Reliable Telemetry in White Spaces using Remote Attestation

Omid Fatemieh
University of Illinois
Urbana-Champaign

Michael LeMay
University of Illinois
Urbana-Champaign

Carl A. Gunter
University of Illinois
Urbana-Champaign

ABSTRACT

We consider reliable telemetry in white spaces in the form of protecting the integrity of distributed spectrum measurements against coordinated misreporting attacks. Our focus is 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 for Southwest Pennsylvania. The results show that with as low as 15% penetration of attestation-capable nodes, more than 94% of the attempts from omniscient attackers can be thwarted.

1. INTRODUCTION

Dynamic spectrum allocation promises to make spectrum use more efficient by enabling opportunistic, unlicensed (secondary) use of 'white-space' frequencies when they are not occupied by licensed (primary) users. This paradigm has gained significant traction due to the increasing demand for wireless services, the limited availability of spectrum, and the FCC's recent ruling that permits operation of unlicensed users in the unused portions of the TV spectrum. This permission is considered the first significant increase in unlicensed spectrum below 5 GHz in over 20 years [2].

Identifying unused portions of spectrum is a key requirement for opportunistic spectrum access. Spectrum availability data is envisioned to be centrally aggregated and consulted to govern the usage of spectrum. At least three scenarios for data collection have been proposed. First, the data may be provided by volunteer white-space devices or deployed sensors to build regional or nationwide spectrum availability databases, or augment the white space geolocation database mandated by the FCC. Second, a white-space service provider may collect spectrum sensing data from white-space devices in its network to determine areas of primary presence [1, 4, 11]. Third, by combining spectrum sensing data from multiple devices (*collaborative sensing*), one can improve the detection accuracy in highly shadowed environments [37].

ACSAC '11 Dec. 5-9, 2011, Orlando, Florida USA

Reliable Telemetry, or reliable central aggregation of sensor data in this context, is threatened by nodes that may report false data with malicious intent. A coordinated group of attacker nodes may aim to *exploit* a spectrum in a given region by falsely reporting that a primary signal is present, or *vandalize* a primary by reporting that its signal is not present and thereby creating interference and chaos. Previous works on this problem have achieved moderate degrees of success by identifying the data from individual or small groups of attackers as abnormal. These approaches suffer from at least two shortcomings. First, their effectiveness is limited by relying solely on sensor data for inferring (dis)trust. Second, they assume limits on the penetration of attackers in an area; attackers either constitute a small fraction of nodes in a small local neighborhood [30], or if they control the majority of nodes in a neighborhood, the preponderance of adjacent neighborhoods must be un-compromised [18].

In this paper, we initiate a new direction in reliable distributed measurement 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. They 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 quantifiably accurate 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 accuracy 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 signal propagation data in the region of interest.

We evaluate our scheme using predicted signal power data obtained from applying empirical signal propagation data to 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. In addition, we systematically enumerate the costs associated with remote attestation and provide detailed data 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. Our evaluation results show that our scheme is highly effective against attacks even in cases where only a small subset of the sensors can be remotely attested. For example, with as low as 15% of nodes being attested,

we show that more than 94% of the attacks can be defended against. The protection gradually improves as the fraction of attested nodes increases.

The contributions of this paper are summarized below:

- A new direction in reliable telemetry against coordinated misreporting attacks that relies on a small subset of attestation-capable sensors.
- A practical framework for using statistical sequential estimation coupled with machine learning classifiers to deter attacks and achieve quantifiably accurate outcome.
- A case study based on real TV transmitter and terrain data from the FCC and NASA in Pennsylvania that includes an evaluation for the proposed scheme, as well as a cost analysis for remote attestation.

2. BACKGROUND AND PROBLEM FORMULATION

In this section we first provide background information on spectrum measurements and remote attestation. Next, we describe our setting and problem statement.

2.1 Spectrum Sensing and Aggregation in White Spaces

The FCC's ruling in November 2008 allows for operation of unlicensed users in the unused portions of the TV spectrum [2]. Wireless communications in this spectrum (below 700 MHz) benefit from favorable signal propagation and penetration properties, which enable long transmission ranges. Access to this spectrum could enable more powerful Internet connections in public areas, campuses, and homes with extended range, fewer dead spots, and improved speeds. Many other applications are envisioned; for example, broadband access for rural areas, extended access to medical care in rural areas, and support for the communications of the advanced meter infrastructure (AMI) [3, 17].

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 [3]: (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.

There exist three scenarios for centrally aggregating spectrum sensing results from sensors in a large region [19]. First, using data from deployed spectrum sensors or volunteer (mobile) white-space devices 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 [6]. Second, 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 prototypes [11]. 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. Third, spectrum sensing results from multiple devices may be combined to improve the detection accuracy at low thresholds in highly shadowed environments [37].

To capture the common aspects 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 report the results to a central server. This form of participatory sensing is also referred to as *crowdsourcing*.
- *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 [16].

2.2 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 [7]. Previous work has also shown that remote attestation can feasibly be implemented in software on-chip for embedded processors such as AVR32 micro-controllers [25].

2.3 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. We refer to these sources as nodes or sensors in the rest of this paper. Due to their widespread adoption, ease of implementation, and small sensing time, we assume that energy detectors will be the only sensors in use [8, 36]. 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 $\times$ 1km square and is the unit in which combining individual results, or *collaborative sensing*, occurs. Within a cell, 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. A common method for combining sensing results within each cell is Equal Gain Combining (EGC), which periodically 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, FCC has mandated -114 dBm as the detection threshold.

We address the problem of performing reliable aggregation of spectrum measurement data contributed by a 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 in any cell, including any number of adjacent cells. For example, they may seek to change the perceived primary signal power for a cell 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 exploitation, the attackers aim to deceive the network to abandon the channel to exclusively use it for themselves, whereas in vandalism the main goal is creating chaos or interference. We focus on canceling the effect of such attackers that have a strong (*e.g. majority*) presence in a cell, and (in the absence of any defense) are able to dominate the cell and flip the detection outcome.

A particularly novel aspect of our work is that we assume that a subset of nodes, for example 20%, are able to perform remote attestation (see Figure 1). For any such *attestation-capable* node, the aggregation server can detect whether it is compromised and thus running illegitimate code. The question that we aim to answer is how 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 (see Section 5).

While some of the nodes may be unreliable or compromised insiders, 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 [32]). 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) Residential Internet access using IEEE 802.22 [36], and (2) AMI communications [17]. In cases where the network contains untrustworthy or mobile devices, secure localization and location verification techniques may ensure nodes' locations are authentic [14, 24, 26, 27]. The above assumptions are common for the type of analysis we perform here [15, 19, 31]; if they are violated then additional protective measures are required.

3. APPROACH

Consider Figure 1 as 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 cell, and compare it to a primary detection threshold to determine whether the spectrum 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 cell A in Figure 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. This approach is safer (in terms of vulnerability to false reports) than one in which the values from the (potentially compromised) regular nodes are also included. 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?

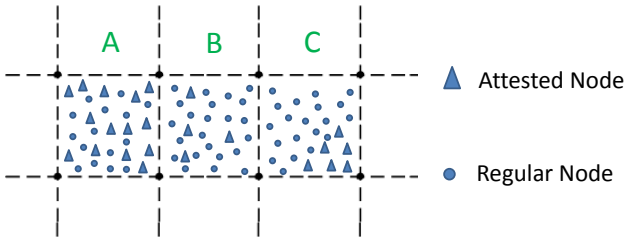

Figure 1: Illustration of a few cells with attested and regular nodes.

Now consider cell B where unlike cell A there are very few trusted 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 cell C; not only there exist very few attested nodes, but their positioning also makes it likely that they do not provide enough diversity. For example, 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?

3.1 Key Issues and Overview

The examples above underline the importance of the following needs. First, there must be a systematic strategy to determine when there is 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.

At a high level, our approach consists of three main phases (summarized in Algorithm 0). First, within each cell we rely on statistical inference and sequential estimation to aggregate data from all of attested sensors as well as 'enough' regular nodes to achieve the application-specified precision goal (Section 3.2). Note that we only include the least required number of regular sensors to limit unnecessary exposure to untrustworthy data. Various inclusion strategies are proposed for this purpose (Section 3.3). The

aggregate is either the mean and median of the data, and is dynamically determined by our algorithm. This choice may change throughout the execution of the algorithm according to pre-specified rules (Section 3.4). Second, the regular nodes that were included in the aggregation process in the cell are compared against the data from the trusted nodes of the 8 neighboring cells. This process involves using machine learning classifiers built from real signal propagation data. The classifier detects irregular signal propagation patterns that most likely represent a coordinated misreporting attack (Section 3.5). Third, after the potentially compromised data is eliminated, we compute the final aggregate.

Algorithm 1 Simplified Approach Overview (for Each Cell)

Input:
(1) *Green Data*: measurements from attested nodes
(2) *Yellow Data*: measurements from regular nodes
(3) *Strategy* $\in$ {*Random, Geo-Diverse, Biased*}: strategy for including data from regular nodes
(4) *Aggregate* $\in$ {*Mean, Median*}: dynamically changes based on the situation

Phase 1: Node Selection
Add *Green Data* to aggregation *Pool*
while ¬SATISFY-PRECISION-REQUIREMENTS(*data* in *Pool*, *Aggregate*) **do**
  if | *Yellow Data* | > 0 **then**
    MOVE-NEXT-ELEMENT-TO-POOL(*Strategy, Yellow Data*)
  else
    Remove all *Yellow Data* from *Pool*
    Go to **Phase 3**
  end if
end while

Phase 2: Attack Detection
Yellow Suspects ← *Yellow Data* in *Pool* from **Phase 1**
Green Neighbors ← averages of *Green Data* in the neighboring cells (*i.e.* 8 numbers)
if SVM-ATTACKER-DETECTION(*Yellow Suspects*, *Green Neighbors*) **then**
  Remove all *Yellow Suspects* from *Pool*
end if

Phase 3: Aggregate Calculation
Compute *Aggregate* from data in *Pool*

3.2 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 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.* $0.95 = 1 - \alpha$).

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 [33]. 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 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 [20].

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 [22]:

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

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

3.3 Intra-cell Inclusion Strategies

We consider three *inclusion strategies* for including regular nodes in the aggregate computation in each cell. The merits and disadvantages of each strategy are discussed in this section and evaluated in Section 4.

Random: Randomly adding data from regular nodes to the data from attested nodes has the advantage that it is in-line with the sampling 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. However, it disallows deploy-

ing targeted inclusion strategies that could potentially lead to lower attacker success rate.

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 [21]. 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 [9]. 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 3.2 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 4).

3.4 Intra-cell Aggregation: Mean or Median?

Within each cell, 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 [18, 38]. 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 (the calculation procedures are presented in Section 3.2). To support our argument about the relatively smaller confidence intervals for mean (with the same number of samples), we generate sample signal propagation data representing a log-normal shadowing model with average power of -95dBm and standard deviation (*a.k.a.* dB-spread) of 4, 6, and 8. Table 1 presents the margins of error achieved using random samples of size 20, 30, 40, and 50 from this distribution.

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 2 illustrates this observation. The attackers' goal is to change the aggregate from a value below the signal threshold of -114 dBm to one above the threshold (*e.g.* -113 dBm). When the median is used (the top picture), the attackers can achieve their desired goal by simply reporting -113 dBm. However, when the average is used (the bottom picture), the attackers need to report an average false report of -105.5 dBm to change the total average to their desired value of -113 dBm. The additional abnormality

Table 1: Margin of error (95% confidence) for randomly generated data of size $|S|$ equal to 20, 30, 40, 50 from log-normal (shadowing) distribution with standard deviation, σ, of 4, 6, and 8.

	$\|S\| = 20$	$\|S\| = 30$	$\|S\| = 40$	$\|S\| = 50$
	$\sigma = 4$			
Mean	1.7	1.4	1.2	1.1
Median	2.3	1.9	1.6	1.5
	$\sigma = 6$			
Mean	2.6	2.1	1.8	1.6
Median	3.5	2.8	2.5	2.3
	$\sigma = 8$			
Mean	3.5	2.8	2.4	2.2
Median	4.7	3.8	3.3	2.6

facilitates detecting them using SVM classifiers (Phase 2 of our approach), and is therefore desirable. 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 strategy helps with reducing the effect of attackers and natural outliers when attackers do not constitute a majority, and makes them more likely to be detected when they do.

We further elaborate on the details of aggregate calculation in Phase 1 with an illustrative example. We start by considering the data from all of the attested nodes in the aggregation pool and initially use *median* as the aggregator. If the margin of error for the median of attested nodes is below the application requirement γ, we stop by declaring the median as the final result. Otherwise, we need more data. Consider a cell with k attested nodes. After adding the k attested nodes, we iteratively add up to $k-1$ additional elements from regular nodes to the aggregation pool.

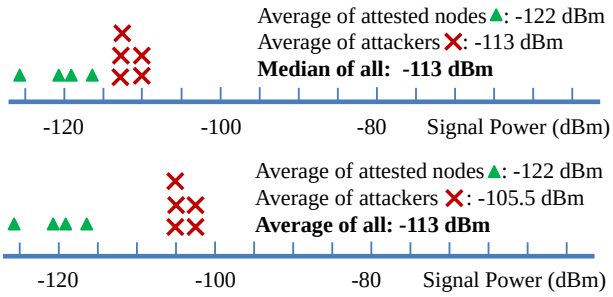

Figure 2: A simplified illustration of why attackers are forced to deviate more when they aim to move the mean (bottom picture) instead of the median (top picture).

The order in which the regular nodes are added to the pool is determined by the chosen inclusion strategy (Random, Geo-Diverse, or Biased). After each addition, if the margin of error for the 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 (using the same inclusion strategy) 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 ignore all the added regular nodes and proceed to Phase 3 where the median of attested nodes is computed as the aggregate.

3.5 Inter-Cell Attacker Detection using Classifiers

When the execution of Algorithm 0 reaches Phase 2, we have obtained an aggregate 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 whose data is included in the calculation are not part of an exploitation or vandalism attack. We first separate the data points from those *regular* nodes that have contributed to the aggregate (*a.k.a.* 'yellow suspects') and compare them to the data from *attested* nodes in the neighboring cells (*a.k.a.* 'green neighbors').

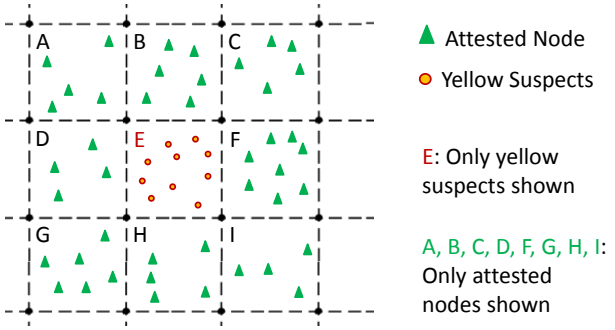

Figure 3: Classification-based attacker detection setting: regular nodes included in the aggregation for cell E and attested nodes from neighboring cells.

To determine if the yellow suspects in a cell represent an attacker-dominated group, we use real signal propagation data in the region to build a *classifier* that is *trained* to differentiate between natural and un-natural signal propagation patterns. The idea is to learn the normal propagation patterns of the signal from the reliable signal propagation data and use it to spot unnatural propagation of signal, which may be caused by malicious false reports.

More specifically, we consider the *local neighborhood* N_E of any cell E to contain E and its 8 neighboring cells (Figure 3). We represent N_E by a 9-element tuple containing the 'average' reported powers from the yellow suspects in E and the average reported power from the green neighbors in a pre-specified order. We call this the *neighborhood representation* of E. For example, for Figure 3, the first element would be the average of yellow suspects in cell E, and the second to ninth elements would be the value in the first element minus the average power of attested nodes in cells A to I (excluding E). Assume for a moment that we have access to reliable power measurements for a subset of the region of interest. This data can be used to create one neighborhood representation for each cell in the area. We refer to each such representation 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, as we will elaborate later, assume we have access to the neighborhood representation for a sufficiently large and diverse set of 'un-natural' (attacker-dominated) cells.

We now cast our problem to a binary *classification* problem. Classification is a machine learning technique that is widely used in domains ranging from spam email detection and unauthorized spectrum usage to fraud detection and speech recognition. 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 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. In our case, the neighborhood representation of a cell serves as its feature vector. The goal is to predict a binary label for a *test* example for which we do not know the label, using the classifier built from training examples [12]. 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.

Now we turn to the problem of obtaining training examples. 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. Having obtained such natural (normal) examples, we modify them to inject *un-natural* training instances to represent attacker-dominated cells.

Building such a classifier from the natural and un-natural examples has been discussed in detail in prior works [19] and has been shown to effectively detect attacker-dominated regions in environments where there is no separation between regular and attestation-capable nodes. By contrast, in our setting, the classifier is applied in a slightly different manner where only the trusted data from the neighbors is used in classification. However, due to potentially low penetration of attested nodes, this translates to less data points being available for classification. This may negatively affect the classification accuracy. We build a similar classifier (using Support Vector Machines (SVM) with quadratic kernels) to detect whether the yellow suspects in a cell look abnormal compared to the green neighbors and evaluate it in Section 4. 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.

4. 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 [37]. For evaluation purposes, we consider this data as the ground truth.

We instantiated our evaluation to an urban/suburban area surrounding Pittsburgh, Pennsylvania. The hilly nature of the area introduces a large amount of legitimate signal variations, which makes the task of precise signal power estimation and attacker detection more challenging (compared to flat areas). Therefore, these experiments should be considered a stress-test for our scheme.

The following points in (*latitude, longitude*) format define the southwest and northeast corners of the considered 20km $\times$ 20km square area in Pennsylvania: $\langle$ (40.35, -80.12), (40.53, -79.884)$\rangle$. Each cell is 1km $\times$ 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

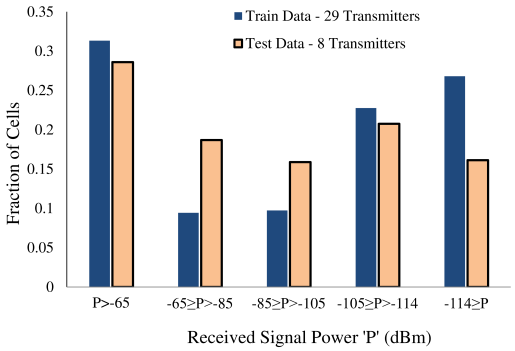

Figure 4: (a) Transmitters in parts of Southwest Pennsylvania / East Ohio. (b) Distribution of received signal for the training and testing data in

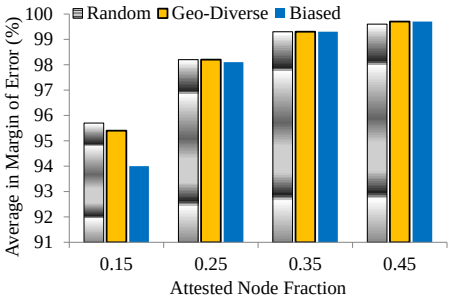

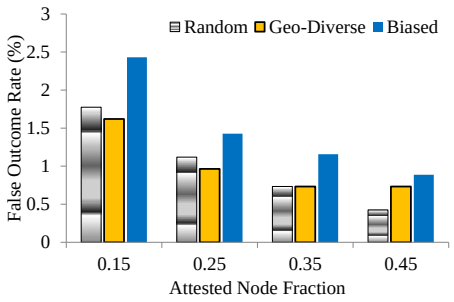

Figure 5: No attack; percentage of cells with ground truth average within the margin of error from the calculated aggregate (left) and false outcome rate (in percentage) as a function of the fraction of attested nodes (right).

transmitters, of which we randomly pick 29 for building the classifier, and 8 for testing it. An illustration of the area, including the location of the majority of DTV transmitters in provided in Figure 4(a). The distribution of the received signal power across all the cells in the region (from all 37 transmitters) is provided in Figure 4(b). Guided by approximate sample size requirements based on methods in Section 3.2, 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. In suburban and rural areas, the densities need to be achieved through provisioning or other means in order for our approach to be fully effective.

4.1 No-Attack Performance

We first evaluate the accuracy 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 5(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. The results show that our approach achieves a high overall success rate in terms of obtaining precise estimates of signal power in a region. They also show that despite Biased's weaker performance in some cases, in most cases the choice of inclusion strategy does not have a significant impact.

As a second performance metric in the absence of attacks, we introduce the *false outcome rate*, representing the fraction of unattacked 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 5(b) represents the false outcome rate as a function of the fraction of attested nodes. The results show that while overall false outcome rates are low, the Biased inclusion strategy is the weakest performer, particularly when the fraction of attested nodes is low. This can be explained by situations in which the few attested nodes are not providing values near the true average power in the cell, and the Biased inclusion strategy aggravates the situation by including similar data that effectively builds up on the already poor samples.

4.2 Performance against 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 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 coordinated 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 [19] and test deterrence of attacks on the

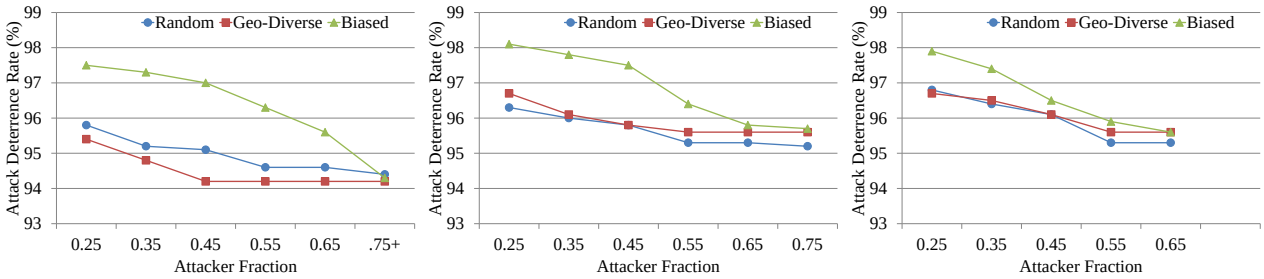

Figure 6: Attack deterrence rate (in percentage) when the average fraction of attested nodes is .15 (left), .25 (center), and .35 (right).

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. 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.

In Figure 6, 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 7 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 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 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.

5. ATTESTATION COSTS

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 increase the amount of *network data* that is transmitted and received, which may also increase the energy consumption of the wireless radio. 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 276 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 [25]. 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 takes about 0.6 seconds to execute. Due to the significantly shorter keys, each attestation operation only generates at least 68 bytes of data.

6. RELATED WORK

Much of the prior work in the context of white space networks uses various abnormality detection techniques to identify individual attackers within a cell as part of collaborative sensing. Such approaches, however, are not capable of detecting cells that form

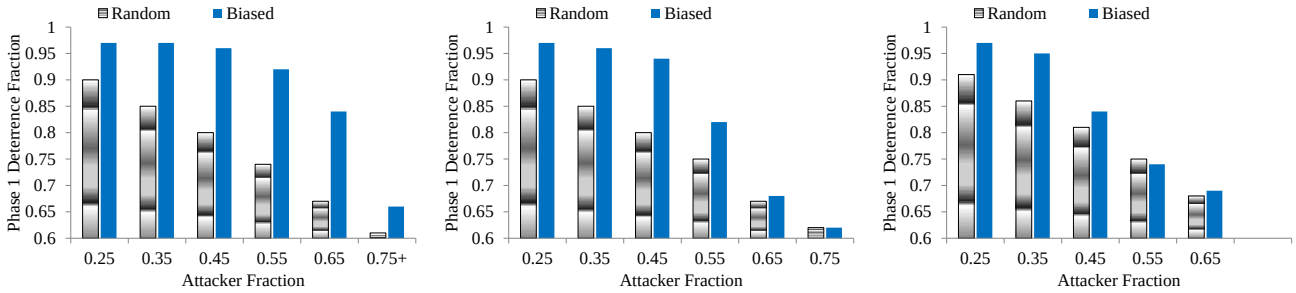

Figure 7: The fraction of attack deterrences in Phase 1. For each bar with value x, $1 - x$ is the fraction deterred in Phase 2. The average fraction of attested nodes is $.15$ (left), $.25$ (center), and $.35$ (right). Results for Geo-diverse (similar to Random) are omitted.

a majority in the cells [15, 31]. For example, Min *et al.*'s approach based on correlation-based filters fails to detect attackers that constitute more than $1/3$ of the population of the nodes in a cell. Fatemieh *et al.* [18, 19] consider detecting attacker-dominated cells by outlier detection and classification techniques, however, their solutions do not consider remote attestation and fall short if a preponderance of neighboring cells are dominated by attackers.

Another body of related work in the context of white space networks considers primary user emulation (PUE) attacks [16, 29]. In a PUE, 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. We consider this problem to be orthogonal to the problem we address.

In the context of sensor networks, Wagner introduced *resilient aggregation* [38], 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 attack detection as we do. Zhang *et al.* [40] propose a framework that identifies readings not statistically consistent with the distribution of readings in a cluster of nearby sensors. Their proposal, however, is not able to handle situations where attacker can compromise a large fraction of the nodes in a cluster. Hur *et al.* [23] propose a trust-based framework in a grid in which each sensor builds trust values for neighbors and reports them to the local aggregator. Their solution, however, does not consider natural uncertainties in the data, does not provide a global view for a centralized aggregator, and cannot identify compromised 'regions.' For a survey on a closely related area of secure data aggregation in wireless sensor networks see [10].

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 [35], detecting compromised nodes [39], and other activities such as secure code update and key establishment [34]. 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.

Insider attacker detection in wireless networks is another area of related work. This problem has been explored in a general setting [13] as well as more specific contexts such as insider jammers. As an illustrative example in the general context of sensor networks, Liu *et al.* [28] propose a solution in which each node builds a distribution of the observed measurements around it and flags deviating neighbors as insider attackers. The solution, however, is local and peer to peer and does not work in areas with more than 25% attackers.

7. CONCLUSIONS

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 nodes. These methods reduce the total cost incurred by attestation. The results show that attestation capability for as low as 15% of the nodes can provide protection against more than 94% of the attacks from omniscient coordinated attackers. The protection improves as the fraction of attested nodes is increased. 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. These are not the only strategies that can be used, and future research should evaluate other strategies. One promising future direction is developing a 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 obtaining robust aggregation results.

8. ACKNOWLEDGEMENTS

We thank Farid Kianifard, Ranveer Chandra, and Ali Farhadi for their valuable comments. This work was supported in part by DOE DE-0000097, HHS 90TR0003-01, NSF CNS 09-64392, NASA 09-VVFCS1-09-0010, NSF CNS 09-17218, NSF CNS 07-16421, and grants from the MacArthur Foundation and Lockheed Martin. The views expressed are those of the authors only.

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

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

[8] 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.

[9] 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.

[10] 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.

[11] 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.

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

[13] 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.

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

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

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

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

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

[23] 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.

[24] L. Lazos and R. Poovendran. Serloc: Robust localization for wireless sensor networks. *ACM Trans. Sen. Netw.*, 1(1):73–100, 2005.

[25] 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.

[26] 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.

[27] 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.

[28] 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.

[29] 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 (Oakland)*, 2010.

[30] A. Min and K. Shin. An optimal sensing framework based on spatial rss-profile in cognitive radio networks. *Sensor, Mesh and Ad Hoc Communications and Networks, 2009. SECON '09. 6th Annual IEEE Communications Society Conference on*, pages 1–9, June 2009.

[31] 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.

[32] 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*, 2004.

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

[34] 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.

[35] A. Seshadri, A. Perrig, L. van Doorn, and P. Khosla. SWATT: SoftWare-based ATTestation for embedded devices. *IEEE Symposium on Security and Privacy (Oakland)*, May 2004.

[36] 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.

[37] 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.

[38] 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.

[39] 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.

[40] 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.