment of this chapter, since that approach only works with the assumption of using the log-distance path loss and log-normal shadow-fading. In order to provide a fair comparison, we evaluate the data-based approach in an environment similar to that of the model-based approach.

The signal power at node N_i is written as $p_i = p_t - (10 \log_{10} r_i^\alpha + S_i)$ where p_t is the transmit power of the primary, r_i is the distance from N_i to the primary transmitter, $10 \log_{10} r_i^\alpha$ represents the path loss with exponent α (typically $2 < \alpha < 4$), and $S_i \sim N(\mu_s, \sigma^2)$ is the loss due to shadow-fading. μ_s is often considered to be 0, and the dB-spread σ independent of the distance to the transmitter (typically $2 \leq \sigma \leq 6$). Therefore we have $p_i \sim N(\mu(r), \sigma^2)$, where $\mu(r) = p_t - (10 \log_{10} r_i^\alpha + \mu_s)$.

Note that the simulation setup and parameters are chosen based on the model-based scheme we simply replicate them here in a larger scale. The simulation environment is an 8192m $\times$ 8192m area in which secondary users are deployed uniformly at random with the density of 0.0008 per square meter. The area is divided into $64 \times 64 = 4096$ square cells of size 128m $\times$ 128m each. Therefore, the expected number of nodes per cell is about 13. Depending on the scenario, primary transmitters with power ranging from 17dBm to 20dBm are placed at different locations in this area to represent wireless microphone primaries. The detection threshold is $\lambda = -74$dBm, $\alpha = 3$ and the standard deviation for the fading and shadowing process, $\sigma = 3$ (in dB scale).

The results are summarized in Table 5.8. It can be seen that the data-based approach outperforms the model-based approach in terms of detection accuracy, however this comes at the cost of moderate false positive rates. Note that data-based approach does not use any information about the nature or specification of signal propagation model, whereas the model-based detection approach requires knowledge of λ, α, and primary powers.

Table 5.8: Model-based vs data-based.

	Fraction of Cells	Model-based		Data-based (Classification)	
		D.A.	F.P.	D.A.	F.P.
$P > -55$	.02	81	0	100	0
$-55 \geq P > -65$	.04	95	0	100	0
$-65 \geq P > -74$	.14	67	0	95	7.4
$-74 \geq P > -80$	.29	85	0	96.7	7.6
$-80 \geq P > -85$	.30	99	0	100	0
$-85 \geq P$	.19	100	0	100	0
Overall	1	89.0	0	98.3	3.5

5.4 Conclusions and Future Work

In this chapter we presented CUSP, a new technique for detecting such attacks while aggregating spectrum sensing data from white space devices spanned over large regions. Our approach uses classification techniques based on SVMs with quadratic kernels to learn to differentiate between natural and un-natural signal propagation patterns in the region of interest. We evaluated the performance of CUSP using real-world transmitter, terrain, and sensor density data from two regions in the US. We showed that CUSP can achieve high detection accuracies even in the most unfavorable situations, *i.e.* hilly urban/suurban areas with significant amounts of additional signal uncertainty.

Multi-Resolution Analysis. In the future, we will enhance the approach to detect attacker-dominated cells at different resolutions. A high-resolution view entails dividing existing cells to smaller cells, whereas a low-resolution view allows for considering a set of neighboring cells as one cell. This enables detecting attackers at a fine level, coarse level, or those that are able to dominate multiple adjacent cells.

Elevation Data as Features. We will add elevation data as features to the training and testing data. This will provide the classifier with more information to learn and decide whether an observed signal propagation pattern is natural. Our preliminary experiments with this approach show improvements of performance in areas with irregular and hilly terrain.

Chapter 6

Trust-Based Protection using Remote Attestation

In this chapter we consider robust radio spectrum telemetry, with focus on the case where a subset of the sensors can be remotely attested. We propose a practical framework for using statistical sequential estimation coupled with machine learning classifiers to deter attacks and achieve quantifiably precise outcome. We provide an application-oriented case study in the context of spectrum measurements in the white spaces. The study includes a cost analysis for remote attestation, as well as an evaluation using real transmitter and terrain data from the FCC and NASA[1].

6.1   Motivation and Approach

Consider Figure 6.1 as a part of the region of interest for performing reliable aggregation of spectrum measurement data. There exist two types of nodes; attestation-capable nodes (triangles), and regular nodes (circles). In any particular cell, the goal is to obtain an estimate of the signal power in that tile, and compare it to a primary detection threshold to determine whether the channel is unused. Assume for now that we have performed remote attestation on all attestation-capable nodes and have excluded those we believe are compromised. Therefore, the remaining attestation-capable nodes are considered *trusted* or *attested*. For regular nodes, however, we do not have any prior information regarding their legitimacy.

Consider tile A in Figure 6.1 in which about half of the nodes are attested. One may argue that the high number of reliable nodes provides enough diversity to absorb the variations due to path loss and shadow-fading, and therefore there is no need to include the results of regular nodes.

---

[1]The majority of the material in this chapter is adopted from Fatemieh, LeMay, and Gunter's recently submitted manuscript [33].

Figure 6.1: Illustration of a few cells with attestation-capable and regular nodes.

This approach is probably safer than using the values from the regular nodes that may include compromised nodes, but what if the rest of the regular nodes are also legitimate? Is the safety worth the reduced precision? How would we determine whether it make sense to rely only on trusted nodes, or we should use the data from regular nodes as well? And if so, which ones?

Now consider tile B where unlike tile A there are very few reliable nodes. Therefore, there is a high chance that aggregating the measurements from such a small number of nodes does not provide enough diversity to obtain a precise measurement (estimate) of the signal power. A similar situation can be seen in tile C; not only do there exist few attested nodes, but their positioning also makes it more likely that they do not provide enough diversity as they may all be behind an obstruction that attenuates the signal. Therefore, it seems necessary to include results from at least some of the regular nodes. But what if some or all of them are compromised, and they skew the results to achieve their malicious goal instead of adding legitimate diversity?

6.1.1 Key Issues and Overview

The examples above underline the importance of addressing the following issues. First, there must be a systematic strategy to determine when we have achieved enough diversity in the results that we can stop collecting additional data within a cell. Second, if we decide we need additional data beyond those from attested nodes, there should exist a strategy to decide which nodes to include. Third, for each cell in which additional regular nodes are added to the data 'pool,' we need a strategy to ensure that the added nodes are not dominated by attackers. The rest of this section aims to address the above concerns.

At a high level, our approach consists of three main phases, as summarized in Algorithm 3. First, within each tile we rely on basic statistical inference to aggregate data from all of attested sensors as well as 'enough' regular nodes to achieve the application-specified precision goal. Note that we only include the least required regular sensors to limit the unnecessary exposure from untrustworthy data. As will be discussed later in this section, various inclusion strategies may be used for this purpose. Second, the regular nodes that were included in the aggregation process in the cell are compared against the data from the reliable nodes of the neighboring cells. This process is performed using machine learning SVM classification to detect irregular signal propagation patterns that most likely represent a coordinated misreporting attack. Third, as will be detailed in the rest of this section, an aggregate (*e.g.*, mean or median) is calculated.

Algorithm 3 Simplified Overview of Approach (for Each Cell)

Input:
(1) *Green Data*: measurements from attested nodes,
(2) *Yellow Data*: measurements from regular nodes,
(3) *Strategy*: strategy for including data from regular nodes

Phase 1: Node Selection
Add *Green Data* to aggregation *pool*
while $\big($!*Satisfy-Precision-Requirements* (*data* in *pool*)$\big)$
   if $\big($*size(Yellow Data)*$> 0\big)$
      Move-Next-Element-To-Pool(*Strategy, Yellow Data*);
   else
      Remove all *Yellow Data* from *pool*; Go to **Phase 3**
end

Phase 2: Attack Detection
Yellow Suspects $\leftarrow$ *Yellow Data* in *pool* from **Phase 1**
Green Neighbors $\leftarrow$ averages of *Green Data* in the neighboring cells (*i.e.* 8 numbers)
attack = SVM-Detection(*Yellow Suspects*, *Green Neighbors*)
if (*attack*)
   Remove all *Yellow Suspects* from *pool*

Phase 3: Aggregate Calculation
Compute aggregate based on data in *pool*

6.1.2 Intra-Cell Node Selection

The two main options for aggregating measurements in a cell are calculating the average (EGC) or median of the data (observations). A collection of observations is referred to as a sample. The goal is to use all of attested nodes plus a dynamically selected set of regular nodes such that we can ensure the computed aggregate is within a pre-defined distance of the real mean or median for the signal in the cell. The median has a key advantage over the mean as an aggregate; it is less vulnerable to natural outliers or attacker nodes that constitute a minority of nodes in a cell [31, 75]. However, computing the sample median with a pre-specified confidence interval requires more data (compared to mean). Or dually, with a fixed number of observations, the confidence intervals achieved for the median are larger than those computed for the mean. We present the corresponding calculation procedures in the next section.

However, if the attackers obtain even a weak majority in a cell, they can move the median to their desired number while being less 'abnormal.' Figure 6.2 illustrates this observation. The additional abnormality facilitates detecting them in Phase 2 of our approach, and therefore may be desired. Hence, we will rely on median when the attested nodes represent the majority of nodes in the cell and rely on the mean otherwise. This approach is detailed in Section 6.1.5.
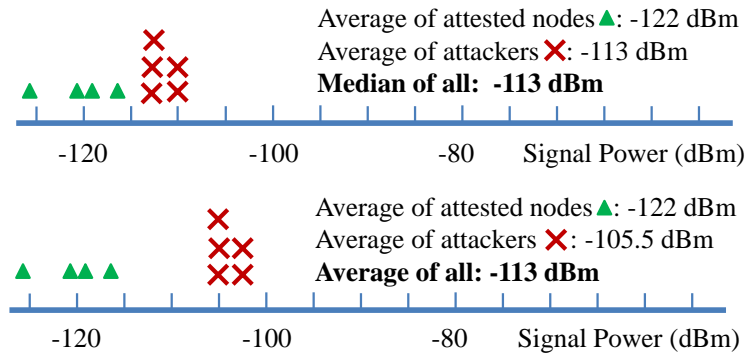

Figure 6.2: A simplified illustration of why attackers are forced to deviate more when they can only affect the mean rather than the median.

6.1.3   Using Statistical Inference to Ensure Precision

For many applications, including aggregation of spectrum sensing data, it is not clear in advance how many sensors (observations) should be used in each aggregation effort in order to achieve the desired precision in the (estimation) outcome. Instead, data is evaluated as it is collected, and further sampling is stopped in accordance with a pre-defined *stopping rule*. This process is also referred to as *sequential estimation*. In our case, we aim to achieve an acceptable precision in the results while using as few data points from regular nodes as possible. We argue that applying basic sequential estimation for achieving fixed width confidence interval for the estimated aggregate is an ideal tool to achieve our goal. By stating the acceptable *margin of error* (half the width of a confidence interval) for the quantity being estimated, the application can ensure with high confidence that the estimated outcome from the sample data is 'close enough' to the true value. In other words, with high confidence (*e.g.* 95%), it can be assured that the true mean (or median) is within a γ margin of error from the estimated value (*e.g.* $\gamma = 3dB$). This is also referred to in the form of a *coverage probability* (*e.g.* $.95 = 1 - \alpha$).

We first focus on a sequential procedure for finding fixed-width confidence intervals for the mean. Let $x_1, x_2, ...$ be a sequence of independent and identically distributed (i.i.d.) random variables having an unknown density function $f(x), x \in R$. The i.i.d. assumption is not absolutely true for sensors that are very close and face correlated shadowing; however in view of practical considerations we proceed with this assumption, which is in-line with the commonly used log-normal shadowing model. Let μ and σ^2 represent the mean and variance of density function $f(x)$. It is known that no fixed-sample size procedure will provide a fixed-width confidence interval for μ having a prescribed coverage probability at the same time. The famous Chow-Robbins procedure for sequential estimation defines the following stopping rule for a confidence interval of size 2γ:

$$N = \inf\{n \geq n_0, n \geq a^2 \gamma^{-2} s_n^2\}$$

where $n_0 \geq 2$ is the initial sample size, $a = z_{(1-\alpha/2)}$ is the $100(1 - \alpha/2)$ percentile of the standard

71

normal distribution $N(0, 1)$ (*e.g.* if $\alpha = .05$ then $a = 1.96$), and s_n is the sample standard deviation of n observations. The Chow-Robbins procedure is asymptotically tight, in the sense that the coverage probability is asymptotically $1 - \alpha$, and is also asymptotically efficient in the sense that the average required number of samples is asymptotically equal to an optimal fixed-sample procedure with *known* σ^2 [37].

Now we turn to the median. We begin by placing the measurements in order, that is: $x_{(1)} < x_{(2)} < ... < x_{(n)}$. The goal is to find an interval $x_{(a)} < m < x_{(b)}$ such that $P(x_{(a)} < m < x_{(b)}) = 1 - \alpha$, where $1 - \alpha$ is the desired probability that the interval captures the median.

In order to have $x_{(a)} < m$, at least a of the observations must fall less than m, and in order to have $m < x_{(b)}$, at most $b - 1$ of the observations must fall less than or equal to m. Since m is the median and since the distribution of the X's is continuous, we have

$$P(X < m) = P(X \leq m) = .5.$$

Assuming independent observations, the probability that at least a and at most $b - 1$ of the observations fall less than m is given by the binomial probability with $p = .5$, that is $\sum_{k=a}^{b-1} \binom{n}{k} (.5)^n$. To construct a $100(1 - \alpha)\%$ confidence interval for m, we choose a and b so that this sum is $1 - \alpha$. For large samples, approximate values of a and b may be found by using the normal approximation to the binomial distribution. We may obtain a and b by solving for them in the following equations [41]:

$$\frac{a - .5n}{\sqrt{.25n}} = -z_{(1-\alpha/2)}, \quad \frac{b - 1 - .5n}{\sqrt{.25n}} = z_{(1-\alpha/2)}$$

Note that both the confidence intervals were calculated by assuming the distribution of the original population is unknown.

6.1.4 Inclusion Strategies

We consider three *inclusion strategies* for including regular nodes in the aggregate computation in each cell.

Random: Randomly adding data from regular nodes to the data from attested nodes has the advantage that it is in-line with the assumptions made in computing the confidence intervals. In addition, the randomness reduces the attacker's chances of selectively compromising nodes and carefully crafting false measurements with minimum abnormality.

Geo-Diverse: By selecting a geographically diverse set of regular nodes, we add diversity to the results and reduce the chances of selecting (regular) nodes that are experiencing similar shadowing effects. To achieve this goal, we use the widely cited Gudmundson shadow correlation model [38]. According to this model, the correlation in shadow-fading in distance Δx is represented as:

$$R(\Delta x) = e^{\frac{-\Delta x}{d_{corr}}}$$

with the correlation length d_{corr} dependent on the environment. Empirical studies suggest values between $25m$ to $120m$ for urban areas [16]. Using this model, we suggest the following greedy approach to adding nodes to the aggregation pool. Before each addition to the pool, we compute the aggregate correlation of all nodes already in the aggregation pool with the candidates to be added to the pool. At each step, we add the node with the least aggregate correlation with existing nodes.

Biased: In this approach, we sort the data from the regular nodes in the increasing order of the absolute value of their difference to the median of the attested nodes. At each step, we move values to the aggregation pool according to their rank in the sorted list. This approach has the disadvantage that creates a 'bias' in the aggregate calculation process, which makes the computations in Section 6.1.3 inaccurate. However, in many cases, this bias effectively works as an implicit weighting mechanism in situations where attackers have only compromised a subset of the regular nodes. In those situations, this approach may limit the number of measurements from compromised nodes that will be included in the final result (see the results in Section 6.2).

6.1.5 Aggregate Selection and Inter-Cell Attacker Detection

Based on the earlier observations regarding outlier resilience, the width of confidence intervals, and ease of detecting coordinated attackers, we finalize Algorithm 3 as follows. In Phase 1, we start by considering the data from all of the attested nodes in the aggregation pool and initially use *median* as the aggregator. Consider a cell with k attested nodes. We iteratively add up to $k - 1$ elements from regular nodes to the aggregation pool according to the desired inclusion strategy. After each addition, if the margin of error for median is reduced to a value lower than γ, we transition to Phase 2. If this condition is not met at any point and there exist additional measurements, we switch to using *mean* as the aggregator. Again, we continue adding new data from the regular nodes to the aggregation pool until the stopping rule is satisfied. If so, we transition to Phase 2. Otherwise, if adding all of the regular nodes does not result in satisfying the stopping rule, we simply ignore all the added regular nodes and compute the median of attested nodes as the aggregator.

In Phase 2, we have an aggregate (either mean or median) from data provided by all of the attested nodes, as well as *some or all* of the regular nodes in the cell. In this phase, we aim to ensure that the regular nodes are not mounting an exploitation or vandalism attack. To that end, we separate the data points from those regular nodes that have contributed to the aggregate ('yellow suspects') and compare them to the data from attested nodes in the neighboring cells. More specifically, we first compute the average of the reports from the yellow suspects. Next, we consider the averages of attested nodes in each of the eight neighboring cells (see Figure 6.3). We refer to these nodes as 'green neighbors.' We use these 9 data points in a 9-element tuple; the first element represents the average power from yellow suspects in the cell under investigation, and the next 8 cells represent the difference between averages of yellow suspects and each of the 8 green neighbors.

Having obtained the 9-element representation for each cell, we feed it to an attacker-detection classifier. The classifier is pre-built in a one-time process using an initial trusted set of data. Building such a classifier has been discussed in detail in earlier work [32], and has proven to effectively detect attacker-dominated regions in regular settings where there is no separation between regular
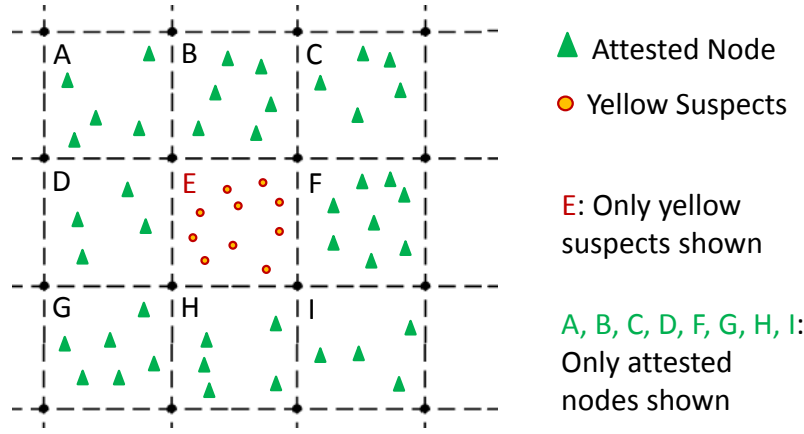
74

Figure 6.3: Classification-based attacker detection setting: regular nodes included in the aggregation for cell E and attested nodes from neighboring cells.

and attestation-capable nodes. In our setting, however, we know that none of the 8 neighboring cells is dominated by attackers. We suggest using the same classifier for detecting whether the yellow suspects in a cell look abnormal compared to the green neighbors. If the classifier considers the data to be anomalous, we only rely on the median of the attested nodes in that cell. Otherwise, the aggregate computed in Phase 1 (using a mix of attested and regular nodes) is valid and should be used as the representative signal power in that cell.

6.2   Evaluation

We evaluate our system using predicted signal propagation data obtained from real transmitters and terrain data. More specifically, the TV transmitter location, signal power, height, and frequency is obtained from FCC databases and terrain (*i.e.* elevation for any given point) is obtained from NASA databases [5]. We choose the FCC-endorsed Longley-Rice empirical outdoor signal propagation model to generate predicted signal power for any location and frequency of interest. Longley-Rice takes into account the effects of terrain as well as transmitter's power, location, frequency, and height. To account for additional uncertainties due to factors such as shadow-fading we add log-normal variations with a mean of zero and a standard deviation (dB-spread) of $\sigma_{dB} = 6$ to the predicted signal power for each point [72]. For evaluation purposes, we consider this data as

the ground truth.

We instantiated our evaluation to a flat rural/suburban area surrounding Champaign, Illinois and a hilly urban/suburban area surrounding Pittsburgh, Pennsylvania. We only provide results for Pittsburgh due to space constraints, since in almost all aspects our approach is challenged more in that area. The following points in (*latitude, longitude*) format define the southwest and northeast corners of the considered 20km × 20km square area in Pennsylvania: ⟨ (40.35, -80.12), (40.53, -79.884)⟩. Each cell is 1km × 1km. We focus on signals from all DTV transmitters within a 150 mile radius of this area with estimated received powers higher than -130dBm. This results in a list of 37 DTV transmitters. Guided by approximate sample size requirements based on methods in Section 6.1.3, we consider nodes to be scattered with an expected density E_d of 50 nodes per cell. To add variation and randomness, we consider the number of nodes to be normally distributed with a mean of E_d, and a standard deviation of 10. Such densities will be easily achievable in urban areas, and need to be achieved through provisioning or other means in suburban areas for our approach to be effective.

6.2.1   No-Attack Performance

We first evaluate the precision of predictions generated by our approach when there is no attack. We compare the aggregate produced by our approach to the ground truth (real average power in the cell). In Figure 6.4(a) we show the percentage of cells for which the real average power is within the chosen margin of error $\epsilon = 3$dB from the calculated aggregate.

As a second performance metric in the absence of attacks, we introduce the *false outcome rate*, representing the fraction of un-attacked cells with ground truth power above (below) the primary detection threshold of -114dBm that due to errors in our approach are mistakenly assigned an aggregate below (above) -114dBm.

Figure 6.4: No attack; percentage of cells with ground truth average within the margin of error from the calculated aggregate (left) and false outcome rate (in percentage) as a function of the fraction of attested nodes (right).

6.2.2 Performance against Omniscient Attackers

To gauge performance in the presence of attacks, we simulate omniscient (and coordinated) attackers that perform exploitation and vandalism attacks. Attacker nodes act in cooperation and know the exact number, measurements, and type of all the other nodes, as well as the inclusion strategy in use (Random, Geo-diverse, or Biased). In cells where the ground truth is below the -114dBm threshold, they cooperate to perform exploitation to change the aggregate to a value above the threshold. Similarly, in cells where the ground truth is above -114dBm, they aim for vandalism by moving the aggregate to a value below the threshold. In the both cases, the attackers minimize the deviation of their false reports from the measurements of un-compromised nodes by choosing to report values that move the aggregate slightly below (above) the threshold (.5 dB here) in order to perform exploitation (vandalism). This maximizes their chances of being included in the aggregate pool in Phase 1 and minimizes their chances of being detected in Phase 2. If the attackers conclude that the protections in Phase 1 do not allow them to 'flip' the aggregate, they refrain from reporting false reports to avoid detection.

To evaluate effectiveness against omniscient attacks, we introduce the *deterrence rate*. This metric represents the fraction of attacks by omniscient attackers that our approach thwarts. Deterrence may occur in phase 1 (by partial or total exclusion from the pool), or in phase 2 where their attack is detected by the classifier. We use data from 29 of the transmitters to build a unified classifier for

the region [32] and test deterrence of attacks on the remaining 8 channels. The deterrence rates for cases with average attested fractions ranging from .15 to .35, and average attacker fraction ranging from .25 to .85 are presented in Figure 6.5. For attested fractions higher than .35, our results (omitted due to space constraints) show that it is more beneficial to avoid the complexities of our approach and only rely on the average of attested nodes.


Figure 6.5: Attack deterrence rate (in percentage) when the average fraction of attested nodes is .15 (left), .25 (center), and .35 (right).

In Figure 6.5, a surprising phenomenon can be seen in the case of Biased attacks. In some cases, when the attested fraction is increased (particularly from .25 to .35), the deterrence rate decreases. While this can be considered a flaw for the biased scheme, it can be described as follows. When the attested fraction is increased, there is less competition from regular un-compromised nodes (for attacker nodes) to report values close to the average of attested nodes and enter the aggregation pool. Therefore, the attackers have a higher chance of entering the pool with false reports, influencing the results, and passing Phase 1. The results in Figure 6.6 show this observation; unlike Random and Geo-diverse cases in which the deterrence at phase 1 does not change or increases as the attested fraction increases, the rate decreases for the Biased strategy.

Overall, the results show the following. (1) All three approaches are highly effective against omniscient attacks, even in cases where a small fraction of nodes (.15) are attested. (2) In terms of attack deterrence, the Biased inclusion strategy outperforms others. This is particularly true with lower attested and attacker fraction. This can be attributed to the difficulty of influencing the aggregate by attackers in these situations, since the attacker has to fulfil two conflicting goals of reporting values close to the attested average (to be included in pool) and at the same time far from

Figure 6.6: The fraction of attack deterrences in Phase 1. For each bar with value x, $1 - x$ is the fraction deterred in Phase 2. The average fraction of attested nodes is .15 (left), .25 (center), and .35 (right). Results for Geo-diverse (similar to Random) are omitted.

the attested average (to move the aggregate and perform attack). (3) The relative outperformance of the Biased approach comes at the price of relatively higher false outcome rates when there is no attack.

6.3 Cost Considerations

Remote attestation can introduce potentially significant additional costs into a system. This section briefly surveys these costs for implementations of two remote attestation architectures. The first uses a TPM, which is a distinct coprocessor, whereas the second is implemented primarily in software, requiring only small hardware adaptations. The TPM-based architecture represents an upper bound on the cost of attestation, since the TPM is intended for use in desktop PCs with practically unlimited power supplies. The software-based architecture represents a low-cost alternative, although hardware and software innovations may result in architectures with even lower costs. The reason we include this section is to emphasize the fact that attestation introduces significant costs, which motivates our approach to leveraging relatively few attested nodes to establish trust in spectrum sensing results. The specific tradeoff between trust and cost can be made on a case-by-case basis.

Costs arise from various sources. Remote attestation support often requires additional hardware resources, which increase *manufacturing* costs. Some schemes involve a coprocessor, and even those primarily implemented in software may necessitate larger memories to store their code and

data. Additional *energy* may be consumed by several components involved in a remote attestation transaction. Coprocessors and processors executing software routines both consume energy. Additionally, coprocessors usually consume some energy when inactive, and enlarged memories may require additional energy. Remote attestation transactions also increase the amount of data that is transmitted and received over the network, which may increase the energy consumption of the wireless radio in addition to its obvious *network bandwidth* cost. Increased network utilization can also impose *time* costs, as can remote attestation transaction processing.

We evaluated an Atmel AT97SC3203 TPM installed in a desktop PC. It imposes a manufacturing cost for the TPM chip itself, and potentially for expanded memories to support interface software installed on the attested processor. We measured its energy consumption using a Digital Multi-Meter (DMM). It draws 10.6mW of power when idle, which is likely to account for the bulk of its total energy consumption. It consumes around 58.9 mJ when an attestation certification is generated. Other operations require some energy, but are unlikely to contribute significantly to total consumption either due to their infrequent invocation or the fact that they do not involve expensive routines such as digital signature generation. Attestation operations require around 1.1 second to execute and generate at least 296 bytes of uncompressed data if the TPM uses a 2048-bit RSA key and the 160-bit SHA-1 hash algorithm, regardless of the specific protocol in use. For reference, we measured the energy consumption of a Digi XBee 802.15.4 radio using an oscilloscope, and determined that transmitting a packet with an x-byte payload consumed about $(0.017x + 1.83)$ mJ of energy at 1mW.

We also evaluated a software-based attestation scheme on an Atmel AVR32 AT32UC3A0512 microcontroller [53]. It only consumes extra energy when it is active. It uses Elliptic-Curve Cryptography (ECC) rather than RSA, which uses shorter keys (192 bits in this prototype) and simpler computations. Thus, although it does not use any hardware accelerators such as those in the TPM, it still consumes similar amounts of energy during attestation operations. Each operation actually only takes about 0.6 seconds to execute. Due to the significantly shorter keys, each attestation operation only generates at least 44 bytes of data.

6.4 Conclusions and Future Work

The use of statistical sequential estimation and classification methods can help evaluate and improve the trustworthiness of spectrum sensing results generated by a network containing a limited number of attested sensors. These methods reduce the total cost incurred by attestation. Our evaluation determined that the Biased node inclusion strategy is the most effective at deterring attacks, but also generates more false positives than Random or Geo-diverse strategies. All three strategies result in substantial attack deterrence. These are not the only strategies that can be used, and future research should evaluate other strategies.

Another direction for further research is developing framework for formulating costs associated with including regular and attested nodes, and systematically striking a balance between the costs (from spectrum data aggregation and remote attestation) and achieving robust aggregation results. Solving the following minimization problem may yield one such solution:

$$\min_{\alpha_\Delta^i, \alpha_\circ^j} \left[DISTANCE\left(\frac{1}{S}\left[\sum_{i=1}^{N} \alpha_\Delta^i \Delta^i + \sum_{j=1}^{M} \alpha_\circ^j \psi(\circ^j)\circ^j\right], P_{real}\right)\right.$$
$$\left. + C_\Delta \sum_{i=1}^{N} \alpha_\Delta^i + C_\circ \sum_{i=1}^{M} \alpha_\circ^i\right],$$
$$S = \sum_{i=1}^{N} \alpha_\Delta^i + \sum_{j=1}^{M} \alpha_\circ^j \psi(\circ^j)$$

where $\alpha_\Delta^i, \alpha_\circ^i \in \{0, 1\}$ denote whether an attestation-capable node i or regular node j will be included in the aggregate calculation, and N, M are the total number of attestation-capable and regular nodes in the cell, respectively. Δ^i is the measurement by the i-th attestation-capable node and $\circ^j$ is the measurement by the j-th regular node in the cell. P_{real} is the ground truth average power in the cell; since it is not known, we will substitute it with a representative value (*e.g.* average of neighboring cells). $\psi(\circ^j)$ denotes the probability that the reading $\circ^j$ by a regular node is legitimate (and will be estimated by an SVM classifier). C_Δ is the cost associated with using each attestation-capable node, and $C_\circ$ is the cost associated with using each regular node. Note

that these costs are uniform within each class of nodes.

Chapter 7

Related Work

This chapter covers the related works in two general categories: white space networks, and sensor and ad-hoc networks. We also briefly discuss related work in the context of remote attestation in distributed sensing, which complements the background material on remote attestation in Section 2.1.3.

7.1 White Space Networks

Most prior work in the context of white space networks considers identifying individual attackers within a cell as part of collaborative sensing. Such approaches are not capable of detecting cells that are dominated by attackers. Min *et al.* [59] group sensors in a neighborhood to clusters (cells), and exclude or minimize the effect of abnormal sensor reports using shadow fading correlation-based filters. However, it fails to detect attackers that constitute more than $1/3$ of the population of the nodes in a cell. Kaligineedi *et al.* [45] address a similar problem by pre-filtering outlying sensing data, and a strategy to assign trust factors to nodes for weighting measurements and potentially omitting some nodes. In addition to the general problems enumerated with outlier-detection techniques, the attacker model is too simplistic and falls short in cases where attackers constitute a large fraction of nodes in a cell, or employ sophisticated misreporting strategies.

also consider malicious false reporting in collaborative sensing. They propose a reputation-based data fusion technique named weighted sequential probability ratio test which is based on the sequential probability ratio test. The proposed scheme, however, heavily depends on apriori knowledge of the reporting values of radios given the true state of the world. It also does not

account for spatial variability of spectrum availability and only focuses on detection in a small region. In addition, the proposed mechanism is limited to hard-combining collaboration techniques, whereas we consider soft-combining techniques to obtain more information from each node.

Chen *et al.* [25, 26] propose a weighted, reputation-based data fusion technique based on the sequential probability ratio test. Their approach only considers hard 0/1 decisions from each node, requires prior knowledge about the false positive and false negative ratios at each node. In addition, it does not account for spatial variability of spectrum availability and only focuses on detection in a small region. Therefore, it cannot detect attacker-dominated regions. Lee *et al.* [52] have considered clustering of nodes and weighted collaborative sensing, but they do not particularly focus on security aspects of collaborative sensing.

Chen *et al.* [27] consider primary user emulation attacks in which an attacker may modify the air interface of a radio to mimic a primary transmitter signal's characteristics, thereby causing legitimate secondary users to erroneously identify the attacker as a primary user. They propose LocDef, which utilizes both signal characteristics and location of the transmitter to verify primary transmitter signals. An alternative is using cryptographic and wireless link signatures to authenticate primary users' signal in presence of attackers that may mimic the same signal. Liu *et al.* [58] achieve this by using a helper node close to a primary user to enable a secondary user to verify cryptographic signatures carried by the helper node's signals and then obtain the helper node's authentic link signatures to verify the primary users signals. We consider this problem to be complementary to the problem we address.

A number of proposals motivate and identify various security issues in cognitive radio networks [22, 28, 67]. Although the attacks we consider, or a variation of them, are mentioned in these works and high-level ideas for mitigating them are proposed, none of them provide detailed solutions for addressing them.

7.2 Sensor and Ad-Hoc Networks

There exists a rich body of related work on this topic in the sensor network security literature. We consider the following to be the most relevant ones. Wagner introduced *resilient aggregation* [75], where he studies resilience of various aggregators to malicious nodes in an analytical framework based on statistical estimation theory and robust statistics. However, his work is limited to small regions and does not consider attacker detection as we do. Zhang *et al.* [77] propose a framework that identifies readings not statistically consistent with the distribution of readings in a cluster of nearby sensors. Their proposal, however, is local, that is only works for a small region. For example, it is not able to handle situations where attacker can compromise a large fraction of the nodes in a cluster. It also assumes the data comes from a distribution in the time domain, which in not a valid assumption in our domain. Hur *et al.* [44] propose a trust-based framework in a grid in which each sensor builds trust values for neighbors and reports them to the local aggregator. Our work is similar to this work in that it is based on a grid. Their solution, however, does not provide a global view for a centralized aggregator, and also cannot identify compromised 'regions.' They also do not consider uncertainties in the data. An avid reader may refer to following list for additional resources in the related are of secure data aggregation in wireless sensor networks [17, 24, 35, 42].

Insider attacker detection in wireless networks is another area of related work. This problem has been explored in a general setting [20, 43, 78] as well as more specific contexts such as insider jammers [50]. As an illustrative example in the general context of sensor networks, Liu *et al.* [56] propose a solution in which each node builds a distribution of the observed measurements around it and flags deviating neighbors as insider attackers. This work is again local and peer to peer and does not work in areas with more than 25% attackers. Krishnamachari *et al.* [49] consider fault tolerant event region detection in sensor networks using a Bayesian framework. This work differs from our work in that it only considers faulty nodes that are not necessarily malicious, the faulty nodes are assumed to be uniformly spread, and the node itself participates in the detection process.

Ganeriwal *et al.* [35] propose a reputation-based trust framework, where each sensor maintains

a local reputation and trust for its neighbors. This work is very general, and is mainly focused on local decision making at each sensor. It is also local and peer to peer, meaning that the reputation is typically considered to be updated based on the quality of pairwise interactions between nodes.

Remote Attestation   There has been a number of works on utilizing remote attestation capability to achieve security in sensor networks. For example, there has been efforts on proposing architectures and building platforms [66], detecting compromised nodes [76], and other activities such as secure code update and key establishment [65]. To the best of our knowledge, no prior work has considered the problem of using attestation to defend against malicious false reports by omniscient attackers in the context of white-space distributed spectrum measurement.

Chapter 8

Conclusions

In this dissertation, we motivated the importance of protection against exploitation and vandalism attacks in the context of spectrum telemetry in white-space networks. This importance is magnified in view of the growing interest in using white-spaces for a wide range of applications. Unlike the thin body of related work on this subject, we focused on a range of attacks, including those originating from omniscient attackers that may constitute the majority of nodes in a small area, and launch carefully crafted attacks. To defend against these attacks, we offered a problem formulation based on a grid of small cells and three general solutions.

The first solution, also called model-based, relies on models for signal propagation and shadow-fading to build an attacker-detection model based on outlier detection. Attackers are either individually detected inside each cell, or detected as a group by corroboration among neighboring cells in a hierarchical structure. This approach, while shown to be effective against attacks, suffers from at least two shortcomings. First, it relies on fairly accurate knowledge about signal propagation and shadow fading models and parameters in order to succeed, which is an unrealistic assumption. This is particularly true in urban areas or hilly terrains. Second, while some of the threshold parameters for attacker detection are automatically derived by the proposed approach, there exists one important threshold that requires conjecturing and manual tuning for each given region, which makes it impractical.

The second solution, CUSP, is a data-based technique that aims to address the limitations of the model-based scheme by solely relying on an initial trusted set of signal propagation data. CUSP uses the data to build SVM classifiers with quadratic kernels that are trained to differentiate between natural and un-natural signal propagation patterns in the region of interest. The resulting

approach is practical and effective for application in all areas and avoids arbitrary assumptions about models, parameters, and thresholds in favor of direct training data. We showed the effectiveness of this approach using a novel evaluation method based on real transmitter and terrain data. We showed that CUSP can achieve high detection accuracies even in the most unfavorable situations, *i.e.*, hilly urban/suburban areas with significant amounts of additional signal uncertainty. We also compared the performance of this technique with the model-based approach and showed that, despite not relying on closed-form formulas, parameters, and manual tuning, it outperforms the model-based technique in terms of attacker detection. A potential drawback is higher false positive rates, which is aggravated in environments with considerable natural variations in signal power within short distances. At a high-level, a remedy would entail modifying the feature space to increase its descriptive power. Adding elevation data to the feature space used in classification was discussed as a promising direction for future work.

In the third solution, we considered the case where the network constitutes a limited number of attestation-capable sensors. We presented a trust-based approach and showed that using statistical sequential estimation and classification methods can help deter attacks, while achieving quantifiably percise spectrum sensing outcome when possible. We showed that these methods reduce the total cost incurred by attestation. Our evaluation determined that the Biased node inclusion strategy is the most effective at deterring attacks, but also generates more false positives than Random or Geo-diverse strategies. All three strategies result in substantial attack deterrence. These are not the only strategies that can be used, and future research should evaluate other strategies. We also discussed another area for further research: developing a framework for formulating costs associated with using regular and attested nodes, and systematically striking a balance between the costs (from spectrum data aggregation and remote attestation) and achieving robust aggregation results.

References

[1] CogNea: Cognitive Networking Alliance. *http://www.cognea.org/*.

[2] FCC, ET Docket No FCC 08-260, November 2008.

[3] FCC, Second Memorandum Opinion and Order, ET Docket No FCC 10-174, September 2010.

[4] IEEE 802.22 WRAN WG on Broadband Wireless Access Standards. *http://www.ieee802.org/22*.

[5] Microsoft Research WhiteFi Service. *http://whitespaces.msresearch.us/*.

[6] Power Harvesting: Induction Magic. *National Decense Education Program, http://www.ndep.us/Power-Harvesting-Induction-Magic*.

[7] Power Harvesting: The Bat Hook. *National Decense Education Program, http://www.ndep.us/Power-Harvesting-The-Bat-Hook*.

[8] S. 649: Radio Spectrum Inventory Act. *http://www.govtrack.us/congress/bill.xpd?bill=s111-649*.

[9] Spectrum Bridge. *http://www.spectrumbridge.com*.

[10] Trusted Computing Group. *http://www.trustedcomputinggroup.org/*.

[11] US Census Bureau. *http://www.census.gov*.

[12] Philadelphia Citywide WiFi Officially Shut Down. *http://www.engadget.com/2008/05/13/philadelphia-citywide-wifi-officially-shut-down/*, 2008.

[13] CTIA: The Wireless Association Files Ex Parte to FCC to Request More Spectrum. *http://www.ctia.org/media/press/body.cfm/prid/1866*, 2009.

[14] Luxembourg: Model of a Successful Muni Wi-Fi Network. *http://www.muniwireless.com/2009/02/15/luxembourg-model-muni-wifi-network/*, 2009.

[15] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty. Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey. *Computer Networks*, 50(13):2127–2159, 2006.

[16] A. Algans, K. Pedersen, and P. Mogensen. Experimental analysis of the joint statistical properties of azimuth spread, delay spread, and shadow fading. *IEEE Journal on Selected Areas in Communications*, 20(3):523 –531, Apr. 2002.

[17] H. Alzaid, E. Foo, and J. G. Nieto. Secure data aggregation in wireless sensor network: a survey. *AISC '08: Proceedings of the sixth Australasian Conference on Information Security*, pages 93–105, 2008.

[18] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pages 27–38, 2009.

[19] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, 2006.

[20] J. Branch, B. Szymanski, C. Giannella, R. Wolff, and H. Kargupta. In-network outlier detection in wireless sensor networks. *Distributed Computing Systems, 2006. ICDCS 2006. 26th IEEE International Conference on*, 2006.

[21] G. Buchwald, S. Kuffner, L. Ecklund, M. Brown, and E. Callaway. The design and operation of the ieee 802.22.1 disabling beacon for the protection of tv whitespace incumbents. *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1 –6, 2008.

[22] J. Burbank. Security in cognitive radio networks: The required evolution in approaches to wireless network security. *CrownCom '08: International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, pages 1–7, May 2008.

[23] S. Capkun and J.-P. Hubaux. Secure positioning in wireless networks. *Selected Areas in Communications, IEEE Journal on*, 24(2):221 – 232, 2006.

[24] H. Chan, A. Perrig, B. Przydatek, and D. Song. Sia: Secure information aggregation in sensor networks. *Journal of Computer Security*, 15:69–102, January 2007.

[25] R. Chen and J.-M. Park. Ensuring trustworthy spectrum sensing in cognitive radio networks. *SDR '06: IEEE Workshop onNetworking Technologies for Software Defined Radio Networks*, pages 110–119, Sept. 2006.

[26] R. Chen, J.-M. Park, and K. Bian. Robust distributed spectrum sensing in cognitive radio networks. *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 1876 –1884, 2008.

[27] R. Chen, J.-M. Park, and J. Reed. Defense against primary user emulation attacks in cognitive radio networks. *IEEE Journal on Selected Areas in Communications*, 26(1):25–37, Jan. 2008.

[28] T. Clancy and N. Goergen. Security in cognitive radio networks: Threats and mitigation. *CrownCom '08: International Conference on Cognitive Radio Oriented Wireless Networks and Communications*, pages 1–8, May 2008.

[29] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, pages 273–297, 1995.

[30] O. Fatemieh, R. Chandra, and C. A. Gunter. Low Cost and Secure Smart Meter Communications using the TV White Spaces. *Proceedings of ISRCS '10: IEEE International Symposium on Resilient Control Systems*, August. 2010.

[31] O. Fatemieh, R. Chandra, and C. A. Gunter. Secure Collaborative Sensing for Crowdsourcing Spectrum Data in White Space Networks. *Proceedings of DySPAN '10: IEEE International Dynamic Spectrum Access Networks Symposium*, April. 2010.

[32] O. Fatemieh, A. Farhadi, R. Chandra, and C. A. Gunter. Using Classification to Protect the Integrity of Spectrum Measurements in White Space Networks. *Proceedings of NDSS '11: 18th Annual Network and Distributed System Security Symposium*, Feb. 2011.

[33] O. Fatemieh, M. LeMay, and C. A. Gunter. Reliable Telemetry in White Spaces using Remote Attestation. *submitted to SECON '11: IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, 2011.

[34] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 99(PrePrints), 2009.

[35] S. Ganeriwal, L. K. Balzano, and M. B. Srivastava. Reputation-based framework for high integrity sensor networks. *ACM Transactions on Sensor Networks*, 4(3):1–37, 2008.

[36] A. Ghasemi and E. Sousa. Spectrum sensing in cognitive radio networks: requirements, challenges and design trade-offs. *Communications Magazine, IEEE*, 46(4):32–39, April 2008.

[37] M. Ghosh, N. Mukhopadhyay, and P. K. Sen. *Sequential Estimation*. John Wiley and Sons, Inc, New York, NY, 1997.

[38] M. Gudmundson. Correlation model for shadow fading in mobile radio systems. *Electronics Letters*, 27(23):2145 –2146, 1991.

[39] I. Guyon, J. Weston, S. Barnhill, and V. Vapnik. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46(1-3):389–422, 2002.

[40] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, San Francisco, CA, second edition, 2006.

[41] J. J. Higgins. *An Introduction to Modern Nonparametric Statistics*. Thomson Learning, Stamford, CT, 2004.

[42] L. Hu and D. Evans. Secure aggregation for wireless networks. *SAINT-W '03: Proceedings of the 2003 Symposium on Applications and the Internet Workshops*, page 384, 2003.

[43] Y.-a. Huang and W. Lee. A cooperative intrusion detection system for ad hoc networks. *SASN '03: Proceedings of the 1st ACM workshop on Security of ad hoc and sensor networks*, pages 135–147, 2003.

[44] J. Hur, Y. Lee, S.-M. Hong, and H. Yoon. Trust management for resilient wireless sensor networks. *Information Security and Cryptology - ICISC 2005*, pages 56–68, 2006.

[45] P. Kaligineedi, M. Khabbazian, and V. Bhargava. Secure cooperative sensing techniques for cognitive radio systems. *ICC '08: IEEE International Conference on Communications*, pages 3406–3410, May 2008.

[46] J. Kennedy and E. Keeping. *Mathematics of Statistics*. Van Nostrand, Princeton, NJ, 3rd edition, 1962.

[47] H. Kim and K. G. Shin. In-band spectrum sensing in cognitive radio networks: energy detection or feature detection? *MobiCom '08: Proceedings of the 14th ACM international conference on Mobile computing and networking*, pages 14–25, 2008.

[48] H.-C. Kim, S. Pang, H.-M. Je, D. Kim, and S. Y. Bang. Pattern classification using support vector machine ensemble. *Pattern Recognition, 2002. Proceedings. 16th International Conference on*, 2002.

[49] B. Krishnamachari and S. Iyengar. Distributed bayesian algorithms for fault-tolerant event region detection in wireless sensor networks. *IEEE Transactions on Computers*, 53(3):241–250, March 2004.

[50] L. Lazos, S. Liu, and M. Krunz. Mitigating control-channel jamming attacks in multi-channel ad hoc networks. *WiSec '09: ACM conference on Wireless network security*, 2009.

[51] L. Lazos, R. Poovendran, and S. Čapkun. Rope: robust position estimation in wireless sensor networks. *Proceedings of the 4th international symposium on Information processing in sensor networks*, 2005.

[52] J. Lee, Y. Kim, S. Sohn, and J. Kim. Weighted-cooperative spectrum sensing scheme using clustering in cognitive radio systems. *Advanced Communication Technology, 2008. ICACT 2008. 10th International Conference on*, 1:786–790, Feb. 2008.

[53] M. LeMay and C. Gunter. Cumulative attestation kernels for embedded systems. volume 5789 of *Lecture Notes in Computer Science*, pages 655–670. Springer Berlin / Heidelberg.

[54] Z. Li, W. Trappe, Y. Zhang, and B. Nath. Robust statistical methods for securing wireless localization in sensor networks. *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, page 12, 2005.

[55] D. Liu, P. Ning, A. Liu, C. Wang, and W. K. Du. Attack-resistant location estimation in wireless sensor networks. *ACM Transactions on Information and System Security (TISSEC)*, 11:22:1–22:39, July 2008.

[56] F. Liu, X. Cheng, and D. Chen. Insider attacker detection in wireless sensor networks. *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 1937 –1945, May 2007.

[57] S. Liu, Y. Chen, W. Trappe, and L. Greenstein. Aldo: An anomaly detection framework for dynamic spectrum access networks. *INFOCOM 2009, IEEE*, pages 675 –683, 2009.

[58] Y. Liu, P. Ning, and H. Dai. Authenticating primary users' signals in cognitive radio networks via integrated cryptographic and wireless link signatures. *IEEE Symposium on Security and Privacy*, 2010.

[59] A. Min, K. Shin, and X. Hu. Attack-tolerant distributed sensing for dynamic spectrum access networks. *ICNP '09: IEEE International Conference on Network Protocols*, pages 294–303, 2009.

[60] S. Mishra, A. Sahai, and R. Brodersen. Cooperative sensing among cognitive radios. *ICC '06: IEEE International Conference on Communications*, 4:1658–1663, June 2006.

[61] R. Murty, R. Chandra, T. Moscibroda, and V. Bahl. Eliminating the need for low threshold spectrum sensing in white space networks. *Microsoft Research Technical Report MSR-TR-2010-127*, September 2010.

[62] J. Newsome, E. Shi, D. Song, and A. Perrig. The sybil attack in sensor networks: analysis & defenses. *IPSN '04: Proceedings of the 3rd international symposium on Information processing in sensor networks*, pages 259–268, 2004.

[63] T. Rappaport. *Wireless Communications: Principles and Practice*. IEEE Press, New York, 1996.

[64] N. Sastry, U. Shankar, and D. Wagner. Secure verification of location claims. *Proceedings of the 2nd ACM workshop on Wireless security*, pages 1–10, 2003.

[65] A. Seshadri, M. Luk, and A. Perrig. SAKE: Software attestation for key establishment in sensor networks. *Proceedings of the 2008 International Conference on Distributed Computing in Sensor Systems (DCOSS)*, June 2008.

[66] A. Seshadri, A. Perrig, L. van Doorn, and P. Khosla. Swatt: software-based attestation for embedded devices. *Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on*, pages 272 – 282, May 2004.

[67] A. Sethi and T. Brown. Hammer model threat assessment of cognitive radio denial of service attacks. *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1–12, Oct. 2008.

[68] N. Shankar, C. Cordeiro, and K. Challapali. Spectrum agile radios: utilization and sensing architectures. *New Frontiers in Dynamic Spectrum Access Networks, 2005. DySPAN 2005. 2005 First IEEE International Symposium on*, pages 160–169, Nov. 2005.

[69] R. Solera-Ure na, D. Martín-Iglesias, A. Gallardo-Antolín, C. Peláez-Moreno, and F. Díaz-de María. Robust asr using support vector machines. *Speech Commun.*, 49(4):253–267, 2007.

[70] C. R. Stevenson, G. Chouinard, Z. Lei, W. Hu, S. J. Shellhammer, and W. Caldwell. Ieee 802.22: the first cognitive radio wireless regional area network standard. *Communications Magazine*, 47(1):130–138, 2009.

[71] A. Taherpour, Y. Norouzi, M. Nasiri-Kenari, A. Jamshidi, and Z. Zeinalpour-Yazdi. Asymptotically optimum detection of primary user in cognitive radio networks. *IET Communications*, 1(6):1138–1145, 2007.

[72] R. Tandra, A. Sahai, and S. Mishra. What is a spectrum hole and what does it take to recognize one? *IEEE Magazine Special Issue on Cognitive Radio*, 97(5):824–848, May 2009.

[73] V. V. Vazirani. *Approximation Algorithms*. Springer, March 2004.

[74] E. Visotsky, S. Kuffner, and R. Peterson. On collaborative detection of tv transmissions in support of dynamic spectrum sharing. *New Frontiers in Dynamic Spectrum Access Networks, 2005. DySPAN 2005. 2005 First IEEE International Symposium on*, pages 338–345, Nov. 2005.

[75] D. Wagner. Resilient aggregation in sensor networks. *SASN '04: Proceedings of the 2nd ACM workshop on Security of ad hoc and sensor networks*, pages 78–87, 2004.

[76] Y. Yang, X. Wang, S. Zhu, and G. Cao. Distributed software-based attestation for node compromise detection in sensor networks. *Reliable Distributed Systems, 2007. SRDS 2007. 26th IEEE International Symposium on*, pages 219 –230, 2007.

[77] W. Zhang, S. Das, and Y. Liu. A trust based framework for secure data aggregation in wireless sensor networks. *Sensor and Ad Hoc Communications and Networks, 2006. SECON '06. 2006 3rd Annual IEEE Communications Society on*, 1:60–69, Sept. 2006.

[78] Y. Zhang and W. Lee. Intrusion detection in wireless ad-hoc networks. *MobiCom '00: 6th annual international conference on Mobile computing and networking*, pages 275–283, 2000.

Author's CV

OMID FATEMIEH

Dept. of Computer Science   Tel: (217) 721 7679

201 N. Goodwin Ave   sfatemi2@illinois.edu

Urbana, IL, 61801   http://cs.uiuc.edu/homes/sfatemi2

EDUCATION

PhD, Computer Science   Aug '04 - Jan '11

University of Illinois at Urbana-Champaign (UIUC), Urbana, IL

Advisor: Prof. Carl. A. Gunter

Thesis Title: Assuring Robustness of Radio Spectrum Telemetry against Vandalism and Exploitation.

GPA: 3.75/4.0

BSc, Computer Engineering with concentration in Software Engineering   Sep '00 - Jul '04

Sharif University of Technology (SUT), Tehran, Iran

- Ranked 3^{rd} in GPA (18.1/20.0) among 85 students in the class of 2004.

- Ranked 123^{th} among 350,398 participants in Iran's annual nationwide university entrance exam in 2000.

RESEARCH AND INDUSTRY EXPERIENCE

Research Assistant   Jun '05 - Present

Illinois Security Lab, UIUC, Urbana, IL

- PhD Thesis: Assuring Robustness of Radio Spectrum Telemetry against Vandalism and Exploitation (collaboration with Microsoft Research).

- Messaging, access control, and confidentiality based on attributes (collaboration with National Center for Supercomputing Applications).

- Analysis and adaptive defense against denial of service attacks in the Internet (collaboration with University of Pennsylvania and Bell Labs).

Technical Program Manager Intern   May '10 - Aug '10

Windows Kernel Group, Microsoft, Redmond, WA

Proposed and wrote the specification for a new feature in the kernel of Windows,
implemented a prototype, and validated it using internally-collected user data.

Technology Intern   May '08 - Aug '08

Morgan Stanley, New York, NY

Analyzed an internal audit system, identified metrics to represent its perfor-
mance, and developed software to compute and output the metrics.

Research Intern   Jun '06 - Aug '06

Bell Labs, Alcatel-Lucent, Murray Hill, NJ

Analyzed and simulated distributed denial of service attacks on the Internet.

Software Engineer Intern   Jun '03 - Aug '03

Hirbod Rayaneh Company, Tehran, Iran

Developed multiple modules of a commercial accounting system.

JOURNAL AND CONFERENCE PAPERS

- O. Fatemieh, A. Farhadi, R. Chandra, C. A. Gunter, Using Classification to Protect the Integrity of Spectrum Measurements in White Space Networks, Proceedings of *Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, Feb 2011 (forthcoming).

- (First three authors alphabetical) R. Bobba, O. Fatemieh, F. Khan, A. Khan, C. A. Gunter, H. Khurana, M. Prabhakaran, Attribute-based Messaging: Access Control and Confidentiality, *ACM Transactions on Information and System Security (TISSEC Journal)*, 2011 (forthcoming).

- O. Fatemieh, R. Chandra, C. A. Gunter, Low Cost and Secure Smart Meter Communications using the TV White Spaces, Proceedings of *IEEE International Symposium on Resilient Control Systems (ISRCS)*, Idaho Falls, Idaho, Aug 2010.

- O. Fatemieh, R. Chandra, C. A. Gunter, Secure Collaborative Sensing for Crowdsourcing Spectrum Data in White Space Networks, Proceedings of *IEEE International Dynamic Spectrum Access Networks Symposium (DySPAN)*, Singapore, Apr 2010.

- S. Khanna, S. Venkatesh, O. Fatemieh, F. Khan, C. A. Gunter, Adaptive Selective Verification. Proceedings of *IEEE Conference on Computer Communications (INFOCOM)*, Phoenix, AZ, Apr 2008.

- M. LeMay, O. Fatemieh, C. A. Gunter, PolicyMorph: Interactive Policy Transformations for a Logical Attribute-Based Access Control Framework. Proceedings of *ACM Symposium on Access Control Models And Technologies (SACMAT)*, Sophia Antipolis, France, Jun 2007.

- (First three authors alphabetical) R. Bobba, O. Fatemieh, F. Khan, C. A. Gunter, and H. Khurana, Using Attribute-Based Access Control to Enable Attribute-Based Messaging, Proceedings of *Annual Computer Security Applications Conference (ACSAC)*, Miami, FL, Dec 2006.

- O. Fatemieh, A Detailed Review of the Book Discovering Knowledge in Data: An Introduction to Data Mining, by D. T. Larose, *Journal of Biopharmaceutical Statistics (JBS)*, Vol. 16, No. 1, Jan 2006.

- O. Fatemieh, Internet2: Introduction, Applications and Methods, *Computer Report Journal of Informatics Society of Iran*, Vol. 25, No. 153, Jul 2003.

WORKING AND SUBMITTED PAPERS

- S. Khanna, S. Venkatesh, O. Fatemieh, F. Khan, C. A. Gunter, An Efficient Adaptive Countermeasure to Thwart DoS Attacks, submitted to *IEEE/ACM Transactions on Networking (ToN)* in Mar 2010.

- O. Fatemieh, M. LeMay, C. A. Gunter, Reliable Telemetry in White Spaces using Remote Attestation, submitted to *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)* in Dec 2010.

- O. Fatemieh, R. Chandra, C. A. Gunter, Secure Spectrum Measurement in White Space Networks, to be submitted to *ACM Transactions on Information and System Security (TISSEC Journal)* in Feb 2011.

WORKSHOPS, TECHNICAL REPORTS, AND POSTERS

- R. Shankesi, O. Fatemieh, C. A. Gunter, Resource Inflation Threats to Denial of Service Countermeasures, *UIUC Technical Report, http://hdl.handle.net/2142/17372* , Urbana, IL, Oct 2010.

- S. Khanna, S. Venkatesh, O. Fatemieh, F. Khan, C. A. Gunter, Nimble Clients Thwart Versatile DDoS Adversaries, *Information Theory and Applications Workshop*, San Diego, CA, Feb 2009.

- O. Fatemieh, F. Khan, M. Greenwald, C. A. Gunter, S. Khanna, J. Messeguer, S. Venkatesh, Adaptive Denial of Service Defense, *Midwest Security Workshop (MSW)*, Purdue University, West Lafayette, IN, Apr 2007.

- O. Fatemieh, R. Bobba, F. Khan, C. A. Gunter, and H. Khurana, Attribute Based Messaging, Poster at *ITI Workshop on Dependability and Security*, Information Trust Institute, Urbana, IL, Dec 2005.

- M. LeMay, O. Fatemieh, S. Katasani, N. Borisov, and C. A. Gunter, Self-Diagnosing Logical Access Controls, Poster at *ITI Workshop on Dependability and Security*, Information Trust Institute, Urbana, IL, Dec 2005.

TEACHING EXPERIENCE

Teaching Assistant
Department of Computer Science, UIUC

- Information Assurance (graduate level)   Fall '09

- Computer Security Architecture (graduate level)   Spring '05

- Information Assurance (graduate level)   Fall '04

Guest Lecturer
Department of Computer Science, UIUC

- Computer Security II - Two lectures on wireless security   Spring '10

- Information Assurance - One lecture on confidentiality policies   Fall '09

- Computer Security II - Two lectures on wireless security   Spring '09

- Computer Security Architecture - Two lectures on cryptography   Spring '05

Organizer and Lecturer, Brief Intro. to Machine Learning and Data Mining   Spring '10
Illinois Security Lab, UIUC
Organized and lectured introductory material on Machine Learning
and Data Mining to members of the Illinois Security Lab in four sessions.
Undergraduate Teaching Assistant
Department of Computer Engineering, SUT

- Computer Networks   Spring '04

- Computer Structure and Language   Fall '03

TECHNICAL SKILLS

Programming: C++, C#, Java, Pascal, Perl, Verilog, Tcl/Tk, Motorola 68000 Assembly.

Software and Standards: NS2, MATLAB, SQLServer, PostgreSQL, XQuery, XACML, XMLSchema.

PROFESSIONAL SERVICE

Journal Reviewer: ACM Transactions on Computer Systems (2009), Journal of Computer Security (JCS - 2005, 2007).

Conference/Workshop Reviewer: ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2010), IEEE International Dynamic Spectrum Access Networks Symposium (DySPAN 2010), IEEE International Symposium on Resilient Control Systems (ISRCS 2010), International Conference on Distributed Computing Systems (ICDCS 2009), IEEE Computer Security Foundations Symposium (CSF 2007), ACM Workshop on Privacy in the Electronic Society (WPES 2007), IEEE International Conference on Network Protocols (ICNP 2006).