The Directional Attack on Wireless Localization

– or –

How to Spoof Your Location with a Tin Can

Kevin Bauer Damon McCoy Eric Anderson Markus Breitenbach
Greg Grudic Dirk Grunwald Douglas Sicker
University of Colorado, Boulder, CO 80309-0430 USA
{bauerk, mccoyd, andersoe, breitenm, grudic, grunwald, sicker}@colorado.edu

***Abstract*—802.11 localization algorithms provide the ability to accurately position and track wireless clients thereby enabling location-based services and applications. However, we show that these localization techniques are vulnerable to non-cryptographic attacks where an adversary uses a low-cost directional antenna to appear from the localization algorithm's perspective to be in another arbitrary location of their choosing. The attacker's ability to actively influence where they are positioned is a key distinguishing feature of the directional attack relative to prior localization attacks that use transmit power control to introduce localization errors. We implement a representative set of received signal strength-based localization algorithms and evaluate the attack in a real office building environment. To mitigate the attack's effectiveness, we develop and evaluate an attack detection scheme that offers a high detection rate with few false positives.**

I. Introduction

Systems that use existing wireless infrastructure to locate and track wireless clients are becoming ubiquitous and have great potential to offer a wide variety of location-aware services. In particular, security services such as rogue device detection [1] and location-based access control [2] are emerging applications for localization systems. As such security services become widely deployed, malicious users will have increased incentive to spoof their locations with sufficient precision to remain hidden or gain unauthorized access to network resources.

To this end, we propose a physical layer localization attack using inexpensive directional antennas. The key distinguishing feature of the directional attack relative to prior localization attacks that use transmit power control to introduce localization errors [3]–[5] is the attacker's ability to actively influence where they are positioned. We call the attacker's ability to appear reliably at a location of their choice teleportation effects. We demonstrate the efficacy of this attack using an extremely low-cost directional antenna built from a tin can, commonly called a "cantenna."

Since most commercial localization systems use the received signal strength (RSS) property of the wireless signals and one of a number of localization algorithms to calculate a location estimate [6]–[12], we limit our evaluation to RSS-based algorithms. While other localization techniques have been proposed using a variety of techniques including angle of arrival (AOA) [13], [14], time of arrival (TOA) [15], and time difference of arrival (TDOA) [16], these systems have not seen the same level of deployment due to the requirement that additional (and often non-commodity) infrastructure must be installed to collect the necessary information to perform the localization.

We implement a representative set of RSS-based localization algorithms and conduct experiments in a real office building environment. The results show that the directional attack enables an adversary to actively influence their location prediction by focusing the directional antenna toward a desired position. In addition, our results indicate that the directional attack is capable of producing significant localization errors that are greater than can be achieved using transmit power control alone. Our results also indicate that even some secure RSS-based localization techniques are vulnerable to the directional attack.

We measure the localization errors and teleportation effects that are produced under the directional attack and also explore how adding transmit power control influences the attack's effectiveness. We find that an attacker who points the directional antenna down long corridors has the ability to appear in the direction of their choice more than 75% of the time. In addition, the results show that these attacks can produce expected localization errors in excess of 18 meters. This is an increase in error of approximately 300% over the expected localization error without attack, and an increase of 200% over previous transmit power control attacks.

Having demonstrated and evaluated the directional attack, we present an attack detection technique based on our empirical data. We show that it achieves a high attack detection rate of over 90% against both directional and transmit power control attacks while maintaining a low false positive rate (at most 10%).

Contributions. This work contributes the following:

1) We propose a physical layer attack against RSS-based 802.11 localization techniques that utilizes an inexpensive directional antenna and gives an attacker significant control over where their position is estimated.
2) We evaluate the directional attack by implementing localization algorithms based upon the k-nearest neighbors and Naïve Bayes classifiers and conducting experiments in a real office building environment.
3) We present and empirically validate a technique to identify directional attackers that offers a high detection rate with few false positives.

II. Background

In this section, we provide a brief introduction to the problem of wireless device localization. We next describe the

localization algorithms used to evaluate the directional attack strategy.

A. RSS-based Localization

Predicting a wireless device's physical location in an indoor environment has been accomplished using techniques based on received signal strength (RSS) [6]–[12], angle of arrival (AoA) [13], [14], time of arrival (ToA) [15], and time difference of arrival (TDoA) [16]. In this paper, we consider only localization techniques that are based on RSS, as these can be constructed with commodity 802.11 hardware and stock drivers.

RSS-based localization refers to the task of estimating an 802.11 device's physical location using only signal strength information. Due to the inherently noisy nature of the RSS measurements, RSS-based localization algorithms typically apply statistical/machine learning techniques, and proceed in two phases:

1) An offline *training phase* is conducted in which several received signal strength indication (RSSI) readings $\vec{t_i} = (r_{i1}, ..., r_{in})$ are collected over a set of n passive receivers and are labeled with the transmitter's true physical location and orientation $p_i = (x_i, y_i, \theta_i)$.
2) During the online *localization phase*, the observed RSSI readings $\vec{o_j} = (r_{j1}, ..., r_{jn})$ are used to produce the device's estimated location $\hat{p_j} = (\hat{x_j}, \hat{y_j})$.

B. k-Nearest Neighbors Localization

RADAR uses the k-nearest neighbors learning algorithm to provide location estimates with minimal localization errors [6]. Using Euclidean distance, the k closest training points to the observed RSS vector in n-dimensional vector space of all signal strength measurements are chosen.[1] The mean of their physical coordinates is computed to produce a location estimate.

In addition to the traditional k-nearest neighbors algorithm, we consider a variant that minimizes the median of the distances in each dimension: $median_{i=1}^{n}(r_i - r_i^{'})^2$. Prior work has suggested that this approach is more resilient to transmit power attacks [17].

C. Naïve Bayes Localization

Localization techniques that use the Naïve Bayes classifier have been proposed in [3], [9], [12], [18]. This approach is based on the application of Bayes' theorem to obtain a position estimate. Using Bayes' theorem, the conditional probability of observing a signal strength vector from the training data at a particular position is computed. During the localization phase, the position estimate is the position that maximizes this probability for the observed signal strength vector.

We also consider a variant of Naïve Bayes called the *Difference Method*, where the mean signal strength at each monitor is computed for a short observation period and the pairwise difference in average signal strengths over the observation period for each monitor is used to train a Bayesian classifier. Prior work has suggested that the difference method is more resilient to signal strength attacks during the localization phase [3].

[1]The specific value of k that produced minimal localization error on non-training data in our experimental deployment was $k = 10$.

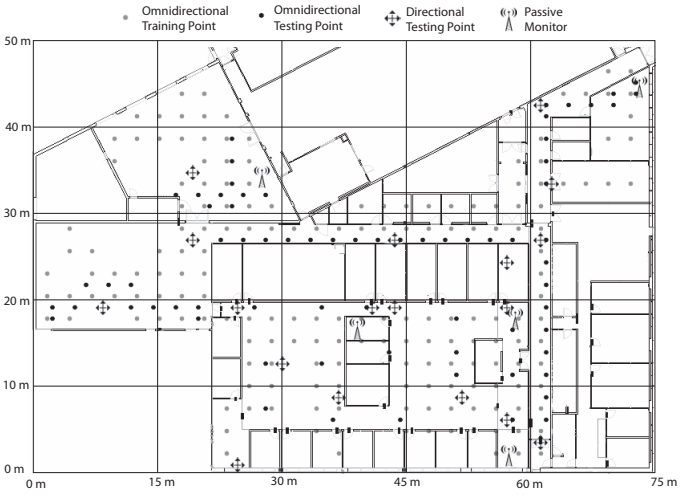

Fig. 1. The layout of the experimental localization deployment

III. Attack Model

Our attack model assumes that there is an indoor 802.11 network and a localization system that attempts to physically localize wireless clients within signal range of the network without any cooperation from the clients. The localization system is composed of passive monitors that collect RSSI information from locations that are different from the monitors' locations.

The adversary's primary goal is to produce significant localization errors focused in a specific direction of their choosing. Their secondary objective is to simply introduce a large amount of error into the system's location prediction, such that the adversary's wireless device appears to be significantly far away from their true location. The adversary utilizes a single commodity 802.11 wireless card and a low-cost directional antenna, such as a "cantenna," constructed from a discarded tin can.

We finally assume that the attacker has no prior knowledge about the placement of the passive monitors, the localization algorithm being used, or the location of the training points.

IV. Experimental Setup and Data Collection

In this section, we describe the experimental process we used to demonstrate and evaluate the effectiveness of the directional attack. We next present the results of our experiments. As the first step toward understanding the consequences of an adversary using a directional antenna, we examine how directional antennas effect localization errors. Here, *localization error* is defined as the Euclidean distance between the user's true physical location and the algorithm's location estimate. We next evaluate the extent to which an adversary can influence their location prediction.

The site for our experiments is a standard office space measuring 75 x 50 meters^2. This indoor environment is sufficiently large and diverse (*i.e.*, it has long hallways, large warehouse-like rooms, and small offices) to provide insight into the attack's behavior in general indoor environments. We deployed five passive monitors with omni-directional antennas throughout the building which recorded the 802.11 traffic within range of them.

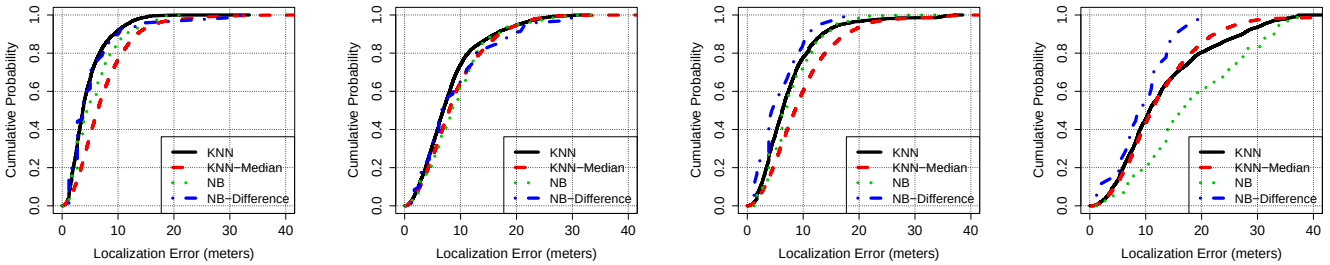

(a) Non-attack data (b) Attenuation attack data (c) Directional attack data (d) Directional and attenuation attack combination data

Fig. 2. Localization error CDFs

To train the RSS-based localization algorithms, we collect packets at 179 training points chosen to provide uniform coverage of the building. For each training point, we broadcast 500 packets using an omni-directional antenna facing each of the four cardinal directions (since the human body tends to attenuate the signal). The packets were generated by a single wireless device transmitting at a constant power level of 16 dBm. These points are tagged with the correct location of the transmitting device.

The adversary's device transmitted packets from 70 different omni-directional testing points, varying the transmission power between 10 dBm and 20 dBm. In addition, the adversary transmitted packets at 18 different testing points using a directional antenna pointed in the four cardinal directions. The cantenna used in these experiments has a 12 dBi gain and a 30 degree half-power beam width.

We assume that the training data is collected by the network operators and localization phase data is generated by an attacker. We further assume that these data sets are physically disjoint, since the attacker has no prior knowledge of where the training points are located. Figure 1 shows a floor plan of the experimental site labeled with monitors, testing, and training points.[2]

V. Experimental Results

In order to establish the accuracy of the localization algorithms in our experimental environment, we first demonstrate that the k-nearest neighbors algorithms (abbreviated KNN and KNN-MED) and Naïve Bayesian techniques (abbreviated NB and NB-DIFF) produce low localization error on non-attack data. We next evaluate the effect that transmit power manipulation has on localization accuracy. We next demonstrate the localization error that can be produced when an adversary uses a directional antenna, and we also characterize the localization error when the attacker's transmit power is varied in combination with a directional antenna. [3] Finally, we evaluate the extent to which an attacker can influence their location prediction.

[2]Our data is available as part of the CRAWDAD wireless trace repository [19]: http://crawdad.cs.dartmouth.edu/cu/rssi.

[3]We compare these attacks using different localization algorithms for the sole purpose of establishing that the attacks have a similar effect on characteristic RSS-based algorithms, not to argue that one algorithm performs better than any other.

A. Localization Without Attack

In order to establish the validity and expected accuracy of the localization algorithms, we first demonstrate that KNN and both Naïve Bayes methods provide reliable room-level localization with relatively low localization error. Using the same omni-directional antenna transmitting at 16 dBm as in the training phase, KNN and both Naïve Bayesian algorithms provided accurate location predictions. Using KNN, the median localization error was 3.6 meters. NB provided location predictions with a median error of 4.3 meters and the NB-DIFF method localized with a median error of 3.7 meters. These localization errors are consistent with prior work [3], [6].[4] The localization error cumulative distribution function (CDF) is provided in Figure 2(a).

Note that the baseline localization error for KNN-MED is significantly higher than the other algorithms. At the median, KNN-MED produced a localization error of 6.3 meters. This is a result of the relatively low number of listening monitors ($n = 5$) that are used in our experimental deployment. Since this algorithm minimizes the median of the distances in all dimensions, when the number of dimensions (*i.e.*, the number of passive listeners n) is small, the median provides a poor distance metric. In fact, Li *et al.* [17] note that additional listening monitors should be deployed in order for this technique to be effective. Unfortunately, this approach may be impractical due to its requirement for many additional monitors.

B. Transmit Power Attacks

The first attack that we consider is where an adversary manipulates their transmit power level in an attempt to introduce additional localization error. Using an omni-directional antenna, the wireless card's power level varies from 10 to 20 dBm. Since training was conducted at a constant 16 dBm transmit power level, we call the scenario in which the adversary decreases the transmit power below 16 dBm an *attenuation* attack. Transmit powers greater than 16 dBm are referred to as *amplification* attacks.

We found that amplification introduced no additional localization errors over those observed without attack (this is

[4]Since we assume a continuous (not discrete) position space, our localization errors are naturally higher for the Bayesian methods, since it requires that the predicted location be a position from the training data.

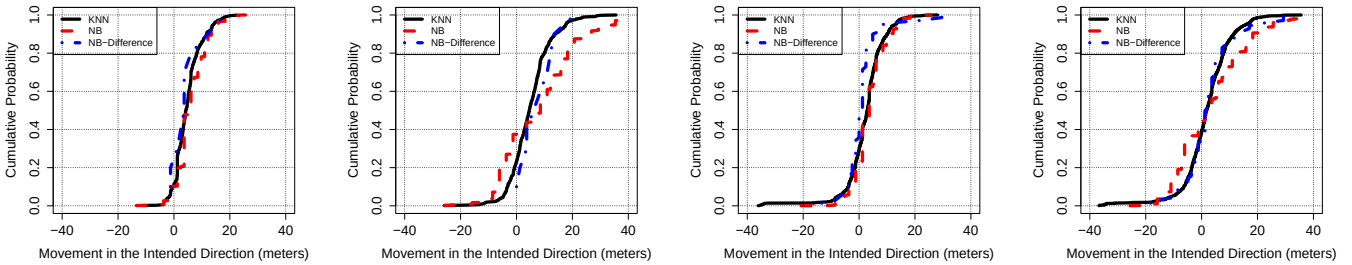

(a) Intelligent directional attacker (b) Intelligent directional attacker with transmit power control (c) Naïve directional attacker (d) Naïve directional attacker with transmit power control

Fig. 3. CDFs of the movement in the general direction in which the directional antenna is focused (Negative values indicate movement in the *opposite* direction of the directional antenna's focus)

consistent with prior work [4]). However, the attenuation effect introduced significant localization errors. This may be because transmit power attenuation reduces the number of sensors that can hear the signal. Transmitting at 10 dBm, at the median KNN produced errors of 6.8 meters and KNN-MED produced errors of 7.9 meters. NB showed errors of 8.1 meters and NB-DIFF had 6.7 meters of localization error. The localization error CDF is given in Figure 2(b). Each algorithm produced a localization error nearly double that which was observed on non-attack data. Transmit power attenuation is a simple technique for an adversary to negatively influence the localization process; however, attenuation alone is not a sufficient tool for allowing an adversary to appear in a direction of their choice.

C. Directional Attacks

We next examine the effect of an attacker with a directional antenna on overall localization error. Testing points are collected using 16 dBm transmit power and a directional antenna. The directional attacker does not intelligently point the directional antenna through hallways or open spaces, but rather points the antenna in all four cardinal directional at each testing point. We call this case the *naïve directional attack*. For KNN, the median localization error was 6.2 meters and KNN-MED produced a median error of 8.5 meters. NB produced a median error of 6.2 meters and NB-DIFF gave 4.7 meters of error at the median. The error CDF for the naïve directional attack is given in Figure 2(c). In general, this attack performs similarly to the attenuation attack.

D. Transmit Power and Directional Attacks

To maximize the expected localization error, an adversary may combine the effects of attenuation attacks with a directional antenna. Here, we still consider the case of the naïve directional attacker. On a localization system based on KNN, this attacker is able to produce an error of 10.6 meters at the median and 11.0 meters for KNN-MED. The NB and NB-DIFF methods produce 16.9 and 9.5 meters of error at the median, respectively. The error CDF for the attenuation and directional attack combination is given in Figure 2(d).

E. Attacker's Ability to Control Position Estimates

We next consider an adversary that uses a directional antenna intelligently, *i.e.*, pointing it down long hallways and through large open spaces. Here, we select directional testing points only where the attacker has a clear line-of-sight view of at least 10 meters before any physical obstruction (such as walls, doors, etc.). We do not consider the directional effects introduced in the KNN-MED algorithm since this technique failed to produce sufficiently low localization errors on non-attack data; thus, we consider the directional effects introduced when the KNN, NB, and NB-DIFF algorithms are used.

Figure 3(a) shows a CDF of movement in the desired direction for the intelligent directional attacker. Intelligently pointing the directional antenna allows the attacker to have more control over their estimated location. In the worst case against the difference method, the smart attacker could move 3 meters in the desired direction half the time and 6 meters in the desired direction 20% of the time. For comparison, Figure 3(c) shows that movement in the intended direction is lower when the directional antenna is pointed arbitrarily.

Figure 3(b) shows the CDF of movement in the attacker's desired direction when an attenuation attack is applied in combination with the directional attack. Since the majority of position estimates exist in the direction of the antenna's focus, the combination attack gives an adversary the ability to reliably teleport in a direction of their choice. For comparison, Figure 3(d) shows that the movement in the direction of the adversary's choice is lower when the directional antenna is pointed arbitrarily.

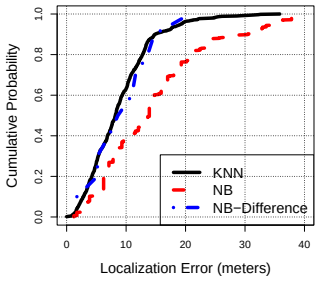

Fig. 4. Localization error for an intelligent attacker with 10 dBm transmit power

Figure 4 shows the localization error CDF for the intelligent directional attacker who attenuated their transmit power. The errors for KNN and the difference method are about the same as for naïve attacker from Section V-D; however, the smart attacker achieves far less localization error from NB. This shows that it is not necessarily the case that using the directional antenna intelligently results in additional overall localization error. However, intelligently focusing the directional antenna

TABLE I
OBSERVED MEDIAN VARIANCE IN RSS

Non-attack	Directional	Attenuation & Directional
4.87 dB	19.67 dB	16.88 dB

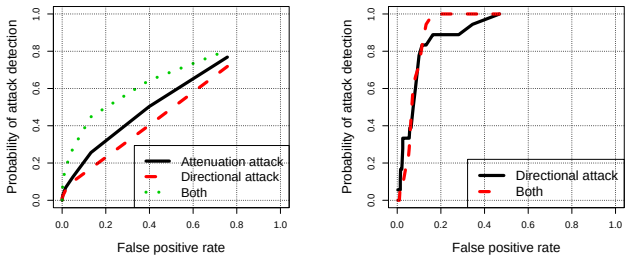

(a) Signal space distance threshold ROC curve for the attenuation attack, directional attack, and combination attack data.

(b) Variance threshold ROC curve for directional and combination attack data sets.

Fig. 5. ROC curves for the distance and variance-based test statistics as τ varies.

in combination with transmit power does offer the ability to reliably appear in a direction of the attacker's choice.

VI. ATTACK DETECTION

Having demonstrated the threat of the directional attack, we next turn our attention to solutions. We first consider attack detection techniques that are agnostic to the localization algorithm applied. These detection techniques work on raw RSS data. We first attempt attack detection by examining the distance between the expected RSS surface derived from training data and an attack RSS vector. Having pointed out the limitations of this method, we propose an improved attack detection technique that is especially tailored to directional attacks.

A. Prior Attack Detection Scheme: Signal Space Distance

Chen *et al.* [4] proposed a localization attack detection framework to detect transmit power manipulation attacks. In this work, the attack detection problem is formulated as statistical significance testing. The detection technique works by first forming an n-dimensional RSS surface S using the training data. Next, during the validation phase an RSS vector s' is observed and the minimum Euclidean distance from this observed vector to the training surface is calculated as follows: $Dist_S = \min\{||s' - s^j|| : \forall s^j \in S\}$. A threshold value τ is derived from training data and the feature vector is classified as an attack if $Dist_S > \tau$.

We applied this detection technique to our attack data sets, but without success. The receiver operating characteristic (ROC) curve — shown in Figure 5(a) — illustrates the relationship between the probability of attack detection and the false positive rate. While the probability of detection becomes sufficiently high (nearly 0.8 for each attack data set), the false positive rate is unacceptably high, approaching 0.8 for each attack data set. This detection scheme produces results that are not significantly different from uniformly random guesses.

The most plausible explanation for the high false positive rate is that Chen *et al.* conduct trace-driven experiments and adjust the transmit power levels significantly more — up to 25 dB — than we consider in our experiments. We assume an attacker that adjusts their transmit power level by *at most* 6 dB. Thus, it is not surprising that this detection technique performs poorly on our data. However, it is important to derive an attack detection scheme that is more reliable under this attack, since we show that even minimal transmit power manipulation in combination with directional antennas can produce significant localization errors.

B. Directional Attack Detection: Minimize RSS Variance

We adopt a similar framework for attack detection as Chen *et al.* [4], but we consider a new test statistic. We assume that the adversary uses a directional antenna as described in Section III. Further, suppose that the attacker's goal is to conceal their true location by attempting to appear in several arbitrary locations. To achieve this effect, the attacker points the directional antenna in all four cardinal directions arbitrarily as their packets are transmitted.

To detect such a directional attack, we first observe that the RSS values across all passive monitors have higher variances than in the non-attack case. Table I shows the expected median in variance for non-attack, directional attack, and directional/attenuation attack data. The attack data sets exhibit more variance on average than the non-attack data set. We leverage this observation in the design of an attack detection scheme.

The attack detection works as follows: For each observation window, the localization system records the RSS of all packets received from each source (identified by a MAC address). For an observation window j, the mean and variance $(\mu_{ij}, \sigma^2_{ij})$ of the RSS values observed is calculated for each monitor $i = 1, ..., n$. We take the median of the RSS variances across all n monitors for the observation window:

$$med_j = \textit{median}^n_{i=1}(\sigma^2_{ij})$$

This single metric captures the degree of variability in the signal strengths during the observation window. We use the median to mitigate the influence of outliers. To determine if an attack is occurring, we derive a variance threshold for detection from our training data τ. If $med_j > \tau$, then the data is classified as an attacks.[5]

The ROC curve for this attack detection method is given in Figure 5(b). Note that this method offers both higher detection rates and fewer false positives than the signal space distance test statistic. This variance-based detection scheme is a viable strategy to detect an attacker who arbitrarily focuses the directional antenna to produce large localization errors because the resulting RSS values tend to vary differently for each orientation of the directional antenna. Thus, such an attack produces greater variance in the observed RSS values over short time periods in comparison to the non-attack case, which tends to produce more stable RSS values with low variances.

[5]It is possible that an attacker could use one MAC address to prepare the attack and another address to launch the attack. More intelligent device fingerprinting methods [20] could be applied to identify the device in this case.

C. Toward Detecting the General Case of the Attack

Having demonstrated how to detect an attacker who points the directional antenna in arbitrary directions, we now focus on detecting the attack in the general case. Suppose that an adversary wishes to appear at a location distant from their true position. While detecting this form of the attack is difficult, we propose using data smoothing techniques to mitigate its effects. This can be done in combination with the distance-based and variance-based detection schemes presented in Sections VI-A and VI-B.

One promising strategy is to compute the location estimates and then filter them through a smoothing mechanism, such as a single exponential smoothing function. Such an iterative smoothing function is defined as follows:

$$\hat{s_i} = \alpha \hat{p}_{i-1} + (1-\alpha)\hat{s}_{i-1}, \text{ for } \alpha \in (0,1]$$

where $\hat{p}_{i-1}$ is the most recent location prediction, α is a smoothing constant that influences how fast the smoothed values change, and $\hat{s}_{i-1}$ is the previous smoothed location prediction. This dampens the location estimate's movement and mitigates the effects of outliers in the location predictions. The impact of smoothing filters on the accuracy of location estimates is left to future work.

VII. Related Work

Previous work has studied the performance of RSS-based localization algorithms in various adversarial scenarios. Chen *et al.* [21] show that physical materials such as foil and even the human body can be used to implement an attenuation attack. Through trace-driven experiments, they evaluate the robustness of a variety of RSS-based localization algorithms under this attack and observe that performance generally degrades with the severity of the attack. We confirm this observation in our study of transmit power attacks in Section IV.

Tao *et al.* [3] study RSS-localization when the assumptions made during the training phase are violated. This includes introducing variation in transmit power level and changing the card itself. A variant of the naïve Bayes classifier is proposed and we confirm its robustness to transmit power control attacks relative to other localization algorithms.

Our work proposes a variation of the transmit power attacks whereby an adversary uses an inexpensive directional antenna to manipulate signal strength properties. While the directional attack can produce significant localization errors, it is the only proposed attack to date that allows an adversary to actively influence *where* they appear to be from the perspective of the localization algorithm.

VIII. Conclusion

We experimentally demonstrated that RSS-based localization techniques are vulnerable to significant localization errors introduced by an adversary with a low-cost directional antenna. In our evaluation, we implement several RSS-based localization algorithms and conduct experiments in an office space environment. In particular, we show that directional attackers not only have the ability to introduce significant localization errors, but also have the ability to reliably *teleport* in a direction of their choice. To address the directional attack, we propose a new test statistic for attack detection based on an analysis of the variance in RSS values. Since RSS-based localization systems rely on training data and ranging, it is unlikely that one can be built to accurately locate a directional attacker.

IX. Acknowledgements

We thank James Martin for granting access to our office building test-bed. We also thank the anonymous reviewers for their helpful comments and suggestions. This work was partially funded by NSF awards ITR-0430593 and CRI-0454404.

References

[1] D. B. Faria and D. R. Cheriton, "Detecting identity-based attacks in wireless networks using signalprints," in *WiSe '06: Proceedings of the 5th ACM workshop on Wireless security*. New York, NY, USA: ACM Press, 2006, pp. 43–52.

[2] I. Ray and L. Yu, "Short paper: Towards a location-aware role-based access control model," in *SECURECOMM '05: Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 234–236.

[3] P. Tao, A. Rudys, A. M. Ladd, and D. S. Wallach, "Wireless LAN location-sensing for security applications," in *WiSe*, 2003.

[4] Y. Chen, W. Trappe, and R. P. Martin, "Attack detection in wireless localization," in *INFOCOM*. IEEE, 2007, pp. 1964–1972.

[5] Y. Chen, K. Kleisouris, X. Li, W. Trappe, and R. P. Martin, "The robustness of localization algorithms to signal strength attacks: A comparative study," in *DCOSS*, 2006, pp. 546–563.

[6] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *INFOCOM*, 2000, pp. 775–784.

[7] D. Madigan, E. Elnahrawy, R. P. Martin, W.-H. Ju, P. Krishnan, and A. Krishnakumar, "Bayesian indoor positioning systems." in *INFOCOM*. IEEE, 2005, pp. 1217–1227.

[8] J. Ash and L. Potter, "Sensor network localization via received signal strength measurements with directional antennas," *Conference on Communication, Control, and Computing*, pp. 1861–1870, Sept 2004.

[9] A. M. Ladd, K. E. Bekris, A. Rudys, L. E. Kavraki, and D. S. Wallach, "Robotics-based location sensing using wireless ethernet," in *MobiCom*, 2002, pp. 227–238.

[10] T. King, S. Kopf, T. Haenselmann, C. Lubberger, and W. Effelsberg, "Compass: A probabilistic indoor positioning system based on 802.11 and digital compasses," in *WiNTECH*, 2006, pp. 34–40.

[11] M. A. Youssef, A. Agrawala, and A. U. Shankar, "Wlan location determination via clustering and probability distributions," in *PERCOM*, Washington, DC, USA, 2003.

[12] E. Elnahrawy, X. Li, and R. P. Martin, "The limits of localization using signal strength: A comparative study," in *SECON*. IEEE, 2004.

[13] E. Elnahrawy, J. Austen-Francisco, and R. P. Martin, "Adding angle of arrival modality to basic RSS location management techniques," in *ISWPC*, 2007, pp. 464–469.

[14] D. Niculescu and B. Nath, "VOR base stations for indoor 802.11 positioning," in *MobiCom*, 2004, pp. 58–69.

[15] Y.-T. Chan, W.-Y. Tsui, H.-C. So, and P. chung Ching, "Time-of-arrival based localization under NLOS conditions," *IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY*, vol. 55, no. 1, January 2006.

[16] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *MOBICOM*, 2000.

[17] Z. Li, W. Trappe, Y. Zhang, and B. Nath, "Robust statistical methods for securing wireless localization in sensor networks," in *IPSN*, 2005.

[18] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki, "Practical robust localization over large-scale 802.11 wireless networks," in *MobiCom*, 2004, pp. 70–84.

[19] "CRAWDAD: A community resource for archiving wireless data at Dartmouth," http://crawdad.cs.dartmouth.edu.

[20] L. C. C. Desmond, C. C. Yuan, T. C. Pheng, and R. S. Lee, "Identifying unique devices through wireless fingerprinting," in *Proceedings of the first ACM conference on wireless network security*, 2008, pp. 46–55.

[21] Y. Chen, K. Kleisouris, X. Li, W. Trappe, and R. P. Martin, "A security and robustness performance analysis of localization algorithms to signal strength attacks," *ACM Trans. Sen. Netw.*, vol. 5, no. 1, pp. 1–37, 2009.