SecureAngle: Improving Wireless Security Using Angle-of-Arrival Information

Jie Xiong
University College London
j.xiong@cs.ucl.ac.uk

Kyle Jamieson
University College London
k.jamieson@cs.ucl.ac.uk

ABSTRACT

Wireless networks play an important role in our everyday lives, at the workplace and at home. However, they are also relatively vulnerable: physically located off site, attackers can circumvent wireless security protocols such as WEP, WPA, and even to some extent WPA2, presenting a security risk to the entire network. To address this problem, we propose SecureAngle, a system designed to operate alongside existing wireless security protocols, adding defense in depth. SecureAngle leverages multi-antenna APs to profile the directions at which a client's signal arrives, using this angle-of-arrival (AoA) information to construct signatures that uniquely identify each client. We identify SecureAngle's role of providing a fine-grained location service in a multi-path indoor environment. With this location information, we investigate how an AP might create a "virtual fence" that drops frames received from clients physically located outside a building or office. With SecureAngle signatures, we also identify how an AP can prevent malicious parties from spoofing the link-layer address of legitimate clients. We discuss how SecureAngle might aid whitespace radios in yielding to incumbent transmitters, as well as its role in directional downlink transmissions with uplink AoA information.

Categories and Subject Descriptors

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

General Terms

Design, experimentation, measurement, security

Keywords

Wireless, 802.11, SecureAngle, Angle of arrival

Hotnets '10, October 20–21, 2010, Monterey, CA, USA.

1. INTRODUCTION

In the past few years, wireless data connectivity has continued its transition into an essential utility. Enterprises, city governments, and homes have rolled out wireless local-area networks across buildings, campuses, and even entire cities. Cellular phones are ubiquitous, with the steady progression in cellular standards driving data rates increasingly higher to meet an increasing demand. By these measures, wireless networks have been a staggering success.

However, from a security perspective, wireless access points present a number of difficulties. To see why, consider a wireless network in the local area. Once an attacker has compromised an access point, she may both eavesdrop on users' traffic and inject traffic into the wireless network. Security protocols such as WEP, WPA, LEAP, and WPA2 (IEEE 802.11i) have been proposed in the past few years, however they have a track record of being compromised [16, 13, 4]. Moreover, once vulnerabilities are discovered, they are slow to be fixed: six years after WEP was known to be insecure, Bittau *et al.* reported that a staggering 76% of secured APs in London still used it [5].

The result is an ongoing competition between new exploits and better wireless security protocols. To fundamentally change this status quo, we propose SecureAngle, an approach to wireless security designed to provide defense in depth, operating alongside and strengthening protocol-based wireless security measures.

The other recent trend in the design of both local- and metropolitan-area wireless access points is the dramatically increasing number of antennas at the AP, mainly to bolster capacity and coverage with multiple-input, multiple-output (MIMO) techniques. IEEE 802.11n and Long Term Evolution (LTE, also known as 4G) are the most recent standards in local-area and cellular networks, respectively, and both exploit MIMO extensively through the use of many antennas at the access point. We believe that in the future, the number of antennas at the access point will increase several-fold, to meet the demand for MIMO links and spatial division multiplexing [15].

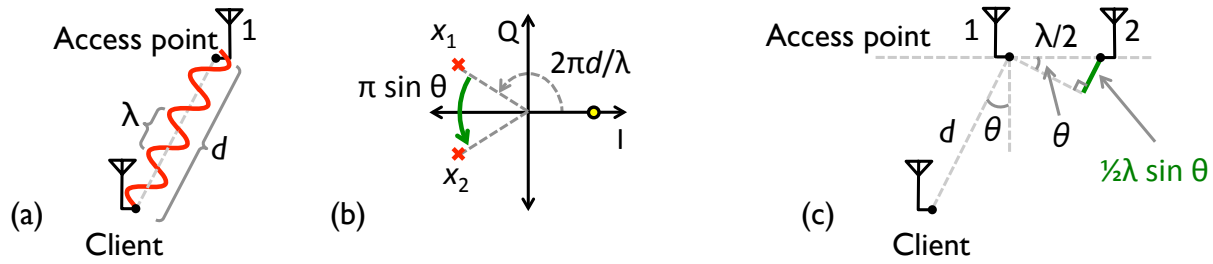

Figure 1: SecureAngle's principle of operation: (a) The phase of the signal goes through a 2π cycle every radio wavelength λ. (b) The complex representation of the sent (filled dot) and received (crosses) signals at both antennas in (c). Phase corresponds to angle measured from the positive in-phase (I) axis. (c) A signal arriving at bearing θ to two antennas.

The key observation we make in this paper is that with the right signal processing at the physical layer and a simple, environment-independent calibration step described in §2.2, we can add a layer of security to an LTE, 802.11n, or other multi-antenna access point. We accomplish this by indirectly measuring the distances between an incoming signal's arrival at each antenna, using these measurements to identify the incoming signal's *angles-of-arrival*, as we describe in §2.1. SecureAngle can then use this AoA information to construct an *AoA signature* that is unique to each client and extremely difficult for an attacker to forge. Even if an attacker has the location information of both client and AP, it's not easy for the attacker to forge the direct path AoA and even more difficult for the attacker to forge all the multipath AoAs. Doing so would require the attacker to know the locations of all obstacles in the vicinity of the AP and client. The combined direct path and reflection path AoAs form the unique signature for each client. With direct path AoA data from two APs, fine indoor location information can also be obtained.

We show that AoA signatures have systems security value with two representative applications. We are not replacing any exiting security protocol. SecureAngle works along with them and adds another layer of security. In both applications, our threat model consists of an attacker equipped with an omnidirectional antenna, directional antenna (as the attackers were equipped in the TJ Maxx attacks of 2006 [8]), or antenna array, and who has successfully penetrated the protocol-based security in use at the access point, whatever that may be. Demonstrating useful applications under this rather strong attacker model substantiates our claim that SecureAngle provides defense in depth against wireless exploits; we now describe the two applications for which we believe SecureAngle can be of most immediate use.

Virtual fences. We investigate restriction of use to the building or room containing the access point. This would be appropriate, for example, in an enterprise setting where the company is contained in a secured building, and it is desired that only clients within the building be allowed wireless access. With direct path AoA information obtained from multiple SecureAngle APs, high-precision indoor location can be determined to enable this service.

Address spoofing prevention. Another application of AoA-based signatures is to aid systems that detect network anomalies and misconfigurations [9, 1], such as when a client spoofs the link-layer address of a legitimate stationary client. Link-layer address spoofing can grant unauthorized access, if the only method of wireless security is an address-based access control list, and spoofing often forms the basis of more sophisticated attacks against the security protocols we mention above. For two clients located at different locations, their AoA signatures are very unlikely to be the same. A fundamental challenge for this application is that AoA signatures change to some degree when obstacles in the environment or clients themselves move, and therefore must be tracked and updated. There also must be a significant difference between the certified signature and an attacker's signature so that they can be discriminated from each other.

We detail these two applications (§2), present experimental evidence for their feasibility (§3), discuss related work (§4), and conclude the paper with a discussion of other more future-looking applications of SecureAngle signatures (§5).

2. DESIGN

We now describe SecureAngle bottom-up, starting with a description from first principles of how we calculate client signatures based on their transmissions' angle of arrival at the access point, and culminating with a system-level sketch.

2.1 Angle of arrival signatures

In both indoor and outdoor wireless channels, a sender's signal reflects off objects in the environment, resulting in multiple copies of the signal arriving at the access point; this phenomenon is known as multipath. For clarity of exposition, we first describe how to compute

angle of arrival when there is just one path from transmitter to access point, then generalize the principles to handle multipath wireless propagation.

Phase of the wireless signal. The key to computing angle of arrival of a wireless signal is to analyze its phase, a quantity that progresses linearly from zero to 2π every radio wavelength λ along the path from client to access point, as shown in Figure 1(a). This means that the access point receives signals with an added phase determined by the path length d from the client. Phase is particularly easy to analyze because software-defined and hardware radios represent the phase of the wireless signal graphically using an *in-phase-quadrature* (I-Q) plot, as shown in Figure 1(b), where angle measured from the I axis indicates phase. Using the I-Q plot, we see that the distance d adds a phase of $2\pi d/\lambda$ as shown by the angle measured from the I axis to the cross labeled x_1 (representing the signal received at antenna one). Furthermore, at the access point with two antennas depicted in Figure 1(c), the distance along a path arriving at bearing θ is a fraction of a wavelength greater to the second antenna than it is to the first, the fraction depending on θ, the angle of arrival.

These facts suggest a particularly simple way to compute θ at a two-antenna access point in the absence of multipath. First, use a software-defined or hardware radio to measure x_1 and x_2 directly, compute the phase of each ($\angle x_1$ and $\angle x_2$), and then solve for θ ($\angle x_1 - \angle x_2$ is between $-\pi$ and π) as:

$$\theta = \arcsin\left(\frac{\angle x_2 - \angle x_1}{\pi}\right) \quad (1)$$

In real-world multipath environments, however, Equation 1 breaks down because multiple paths' signals sum in the I-Q plot, breaking the simple two-antenna exposition above. However, adding antennas can resolve the ambiguity. The best known AoA estimation algorithms are based on eigenstructure analysis of a *correlation matrix* formed by samplewise-multiplying the raw signal from the lth antenna with the raw signal from the mth antenna, then computing the mean of the result. The lth column and mth row of this $m \times l$ matrix is therefore the mean correlation between the lth and mth antennas' signals. The principles of these algorithms are the same as in the two antenna case, but the mathematics are more involved; we refer the interested reader to the MUSIC [12] algorithm for more information.

The output of such AoA estimation algorithms, shown below in §3, is a *pseudospectrum*: a continuous plot of likelihood versus angle. We use the pseudospectrum as our client signature.

2.2 Access point calibration

Equipping the access point with multiple antennas is necessary for SecureAngle, but does not suffice to calculate angle of arrival as described in the preceding section. As we see in the right-hand section of Figure 2 labeled "WARP," each radio receiver incorporates a 2.4 GHz oscillator whose purpose is to convert the incoming radio frequency signal to its representation in I-Q space shown, for example, above in Figure 1(b). An undesirable consequence of this *downconversion* step is that it introduces an unknown phase offset to the resulting signal in I-Q space, rendering our proposed method of measuring angle of arrival (and MIMO) inoperable.

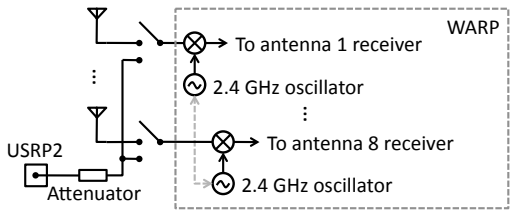

Figure 2: SecureAngle physical layer design. All 2.4 GHz oscillators are synchronized, and inputs to the eight receivers switch between the antennas and the signal generator (labeled "USRP2"); we use the latter for calibration.

To remedy this, MIMO systems *phase lock* each radio's oscillator together, so that they run at exactly the same frequency. We represent this by the dotted line between oscillators in Figure 2. This suffices for MIMO, but not for our application, because the downconverters of even phase-locked systems introduce an unknown but constant phase difference to each receiver, which manifests as an unknown phase added to the constellation points in Figure 1(b).

Our solution is to calibrate the array, measuring each phase offset directly. The USRP2 in Figure 2 transmits a continuous 2.4 GHz carrier through a 36 dB attenuator, which we split into eight signals and feed into the radio front ends. Since each of the eight paths from the USRP2 to a radio receiver is of equal length, the signals we measure when the switches in Figure 2 are each in the lower position yield seven relative phase offsets for antennas 2–8, relative to antenna one. Subtracting these relative phase offsets from the incoming signals over the air then cancels the unknown phase difference, and the methods of §2.1 become applicable.

2.3 System design

In this section, we sketch how the above techniques can be integrated with a functioning wireless network, for both our representative applications.

2.3.1 Virtual fences

Using the SecureAngle signatures described above, we demonstrate in Section 3 that after overhearing just one packet, it is possible to measure approximately three quarters of our clients' bearings to the access point to

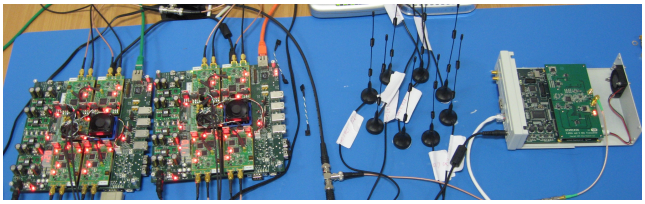

Figure 3: The SecureAngle prototype. Two Rice WARP platforms (left) provide a total of eight antennas and radio chains, while a cable-connected USRP2 software-defined radio (right) calibrates the array.

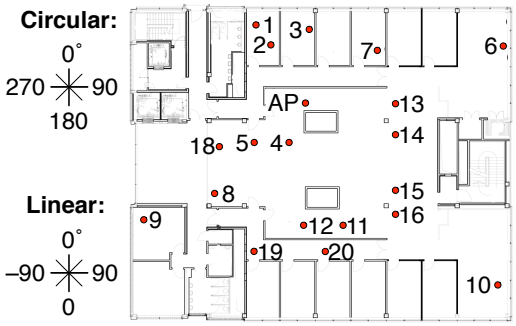

Figure 4: Testbed environment: Soekris clients are numbered, and the WARP access point is labeled "AP."

within 2.5° and all clients' bearings to within 14° with 95% confidence. We obtain this result even in an indoor office environment where multipath reflections are highly prevalent. In an environment where more than two access points are computing this bearing information, the intersection point of the direct path AoA is identified as the location of client.

2.3.2 *Address spoofing prevention*

To prevent address spoofing, SecureAngle records a legitimate client's signature S_{c_l} during the initial training stage and associates this signature with the MAC address. For all the incoming packets associated with this MAC address, signatures will be compared with S_{c_l}, the experimental hypothesis being that there is a significant difference between S_{c_l} and an attacker's signature, so that they can be discriminated from each other. Since S_{c_l} changes when the client or nearby obstacles move, the AP needs to track and update S_{c_l}. We can accomplish this using uplink traffic that the clients send to the AP. If a malicious client injects traffic into the network , the AP can detect the consequent change of signature and flag the injection event.

3. EVALUATION

In this section, we provide some empirical evidence to support SecureAngle's utility for the two applications outlined above.

Prototype implementation. The prototype in Figure 3 uses two WARP FPGA-based wireless platforms, each equipped with four radio front ends and four antennas. The WARPs run a custom hardware design of WARPLab calibrated as described in §2.2. The two WARP boards are also modified to share the same sampling clocks to remove frequency offset. The eight antennas attached to WARP are placed in linear or circular arrangements.[1] In the linear arrangement, they are spaced at a half wavelength distance (6.13 cm). The circular arrangement is actually an octagon with 4.7 cm sides and an antenna at each corner. We place the prototype access point at the point marked "AP" in our testbed, shown in Figure 4.

To capture traffic, the WARPs sample and buffer 20 MHz of signal bandwidth over time periods of 0.4 ms in length, then transfer the buffered samples over Ethernet to a computer running Matlab, on which we realize the Schmidl-Cox [11] OFDM packet detection algorithm to locate packets in the raw samples. In real wireless networks, measurements based on just one signal sample as described above are sensitive to background noise and interference from other senders. We therefore detect individual packets in the incoming stream of samples, and compute the correlation matrix to obtain mean phase differences with each entire packet.

3.1 Measurement accuracy

The first question we ask is: how accurately can SecureAngle determine the bearing to a client? To answer this question, we examine pseudospectra from the 20 Soekris clients shown in Figure 4. We are able to obtain very accurate bearing results for all the clients regardless of proximity to the AP or location inside or outside the AP's room. In an indoor environment with strong multipath propagation, reflections may generate false positive direct path AoA results. It is critical for us to eliminate these false positive AoAs if we are to determine the true bearing of clients.

Figure 5 shows the bearing results for all the clients with a circular AP antenna arrangement. With the proposed scheme, we are able to obtain the pseudospectrum graph which indicates the likelihood of energy received at each angle from 0 to 360° (examples of the pseudospectrum are shown in Figs. 6 and 7). We compute the bearing of each client as the angle corresponding to the maximum point on its pseudospectrum. We compute 10 pseudospectra for each client, each from a different packet, and plot the mean obtained bearing as well as 99% confidence interval in Figure 5. Client 6

[1]The AoA range for the linear arrangement is between −90 and 90°, since clients on the two sides of the line formed by the antennas are not differentiable. The circular arrangement solves this problem effectively.

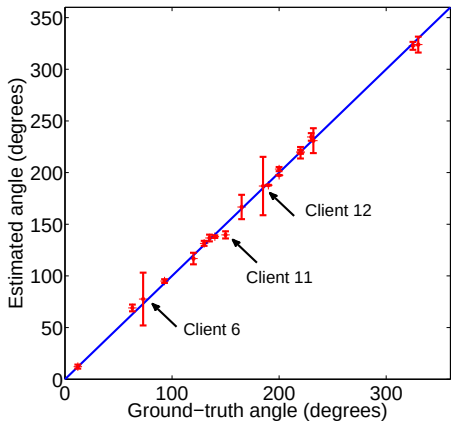

Figure 5: Measured versus ground truth bearing estimation for Soekris clients in the office environment shown in Figure 4 (error bars indicate 99% confidence intervals).

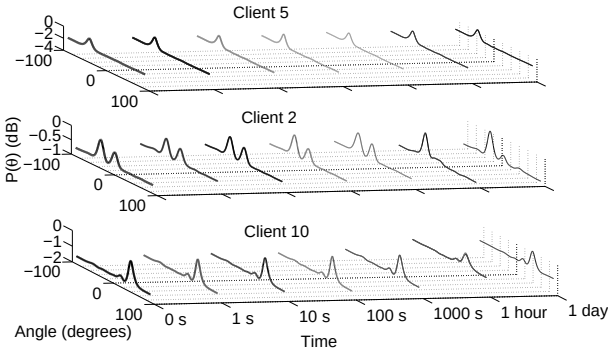

Figure 6: Stability of AoA signatures for three clients: each curve on each subplot shows the client's pseudospectrum at logarithmically-spaced time intervals.

and 12 have greater variance as Client 12 is blocked partially by a large cement pillar while client 6 is far away and multipath reflections are strong. Client 11 is completely blocked by the pillar and gives a little bit smaller value close to the true angle. The mean 99% confidence interval for all the clients is as small as 7°. Even with very strong multipath reflections, the direct path bearing corresponds to the highest peak in the pseudospectrum most of the time. For those few scenarios which the reflection paths are much stronger than the direct path, multiple APs can be applied to remove the false positive direct path AoA as those false positive AoAs obtained from different APs may not intersect with each other. Furthermore, direct path AoA is relatively stable compared with reflection path AoAs.

3.2 Measurement stability

SecureAngle depends not just on the stability of single antenna channels, but also specifically on the stability, or *coherence time* of the multiple antenna channel shown in Figure 1(c). Studies of a 4x4 MIMO channel at 2 GHz find median coherence times of between 25 ms for a walking-speed receiver and 125 ms for a stationary receiver [3] outdoors, where more and faster motion in the environment shortens coherence times [6]. Figure 6 shows AoA pseudospectra from three different clients with a linear AP antenna arrangement. To show that the AoA signatures are stable with time, each subplot of Figure 6 is composed of pseudospectra generated from packets recorded zero, one, 10, 100 and 1000 seconds, as well as one hour and one day later, all from the same client. We choose three representative clients: one (Client 2) located in another room nearby the access point. The other two clients are located in the same room as the access point but one (Client 5) is near by and the other (Client 10) is far away. The highest peak on the pseudospectrum graph usually corresponds to the direct path between the client and access point. We see clearly on the figure that the direct-path peak is quite stable while the multipath reflection peaks (smaller peaks) sometimes vary. From minute to minute, pseudospectra are quite stable, suggesting that for the purpose of protection against address spoofing, they can be tracked in the presence of normal indoor motion.

3.3 Measurement resolution

We show the effect of number of antennas on the AoA results in this section. When there are not many multipath reflections and the direct path signal is relatively strong, even two and four antennas can generate quite accurate results. However, in a more challenging environment when the direct path is relatively weak and multipath reflections are strong, increasing the number of antennas will result in more accurate bearing estimation. Another benefit of having more antennas is that the resolution of the pseudospectrum graph improves. We take Client 12 blocked by the pillar which has strong multipath reflections as the example. In Figure 7, we show the AoA pseudospectrum plot for the same packet with 2, 4, 6 and 8 antennas in linear arrangement.

A two-antenna arrangement generates one peak. Four antennas yield better resolution than two antennas with the measured bearing closer to the true bearing. However, with four antennas, it is not possible to differentiate two incoming signals within a 45° range. If the direct path and reflection path are within a 45° range, four antennas will not be able to generate two peaks for each of them, instead, one peak with an angle between the two incoming signals will be seen. However, this bearing is usually close to the true bearing. Once six antennas are used, we find that both the direct path and multipath components are visible. With eight antennas, we have even better resolution and more accurate results. With future wireless acess point designs

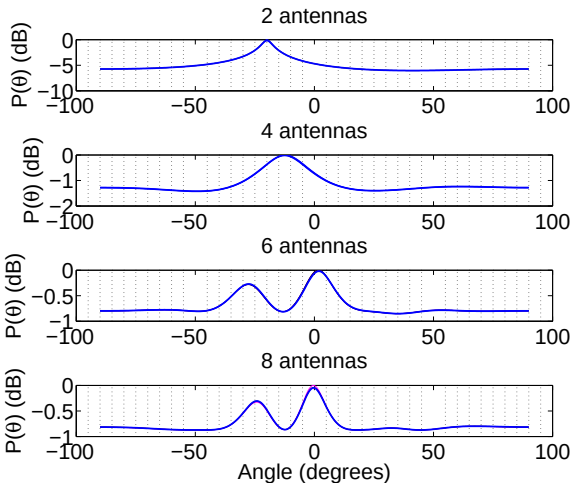

Figure 7: More antennas improve resolution and accuracy, resulting in a more specific signatures.

(such as SAM [15]) scaling up the number of antennas at the access point, the trend favors our design.

4. RELATED WORK

Many schemes have been proposed to capture signatures that characterize wireless clients' identity and location. The most widely used physical layer information is received signal strength (RSS) [2, 7]. While readily available from commodity hardware RSS is very coarse compared to physical-layer information, so is prone to error if few packets are available. Furthermore, attackers with directional antennas can subvert RSS-based systems [10]. Wong et al. [17] investigate the use of AoA information for localization using high-bandwidth (1 GHz) sampling rate. They apply a channel impulse response method which requires a probe packet with a long training sequence and a very high SINR (60 dB in their experiemnts) which is not realistic in real life. Geo-fencing [14] utilizes directional antennas and frames coding approach to control the indoor coverage boundary. Each AP sends out partial frames and only the clients located in the overlapping region covered by multiple APs can decode all the packets. Compared to Geo-fencing, SecureAngle does not mandate the rearrangement of traffic in the wireless network, but has not yet been shown to offer location-based security guarantees in the presence of adversaries with directional antennas – we leave this for future work.

5. CONCLUSION AND FUTURE WORK

We have described SecureAngle, a system that identifys angle-of-arrival data from the physical layer to fingerprint clients, adding a layer of wireless security.

In future work, we plan to integrate packet detection and AoA algorithms into the FPGA. We also plan to test our applications with client mobility and track the mobility trace with multiple APs. 3D location tracking will also be considered. Seamless location service for mobile devices using both GPS and SecureAngle's bearing estimation will be possible. This would cover both indoors where GPS signals do not propagate and outdoors. Based on our experiments, we believe SecureAngle could be as accurate as GPS. With AoA information obtained, high efficiency downlink directional transmission will also be feasible resulting in higher throughput and better reliability.

6. REFERENCES

[1] P. Bahl, R. Chandra, A. Greenberg, S. Kandula, D. Maltz, and M. Zhang. Towards highly reliable enterprise network services via inference of multi-level dependencies. In *Proc. of ACM SIGCOMM*, 2007.

[2] P. Bahl and V. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proc. of IEEE Infocom*, pages 775–784, 2000.

[3] M. A. Beach, M. Hunukumbure, and M. W. Webb. Dynamics of spatial eigenmodes in measured MIMO channels with different antenna modules. In *Proc. of the IET European Conf. on Antennas and Propagation*, 2007.

[4] M. Beck and E. Tews. Practical attacks against WEP and WPA. In *Proc. of ACM WiSec*, pages 79–86, 2009.

[5] A. Bittau, M. Handley, and J. Lackey. The final nail in WEP's coffin. In *Proc. of the IEEE Symp. on Security and Privacy*, pages 386–400, May 2006.

[6] J. Camp and E. Knightly. Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation. In *Proc. of ACM MobiCom*, pages 315–326, Sept. 2008.

[7] D. Faria and D. Cheriton. Radio-layer security: Detecting identity-based attacks in wireless networks using signalprints. In *Proc. of ACM WiSe*, 2006.

[8] D. Goodin. Lax security led to TJX breach. *The Register*, May 2007. `http://www.theregister.co.uk/2007/05/04/txj_nonfeasance`.

[9] S. Kandula, R. Mahajan, P. Verkaik, S. Agarwal, J. Padhye, and P. Bahl. Detailed diagnosis in enterprise networks. In *Proc. of ACM SIGCOMM*, 2009.

[10] N. Patwari and S. Kasera. Robust location distinction using temporal link signatures. In *Proc. of the ACM MobiCom Conf.*, pages 111–122, Sept. 2007.

[11] T. M. Schmidl and D. C. Cox. Robust Frequency and Timing Synchroniation for OFDM. *IEEE Trans. on Communications*, 45(12):1613–1621, Dec. 1997.

[12] R. Schmidt. Multiple emitter location and signal parameter estimation. IEEE *Trans. on Antennas and Propagation*, AP-34(3):276–280, Mar. 1986.

[13] B. Schneier, Mudge, and D. Wagner. Cryptanalysis of Microsoft's PPTP authentication mechanisms (MS-CHAPv2), Oct. 1999.

[14] A. Sheth, S. Seshan, and D. Wetherall. Geo-fencing: Confining Wi-Fi Coverage to Physical Boundaries. In *Proceedings of the 7th International Conference on Pervasive Computing*, 2009.

[15] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. Voelker. SAM: Enabling practical spatial multiple access in wireless LAN. In *Proc. of ACM MobiCom*, 2009.

[16] E. Tews, R.-P. Weinmann, and A. Pyshkin. Breaking 104-bit WEP in less than 60 seconds. *Springer Lecture Notes in Computer Science*, 4867:188–202, Jan. 2008.

[17] C. Wong, R. Klukas, and G. Messier. Using WLAN infrastructure for angle-of-arrival indoor user location. In *Proc. of the IEEE VTC Conf.*, pages 1–5, Sept. 2008.