

# A Security and Robustness Performance Analysis of Localization Algorithms to Signal Strength Attacks

YINGYING CHEN

Stevens Institute of Technology

KONSTANTINOS KLEISOURIS

Rutgers University

XIAOYAN LI

Lafayette College

and

WADE TRAPPE and RICHARD P. MARTIN

Rutgers University

---

Recently, it has been noted that localization algorithms that use signal strength are susceptible to noncryptographic attacks, which consequently threatens their viability for sensor applications. In this work, we examine several localization algorithms and evaluate their robustness to attacks where an adversary attenuates or amplifies the signal strength at one or more landmarks. We study both point-based and area-based methods that employ received signal strength for localization, and propose several performance metrics that quantify the estimator's precision, bias, and error, including Hölder metrics, which quantify the variability in position space for a given variability in signal strength space. We then conduct a trace-driven evaluation of a set of representative algorithms, where we measured their performance as we applied attacks on real data from two different buildings. We found the median error degraded gracefully, with a linear response as a function of the attack strength. We also found that area-based algorithms experienced a decrease and a spatial-shift in the returned area under attack, implying that precision increases though bias is introduced for these schemes. Additionally, we observed similar values for the average Hölder metric across most of the algorithms, thereby providing strong experimental evidence that nearly all the algorithms have similar average responses to signal strength attacks with the exception of the Bayesian Networks algorithm.

---

Preliminary results of this article have been presented in part in DCOSS 2006 [Chen et al. 2006b]. Authors' addresses: Y. Chen, Stevens Institute of Technology, Hoboken, NJ. 07030; email: yingying.chen@stevens.edu; K. Kleisouris, W. Trappe, and R. P. Martin, Rutgers University, Piscataway, NJ; email: {kkonst, martin}@cs.rutgers.edu, trappe@winlab.rutgers.edu; X. Li, Lafayette College, Easton, PA; email: xli@cs.lafayette.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org). © 2009 ACM 1550-4859/2009/02-ART2 \$5.00 DOI 10.1145/1464420.1464422 <http://doi.acm.org/10.1145/1464420.1464422>

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms: Security, Performance, Algorithms, Design, Measurement, Experimentation

Additional Key Words and Phrases: Localization, robustness, sensor networks, signal strength attacks, susceptibility

**ACM Reference Format:**

Chen, Y., Kleisouris, K., Li, X., Trappe, W., and Martin, R. P. 2009. A security and robustness performance analysis of localization algorithms to signal strength attacks. *ACM Trans. Sen. Netw.*, 5, 1, Article 2 (February 2009), 37 pages. DOI = 10.1145/1464420.1464422 <http://doi.acm.org/10.1145/1464420.1464422>

---

## 1. INTRODUCTION

Accurately localizing sensor nodes is a critical function for many higher level applications such as health care monitoring, wildlife animal habitat tracking, emergency rescue and recovery, location-based access control, and location-aware content delivery. Over the past few years, many localization algorithms have been proposed to localize wireless devices and sensors and provide location information to new classes of location-oriented applications. Out of the myriad of localization methods, algorithms that use received signal strength (RSS) as the basis of localization are very attractive options as using RSS allows the localization system to reuse the existing communication infrastructure, rather than requiring the additional cost needed to deploy specialized localization infrastructure, such as ceiling-based ultrasound, GPS, or infrared methods [Hazas and Ward 2003; Priyantha et al. 2000; Savvides et al. 2001]. In particular, all commodity radio technologies, such as 802.11, 802.15.4, and Bluetooth provide RSS values associated with packet reception, and thus localization services can easily be built for such systems. Further, RSS-based localization is attractive as the techniques are technology-independent: an algorithm can be developed and applied across different platforms, whether 802.11 or Bluetooth. In addition, it provides reasonable accuracy with median errors of one to five meters [Elnahrawy et al. 2004]. However, as more location-dependent services are deployed, they will increasingly become tempting targets for malicious attacks. Adversaries may alter signal strength measurements for the purpose of accessing services that are based on location information (e.g. WLAN access may only be granted to devices inside of a building.).

Unlike traditional systems, the localization infrastructure is sensitive to a variety of attacks, ranging from conventional to noncryptographic, that can subvert the utility of location information. Conventional attacks, where an adversary injects false messages, can be isolated and protected against using traditional cryptographic methods, such as authentication. However, there is a completely orthogonal set of attacks that are noncryptographic, where the measurement process itself can be corrupted by adversaries. Unfortunately, these noncryptographic attacks cannot be addressed by traditional security services. Thus, it is desirable to study the impact of these attacks on localization algorithms and explore methods to detect and further to eliminate these attacks from the network. Although there has been recent research on securing

localization [Brands and Chaum 1994; Capkun and Hubaux 2005; Capkun and Hubaux 2006; Li et al. 2005; Liu et al. 2005; Sastry et al. 2003], to date there has been no study on the robustness of the existing generation of RSS-based localization algorithms to physical attacks.

Rather than jumping to the immediate conclusion that all RSS-based localization systems are vulnerable, we believe that a thorough performance evaluation of existing RSS-based localization schemes is warranted. Such an evaluation would represent a valuable contribution to a wireless sensor network designer as it would help drive protocol decisions, and allow the engineer to decide whether more complicated secure localization algorithms are truly necessary. In this article, we expand on our study [Chen et al. 2006b] and detail an investigation into the susceptibility of a wide range of signal strength localization algorithms to attacks on the Received Signal Strength (RSS). Specifically, we examine the response of several localization algorithms to unanticipated power losses and gains—attenuation and amplification attacks. In these attacks, the attacker modifies the RSS of a sensor node or landmark, for example, by placing an absorbing or reflecting material around the node or landmark. Notably, we expand the set of attack scenarios to include amplification or attenuation attacks on combinations of landmarks, as well as analyze the results of simultaneous amplification and attenuation on multiple landmarks. We investigate both point-based and area-based algorithms that utilize RSS to perform localization. In order to evaluate the robustness of these algorithms, we provide a generalized characterization of the localization problem, and then present several performance metrics suitable for quantifying performance, including estimator angle bias, estimator distance error, and estimator precision. Additionally, an essential contribution of our work is the introduction of a new family of localization performance metrics, which we call Hölder metrics. These metrics quantify the susceptibility of localization algorithms to perturbations in signal strength readings. We use worst-case and average-case versions of the Hölder metric, which describe the maximum and average variability as a function of changes in the RSS. We then experimentally evaluate the performance of a wide variety of localization algorithms after applying attenuation and amplification attacks to real data measured from two different office buildings.

Using experimentally observed localization performance, we found that the errors for a wide variety of algorithms scaled with surprising similarity under attack. The single exception was the Bayesian Networks algorithm, which degraded slower than the others in response to attacks against a single landmark and was attack resistant when simultaneously localizing multiple devices without using training data under all-landmark attacks. In addition to our experimental observations, we found a similar average-case response of the algorithms using our Hölder metrics. However, we observed that methods which returned an average of likely positions had less variability and are thus less susceptible than other methods.

We also observed that all algorithms, except Bayesian Networks without using training data, degraded gracefully, experiencing linear scaling in localization error as a function of the amount of loss or gain (in dB) introduced by

an attack. This observation applied to various statistical descriptions of the error, led us to conclude that no algorithm collapses in response to an attack. This is important because it means that, for all the algorithms we examined, there is no tipping point at which an attacker can cause gross errors. In particular, we found the mean error of most of the algorithms for both buildings scaled between 1.3–1.8 ft/dB when all the landmarks were attacked simultaneously, and 0.5–0.8 ft/dB when a single landmark was attacked. Additionally, the performance of the mean response of algorithms with multiple landmarks under attack is between the all-landmark attack and the single landmark attack, which scaled at 0.4–1.4 ft/dB. Further we observed that mixed attacks with simultaneous attenuation and amplification cause the mean response of algorithms to move faster, ranging from 0.2–2.3 ft/dB. More powerful effects were witnessed when the mixed attack was applied to landmarks that were further apart from each other. We also showed experimentally that RSS can be easily attenuated by 15 dB, and that, as a general rule of thumb, very simple signal strength attacks can lead to localization errors of 20–30 ft.

Finally, we conducted a detailed evaluation of area-based algorithms as this family of algorithms returns a set of potential locations for the transmitter. Thus, it is possible that these algorithms might return a set with a larger area in response to an attack and could have less precision (more uncertainty) under attack. However, we found all three of our area-based algorithms shifted the returned areas rather than increasing the returned area. Further, one of the algorithms, the Area Based Probability (ABP) scheme, significantly shrank the size of the returned area in response to very large changes in signal strength.

The rest of this article is organized as follows. We begin, in Section 2, by giving an overview of the algorithms used in our performance study and discuss how signal strength attacks can be performed in Section 3. In Section 4, we provide a formal model of the localization problem as well as introduce the metrics that we use in this article. We then examine the performance of the algorithms through an experimental study in Section 5, and discuss the Hölder metrics for these algorithms in Section 6. We wrap up our article by providing a discussion of related work in Section 7. Finally, we conclude in Section 8.

## 2. LOCALIZATION ALGORITHMS

Signal strength is a common physical property used by a widely diverse set of algorithms. For example, most fingerprinting approaches utilize the RSS, for example, Bahl and Padmanabhan [2000] and Battiti et al. [2002], and many multilateration approaches [Madigan et al. 2005] use it as well. Although these algorithms provide several-meter level accuracy, using the RSS is an attractive approach, because the existing wireless infrastructure can be reused—this feature presents a tremendous cost savings over deploying localization-specific hardware. In this article we thus focus on localization algorithms that employ signal strength measurements. In this section, we provide an overview of a representative set of algorithms selected for conducting performance analysis under attack. These algorithms use either deterministic or probabilistic methods for location estimation.

There are several ways to classify localization schemes that use signal strength: range-based schemes, which explicitly involve the calculation of distances to landmarks; and RF fingerprinting schemes whereby a radio map is constructed using prior measurements, and a device is localized by referencing this radio map. In this work, we focus on indoor signal-strength-based localization algorithms utilizing these approaches. We can further break down the algorithms into two main categories: point-based methods, and area-based methods.

## 2.1 Point-Based Algorithms

Point-based methods return an estimated point as a localization result. Here we describe a few representative point-based schemes for our study.

*RADAR (R1).* A primary example of a point-based method is the RADAR scheme [Bahl and Padmanabhan 2000]. In R1, multiple base stations are deployed to provide overlapping coverage of an area, such as an office building. During set up, a mobile host with known position broadcasts beacons periodically, and the signal strength readings are measured at a set of fixed landmarks. Collecting together the averaged signal strength readings from each of the landmarks for different transmitter locations provides a radio map. After training, localization is performed by measuring a wireless device's RSS at each landmark, and the vector of RSS values is compared to the radio map. The record in the radio map whose signal strength vector is closest in the Euclidean sense to the observed signal strength vector is declared to correspond to the location of the transmitter. Variations of RADAR, such as *Averaged RADAR (R2)*, which returns the average of the closest two fingerprints and *Gridded RADAR (GR)*, which uses the Interpolated Map Grid (IMG) as a set of additional fingerprints over the basic RADAR have been proposed in Elnahrawy et al. [2004].

*Highest Probability (P1).* The P1 method uses a probabilistic approach by applying the statistical Bayes' rule to return the point with the highest probability in the preconstructed radio map as the location estimation result [Roos et al. 2002]. There are variations of Highest Probability. *Averaged Highest Probability (P2)* returns the midpoint of the top two training fingerprints. Like GR, *Gridded Probability (GP)* uses fingerprints based on an IMG [Elnahrawy et al. 2004].

## 2.2 Area-Based Algorithms

On the other hand, area-based algorithms return a *most likely* area in which the true location resides. One of the major advantages of area-based methods compared to point-based methods is that they return a region, which has an increased chance of capturing the transmitter's true location. We study three area-based algorithms [Elnahrawy et al. 2004; Madigan et al. 2005], two of them, Simple Point Matching (SPM) and Area Based Probability (ABP), use an Interpolated Map Grid (IMG) and perform scene matching (fingerprint matching) for localization; and the other, Bayesian Networks (BN), is a multilateration algorithm.

*Simple Point Matching (SPM).* In SPM, the floor is divided into small tiles. The strategy behind SPM is to find a set of tiles that fall within a threshold of the RSS for each landmark independently, then return the tiles that form the intersection of each landmark's set. We define the threshold as:

$$s_i \pm q, \quad (1)$$

where  $s_i$  is the expected value of the RSS reading from Landmark  $i$ , and  $q$  is an expected noise level. One way to choose  $q$  is to use the maximum of the standard deviation  $\sigma$  with:

$$\sigma = \max\{\sigma_{ij}; i \in \{1..number\ of\ landmarks\}, j \in \{1..number\ of\ points\}\}. \quad (2)$$

SPM [Elnahrawy et al. 2004] is an approximation of the Maximum Likelihood Estimation (MLE) method.

*Area Based Probability (ABP).* ABP returns a set of tiles bounded by a probability that the transmitter is within the returned tile set. The probability is called the confidence,  $\alpha$ , and is adjustable by the user. ABP assumes the distribution of RSS for each landmark follows a Gaussian distribution with mean as the expected value of RSS reading vector,  $\mathbf{s}$ . The Gaussian random variable from each landmark is independent. ABP then computes the probability of the transmitter being at each tile  $L_i$ , with  $i = 0 \dots L$ , on the floor using Bayes' rule:

$$P(L_i | \mathbf{s}) = \frac{P(\mathbf{s} | L_i) \times P(L_i)}{P(\mathbf{s})}. \quad (3)$$

Given that the transmitter must be at exactly one tile, satisfying  $\sum_{i=1}^L P(L_i | \mathbf{s}) = 1$ , ABP normalizes the probability and returns the most likely tiles up to its confidence  $\alpha$  [Elnahrawy et al. 2004].

*Bayesian Networks (BN).* BN localization is a multilateration algorithm that encodes the signal-to-distance propagation model into the Bayesian Graphical Model for localization [Madigan et al. 2005]. In BN, the overall joint density of  $x \in X$ , where  $x$  is a random variable, only depends on the parents of  $x$ , denoted  $pa(x)$ :

$$p(X) = \prod_{x \in X} p(x | pa(x)). \quad (4)$$

Once  $p(X)$  is computed, the marginal distribution of any subset of the variables of the network can be obtained as it is proportional to the overall joint distribution.

Figure 1 presents two Bayesian Network algorithms,  $M_1$  and  $M_2$ , that we used for our analysis. Each rectangle is a plate, and shows a part of the network that is replicated; in our case, the nodes on each plate are repeated for each of the  $n$  landmarks whose locations are known. The vertices  $X$  and  $Y$  represent location; the vertex  $S_i$  is the RSS reading from the  $i$ th landmark; and the vertex  $D_i$  represents the Euclidean distance between the location specified by  $X$  and  $Y$ , and the  $i$ th landmark. The value of  $S_i$  follows a signal propagation model  $S_i = b_{i0} + b_{i1} \log D_i$ , where  $b_{i0}, b_{i1}$  are the parameters specific to the  $i$ th landmark. The distance  $D_i = \sqrt{(X - x_i)^2 + (Y - y_i)^2}$  in turn depends on






Fig. 1. Bayesian Networks localization algorithm: Bayesian Graphical Models using plate notation.

the location  $(X, Y)$  of the measured signal and the coordinates  $(x_i, y_i)$  of the  $i$ th landmark. The networks model noise and outliers by modeling the  $S_i$  as a Gaussian distribution around the aforementioned propagation model, with variance  $\tau_i$ :

$$S_i \sim N(b_{i0} + b_{i1} \log D_i, \tau_i). \quad (5)$$

The initial parameters  $(b_{i0}, b_{i1}, \tau_i)$  of the model are unknown; usually the training data is used to adjust the specific parameters of the model according to the relationships encoded in the network. Through Markov Chain Monte Carlo (MCMC) simulation, BN returns the sampling distribution of the possible locations of  $X$  and  $Y$  as the localization result.

$M_1$  utilizes a simple Bayesian Network model, as depicted in Figure 1(a), and requires location information in the training set in order to give good localization results. The  $M_2$  model is hierarchical as shown in Figure 1(b), by making the coefficients of the signal propagation model have common parents. The BN  $M_2$  algorithm can localize multiple devices simultaneously with no training set, leading to a zero-profiling technique for location estimation.

The algorithms we have described in this section are summarized in Table I. Although there are a variety of other signal-strength-based localization algorithms that may be studied, we believe that our results are general and can be applied to other point-based and area-based methods.

### 3. CONDUCTING SIGNAL STRENGTH ATTACKS

In this section, we study the feasibility of conducting signal strength attacks. We first discuss the possible attacks on signal strength. We then provide experimental results for signal strength going through various materials. Finally, we derive an attack model for our performance analysis of the robustness of localization algorithms.

Table I. Algorithms Under Study

Algorithm	Abbreviation	Description
Area-Based		
Simple Point Matching	SPM	Maximum likelihood matching of the RSS to an area using thresholds.
Area Based Probability	ABP- $\alpha$	Bayes rule matching of the RSS to an area probabilistically bounded by the confidence level $\alpha\%$ .
Bayesian Network ( $M_1, M_2$ )	BN	Returns the most likely area using a Bayesian network approach.
Point-Based		
RADAR	R1	Returns the closest record in the Euclidean distance of signal space.
Averaged RADAR	R2	Returns the average of the top 2 closest records in the signal map.
Gridded RADAR	GR	Applies RADAR using an interpolated grid signal map.
Highest Probability	P1	Applies maximum likelihood estimation to the received signal.
Averaged Highest Probability	P2	Returns the average of the top 2 likelihoods.
Gridded Highest Probability	GP	Applies likelihoods to an interpolated grid signal map.

### 3.1 Signal Strength Attacks

The first step to tackle a security problem is to put oneself in the role of the adversary and attempt to understand the attacks. To attack signal-strength based localization systems, an adversary must attenuate or amplify the RSS readings. This can be done by applying the attack at the transmitting device, for example simply placing foil around the 802.11 card; or by directing the attack at the landmarks. For example, we may steer the lobes and nulls of an antenna to target selected landmarks. A broad variety of attenuation attacks can be performed by introducing materials between the landmarks and sensors Li et al. [2005].

In order to support the claim that physical attacks on received signal strength are feasible and capable of significantly affecting the results of a localization algorithm, we first examined the possibility of signal strength attacks. Next, we report results of actual experiments to quantify the effectiveness of various ways of attenuating/amplifying signal strength.

### 3.2 Experimental Results of Attacks

Our experiments were performed in our laboratory on the third floor of the CoRE building at Rutgers University, as shown in Figure 4(a). There are four



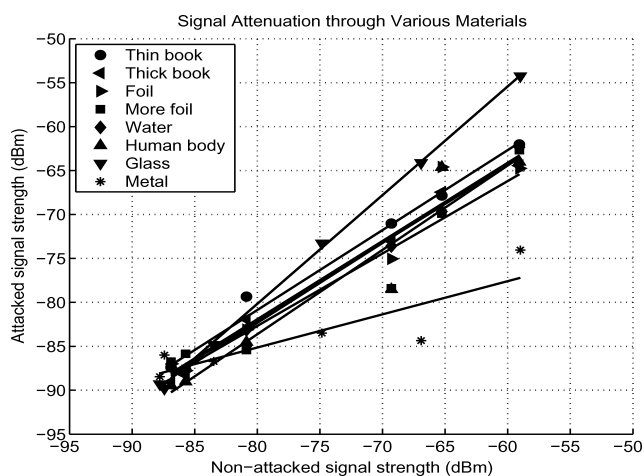


Fig. 2. Signal strength when going through a barrier.

landmarks deployed on the third floor of CoRE. We measured the RSS of beacon signals coming from each of them. The RSS readings were collected using a laptop with an Orinoco Silver wireless card, using `iwlist` to sample the signal strength. In order to mitigate the effect of fluctuations, we collected samples once every second for 10 minutes, and averaged the signal strength over 600 samples.

As noted earlier, an adversary may attack the signal strength by attenuating or amplifying the RSS readings. This can be done either at the receiver or at the transmitter. Our aim is to find the results of power loss in dB by simple attacks. Therefore, in the experiments, we placed various obstruction materials close to the laptop's wireless card and measured the RSS values from each landmark at the laptop. The following obstructions were used: a thin book, a thick book, a layer of metal foil, three layers of foil (referred to as more foil), a mug filled with water (referred to as water), a glass mug (referred to as glass), a metal cabinet (referred to as metal), and a human body. These materials are easy to access and attacks utilizing these kind of materials can be simply performed with low cost. The original signal strength values, together with the signal strength measurements in the presence of these objects, are provided in Figure 2. The points represent the measured data from experimental results for various materials, while the lines are the linear least-squares fitting. The results are intended to show the feasibility of using such materials for attacks. As we would expect, highly attenuating materials such as the metal box or foil have a large impact on signal strength, whereas other materials do not affect the signal much. A more comprehensive study of propagation loss through common materials can be found in Wilson [2002]. We note that more powerful attenuation loss is possible by using more advanced materials (such as RF-absorptive carbon fabric). Finally, we note that these results also imply that amplification is possible by removing a barrier (e.g. a door) of the corresponding material or through antenna-based methods.

### 3.3 Attack Model

Based upon the results in Figure 2, we further see that there is a linear relationship between the unattacked signal strength and the attacked signal strength in dB for various materials. The linear relationship implies that there is an easy way for an adversary to perform and control the effect of an attack on the observed signal strength by appropriately selecting different materials. Specifically, we envision that an adversary may suitably introduce and/or remove barriers of appropriate materials so as to attenuate and amplify the signal strength readings at one or more landmarks. Due to the observed linear relationship illustrated in Figure 2, we refer to this as the linear attack model.

In the rest of this article, we will use the linear attack model to describe the effect of an attack on the RSS readings at one or more landmarks. The resulting attacked readings are then used to study the consequent effects on localization for the surveyed algorithms. In particular, in this study, we apply our attacks to individual landmarks, which might correspond to placing a barrier directly in front of a landmark, as well as to the entire set of landmarks, which corresponds to placing a barrier around the transmitting device. Similar arguments can be made for amplification attacks, whereby usually barriers are removed between the source and receivers. Moreover, we apply attenuation, amplification, or a mixture of simultaneous attenuation and amplification attacks to multiple landmarks and study the performance of localization algorithms. The broad collection of our attack scenarios has covered the set of possibilities that an adversary could attempt to accomplish. Although there are many different and more complex signal strength attack methods that can be used, we believe their effects will not vary much from the linear signal strength attack model we use in this article, and note that such sophisticated attacks could involve much higher cost to perform.

## 4. MEASURING ATTACK SUSCEPTIBILITY

The aim of a localization attack is to perturb a set of signal strength readings in order to have an effect on the localization output. When selecting a localization algorithm, it is desirable to have a set of metrics by which we can quantify how susceptible a localization algorithm is to varying levels of attack by an adversary. In this section, we shall provide a formal specification for an attack, and present several measurement tools for quantifying the effectiveness of an attack.

### 4.1 A Generalized Localization Model

In order to begin, we need to specify a model that captures a variety of RF-fingerprinting localization algorithms. Let us suppose that we have a domain  $D$  in two-dimensions, such as an office building, over which we wish to localize transmitters. Within  $D$ , a set of  $n$  landmarks has been deployed to assist in localization. A wireless device that transmits with a fixed power in an isotropic manner will cause a vector of  $n$  signal strength readings to be measured by the  $n$  landmarks. In practice, these  $n$  signal strength readings are averaged over

a sufficiently large time window to remove statistical variability. Therefore, corresponding to each location in  $D$ , there is an  $n$ -dimensional vector of signal readings  $\mathbf{s} = (s_1, s_2, \dots, s_n)$  that resides in a range  $R$ .

This relationship between positions in  $D$  and signal strength vectors defines a fingerprint function  $F : D \rightarrow R$  that takes our real world position  $(x, y)$  and maps it to a signal strength reading  $\mathbf{s}$ .  $F$  has some important properties. First, in practice,  $F$  is not completely specified, but rather a finite set of positions  $(x_j, y_j)$  is used for measuring a corresponding set of signal strength vectors  $\mathbf{s}_j$ . Additionally, the function  $F$  is generally one-to-one, but is not onto. This means that the inverse of  $F$  is a function  $G$  that is not well-defined: there are holes in the  $n$ -dimensional space in which  $R$  resides for which there is no well-defined inverse.

It is precisely the inverse function  $G$ , though, that allows us to perform localization. In general, we will have a signal strength reading  $\mathbf{s}$  for which there is no explicit inverse (e.g. perhaps due to noise variability). Instead of using  $G$ , which has a domain restricted to  $R$ , we consider various pseudo-inverses  $G_{alg}$  of  $F$  for which the domain of  $G_{alg}$  is the complete  $n$ -dimensional space. Here, the notation  $G_{alg}$  indicates that there may be different *algorithmic* choices for the pseudo-inverse. For example, we shall denote  $G_R$  to be the RADAR localization algorithm. In general, the function  $G_{alg}$  maps an  $n$ -dimensional signal strength vector to a region in  $D$ . For point-based localization algorithms, the image of  $G_{alg}$  is a single point corresponding to the localization result. On the other hand, for area-based methods, the localization algorithm  $G_{alg}$  produces a set of likely positions.

An attack on the localization algorithm is a perturbation to the correct  $n$ -dimensional signal strength vector  $\mathbf{s}$  to produce a corrupted  $n$ -dimensional vector  $\tilde{\mathbf{s}}$ . Corresponding to the uncorrupted signal strength vector  $\mathbf{s}$  is a correct localization result,  $\mathbf{p} = G_{alg}(\mathbf{s})$ , while the corrupted signal strength vector produces an attacked localization result  $\tilde{\mathbf{p}} = G_{alg}(\tilde{\mathbf{s}})$ . Here,  $\mathbf{p}$  and  $\tilde{\mathbf{p}}$  are set-valued and may either be a single point or a region in  $D$ .

#### 4.2 Attack Susceptibility Metrics

We wish to quantify the effect that an attack has on localization by relating the effect of a change in a signal strength reading,  $\mathbf{s}$ , to the resulting change in the localization result,  $\mathbf{p}$ . We shall use  $\mathbf{p}_0$ , to denote the correct location of a transmitter,  $\mathbf{p}$ , to denote the estimated location (set) when there is no attack being performed, and  $\tilde{\mathbf{p}}$  to denote the position (set) returned by the estimator after an attack has affected the signal strength. Figure 3 illustrates the relationship between the true location and the estimated locations. There are several performance metrics that we will use:

—*Estimator angle bias.* The perturbation on the signal strength vector caused by an attack will result in the variability of location estimation in the physical space. We want to investigate the bias along the angular dimension. That is, if we plot the relative error position in polar coordinates, for an unbiased estimator the error would have an equal probability of falling along any angle. However, when attacking a single landmark, we may expect an angular bias to be introduced. The estimation angle bias is studied by calculating the

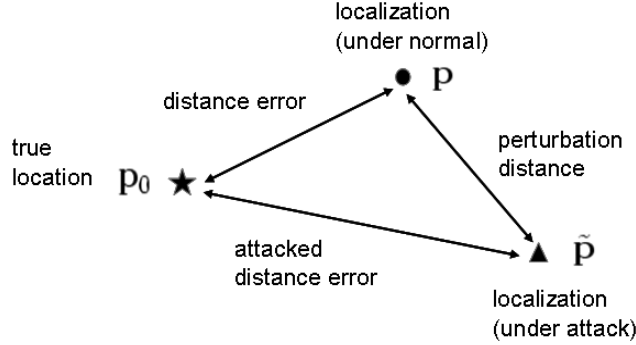


Fig. 3. Interpretation of distances in location estimation.

estimated position for different experimental trials, and comparing these results, in a spatial sense, to the true position. An angularly-unbiased algorithm should uniformly cover the 360 degrees around the true location. For area-based methods, we replace  $\tilde{\mathbf{p}}$ , which is a set, with its median (along the  $x$  and  $y$  dimensions separately) to get a point. The angular bias is an important metric since it can serve as an indication as to whether an attacker can skew the localization result in a specific direction—algorithms with more angular bias are more skewable and hence worse choices for deployment since an adversary can use this knowledge to its advantage.

—*Estimator distance error.* An attack will cause the magnitude of  $\mathbf{p}_0 - \tilde{\mathbf{p}}$  to increase. For a particular localization algorithm  $G_{alg}$ , we are interested in the statistical characterization of  $\|\mathbf{p}_0 - \tilde{\mathbf{p}}\|$  over all possible locations in the building. The characterization of  $\|\mathbf{p}_0 - \tilde{\mathbf{p}}\|$  depends on whether a point-based method or an area-based method is used, and can be described via its mean and distributional behavior. For a point-based method, we may measure the cumulative distribution (cdf) of the error  $\|\mathbf{p}_0 - \tilde{\mathbf{p}}\|$  over the entire building. For area-based metrics, we calculate the CDF of the distance between the median of the estimated locations  $\tilde{\mathbf{p}}_{med}$  and the true location, i.e.  $\|\mathbf{p}_0 - \tilde{\mathbf{p}}_{med}\|$ .

The CDF provides a complete statistical specification of the distance errors. It is often more desirable to look at the average behavior of the error. For point-based methods, the average distance error is simply  $E[\|\mathbf{p}_0 - \tilde{\mathbf{p}}\|]$ , which is just the average of  $\|\mathbf{p}_0 - \tilde{\mathbf{p}}\|$  over all locations. Area-based methods allow for more options in defining the average distance error. First, for a particular value of  $\mathbf{p}_0$ ,  $\tilde{\mathbf{p}}$  is a set of points. For each  $\mathbf{p}_0$ , we get a collection of error values  $\|\mathbf{p}_0 - \mathbf{q}\|$ , as  $\mathbf{q}$  varies over points in  $\tilde{\mathbf{p}}$ . For each  $\mathbf{p}_0$ , we may extract the minimum, 25th percentile, median, 75th percentile, and maximum. These quartile values of  $\|\mathbf{p}_0 - \mathbf{q}\|$  are then averaged over the different positions,  $\mathbf{p}_0$ .

—*Estimator Precision.* An area-based localization algorithm returns a set  $\mathbf{p}$ . For localization, precision refers to the size of the returned estimated area. This metric quantifies the average value of the area of the localized set  $\mathbf{p}$  over different signal strength readings,  $\mathbf{s}$ . Generally speaking, the smaller the size of the returned area, the more precise the estimation. When an attack is conducted, it is possible that the precision of the answer  $\tilde{\mathbf{p}}$  is affected.

—*Precision vs. Perturbation Distance.* The perturbation distance is the quantity  $\|\mathbf{p}_{med} - \tilde{\mathbf{p}}_{med}\|$ . The precision vs. perturbation distance metric depicts the functional dependency between precision and increased perturbation distance.

—*Hölder Metrics.* In addition to error performance, we are interested in how dramatically the returned results can be perturbed by an attack. Thus, we wish to relate the magnitude of the perturbation  $\|\mathbf{s} - \tilde{\mathbf{s}}\|$  to its effect on the localization result, which is measured by  $\|G_{alg}(\mathbf{s}) - G_{alg}(\tilde{\mathbf{s}})\|$ . In order to quantify the effect that a change in the signal strength space has on the position space, we borrow a measure from functional analysis Lang [1993], called the Hölder parameter (also known as the Lipschitz parameter) for  $G_{alg}$ . The Hölder parameter,  $H_{alg}$  is defined via:

$$H_{alg} = \max_{\mathbf{s}, \mathbf{v}} \frac{\|G_{alg}(\mathbf{s}) - G_{alg}(\mathbf{v})\|}{\|\mathbf{s} - \mathbf{v}\|}, \quad (6)$$

where  $\mathbf{s}$  and  $\mathbf{v}$  are all the possible combinations of signal strength vectors in signal space. For continuous  $G_{alg}$ , the Hölder parameter measures the maximum (or worst-case) ratio of variability in position space for a given variability in signal strength space. Since the traditional Hölder parameter describes the worst-case effect an attack might have, it is natural to also provide an average-case measurement of an attack, and therefore we introduce the average-case Hölder parameter:

$$\bar{H}_{alg} = \text{avg}_{\mathbf{s}, \mathbf{v}} \frac{\|G_{alg}(\mathbf{s}) - G_{alg}(\mathbf{v})\|}{\|\mathbf{s} - \mathbf{v}\|}. \quad (7)$$


These parameters are only defined for continuous functions  $G_{alg}$ , but many localization algorithms are not continuous. For example, if we look at  $G_R$  for RADAR, the result of varying a signal strength reading is that it will yield a *stair-step* behavior in position space: small changes will map to the same output and then suddenly, as we continue changing the signal strength vector, there will be a change to a new position estimate (we have switched over to a new Voronoi cell in signal space). In reality, this behavior does not concern us too much, as we are merely concerned with whether adjacent Voronoi cells map to close positions. We will revisit this issue in Section 6. Finally, we emphasize that Hölder metrics measure the perturbability of the returned results, and do not directly measure error.

## 5. EXPERIMENTAL RESULTS


In this section we present our experimental results. We first describe our method. Next, we examine the impact of attacks on the RSS to localization bias and localization error under different attacking scenarios. We then quantify the algorithms' linear responses to RSS changes. Finally, we present a precision study that investigates the impact of attacks on the returned areas for area-based algorithms.

### 5.1 Experimental Setup

Figure 4 shows our experimental setup. The floor map on the left, (a) is the third floor of the CoRE building at Rutgers, which houses the computer science



(a) CoRE Building



(b) Industrial Lab

Fig. 4. Deployment of landmarks and training locations on the experimental floors.

department and has an area of  $200 \times 80\text{ft}$  ( $16000\text{ft}^2$ ). The other floor, shown in (b), is an industrial research laboratory (we call it the Industrial Lab), which has an area of  $225 \times 144\text{ft}$  ( $32400\text{ft}^2$ ). The stars are the training points, the small dots are testing points, and the larger squares are the landmarks, which are 802.11 access points. Notice that the four CoRE landmarks are more colinear than the five landmarks in the Industrial Lab. Next, we perform a trace-driven simulation study to apply our linear attack model to the experimental data collected from these two buildings.

For both attenuation and amplification attacks, we ran the algorithms but modified the measured RSS of the testing points collected from these two office buildings. Specifically, we altered the RSS by  $\pm 5\text{dB}$  to  $\pm 25\text{dB}$ , in increments of  $5\text{dB}$ . We experimented with different ways to handle signals that would fall below the detectable threshold of  $-92\text{dBm}$  for our wireless cards. We found that substituting the minimal signal ( $-92\text{dBm}$ ) produced about the



same localization results and did not require changing the algorithms to special case missing data.

We experimented with different training set sizes, including 20, 35, 60, 85, 115, 145, 185, 215, 245, 253, and 286 points. Experimental data was collected at a total of 286 locations in the CoRE building and at a total of 253 locations in the Industrial Lab. Although there are some small differences, we found that the behavior of the algorithms matches previous results [Elnahrawy et al. 2004] and varied little after using 115 training points. We therefore chose to use a training set size of 115 for this study.

## 5.2 Localization Angle Bias

In this section, we study the angular bias of the localization schemes introduced by signal strength attacks. For the Industrial Lab, Figure 5(a) shows the localization result of ABP under no attack for the relative estimation positions to the true locations, setting as the origin, over all the localization attempts. The normal performance of the algorithms is unbiased, with the localization results uniformly distributed around the true locations.

Figure 5(b) shows the relative position estimation results under 25dB attenuation attack on all landmarks, while Figure 5(c) and Figure 5(d) show the attacked results on single landmarks, landmark 1 and landmark 3, respectively. Figure 4(b) shows that landmark 1 and landmark 3 are placed in diagonal positions across the Industrial Lab. We have observed that signal strength attacks have affected the localization schemes by introducing angular bias on the results with the location estimation more likely to be in the fourth quadrant relative to the true location when landmark 1 is attacked, as shown in Figure 5(c). Because landmark 1 is placed in the upper left corner in the building floor map shown in Figure 4, signal attenuation on landmark 1 made the localization system think the sensor node is farther away from landmark 1, and thus the resulting localization results under attack have been pushed into the fourth quadrant. This effect has been proved by examining the localization results when landmark 3 is under attack. As presented in Figure 5(d), the relative localization results are mostly in the second quadrant since landmark 3 is placed in the lower right corner of the building floor map. Further, as expected, for simultaneous landmark attacks, the localization results are distributed around the true locations randomly, but with much larger estimation errors as presented in Figure 5(b). We have observed similar effects for the other algorithms in the Industrial Lab and the CoRE building.

## 5.3 Localization Error Analysis

In this section, we analyze the estimator distance error through the statistical characterization of  $\|\mathbf{p}_0 - \tilde{\mathbf{p}}\|$  by presenting the error CDFs of all the algorithms as a function of attenuation and amplification attacks. The CDF provides a complete statistical specification of the distance errors. Specifically, we study the localization error under four attack scenarios: an all-landmark attack; a single landmark attack; attacks involving multiple landmarks; and attacks involving simultaneous amplification and attenuation on multiple landmarks.

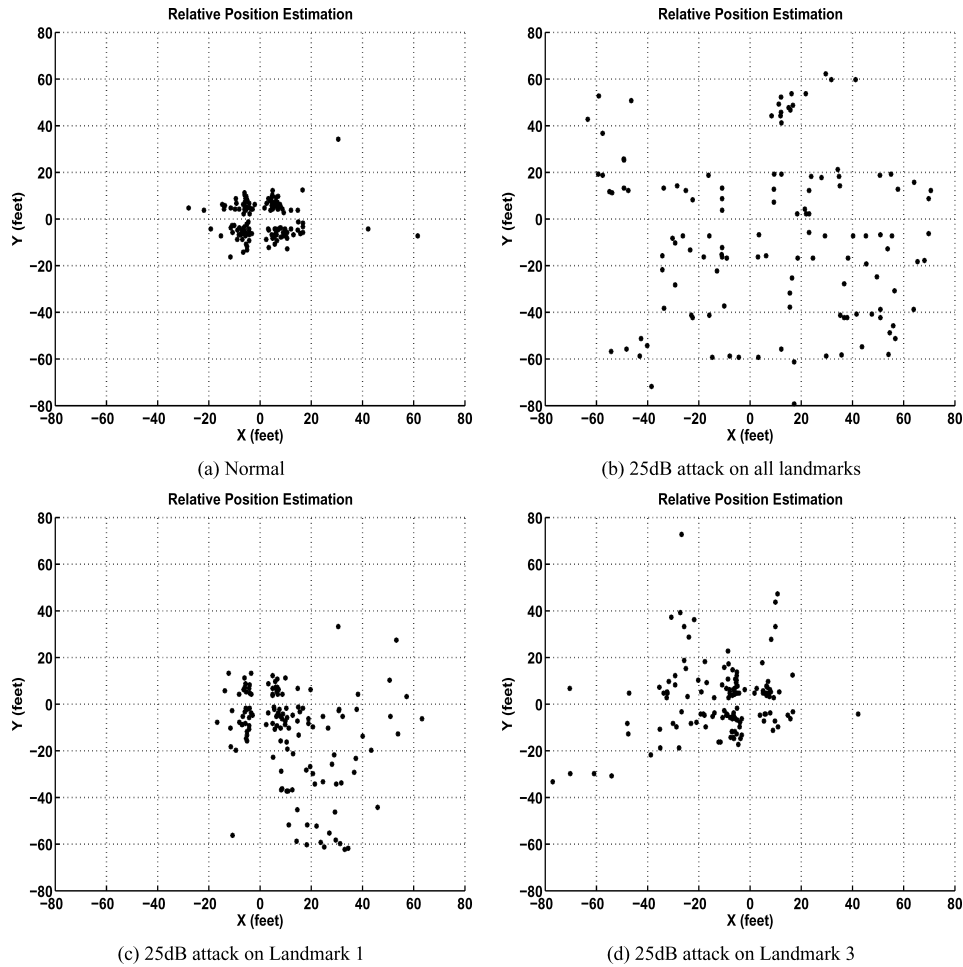


Fig. 5. ABP: localization estimation relative to the true locations for the Industrial Lab.

As a baseline, Figure 6(a) shows the normal performance of the algorithms for the CoRE building and (e) shows the results for the Industrial Lab. For the area-based algorithms, the median tile error is presented, as well as the minimum and maximum tile errors for ABP-75. For BN, we present the results using the simple Bayesian Network  $M_1$  algorithm, denoted as BNmed in the plot. Note that the results from Bayesian Network  $M_2$  are, in fact, better than  $M_1$ , and are comparable to the results for the RADAR scheme R1. However, for the sake of clarity of the plot, we have chosen to only present the results of  $M_1$ . As in previous work [Elnahrawy et al. 2004], the algorithms all obtain similar performance, with the exception of BN, which slightly under-performs the other algorithms.

First, we look at the performance of localization algorithms under an all-landmark attack. Figures 6(b) and 6(c) show the error CDFs under simultaneous landmark attenuation attacks of 10 and 25 dB for CoRE, respectively,

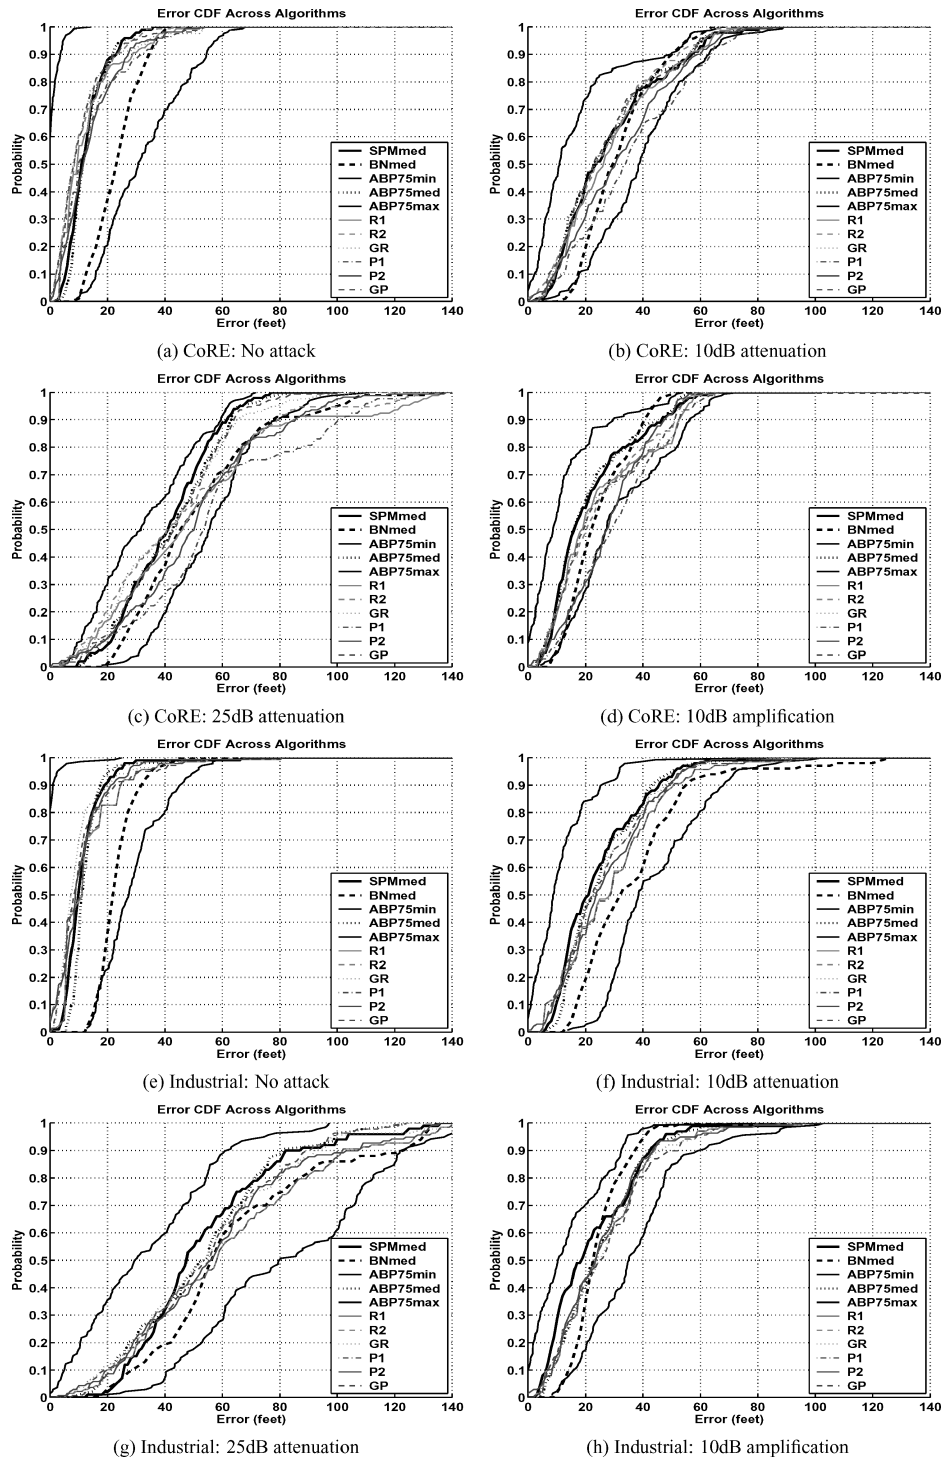


Fig. 6. Error CDF across localization algorithms when attacks are performed on all landmarks.

while Figures 6(f) and 6(g) show similar results in the industrial lab. First, the bulk of the curves shift to the right by roughly equal amounts: no algorithm is qualitatively more robust than the others. Comparing the two buildings, the results show that the industrial lab errors are slightly higher for attacks at equal dB, but again, qualitatively the impact of the building environment is not very significant.

Figures 6(d) and 6(h) show the error CDFs for the CoRE and Industrial Lab under a 10 dB amplification attack. The results are qualitatively symmetric with respect to the outcomes of the 10dB attenuation attack. We found that, in general, comparing amplifications to attenuations of equal dB, the errors were qualitatively the same.

An interesting feature is that in CoRE the minimum error for ABP-75 also shifts to the right by roughly the same amount as the other curves. Figures 6(a) and 6(e) show that, in the non-attacked case, the minimum tile error for ABP-75 is quite small, meaning that the localized node is almost always within or very close to the returned area. However, under attacks, the closest part of the returned area moves away from the true location at the same rate as the median tile. We observed similar effects for the SPM and BN algorithms. We noticed that under large attacks around 25 dB, the median error CDF curves in the Industrial Lab have similar performance to those from the CoRE building, but there are two curves that seem to be outliers, namely ABP75 min and ABP75max. These two curves represent the best and the worst cases from the ABP algorithm. We see that they are not moving at the same speed as the median errors, when compared with the results of the CoRE building. This tells us that the variance/spread of the performance of area-based algorithms in the Industrial Lab has increased under an all-landmark attack, but that the average behavior is consistent across the two buildings.

We then examine attacks against a single landmark. We found attacks against certain landmarks had a much higher impact than against others in the CoRE building. Figures 7(a) and 7(b) show the difference in the error CDF by comparing attacks of landmarks 1 and 2. Figure 4(a) shows that landmark 1 is at the left end of the building, while landmark 2 is in the center and is close to landmark 4. The tails of the curves in Figure 7(a) are much worse than for 7(b), showing that when landmark 1 is attacked, significantly more high errors are returned. Figures 7(c) and 7(d) show a similar effect for amplification attacks. This is because landmark 1 is at one end of the building alone. The contribution of the signal strength reading from landmark 1 plays an important role in localization, while the contribution of landmark 2 can be reduced by the contribution from the nearby landmark 4 when under attack.

The Industrial Lab results in Figures 7(e)–(h) show much less sensitivity to landmark placement compared to the CoRE building. Figure 4(b) shows that landmark 5 is centrally located and we initially suspected this would result in increased attack sensitivity. However, the error CDFs show that the remaining four landmarks provide sufficient coverage: as landmark 5 is attacked, the error CDFs are not much different from attacking landmark 4. The landmark placement in the CoRE building is colinear (to maximize the signal coverage

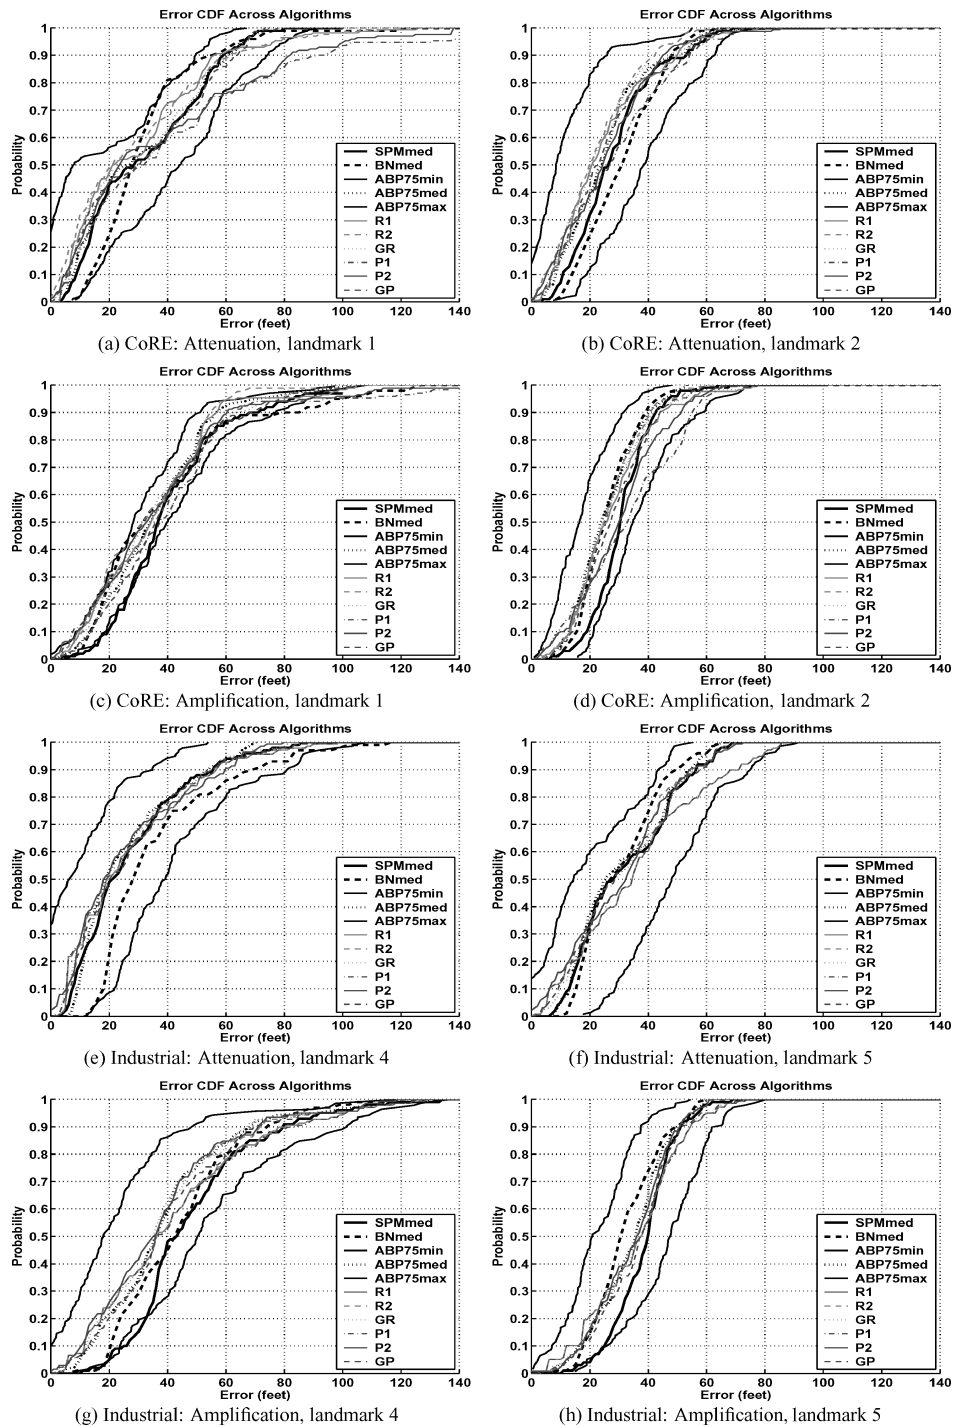


Fig. 7. Error CDF across localization algorithms when attacks are performed on an individual landmark. The attack is 25dB of signal attenuation and signal amplification respectively.

on the floor), while the landmark placement in the Industrial Lab is closer to the optimal landmark placement for location accuracy. We believe that better landmark placement for localization [Chen et al. 2006a] in the Industrial Lab can account for the localization performance being less sensitive to landmark placement under attack.

Next, we study attacks on more than one landmark, but not on all landmarks. Figure 8 presents the localization results in the CoRE building when attenuation attacks are performed on multiple landmarks, specifically on landmark pairs, 1 and 2, 1 and 3, and 2 and 4. We found that attacks on landmark pair 1 and 3 shown in Figure 8(d) cause larger errors compared to results in Figures 8(b) and 8(f) when attacking landmark pairs 1 and 2, and 2 and 4. Since landmarks 1 and 3 are placed at two ends of the building alone, the contribution of the RSS reading from these two landmarks is significant compared to the readings from landmarks 2 and 4, which are closely placed and can cover each other. In general, the impact of multiple landmark attacks on localization performance is between the performance of a single landmark attack and an all-landmark attack.

Fourth, we look at the attack scenario in which the adversary simultaneously performs both amplification and attenuation attacks on multiple landmarks. The localization results are presented in Figure 9 for the CoRE building. For a direct comparison, we present results when mixed attacks are applied on landmark pairs, 1 and 2, 1 and 3, and 2 and 4. We should expect that such an attack would be more effective in falsifying the location results, and this is what we observe. But beyond this, we observe that the performance depends heavily upon which landmarks are attacked. We found that if the attacked landmarks are close to each other, for example landmarks 2 and 4, which are located in the center of the building, the effects of amplification and attenuation attacks are canceled out. Thus the impact of mixed attacks does not lead to significant perturbation in the localization results, as shown in Figure 9(f), which is about the same as under the single landmark attacks displayed in Figure 7. However, if the attacked landmarks are farther away from each other, such as landmarks 1 and 3, which are located at opposite ends of the building, the simultaneous amplification and attenuation attacks can be very harmful and cause larger localization errors for all the algorithms presented in Figure 9(d). The behavior of the error CDFs in Figure 9(d) is qualitatively different from others with very long tails. The effects of the amplification attack on landmark 1 and the attenuation attack on landmark 3 pushed the localization results further in one direction, and thus introduced large localization bias.

The four attack scenarios we studied covered a broad collection of possible combinations of signal strength attacks. We found that simultaneously attacking all landmarks has more impact on localization performance than attacking an individual landmark. Further, simultaneous amplification and attenuation attacks on certain landmarks can cause qualitatively larger errors than other kinds of attacks. Most importantly, we observed that none of the localization algorithms outperforms the others for the attacks we examined.



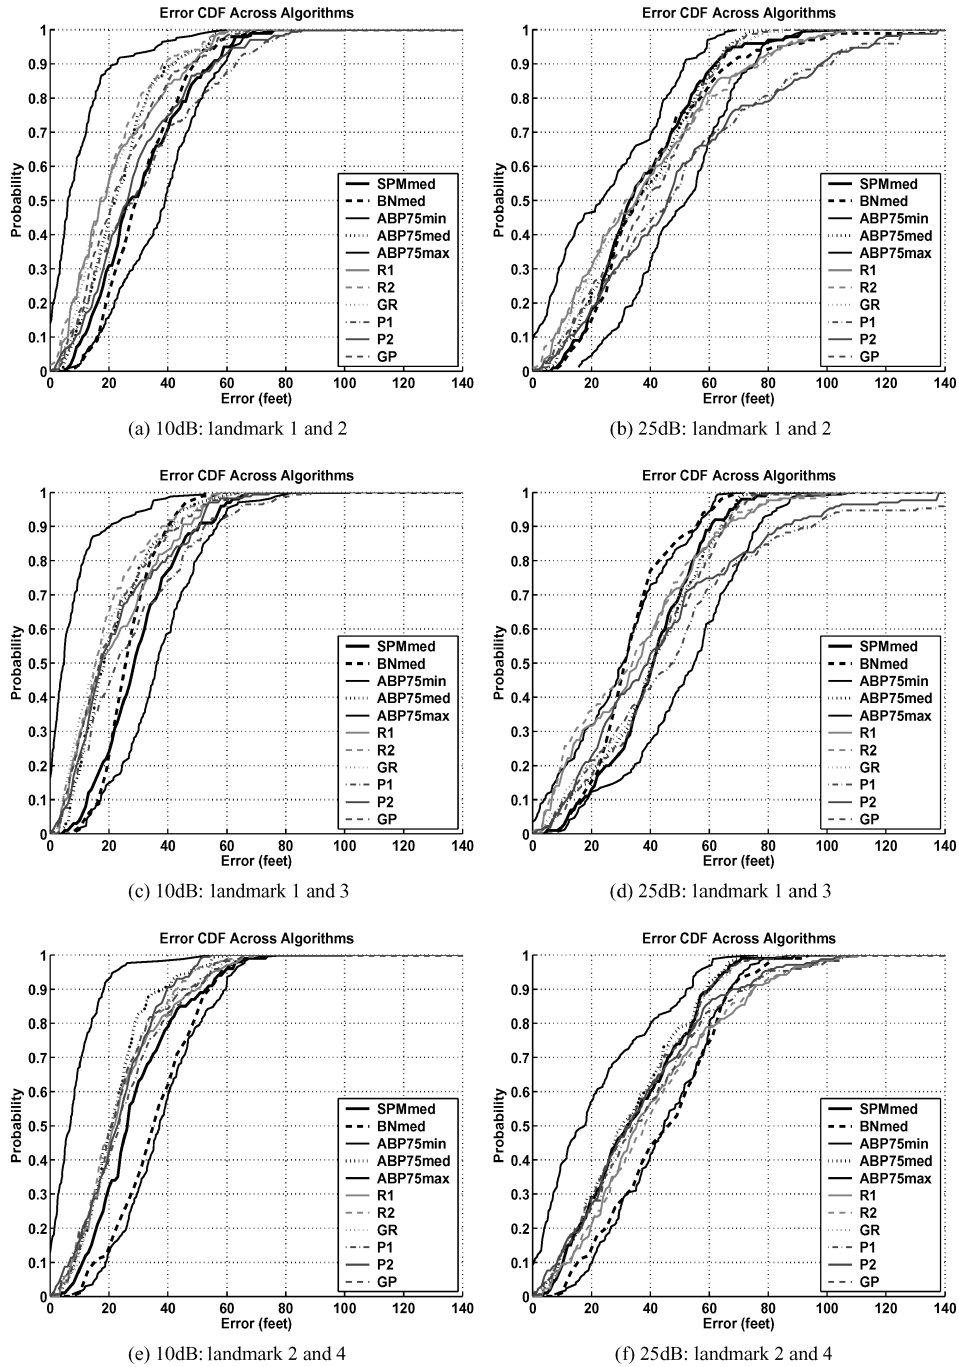


Fig. 8. CoRE: Error CDF across localization algorithms when attenuation attacks are performed on multiple landmarks.

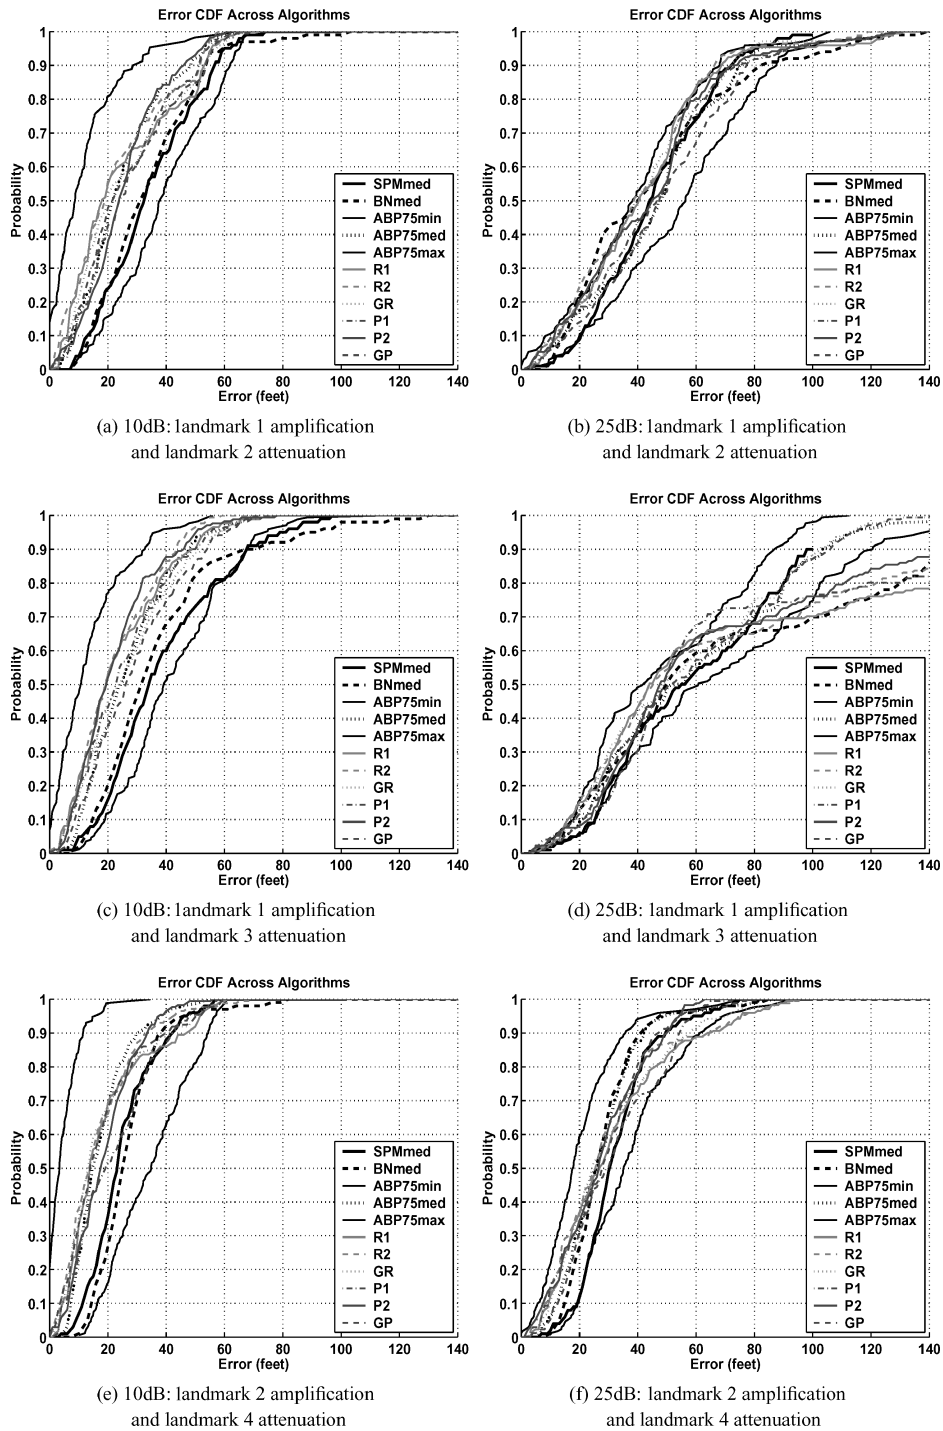


Fig. 9. CoRE: Error CDF across localization algorithms when amplification and attenuation attacks are simultaneously performed on multiple landmarks.

## 5.4 Linear Response

In this section, we show that the average distance error,  $E[\|\mathbf{p}_0 - \hat{\mathbf{p}}\|]$ , of all the algorithms scales in a linear way to attacks. That is, the mean localization error changes linearly with respect to the size of the signal strength change introduced in dB (recall dB is a log-scaled change in power).

Median error versus RSS attenuation for an all-landmark attack is plotted in Figures 10(a) and 10(e), and for individual landmarks in the other figures. Figure 11 plots the median localization error under simultaneous signal strength attenuation and amplification attacks on multiple attacks. Points are data derived from experimental results, and the lines are linear least-squares fits. For BN, in this section we only present the results from Bayesian Network  $M_1$  to compare with other algorithms. But note that the performance for the  $M_2$  algorithm is comparable. The most important feature is that, in all cases, the median responses of all the algorithms fits a line extremely well, with an average  $R^2$ -statistic of 0.97 for both the CoRE and Industrial Lab. The mixed attacks, with amplification attack on landmark 1 and attenuation attack on landmark 3 in CoRE, shown in Figure 11(d), is an exceptional case with  $R^2$  of 0.86 as the worst case.

Comparing the slopes across all the algorithms presented in Tables II, III, and IV, we found a mean change in positioning error versus signal attenuation of 1.55 ft/dB under an all-landmark attack with a minimum of 1.3 ft/dB and maximum of 1.8 ft/dB. For the single landmark attack, the slope was substantially less, 0.64 ft/dB, although BN degrades consistently less than the other algorithms, at 0.44 ft/dB. Under attenuation attacks on multiple landmarks, the localization algorithms move at the speed of 0.9 ft/dB to 1.4 ft/dB, which is between the results of a single landmark attack and an all-landmark attack. However, the median error moves faster under simultaneous amplification and attenuation attacks on landmarks 1 and 3, at the speed of 1.8–2.2 ft/dB, as shown in Table IV. We note the mean error tops out when the attack strength is 25 dB. This confirms our analysis in Figure 9(d) that applying simultaneous amplification and attenuation attacks on landmarks that are farther apart causes larger impacts on the performance of localization schemes, although in practice it is hard for an adversary to conduct simultaneous amplification and attenuation attacks without using sophisticated equipment. In general, the linear fit results are quite important, as it means that no algorithm has a cliff where the average positioning error suffers a catastrophic failure under attack. Instead, it remains proportional to the severity of the attack.

While the median error characterizes the overall response to attacks, it does not address whether an attacker can cause a few, large errors. We examined the response of the maximum error as a function of the strength of the attack on an all-landmark attack: how the 100<sup>th</sup> percentile error scales as a function of the change in dB under an all-landmark attack. The all-landmark attack corresponds to a common attack scenario. It is thus desirable to study the worst-case situation under an all-landmark attack. We note that this characterization is not the same as, nor is directly related to, the Hölder metrics. Those metrics define the rates of change between physical and signal space within

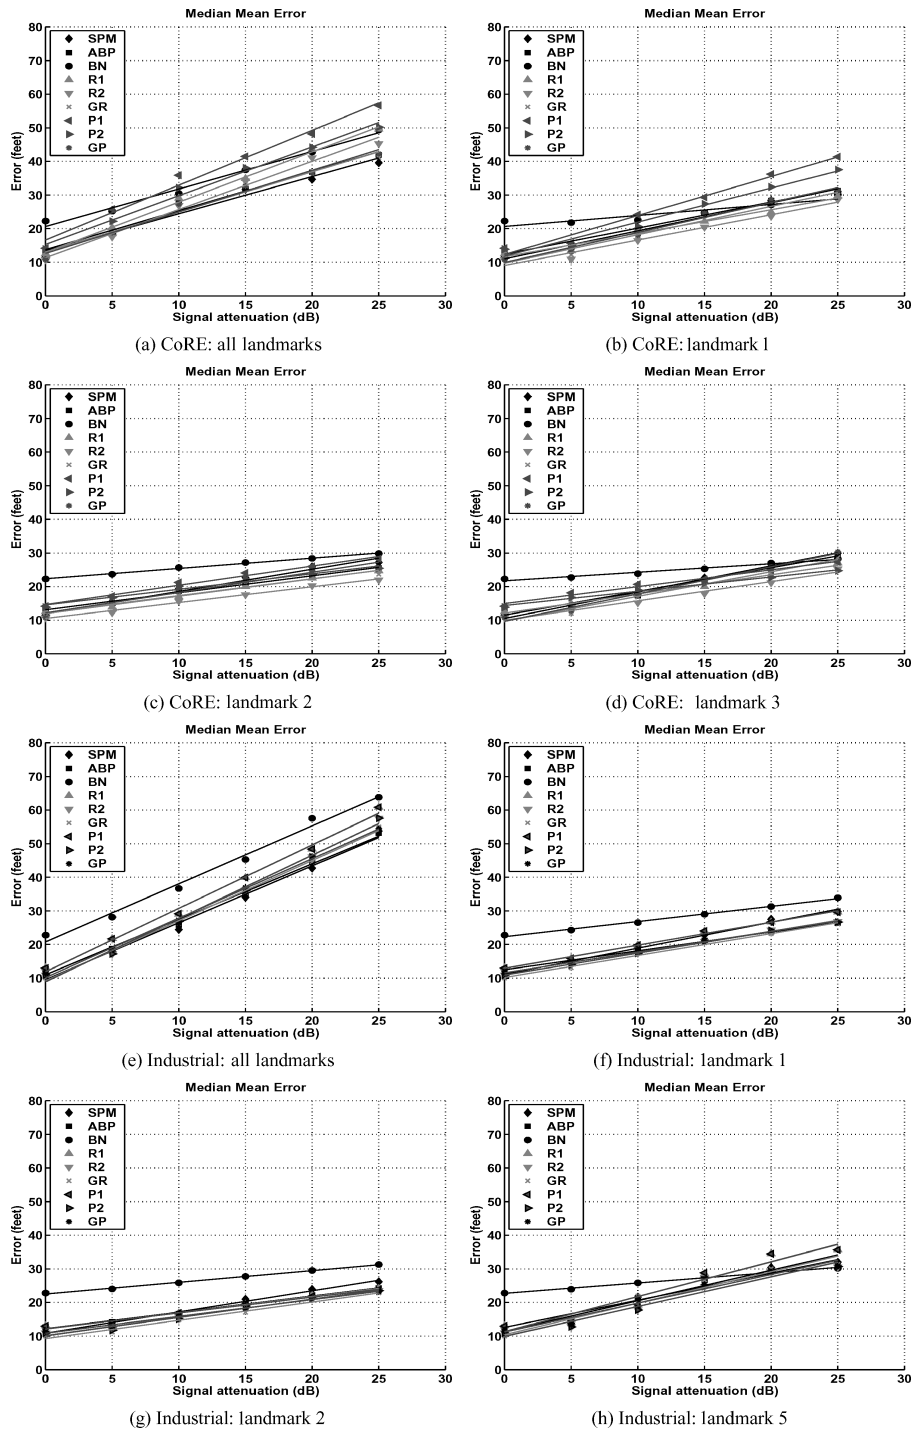


Fig. 10. Average location estimation error across localization algorithms under signal strength attenuation attack.

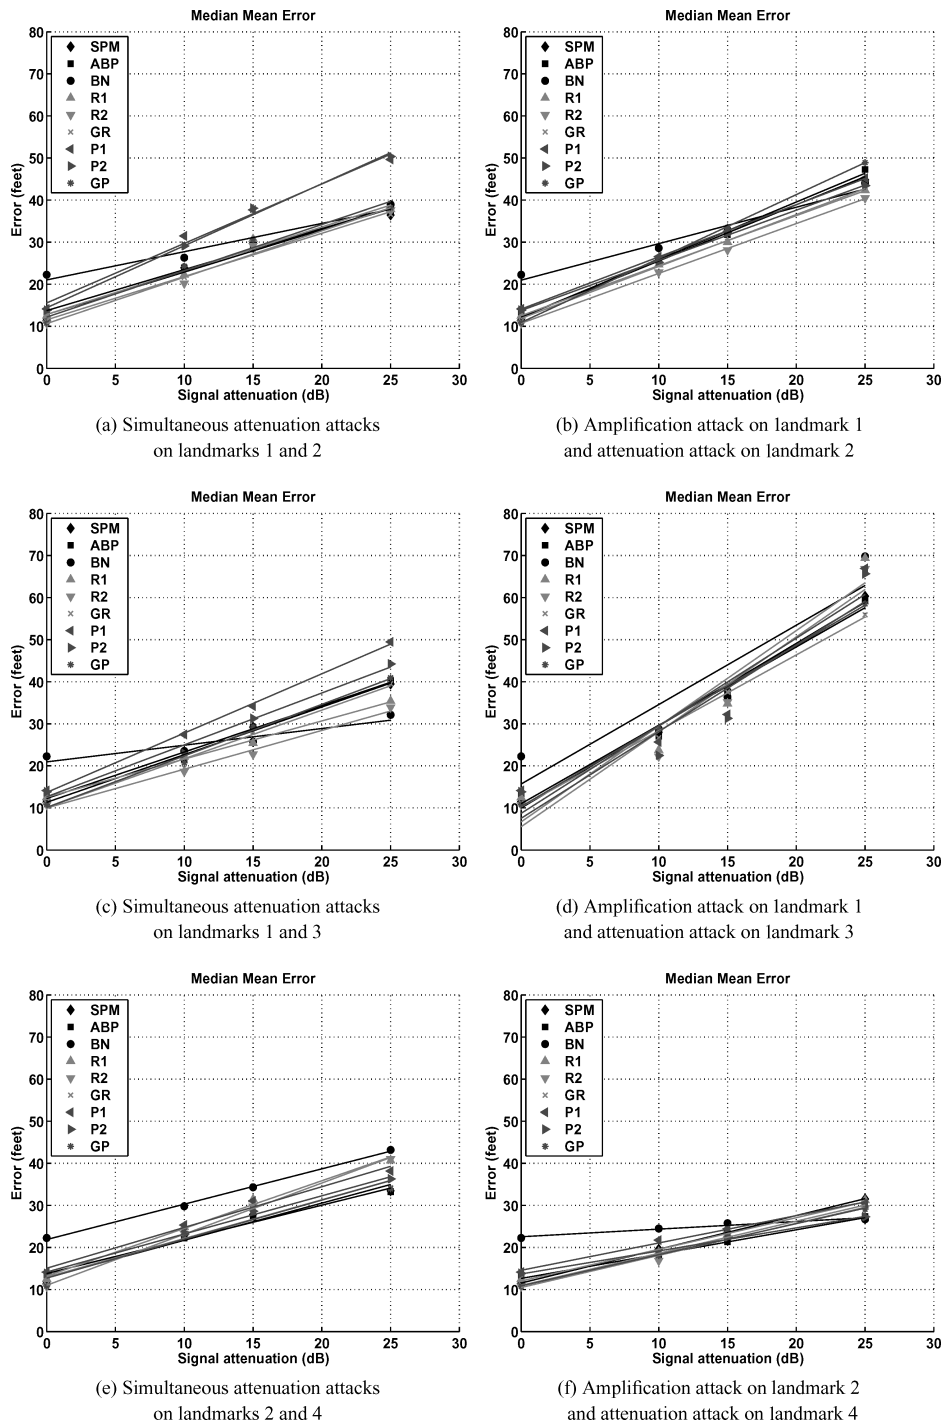


Fig. 11. CoRE: Average location estimation error across localization algorithms under simultaneous signal strength attenuation and amplification attacks on multiple landmarks.

Table II. CoRE: Slopes of Average Error from Linear Regression for Attenuation Attacks on all Landmarks and Individual Landmark

Buildings	CoRE: Attenuation Attack				
Landmarks	All	1	2	3	4
Area-Based					
SPM	1.1048	0.8331	0.662	0.7816	0.6244
ABP-75	1.1656	0.7783	0.5049	0.7052	0.384
BN	1.1157	0.3287	0.3065	0.2544	0.493
Point-Based					
R1	1.4922	0.7006	0.5151	0.5702	0.7941
R2	1.4327	0.7534	0.4687	0.5732	0.7425
GR	1.1896	0.8440	0.5033	0.7357	0.7124
P1	1.6306	1.1597	0.5728	0.5026	0.3644
P2	1.4505	1.0123	0.464	0.4251	0.3063
GP	1.2359	0.8915	0.6028	0.8103	0.4595
Average	1.3131	0.8113	0.5111	0.5954	0.5423

Table III. Industrial: Slopes of Average Error from Linear Regression for Attenuation Attacks on all Landmarks and Individual Landmark

Buildings	Industrial Lab: Attenuation Attack					
Landmarks	All	1	2	3	4	5
Area-Based						
SPM	1.6901	0.7753	0.6283	0.5485	0.6455	0.9103
ABP-75	1.6479	0.5615	0.4852	0.4146	0.5469	0.8072
BN	1.7249	0.4528	0.3487	0.5215	0.5615	0.3094
Point-Based						
R1	1.8823	0.6827	0.4837	0.4286	0.5867	1.0356
R2	1.8816	0.6524	0.5394	0.4000	0.5861	0.8800
GR	1.7860	0.6514	0.5410	0.4668	0.6331	0.9358
P1	1.8854	0.6856	0.4710	0.4532	0.5881	1.0390
P2	1.8802	0.6448	0.5431	0.4023	0.5875	0.8861
GP	1.7666	0.6148	0.4976	0.4800	0.6213	0.8553
Average	1.7939	0.6357	0.504	0.4573	0.5952	0.8510

Table IV. CoRE: Slopes of Average Error from Linear Regression for Mixed Attacks of Signal Attenuation and Amplification on Multiple Landmarks

Buildings	Attenuation Attacks			Amplification and Attenuation Attacks		
Landmarks	1 and 2	1 and 3	2 and 4	1 and 2	1 and 3	2 and 4
Area-Based						
SPM	1.0054	1.1328	0.8836	1.3358	1.9556	0.8018
ABP-75	0.9740	1.1050	0.8125	1.3670	1.8628	0.5778
BN	0.6716	0.3965	0.8401	0.8665	1.8868	0.1812
Point-Based						
R1	1.0392	0.9069	1.1326	1.1895	2.2731	0.7522
R2	1.1013	0.9222	1.2148	1.1841	2.2552	0.7633
GR	1.0276	1.1559	0.9196	1.2337	1.8046	0.7642
P1	1.4142	1.4104	0.9683	1.2414	2.0808	0.6492
P2	1.4735	1.2330	0.9054	1.1921	2.0606	0.5472
GP	1.1003	1.2246	0.9271	1.5197	1.9138	0.7387
Average	1.0897	1.0541	0.9560	1.2367	2.0104	0.6417



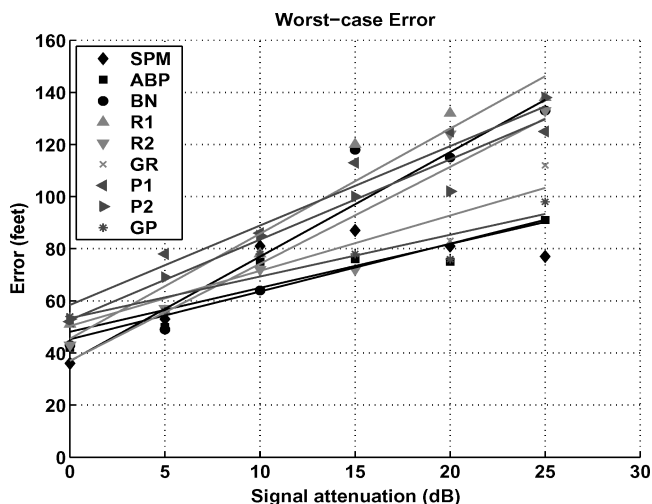


Fig. 12. CoRE: Maximum error as a function of attack strength from an all-landmark attack.

the localization function itself, while here we characterize the change in the estimator error to the change in signal:  $\|\mathbf{p}_0 - \tilde{\mathbf{p}}\|/\|\mathbf{s} - \mathbf{v}\|$ .

Figure 12 plots the worst-case error for each algorithm as a function of signal dB for the CoRE building under an all-landmark attack. The figure shows that almost all the responses are again linear, with least-squares fits of  $R^2$  values of 0.84 or higher, though SPM does not have a linear response. The second important point is the algorithms' responses vary, falling into three groups. BN, R1, and R2 are quite poor, with the worst case error scaling at about 4 ft/dB. P1 and P2, are in a second class, scaling at close to 3 ft/dB. The gridded algorithms, GP and GR, as well as ABP-75 fair better, scaling at 2 ft/dB or less. Finally, SPM is in a class by itself, with a poor linear fit ( $R^2$  of 0.61) and the maximum error topping out at about 85 ft after 15 dB of attack.

Examining the error CDFs and the maximum errors, we can see that most of the localizations move fairly slowly in response to an attack, at about 1.5 ft/dB. However, for some of the algorithms, particularly BN, R1, and R2, the top part of the error CDF moves faster, at about 4 ft/dB. What this means is that, for a few selected points, an attacker can cause more substantial errors of over 100 ft. However, at most places in the building, an attack can only cause much smaller errors.

Figure 10 shows that BN is more robust compared to other algorithms for individual landmark attacks. Recall BN uses a Monte-Carlo sampling technique (Gibbs sampling) to compute the full joint-probability distribution for not just the position coordinates, but also for every node in the Bayesian network. Under a single landmark attack we found the network reduces the contribution of network nodes directly affected by the attacked landmark to the full joint-probability distribution while increasing other landmarks' contributions. In effect, the network discounts the attacked landmark's contribution to the overall joint-density because the attacked data from that landmark is highly unlikely, given the training data.

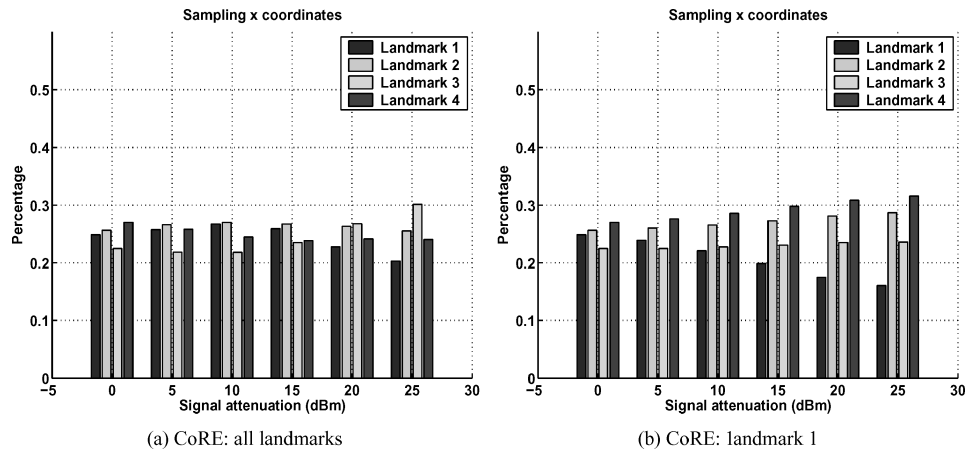


Fig. 13. Contribution of each landmark during sampling in the BN algorithm under attenuation attacks.

To show this effect we developed our own Gibbs sampler so that we could observe the relative contributions of each node in the Bayesian network to the final answer. Figure 13 shows the percentage contribution for each landmark to the overall joint-density. For instance, in CoRE, the contribution of each landmark starts almost uniformly. When landmark 1 is under attack, its contribution goes from 0.25 down to 0.15.

### 5.5 Precision Study

In this section, we examine the area-based algorithms' precision in response to attacks. Figure 14 shows a localization example of the area-based algorithms in the CoRE building. The actual point is shown as a big dot and the convex hulls of the returned areas are outlined. Normally, the SPM and ABP algorithms perform similarly, while the BN algorithm has a much different profile: returning the sampling distribution of the possible estimation. Under signal strength attacks, we observed that the returned areas are reduced and shifted from the true location.

Figure 15 shows the CDF of the precision (size of the returned area) for different area-based algorithms under attack for all the landmarks in CoRE and Industrial Lab. We found that overall the algorithms did not become less precise in response to attacks, but rather, the algorithms tended to shift and shrink the returned areas. Figure 15(a) shows a small average shrinkage for SPM in the CoRE building, and likewise, 15(b) shows a similar effect for BN.

ABP-75 had the most dramatic effect. Figures 15(c) and 15(d) show the precision versus the attack strength for both buildings. The shrinkages are quite substantial. We found that, under attack, the probability densities of the tiles shrank to small values that were located on a few tiles—reflecting the fact that an attacks decreases the likelihood of a position to localize a node. We also found that this effect held for amplification attacks, as is shown in Figure 15(d).

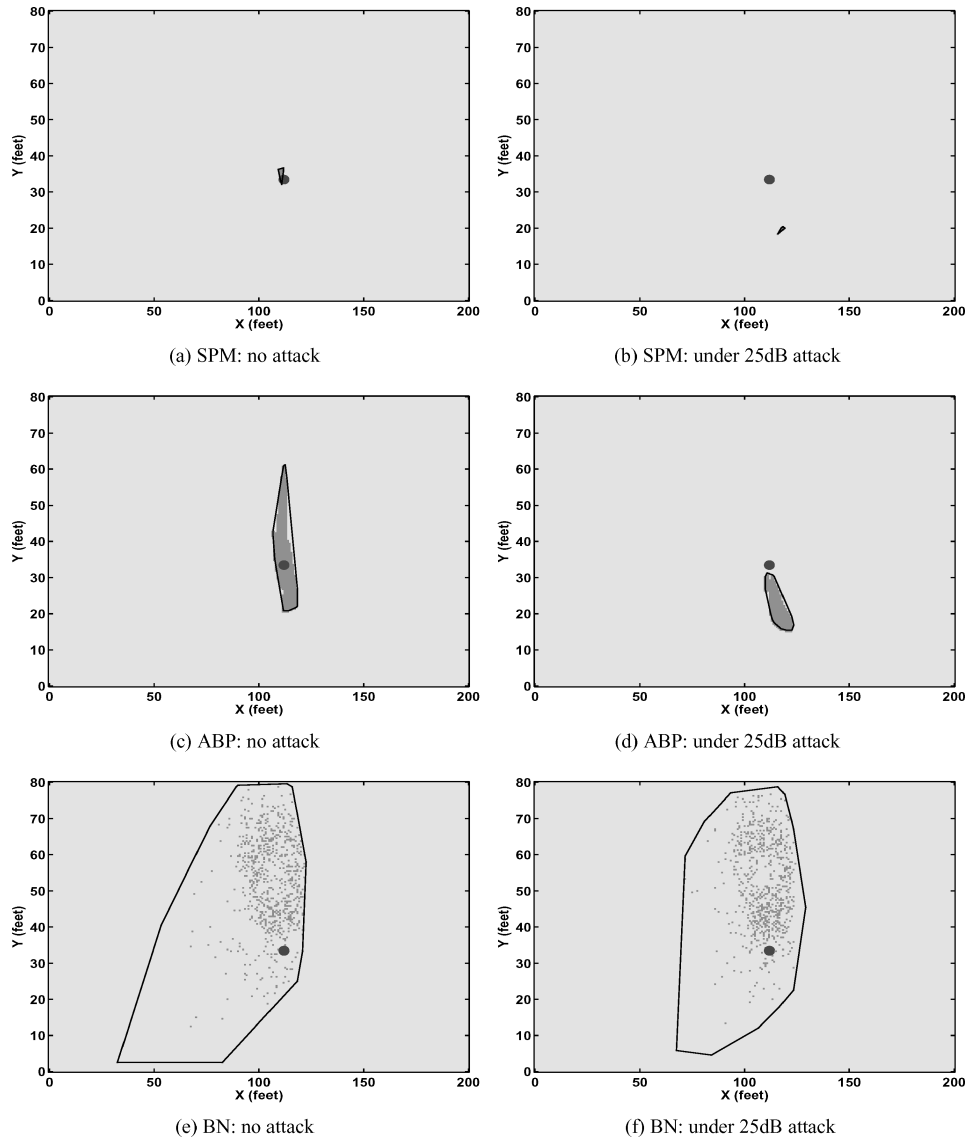


Fig. 14. CoRE: comparison of localization results from the area-based algorithms for a testing point.

The shrinking precision behavior may be useful for attack detection, although a full characterization of how this effect occurs remains for future work.

Examining this effect further, Figure 16 presents the precision versus the perturbation distance  $\|\mathbf{p}_{med} - \tilde{\mathbf{p}}_{med}\|$ , with a least squares line fit. Figure 16(a) shows the effect when attacking all landmarks on the CoRE building. Figure 16(b) shows a downward trend, but much weaker, when one landmark is under attack. We observed similar results for the Industrial Lab. We see mostly linear changes in precision in response to attacks, although with great

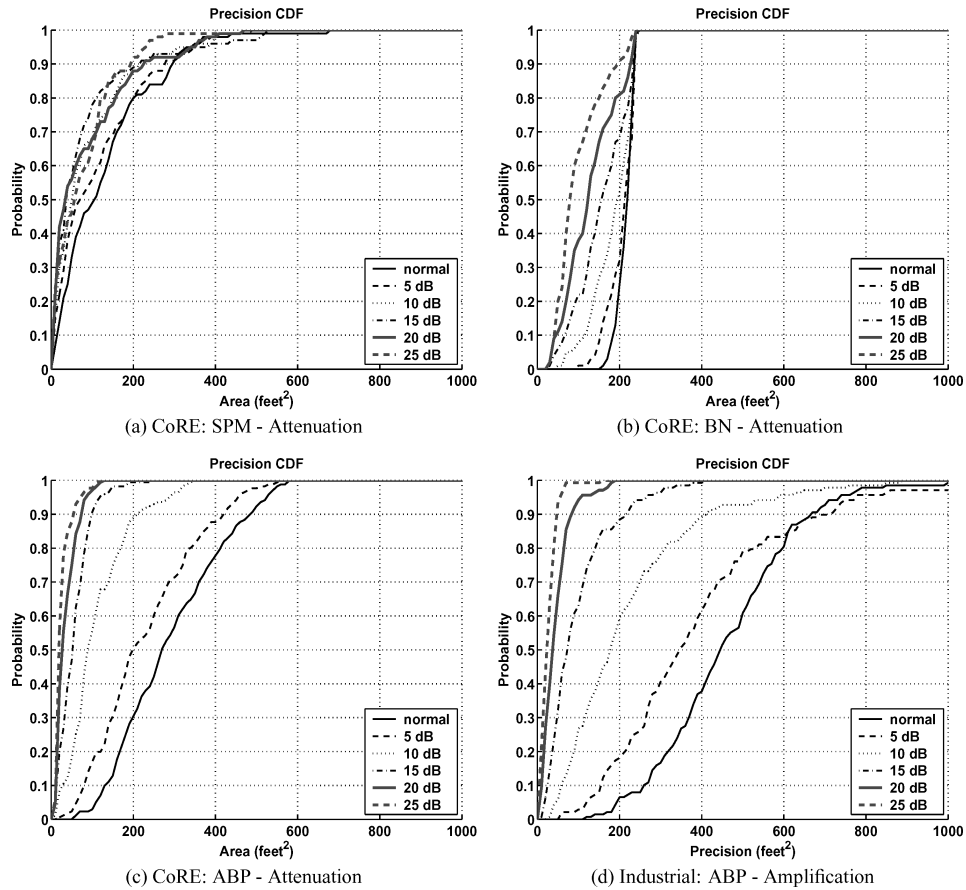


Fig. 15. Analysis of precision CDF across area-based algorithms. The attack is performed on all the landmarks.

differences among the algorithms. The figures show that the decrease in precision as a function of dB is particularly strong for ABP-75.

## 5.6 Robust Multidevice Localization

As presented in Section 2, the Bayesian Network  $M_2$  algorithm can simultaneously localize multiple devices with no training data. Figure 17 presents the error CDFs of  $M_2$  when simultaneously positioning 171 devices under a normal operational situation and with a 25 dB attack applied to all landmarks (i.e., for all signals coming into each landmark) respectively. As shown in Figure 17(a), the performance when simultaneously localizing multiple devices (No Training, Testing = 171) is very similar to that of positioning only one device at a time (Training = 115, Testing = 171). However, we found that under an all-landmark attack with 25 dB severity, the CDF curve of localizing a single device with training data has a large shift to the right, with the same trend as presented in Figure 6; but the curve of localizing multiple devices without training data moves at a much slower speed. This indicates that BN  $M_2$  is more robust

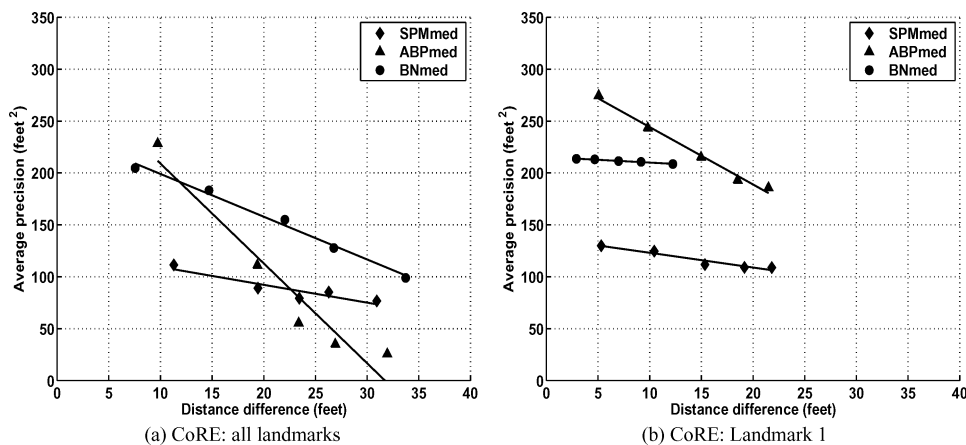
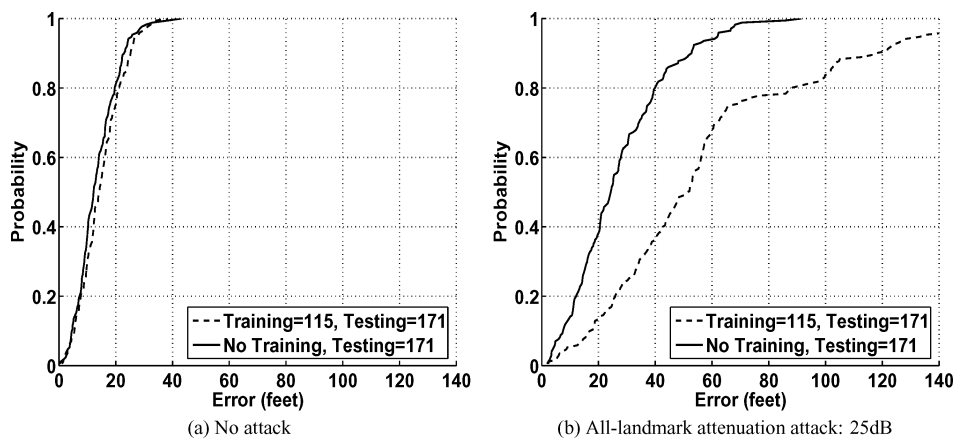


Fig. 16. Precision vs. perturbation distance under attenuation attack.


 Fig. 17. Localization error CDFs using Bayesian Network  $M_2$  algorithm.

than all the other algorithms under attacks, especially under realistic operating conditions involving multiple, simultaneous localization tasks.

Usually Bayesian Networks utilize the training data to predict their model parameters. When no training data is used for localizing multiple transmitting devices,  $M_2$  relies on all the RSS readings from multiple devices to adjust the parameters specified in the model. Under the case of all-landmark attacks, The RSS readings of multiple devices are all corrupted and shifted by a constant. Thus the model parameters predicted by  $M_2$  should be qualitatively similar to those predicted under the normal situation without attack. Therefore, the BN  $M_2$  algorithm is attack-resistant if there is a massive attack, where all signals coming into each landmark are being attacked by adversaries.

Table V. Analysis of Worst-Case  $H$  and Average-Case  $\bar{H}$ 

Algorithms	CoRE: $H$	LAB: $H$	CoRE: $\bar{H}$	LAB: $\bar{H}$
Area-Based				
SPM	23.7646	11.0659	1.8856	2.3548
ABP-75	20.0347	23.0652	1.8548	2.3424
BN	31.7324	14.9168	2.0595	2.5873
Point-Based				
R1	36.2400	20.7846	1.9750	2.3677
R2	19.8586	8.7313	1.9138	2.3058
GR	35.9880	20.6886	1.9691	2.3628
P1	20.8832	20.7846	1.9793	2.3683
P2	19.8586	8.7313	1.9178	2.3058
GP	21.8303	20.6886	1.9649	2.2882

## 6. DISCUSSION ABOUT HÖLDER METRICS

In the previous section we examined the experimental results, and looked at the performance of a set of representative localization algorithms in terms of error and precision. We now focus on the performance of these localization algorithms in terms of the Hölder metrics. The Hölder metrics measure the variability of the returned answer in response to changes in the signal strength vectors.

We first discuss the practical aspects of measuring  $H$  and  $\bar{H}$  for different algorithms. In Section 4, the Hölder parameters are defined by calculating the maximum and average over the entire  $n$ -dimensional signal strength space. In practice, it is necessary to perform a sampling technique to measure  $H$  and  $\bar{H}$ . Additionally, as noted earlier, the definitions of  $H$  and  $\bar{H}$  are only suitable for Hölder continuous functions,  $G_{alg}$ . In reality, several localization algorithms, such as RADAR, are not continuous and involve the tessellation of the signal strength space into Voronoi cells  $V_j$ , and thus only a discrete set of localization results are produced (image of  $V_j$  under  $G_{alg}$ ). Hence, for any  $\mathbf{s} \in V_j$  we have  $G_R(\mathbf{s}) = (x_j, y_j)$ . Unfortunately, for neighboring Voronoi cells, we may take  $\mathbf{s} \in V_j$  and  $\mathbf{v} \in V_i$  such that they are arbitrarily close ( $\|\mathbf{s} - \mathbf{v}\| \rightarrow 0$ ), while  $\|G_R(\mathbf{s}) - G_R(\mathbf{v})\| \neq 0$ . In such a case, the formal calculation of  $H$  and  $\bar{H}$  is not possible. However, for our purposes, we are only interested in measuring the notion of adjacency of Voronoi cells in signal space yielding close localization results. Thus, our calculation of  $H$  and  $\bar{H}$  is only performed over the centroids of the various Voronoi cells for localization algorithms that tessellate the signal strength space.

The Hölder parameters for the different localization algorithms are presented in Table V. Examining these results, there are several important observations that can be made. First, if we examine the results for  $\bar{H}$  we see that, for each building, all of the algorithms have very similar  $\bar{H}$  values. Hence, we may conclude that the average variability of the returned localization result with respect to a change in the signal strength vector is roughly the same for all algorithms. This is an important result as it means, regardless of which RF fingerprinting localization system we deploy, the average susceptibility of the returned results to an attack is essentially identical.



However, if we examine the results for  $H$ , which reflects the worst-case susceptibility, then we see that there are some differences across the algorithms. First, comparing  $H$  and  $\overline{H}$  for both point-based and area-based algorithms, we see that the worst-case variability can be much larger than the average variability. Additionally, the point-based methods appear to cluster. Notably, RADAR (R1) and Gridded Radar (GR) have similar performance across both CoRE and the Industrial Lab, while averaged RADAR (R2) and averaged Highest Probability (P2) have similar performance across both buildings. A very interesting phenomenon is observed by looking at the algorithms that returned an average of likely locations (R2 and P2). Across both buildings, these algorithms exhibited less variability compared to other algorithms. This is to be expected, as averaging is a smoothing operation, which reduces variations in a function. This observation suggests that R2 and P2 are more robust from a worst-case point-of-view than other point-based algorithms.

## 7. RELATED WORK

There have been many active research efforts developing localization systems for wireless and sensor networks. Here, we give a short overview of the different localization strategies. We next discuss the work in secure localization, and then describe the works most closely related to ours.

In general, localization algorithms can be categorized as: range-based versus range-free, scene matching (fingerprint matching), and aggregate or singular. The range-based algorithms involve distance estimation to landmarks using the measurement of various physical properties like RSS [Hightower et al. 2000], Time Of Arrival (TOA) [Enge and Misra 2001] and Time Difference Of Arrival (TDOA) [Priyantha et al. 2000]. Rather than use precise physical property measurements, range-free algorithms use coarser metrics like connectivity [Shang et al. 2003] or hop-counts [Niculescu and Nath 2001] to landmarks to place bounds on candidate positions. In scene matching approaches, a radio map of the environment is constructed, either by measuring actual samples, using signal propagation models, or some combination of the two. A node then measures a set of radio properties (often just the RSS of a set of landmarks), the fingerprint, and attempts to match these to known location(s) on the radio map. These approaches are almost always used in indoor environments because signal propagation is extensively affected by reflection, diffraction and scattering, and thus ranging or simple distance bounds cannot be effectively employed. Matching fingerprints to locations can be cast in statistical terms [Youssef et al. 2003; Roos et al. 2002], as a machine-learning classifier problem [Battiti et al. 2002], or as a clustering problem [Bahl and Padmanabhan 2000]. Finally, a third dimension of classification extends to aggregate or singular algorithms. Aggregate approaches use collections of many nodes in the network in order to localize (often by flooding), while localization of a node in singular methods only requires it to communicate to a few landmarks. For example, algorithms using optimization [Doherty et al. 2001] or multidimensional scaling [Shang et al. 2003] require many estimates between nodes.

Recently, it has been recognized that there are many noncryptographic attacks that can affect localization performance. For example, wormhole attacks tunnel through a faster channel to shorten the observed distance between two nodes [Hu et al. 2003]. Compromised nodes may delay response messages to disrupt distance estimation [Li et al. 2005] and compromised landmarks may even broadcast completely invalid information [Liu et al. 2005]. Physical barriers can directly distort the physical property used by localization. Li et al. [2005] provided a thorough survey of potential attacks to various localization algorithms based on their underlying physical properties.

Secure localization algorithms have been proposed to address these attacks. Capkun and Hubaux [2005] use a distance bounding protocol [Brands and Chaum 1994; Sastry et al. 2003] to upperbound the distance between two nodes. Location estimation (via multilateration) with distances from the bounding protocol can be verified against these bounds and any inconsistency will then indicate attack. Capkun and Hubaux [2006] uses hidden and mobile base stations to localize and verify location estimates. Since such base station locations are hard for attackers to infer, it is hard to launch an attack, thereby providing extra security. Lazos et al. [2005] uses both directional antenna and distance bounding to achieve security. Compared to all these methods, which employ location verification and discard location estimates that indicate attack, Liu et al. [2005] and Li et al. [2005] try to eliminate the effects of attack and still provide good localization. Li et al. [2005] make use of data redundancy and robust statistical methods to achieve reliable localization in the presence of attacks. Liu et al. [2005] propose detecting attacks based on data inconsistency from received beacons and using a greedy search or voting algorithm to eliminate the malicious beacon information.

Lim et al. [2006] proposed a localization algorithm that used the truncated singular value decomposition (SVD) technique and a lateration method for building a zero-configuration, robust, indoor localization system. Based on its theoretical analysis and system development, it is not clear whether this algorithm will be affected by signal strength attacks including an all-landmark attack, a single landmark attack, and mixed attacks of signal attenuation and amplification. Since this algorithm calibrates RSS at known transmission power levels for a wireless device only once at the deployment phase, we suspect that it cannot handle systematic signal strength bias at any dB levels caused by signal attacks, especially when the attack is only performed on a single landmark or multiple landmarks with simultaneous signal amplification and attenuation. Further, Tao et al. [2003] presented an interesting mechanism using the difference of signal strength readings to maintain robustness of localization when a malicious node operates at different power levels or uses different WLAN cards. Although this technique should still work when a node is being attacked, which corresponds to an all-landmark attack scenario in our study, because it takes the differences of RSS readings as localization inputs, this approach cannot address adversarial situations with mixed attacks on multiple landmarks and single landmark attacks. In this work, we studied a representative set of localization algorithms employing signal strength for localization instead of just two algorithms as discussed in Tao et al. [2003]. In previous work, Li et al. [2005]

proposed a possible solution to triangulation-based and fingerprint-based localization, but the susceptibility and performance of various localization methods were not completely investigated. In addition, we have experimentally performed our analysis using real networks deployed in two different buildings, hence our analysis supports the case for robustness in practical system deployments.

## 8. CONCLUSION

In this article, we provided a performance analysis of the robustness of RF-based localization algorithms to attacks that target signal strength measurements. We first examined the feasibility of conducting signal amplification and attenuation attacks, and observed a linear dependency between non-attacked signal strength and attacked signal strength readings for different barriers placed between the transmitter and a landmark receiver. We then provided a set of performance metrics for quantifying the effectiveness of attenuation/amplification attacks and their impacts on localization. Our metrics included localization angular bias, localization error, the precision of area-based algorithms, and a new family of metrics, called Hölder metrics, that quantify the variability of the returned location results versus change in signal strength vectors.

We conducted a trace-driven evaluation of a representative set of point-based and area-based localization algorithms where the linear attack model was applied to data measured in two different office buildings. We investigated the impact of signal attenuation as well as signal amplification attacks on a sensor node or landmark by applying signal perturbations to individual landmarks, multiple landmarks, and all landmarks. We found that the localization error scaled similarly for all algorithms under attack, except for the Bayesian Networks algorithm. Large localization errors are introduced under severe attacks, resulting in 20–30 feet location perturbation under an attack strength of 15 dB. Further, we found that, when attacked, area-based algorithms did not experience a degradation in precision although they experienced degradation in accuracy and more uncertainty in location estimation. One important observation is that Bayesian Networks are more robust under both an individual landmark attack when positioning a single device, as well as an all-landmark attack when localizing multiple devices simultaneously.

We then examined the variability of the localization results under attack by measuring the Hölder metrics. We found that most algorithms had similar average variability, but those methods that returned the average of a set of most likely positions exhibited less variability. This result suggests that the average susceptibility of the returned results to an attack is essentially identical across point-based and area-based algorithms, though it might be desirable to employ either area-based methods or point-based methods that perform averaging in order to lessen the worst-case effect of a potential attack. Additionally this investigation indicates that the performance of most of the RSS-based localization algorithms degrades significantly under signal strength attacks, and consequently that network designers need to resort to more complicated

secure localization algorithms for dealing with potential attacks in an uncontrolled environment.

#### ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their helpful comments.

#### REFERENCES

- BAHL, P. AND PADMANABHAN, V. N. 2000. Radar: An in-building rf-based user location and tracking system. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*. 775–784.
- BATTITI, R., BRUNATO, M., AND VILLANI, A. 2002. Statistical learning theory for location fingerprinting in wireless LANs. Tech. Rep. DIT-02-086, University of Trento, Informatica e Telecomunicazioni.
- BRANDS, S. AND CHAUM, D. 1994. Distance-bounding protocols. In *Proceedings of the Workshop on the Theory and Application of Cryptographic Techniques on Advances in Cryptology*. 344–359.
- CAPKUN, S. AND HUBAUX, J. 2006. Securing localization with hidden and mobile base stations. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*.
- CAPKUN, S. AND HUBAUX, J. P. 2005. Secure positioning of wireless devices with application to sensor networks. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*. 1917–1928.
- CHEN, Y., FRANCISCO, J., TRAPPE, W., AND MARTIN, R. P. 2006a. A practical approach to landmark deployment for indoor localization. In *Proceedings of the Third Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*.
- CHEN, Y., KLEISOURIS, K., LI, X., TRAPPE, W., AND MARTIN, R. P. 2006b. The robustness of localization algorithms to signal strength attacks: a comparative study. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems (DCOSS)*. 546–563.
- DOHERTY, L., PISTER, K. S. J., AND ELGHAOUI, L. 2001. Convex position estimation in wireless sensor networks. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*. 1655–1663.
- ELNAHRAWY, E., LI, X., AND MARTIN, R. P. 2004. The limits of localization using signal strength: A comparative study. In *Proceedings of the First IEEE International Conference on Sensor and Ad hoc Communications and Networks (SECON)*. 406–414.
- ENGE, P. AND MISRA, P. 2001. *Global Positioning System: Signals, Measurements and Performance*. Ganga-Jamuna Press.
- HAZAS, M. AND WARD, A. 2003. A high performance privacy-oriented location system. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications (PerCom)*.
- HIGHTOWER, J., BORRIELLO, G., AND WANT, R. 2000. Spoton: An indoor 3d location sensing technology based on rf signal strength. Tech. Rep. 00-02-02, University of Washington, Department of Computer Science and Engineering.
- HU, Y., PERRIG, A., AND JOHNSON, D. 2003. Packet leashes: a defense against wormhole attacks in wireless networks. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*. 1976–1986.
- LANG, S. 1993. *Real and Functional Analysis*. Springer.
- LAZOS, L., POOVENDRAN, R., AND CAPKUN, S. 2005. Rope: robust position estimation in wireless sensor networks. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*. 324–331.
- LI, Z., TRAPPE, W., ZHANG, Y., AND NATH, B. 2005. Robust statistical methods for securing wireless localization in sensor networks. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*. 91–98.
- LIM, H., KUNG, L., HOU, J., AND LUO, H. 2006. Zero-configuration, robust indoor localization: Theory and experimentation. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*.

- LIU, D., NING, P., AND DU, W. 2005. Attack-resistant location estimation in sensor networks. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*. 99–106.
- MADIGAN, D., ELNAHRAWY, E., MARTIN, R., JU, W., KRISHNAN, P., AND KRISHNAKUMAR, A. S. 2005. Bayesian indoor positioning systems. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*. 324–331.
- NICULESCU, D. AND NATH, B. 2001. Ad hoc positioning system (APS). In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*. 2926–2931.
- PRIYANTHA, N., CHAKRABORTY, A., AND BALAKRISHNAN, H. 2000. The cricket location-support system. In *Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom)*. 32–43.
- ROOS, T., MYLLYMAKI, P., AND TIRRI, H. 2002. A statistical modeling approach to location estimation. *IEEE Trans. Mobile Comput.* 1, 1 (Jan–March), 59–69.
- SASTRY, N., SHANKAR, U., AND WAGNER, D. 2003. Secure verification of location claims. In *Proceedings of the ACM Workshop on Wireless Security*. 1–10.
- SAVVIDES, A., HAN, C.-C., AND SRIVASTAVA, M. 2001. Dynamic fine-grained localization in ad hoc networks of sensors. In *Proceedings of the Seventh Annual ACM International Conference on Mobile Computing and Networking (MobiCom)*.
- SHANG, Y., RUMML, W., ZHANG, Y., AND FROMHERZ, M. P. J. 2003. Localization from mere connectivity. In *Proceedings of the Fourth ACM International Symposium on Mobile Ad-Hoc Networking and Computing (MobiHoc)*. 201–212.
- TAO, P. AND RUDYS, A., AND LADD, A. AND WALLACH, D. 2003. Wireless lan location-sensing for security applications. In *Proceedings of the Second ACM Workshop on Wireless Security (WiSe)*.
- WILSON, R. 2002. Propagation loss through common building materials, 2.4GHz vs. 5GHz. *White paper available at <http://www.magisnetworks.com>*.
- YOUSSEF, M., AGRAWAL, A., AND SHANKAR, A. U. 2003. WLAN location determination via clustering and probability distributions. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 143–150.

Received September 2006; revised May 2007, December 2007; accepted December 2007