

Detecting Identity-Based Attacks in Wireless Networks Using Signalprints

Daniel B. Faria
Computer Science Department
Stanford University
dbfaria@cs.stanford.edu

David R. Cheriton
Computer Science Department
Stanford University
cheriton@cs.stanford.edu

ABSTRACT

Wireless networks are vulnerable to many identity-based attacks in which a malicious device uses forged MAC addresses to masquerade as a specific client or to create multiple illegitimate identities. For example, several link-layer services in IEEE 802.11 networks have been shown to be vulnerable to such attacks even when 802.11i/1X and other security mechanisms are deployed. In this paper we show that a transmitting device can be robustly identified by its *signalprint*, a tuple of signal strength values reported by access points acting as sensors. We show that, different from MAC addresses or other packet contents, attackers do not have as much control regarding the signalprints they produce. Moreover, using measurements in a testbed network, we demonstrate that signalprints are strongly correlated with the physical location of clients, with similar values found mostly in close proximity. By tagging suspicious packets with their corresponding signalprints, the network is able to robustly identify each transmitter independently of packet contents, allowing detection of a large class of identity-based attacks with high probability.

Categories and Subject Descriptors

C.2.3 [Computer Communication Networks]: Network Operations - *Network Monitoring*; C.2.5 [Computer Communication Networks]: Local and Wide-Area Networks

General Terms

Design, Measurement, Security.

Keywords

Denial-of-Service Attacks, Security, Wireless LANs, Location-Based Services, IEEE 802.11.

1. INTRODUCTION

Denial-of-service (DoS) attacks can bring networks to a halt by saturating communication links or by flooding hosts with requests that induce computationally expensive operations or unnecessary allocation of resources. Wireless LANs

(WLANs) are yet another scenario for DoS attacks, though with the added complication that the wireless medium makes it easier for the injection of attack traffic.

Several DoS attacks in wireless LANs are possible because these networks lack reliable client identifiers before upper-layer authentication mechanisms are evoked and user credentials are securely established. After a client authenticates successfully and session keys are used to encrypt and authenticate packets sent over wireless links, the network can securely verify if the source MAC address in a packet is correct. Without this mechanism, however, wireless installations have to rely solely on MAC addresses for client identification: two devices in a network using the same address are treated as a single client, even if they generate conflicting or inconsistent requests.

As MAC addresses can be easily changed through device drivers, simple yet effective identity-based attacks can be implemented with off-the-shelf equipment against multiple link-layer services. IEEE 802.11 networks, for instance, have been shown to be vulnerable to a class of attacks we refer to as *masquerading attacks*, in which a malicious device targets a specific client by spoofing its MAC address or the address of its current access point. Bellardo and Savage have demonstrated that a 10-second deauthentication attack can immediately knock a client off the network and possibly incur minute-long outages given the interaction between 802.11 and TCP [5]. With such tools, a malicious user could render a WiFi hotspot unusable by targeting all active clients or simply maximize the throughput achieved by his own laptop by periodically deauthenticating devices using the same access point as him. These attacks can be currently implemented even if networks deploy recent security standards such as IEEE 802.11i [2].

Another class of identity-based attacks target *resource depletion*: an attacker can generate high rates of requests with random MAC values in order to consume shared resources. For example, authentication protocols such as TLS (popular with 802.11i/802.1X) demand milliseconds of processing time, making servers vulnerable to attacks that consume in the order of 200 Kbps of attack bandwidth [7]. As another example, the attack could target a DHCP server in a publicly available part of the network and consume all IP addresses reserved for visitors. A PDA device left behind inside a corporation could act as a “wireless grenade”, going off at a programmed time and flooding the authentication server with random requests, possibly affecting clients well beyond its communication range.

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WiSe'06, September 29, 2006, Los Angeles, California, USA.
Copyright 2006 ACM 1-59593-557-6/06/0009 ...\$5.00.

In this paper we show that reliable client identifiers, which we call *signalprints*, can be created using signal strength information reported by access points and used to detect misbehaving devices. As a packet of interest (e.g. a deauthentication request) is transmitted over the wireless link, it is sensed by access points within range, which report signal strength measurements (a.k.a RSSI levels) to a centralized server. The request is then “tagged” with a signalprint, a tuple constructed by aggregating all measurements reported. Transmitters at different locations produce distinct signalprints because signal decays with distance, allowing the system to robustly distinguish clients located geographically apart. We present measurements performed within an office building with an IEEE 802.11 network that demonstrate that signalprints can be used to detect masquerading and resource-depletion attacks with high probability.

2. ATTACK MODEL

We assume that malicious clients are provided with standard wireless transmitters. First, we assume they employ omnidirectional antennas, much like most portable wireless devices. The use of directional antennas is discussed in section 6.5. Second, we assume they are able to modify the contents of each outgoing packet. This allows them, among other things, to change source and destination MAC addresses, a capability needed to implement the attacks we are interested in. For example, Bellardo and Savage have described a mechanism that can be used to accomplish this [5]. Finally, we assume that they are provided with multiple transmission power levels and that they can also change that setting on a per-packet basis. In this paper we restrict ourselves to 802.11 networks, but the ideas presented can be equally applied to other wireless LAN technologies.

In terms of their physical location, we assume attackers can move freely around the area covered by the wireless network. Note that in practice, this is only possible in environments with little or no physical security, such as in cafeterias and other hotspots. The probability of mounting successful attacks would be lower in environments with tighter security measures, such as in enterprise installations.

In this paper we focus on the two classes of attacks already mentioned: masquerading and resource depletion attacks.

3. SIGNALPRINTS

3.1 Network Architecture.

We assume a network architecture as shown in figure 1, composed of multiple access points (APs) distributed across the environment that feed traffic information to a centralized server, which we call a *wireless appliance* (WA). We focus on the access points deployed as sensors: they observe the traffic on a channel specified by the WA and collect information such as the received signal strength level for each packet successfully received. This information is then forwarded to the WA, which is able to create a signalprint for each packet of interest.

Our proposed mechanism can be readily deployed. For instance, the architectural requirements just presented are currently satisfied by some 802.11-based wireless intrusion detection systems (WIDSs) and network installations that employ lightweight access points. Some WIDSs work exactly as described above, with a server aggregating infor-




Figure 1: Signalprint creation.

mation from all APs in order to detect security events of interest. Networks with lightweight APs require a central point of control, a device similar in function to what we call a WA. In this case, access points implement minimum functionality – sometimes acting simply as remote radio interfaces – and delegate all other functions to the WA, which is computationally more powerful. An example is the architecture being currently standardized by the CAPWAP Working Group at IETF [6], which should allow installations to scale to large numbers of access points by simplifying network management.

3.2 Signalprint Representation

Conceptually, a signalprint is the signal strength characterization of a packet transmission. Each signalprint is represented as a vector of signal strength measurements, with one entry for each access point acting as sensor. Values in signalprints always appear in the same order, i.e., position i always contains the signal strength level (in dBm) reported by the i^{th} AP. We use the notation $S[i]$ to refer to the i^{th} entry in a signalprint. If an access point does not report an RSSI level for a given packet, a default value equal to its *sensitivity* is used. (The sensitivity of a receiver with respect to a given data rate is defined as the minimum signal strength level needed to achieve a target packet error rate.)

The size of a signalprint is the number of non-default elements it contains, i.e., the number of entries created from actual RSSI measurements. For instance, figure 2(a) shows two signalprints, S_1 and S_2 , both with 7 entries (the number of APs in the network) but with sizes 5 and 6, respectively. (In this case, default values of -95 dBm were used.) Signalprint S_1 was created using RSSI levels reported by APs 1, 3, 4, 5, and 7, while S_2 has values from APs 1, 2, 4, 5, 6, and 7. As an alternative notation, S_1 can also be written as $S_1 : (-50, -, -80, -73, -88, -, -60)$, where default values are omitted.

3.3 Signalprint Generation

Figure 1 illustrates how signalprints are created for wireless transmissions. A client (Client1) is shown transmitting an authentication request through its current access point

(solid line). Before forwarding the packet to the WA, the AP tags it with the RSSI level measured during reception. (Signal strength estimates are commonly made available by IEEE 802.11 device drivers for each packet received.) The other two APs shown in the figure are also configured as sensors and tuned to the same channel. As `Client1` is also within their ranges, they send similar reports to the WA with their own RSSI measurements. As shown at the top of the figure, the WA aggregates all reports and creates the following signalprint for `Client1`: $S_{C1} : (-73, -51, -67)$. A signalprint for a second client, `Client2`, is also shown at the WA. The signalprints produced by both clients are quite different – for example the clients could be located in different offices within a building.

The WA can identify identity-based attacks by comparing signalprints produced by multiple packets. For example, if `Client1` submits a high rate of requests trying to clog the authentication server, the WA can detect it given that many of `Client1`'s transmissions produce similar signalprints. Likewise, the WA can detect if `Client2` mounts a DoS attack against `Client1` by sending 802.11 deauthentication requests with `Client1`'s MAC addresses, as the signalprints produced by the two devices are different.

We assume that a subset of the deployed access points report RSSI measurements to the WA for all transmissions they can detect. In the case of a WIDS that relies on a separate wireless infrastructure, some APs are already permanently configured as sensors. In a CAPWAP network, the WA is responsible for selecting the APs for signalprint processing. In an over-provisioned installation, the WA can select the access points that are not actively serving clients.

Signalprint-based attack detection should be implemented as a reactive mechanism whenever the number of sensors is not sufficient to cover all active channels. For instance, dense 802.11 deployments have at least 15 non-overlapping channels available across both 2.4 and 5 GHz frequency bands. The objective is to maximize the size of the signalprints produced: the more measurements are received for a packet transmission, the more accurate is the information gathered about the location of the corresponding device. For that to be possible, sensor APs need to be listening simultaneously to the proper channel. In large networks, where more channels are required to serve active clients, dividing the sensors across all channels to be monitored would produce short, inaccurate signalprints. For this reason, a two-step monitoring process should be implemented in these situations. First, the WA identifies any abnormal behavior using both active and sensor APs, which are scattered across all channels. When abnormal behavior is detected – such as a surge in the number of 802.1X authentication requests or a high number of association events related to a single client – the WA sets enough sensors to the proper channel to create signalprints for the relevant packets.

3.4 Signalprint Properties

Three properties concerning signalprints enable their use as reliable client identifiers:

Signalprints are hard to spoof. Signal attenuation is a function of the distance between clients and access points, with a strong dependence on environmental factors such as construction materials and obstacles such as furniture items [13, 18]. Consequently, transmitters have little or no control over signal attenuation within the environment,

being unable to considerably change the signalprints they produce. We show that the use of differential signalprints makes the system robust against devices that employ multiple transmission power levels, further decreasing their control over the signalprints generated.

Signalprints are strongly correlated with the physical location of clients, with similar signalprints found mostly in close proximity. In our measurements, performed within a 45m×24m office environment with a total of 12 802.11 access points, devices need to be as close as 5 meters in order to generate similar signalprints with high probability, even when only 6 APs are used. This allows the detection of masquerading attempts when attacker and victim are not in close proximity. If an attacker aims to DoS a specific client and avoid detection, he is forced to move closer to the infrastructure, thus risking exposure.

This property has also been demonstrated by WLAN localization systems that employ an offline training phase where signal strength patterns (essentially signalprints) are created for a set of selected locations (usually called a signal map, or radio map). These systems have consistently achieved average localization errors below 3 meters, mapping areas as large as 19,000 s.f. and with numbers of access points varying between 4 and 20 [4, 20, 17, 24].

Packet bursts transmitted by a stationary device generate similar signalprints with high probability.

Our measurements show that while RSSI levels for a stationary device do oscillate over time due to multiple factors, over 90% of variations are within 5 dB from the median RSSI level. This correlation between consecutive samples has also been reported by other researchers [24]. Consequently, an attacker that mounts a resource depletion attack using random MAC addresses can be easily spotted. While not all signalprints may match each other, the network would still be able to detect that a single transmitter is responsible for a high rate of requests.

Signalprints allow a centrally controlled WLAN to reliably single out clients. Instead of identifying them based on MAC addresses or other data *they* provide, signalprints allow the system to recognize them based on what they look like in terms of signal strength levels.

4. MATCHING SIGNALPRINTS

In this section we demonstrate how matching rules are specified to detect identity-based attacks. In section 4.1 we describe the use of differential signal strength values during matching. In sections 4.2 and 4.3 we describe how values within signalprints are compared using max-matches and min-matches. In section 4.4 we describe how matching rules are specified in terms of these operations.

4.1 Differential Values

Values within a signalprint can be written as absolute values (e.g. RSSI levels in dBm) or as relative values (e.g. with respect to its higher or lower value). We use the term *differential signal strength* to refer to the difference between the value at a given position and the maximum value found in that signalprint. Signalprints are either written with absolute or differential values: for example, a signalprint $S : (-50, -62, -76)$ written using differential signal strength becomes $S : (0, -12, -26)$. Figure 2(b) shows S_1 and S_2




Figure 2: Signalprint matching examples. Figure 2(a) shows two signalprints and their corresponding sizes. Figures 2(b) and 2(c) demonstrate how max-matches and min-matches are computed.

written with both absolute and differential values (the latter shown respectively above and below S_1 and S_2). When matching two signalprints, both need to be written in either absolute or differential values.

The use of differential values increases the robustness of signalprint operations against devices (possibly malicious) that vary their transmission power levels between frames. It is a trick borrowed from differential GPS, where a second, stationary receiver is used to remove timing errors that occur in both paths, between a satellite and each one of the receivers. In our case, this error or unknown quantity is the power level used by a transmitter. With absolute values, changes in transmission power create similar changes in the detected RSSI, which could cause the system to attribute multiple packets sent by a single client to multiple devices. Using differential values, transmissions performed by a stationary transmitter generate similar signalprints, increasing changes of attack detection.

4.2 Max-Matches

Matches are found by comparing values at the same position in two different signalprints. A “max-match” of ϵ dB is found whenever values differ by *at most* ϵ dB. I.e., a 10-dB max-match is found at position i if $abs(S_1[i] - S_2[i]) \leq 10$ and both $S_1[i]$ and $S_2[i]$ are non-default values. The total number of ϵ -dB max-matches found by comparing signalprints S_1 and S_2 is denoted by $maxMatches(S_1, S_2, \epsilon)$.

We decided to remove default values from match computations because they can arise from two distinct scenarios. On one hand, a client can be simply outside the range of an access point, in which case its packets are not detected and RSSI measurements are simply not reported. On the other hand, many events may cause an AP to fail to receive packets independently of signal quality. For instance, two packets sent on the same channel but on different cells may overlap in time, in which case both packets might be incorrectly decoded and dropped by the AP.

In this paper matches are always computed using differential signal strength values. Figure 2(b) shows that 3 10-dB max-matches are found when comparing S_1 and S_2 , i.e. $maxMatches(S_1, S_2, 10) = 3$. Signalprints are shown with both original and differential signal strength values, with matches found at positions 1, 4, and 7. Note that position 5 does not yield a 10-dB max-match when using differential values: the difference equals 21 dB instead of the 8 dB when absolute values are used.

Max-matches are especially useful when looking for signalprints produced by the same transmitter. As we show in section 6, RSSI values produced by a stationary client tend to oscillate within 5 dB from its median value. As a result,

high numbers of max-matches with low values of ϵ (e.g. 5 dB) are likely to occur for a pair of signalprints sent by the same device.

4.3 Min-Matches

Analogous to a max-match, a “min-match” of ϵ dB is found whenever values differ by *at least* ϵ dB. A 10-dB min-match is found at position i if $abs(S_1[i] - S_2[i]) \geq 10$ and both $S_1[i]$ and $S_2[i]$ are non-default values. The total number of ϵ -dB min-matches found when comparing signalprints S_1 and S_2 is denoted by $minMatches(S_1, S_2, \epsilon)$. As shown in figure 2(c), a single 20-dB min-match is found when comparing S_1 and S_2 , at position 4.

Min-matches allow the system to identify, with high probability, when two packets are sent by distinct devices. While small variations in received signal strength occur even for a stationary client, rarely does it change by more than 10 or 15 dB. Consequently, the system can classify two packets as coming from different devices with high confidence if large differences are seen in a signalprint.

4.4 Matching Rules

We say that a pair of signalprints “match” if they satisfy a specified *matching rule*, a boolean expression involving numbers of max-matches and min-matches, and possibly signalprint properties such as size. The matching rule $maxMatches(S_1, S_2, 5) \geq 4$ requires two signalprints to have RSSI values within 5 dB of each other in at least 4 positions.

When specifying matching rules, it is important to account for both signal strength oscillation and lack of feedback from access points. Constant RSSI oscillation makes it unlikely that even signalprints produced by the same stationary device have the exact same RSSI values in multiple positions. Consequently, we usually write max-match clauses with values of ϵ of at least 5 dB. The lack of feedback from some APs prevents matches in all signalprint positions.

As with intrusion detection systems, matching rules are specified with the objective of minimizing false positives, i.e., we want a match to be a strong indication that an attack is taking place. The reason is cost: a match raises an alarm that is likely to be handled by the network administrator. Rules can be made more precise (fewer false positives) by increasing the minimum number of matches and changing the value of ϵ .

5. ATTACK DETECTION

Three attack properties are important to our analysis: R denotes the rate in packets per second (pps) required for a given DoS attack to be effective, S denotes the speed of

the device, while A denotes the number of antennas under the control of the attacker. In this section we assume that devices are stationary ($S = 0$) and provided with a single omni-directional antenna ($A = 1$). In section 6.4 we address the effects of moving devices, while in section 6.5 we assume attackers with directional antennas. Finally, we discuss attacks with multiple antennas in section 7.

5.1 Resource Depletion Attacks

In this scenario, an attacker sends high rates of request messages using random MAC values in order to emulate a high number of clients and consume scarce resources in the network. For example, an attacker can send enough DHCP requests in a hotspot as to consume all available IP addresses, flood access points with association requests in the hopes of exceeding allowed limits, or send high rates of authentication requests to slow down or even disable a shared authentication server.

As an example, Dean *et al.* [7] have also shown that effective low-bandwidth DoS attacks can be mounted against TLS, one of the preferred authentication methods to be used in 802.11i/802.1X [2, 1]. TLS requires cryptographic operations (e.g. RSA and Diffie-Hellman) that when executed in software demand tens of milliseconds even on a dedicated processor. In situations where a dedicated server is not available – some lightweight AP architectures perform authentication at the WA – the overhead imposed by each request could be much higher. A server could therefore be overloaded by a device that generates 100 requests per second, which can be injected into the network while demanding far less than 1Mbps of attack bandwidth.

In this case, the input to the signalprint matching process is a set of packets (e.g. authentication requests) with distinct MAC addresses and their corresponding signalprints. Effective DoS attacks in this category require high packet rates ($R \gg 1$ pps), so many signalprints should be available for processing. By comparing pairs of signalprints, the system can identify subsets generated by the same device.

Matching rules should require multiple max-matches with low values of ϵ because we are looking for signalprints that were generated by the same device and therefore expected to have similar RSSI values in multiple positions. Using 6 APs as sensors, the first rule we evaluate in section 6 for this purpose is $maxMatches(S_1, S_2, 5) \geq 4$. We can decrease the probability of false positives by increasing the required number of max-matches or decreasing the value of ϵ . The second rule we evaluate – $maxMatches(S_1, S_2, 5) \geq 5$ – tends to be satisfied by signalprints generated at locations that are physically closer to each other.

In order to further decrease the probability of false positives, these rules can be extended with min-match clauses. For instance, consider two signalprints that satisfy the second matching rule above by having similar RSSI levels in 5 positions. Now consider the single position that did not produce a max-match. If one of the signalprints has a default value at that position, the likelihood of these signalprints being from the same device does not change much. However, if values are defined in both signalprints and differ by 8 or 10 dB, this likelihood decreases substantially. Therefore, we evaluate a third matching rule that extends the second rule above with a min-match clause: $maxMatches(S_1, S_2, 5) \geq 5 \wedge minMatches(S_1, S_2, 8) = 0$.

5.2 Masquerading Attacks

In masquerading attacks, an attacker targets a specific client by cloning its MAC address or the address of its access point. For instance, Bellardo *et al.* have shown that deauthentication and disassociation attacks can be easily mounted in 802.11 networks and are very effective [5]. Before a client can send packets over the wireless link, it needs to authenticate and associate itself with an AP. In a deauthentication attack, deauthentication requests are sent by an attacker with the MAC address of the victim. The access point, after granting the attacker’s request, removes the victim from the authenticated state and drops all its packets until association is reestablished. Bellardo *et al.* discuss other equally effective masquerading attacks that exploit the association service and the power saving mechanism [5].

In normal situations, 802.11 devices are not expected to generate high rates of authentication or association messages. However, there *are* situations in which well-behaved clients switch between access points with a frequency that is abnormally high. For example, in their study of a large-scale 802.11 network, Kotz *et al.* showed that clients sometimes are overly aggressive when selecting the best access point, which causes them to reassociate more often than necessary [16]. In these cases, multiple APs are within the client’s range with comparable RSSI levels, which may cause it to change APs with small variations in signal strength.

So the WA can detect an unusual traffic pattern, but is an attack really happening? Signalprints can be used to detect attacks with high probability, providing a level of assurance that cannot be achieved by only looking at packet contents.

The input now consists of two sets of packets that represent conflicting requests (e.g. authentication vs. deauthentication messages), all transmitted with the same MAC address. An attack is detected by comparing pairs of signalprints, one from each set. Given that continuous attacks are needed to severely affect a victim’s throughput, large input sets are also expected in this case. For example, to keep a victim off the network, Bellardo *et al.* used up to 10 deauthentication frames per second in their experiments [5].

To detect these attacks, matching rules should require min-matches with large values of ϵ , because we are looking for considerable differences in RSSI that would indicate two (or more) distinct transmitters. In this case, rules can be more precise by either increasing the number of min-matches or increasing the value of ϵ . In our evaluation section, using 6 access points, we look for 10-dB min-matches. We evaluate the performance of two matching rules for this purpose: $minMatches(S_1, S_2, 10) \geq 1$ and $minMatches(S_1, S_2, 10) \geq 2$.

6. EVALUATION

In this section we show that signalprints are strongly correlated with the physical locations within an environment, which allows them to be used as robust, location-dependent client identifiers.

6.1 Testbed

Our testbed consists of a 45×24m (147×78 ft) section of an office environment (the 4A Wing of the Gates Building at Stanford University). As shown in figure 3, it contains a mix of offices (most 3×6m), large labs (at least 8×4.5m), and long corridors. We have installed a total of 12 IEEE 802.11b/g access points, which are mounted at the ceiling

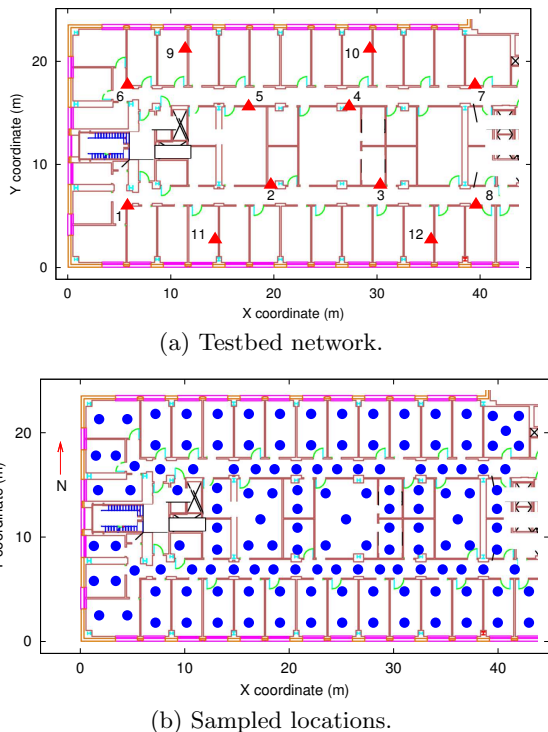


Figure 3: Our testbed and locations sampled.

(height of 2.5 m) and connected to a centralized server. The access points are off-the-shelf Linksys WRT54G units running a Linux distribution (OpenWrt rc3), their exact placement depicted by triangles in figure 3(a). An application running at each AP monitors the wireless channels and reports signal strength information back to the server.

In order to evaluate our mechanism, we created a test dataset by manually sampling 135 locations across the floor, as shown in figure 3(b). At each location a laptop provided with a Cisco 340 PCMCIA card transmitted ping packets at rates between 10 and 20 packets per second for approximately one minute, as our access points had to hop through the three 802.11b/g non-overlapping channels. The client annotated each packet with the location ID, while each access point in range tagged it with an RSSI level and forwarded it to the localization server, which logged all measurement traffic. At all locations the laptop was held at waist level and with the same orientation, with the user facing North as indicated in figure 3(b). A total of over 420,000 signal strength samples were collected, and we created a distinct signalprint for each location by using the median RSSI level with respect to each access point.

6.2 Signal Strength Oscillation

We first demonstrate that while stationary, a wireless client tends to create similar signalprints, despite the inherently unpredictable nature of wireless propagation. This is an important property when we consider the performance of our system against DoS attacks with high rates of requests. For example, even a small fraction of matching signalprints would allow the network to detect a malicious device that sends over 100 authentication requests per second.

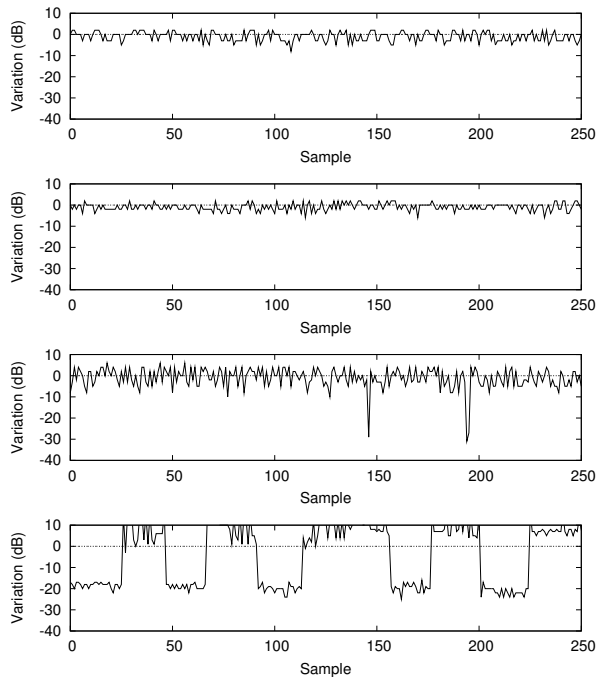


Figure 4: RSSI oscillation for a stationary device. Each graph was created by choosing one of the locations sampled and one access point within its range. It shows the variation in signal strength for consecutive frame transmissions relative to the median RSSI (shown as 0 dBm) for that location-AP pair.

Our measurements suggest that most signal strength oscillations are small, within 5 dB from the median RSSI level. Each graph in figure 4 was created by choosing one of our sampled locations and one of the access points within its range. Each graph shows the variation in signal strength for successive transmissions: the difference between the detected RSSI and the median value for the corresponding location-AP pair (the median, or base level, is shown as 0 dBm). The top two graphs in figure 4 are examples of the behavior detected for most locations: the majority of RSSI oscillations are within 5 dB from the median. Aggregating all the measurements in our dataset – all locations with respect to all access points – we have that over 71%, 90%, and 93% of RSSI oscillations are respectively within 2, 5, and 10 dB from the median RSSI levels.

However, the tail of the distribution is long, and some strong, mostly destructive oscillations do occur. The bottom two graphs in figure 4 are examples of this case. In one of them, the RSSI level is somewhat stable, with a couple of strong oscillations (>25 dB). In the other example, periods with strong RSSI degradation seem to happen with a certain frequency, with a difference in signal strength between the two levels of over 30 dB.¹

While strong RSSI oscillations do occur in our measurements, the path loss between a stationary client and each AP is stable most of the time. In theory, multipath prop-

¹We do not have a definitive explanation for these signal strength variations we have observed.

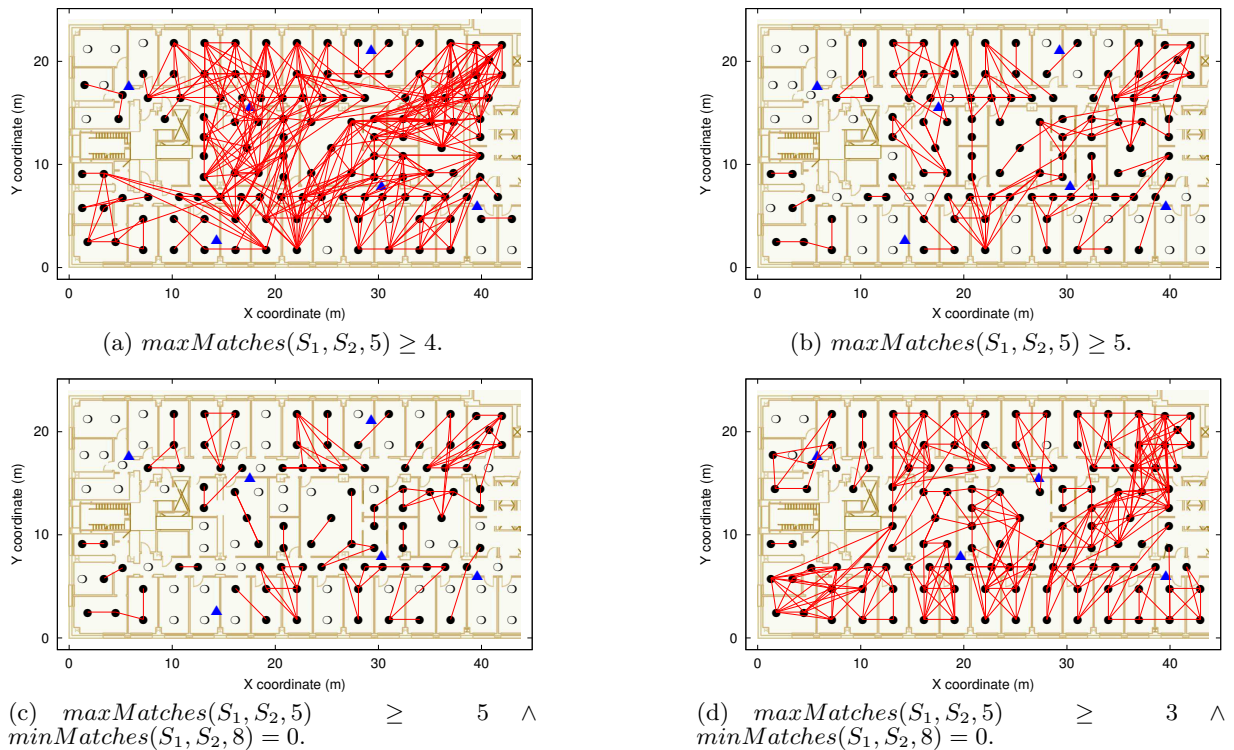


Figure 5: Location pairs satisfying multiple matching rules. Figures 5(a)-5(c) use the 6-AP configuration (APs shown as triangles), while figure 5(d) uses the setup with 4 access points.

agation and other phenomena generate small-scale fading, with possibly strong RSSI variations over time caused by people walking by, doors being closed, and other changes in environment that affect any of the multiple paths taken by transmissions between two devices. In practice, however, these events do not seem to happen often. Perhaps these results are due to techniques developed to decrease the effects of small-scale fading in wireless systems, such as antenna diversity, implemented in most 802.11 devices.

6.3 Signalprints and Physical Proximity

In this section we explore the relationship between signalprints and physical proximity between transmitters in order to detect identity-based attacks. Despite having 12 access points deployed in our testbed, all the results presented in this paper use two AP configurations: one with 6 access points (numbers 3, 5, 6, 8, 10, and 11 in figure 3(a)) and the other with 4 APs (numbers 2, 4, 6, and 8). The configuration with 4 access points is used to evaluate the loss in accuracy when using fewer sensors. Each figure shows the APs being used as triangles and omits the others.

6.3.1 Detecting Packets From a Single Device

As discussed in section 5.1, matching rules that detect when packets are generated by the same device are useful to detect high-rate DoS attacks. By requiring multiple max-matches with low ϵ values, we show that matching similar signalprints are found mostly in close proximity.

Locations that produce signalprints with similar values in multiple positions tend to be physically close. Using the 6-

AP configuration, figure 5(a) shows all location pairs that satisfy the matching rule $\maxMatches(S_1, S_2, 5) \geq 4$ connected by a line segment. Even though many matches are produced, most of them involve locations that are close to each other. Overall, 430 matches were found (4.8% of all pairs), with 51%, 74%, and 91% of them found respectively for locations within 5, 7, and 10 meters from each other. However, there are still many matches found for locations more than 15 meters from each other. All matching results presented in this section are summarized in table 1.

As discussed in section 5.1, matching results are improved if we increase the number of max-matches required. Still using 6 sensors, figure 5(b) shows the locations whose signalprints satisfy the rule $\maxMatches(S_1, S_2, 5) \geq 5$. Compared to figure 5(a), there is a significant reduction in the number of long-distance matches. A total of 150 matches were found (1.7% of all pairs), with respectively 64%, 88%, and 98% of them found for locations within 5, 7, and 10 meters from each other. In this case, there are no matches for locations more than 15 meters apart.

Using min-matches, matching rules can be made even more precise. Figure 5(c) shows that the matching rule $\maxMatches(S_1, S_2, 5) \geq 5 \wedge \minMatches(S_1, S_2, 8) = 0$ further reduces the number of long-distance matches. Like in figure 5(b), this rule still requires a minimum of 5 max-matches of 5 dB, but now rejects all the location pairs for which any difference larger than 8-dB is found. As shown, this rule produces only 97 matches (1.1% of all pairs), with respectively 72%, 91%, and 99% of them found for locations within 5, 7, and 10 meters from each other. In this case

Matching Rule	Figure	# APs	# Matches	$\leq 5\text{m}$	$\leq 7\text{m}$	$\leq 10\text{m}$
$\text{maxMatches}(S_1, S_2, 5) \geq 4$	fig. 5(a)	6	430 (4.8%)	50.9%	74.4%	91.2%
$\text{maxMatches}(S_1, S_2, 5) \geq 5$	fig. 5(b)	6	150 (1.7%)	64.0%	88.0%	98.0%
$\text{maxMatches}(S_1, S_2, 5) \geq 5 \wedge \text{minMatches}(S_1, S_2, 8) = 0$	fig. 5(c)	6	97 (1.1%)	72.2%	90.7%	99.0%
$\text{maxMatches}(S_1, S_2, 5) \geq 3 \wedge \text{minMatches}(S_1, S_2, 8) = 0$	fig. 5(d)	4	317 (3.5%)	62.2%	86.4%	99.1%
$\text{minMatches}(S_1, S_2, 10) \geq 1$	(no figure)	6	8643 (95.6%)	4.6%	10.1%	21.9%
$\text{minMatches}(S_1, S_2, 10) \geq 2$	(no figure)	6	7768 (85.9%)	2.6%	6.9%	17.9%

Table 1: Matching results. Each row shows a matching rule, the figure (if any) containing the signalprints created from our measurements that satisfy that rule, the number of access points used as sensors, the number of matches produced, and the percentages of matches created by locations within 5, 7, and 10 meters from each other. The first four rules are used to detect packets transmitted by the same device, while the last two detect packets sent by distinct devices.

there is a single match for locations more than 10 meters from each other.

Finally, figure 5(d) shows that performance degrades if we decrease the number of access points used to 4, but that results are still satisfactory due to the use of min-matches. With 4 APs, this matching rule requires at least 3 5-dB max-matches and no 8-dB min-matches. It produces 317 matches, but with respectively 62%, 86%, and 99% of matches found for locations within 5, 7, and 10 meters from each other. Note that these numbers are better than the ones related to figure 5(a) even though there are two fewer access points.

These results show that resource depletion attacks can be detected with high probability, as matching signalprints are found mostly for locations that are near each other. Therefore, a large number of matching requests means they are being transmitted from a specific location or area, which could be found by coupling our mechanism with a localization system. Some signalprints produced at the same location may not match due to RSSI oscillations, but this does not prevent the WA from detecting high-rate DoS attacks.

6.3.2 Detecting Packets From Distinct Devices

In this section we evaluate matching rules specified to decrease the probability of false positives when looking for masquerading attacks. We want signalprints to match only if there is a high probability that they were indeed produced by distinct devices. In this case, detecting large RSSI differences is more important than finding similar values, so min-matches play a more important role in these situations.

Most location pairs in our dataset generate signalprints that satisfy the matching rule $\text{minMatches}(S_1, S_2, 10) \geq 1$, i.e., values in at least one position differ by 10 dB or more. As shown in table 1 (5th row) over 95% of all location pairs satisfy this rule. Even a large number of locations that are physically close can be distinguished, with over 400 matches produced for locations less than 5 meters from each other. Overall, these results show that masquerading attacks can be detected with high probability, as at least one access point can tell the two locations apart.

We can decrease the probability of false positives by increasing the minimum number of 10-dB min-matches to 2. As shown in the 6th row in the table, over 85% of all location pairs still produce a match. In this case, a match is an even stronger indication that an attack is taking place, as signalprints differ substantially relative to at least two access points.

6.4 Moving Devices

We do not expect legitimate clients on the move to generate false alarms because they send requests at rates much lower than required by most attacks. For example, consider an 802.11 client that associates with an access point and after some time moves to a different location and requests disassociation. Despite the fact that the two signalprints generated can be quite different, an alarm should not be raised in this situation. An effective disassociation attack requires higher rates of deauthentication requests to keep a client off the network, so only a larger number of matching signalprints detected during a short period of time (e.g. tens of seconds) should generate an alarm.

Unless an attacker moves towards the victim, changing his location does not increase the chances of having a successful masquerading attack. What matters is not how the signalprints he produces compare to each other – for this matter they could be all different – but how similar they are to the one produced by the victim. Attacks are detected as long as there are considerable RSSI differences, which only cease to exist if the attacker moves close to his victim.

Whether an attacker can disguise a resource depletion attack by changing his location over time depends on his speed and the required packet rate. Let us assume that an attacker moves at pedestrian speeds and consider an attack requiring $R > 10$ pps (such as the attack against TLS). In this case, attacks are still detected with high probability. If he transmits at a uniform rate, which has to be close to R pps, he continuously provides the system with information about his location. Packets transmitted close in time generate similar signalprints, allowing the system to track his location if a localization system is available. To avoid being tracked, an attacker needs to alternate periods of packet transmissions and radio silence. During such transmission bursts, however, he needs to send packets at rates *higher* than R pps in order to compensate for the periods of silence. This attack would be also detected because signalprints generated during each burst should match each other with high probability. However, *tracking* the attacker becomes more challenging because these bursts produce location estimates that are further apart.

6.5 Directional and Beamforming Antennas

A single directional or beamforming antenna would be more helpful to an attacker implementing a resource depletion attack than a masquerading attack. In a masquerading attack, it is still hard for an attacker to clone the exact signalprint produced by his intended victim from a large

distance. In close range, an omni-directional transmitter would also be effective while being easier to conceal. During resource depletion attacks, changing the transmission beam allows an attacker to change his signalprint, which decreases the number of matching requests. The probability of detection depends on the number of distinct patterns a transmitter is able to create and the packet rate required by the attack. If an attacker is only able to produce a small number of antenna patterns and an attack requires high packer rates (tens or hundreds of packets per second), some of the signalprints produced are still associated with a large number of requests, allowing detection with high probability.

7. LIMITATIONS

Due to the use of RSSI levels to characterize wireless clients, one inherent limitation of our mechanism is that it may be unable to distinguish two devices located physically close to each other. Masquerading attempts can be detected if there is a noticeable difference in RSSI with respect to at least one access point. As shown in section 6.3.2, this happens even for some locations in close range, possibly due to obstacles that affect one location more than the other. In some situations – such as multiple clients in a conference room – the system may not have compelling evidence that packets are coming from different devices, making masquerading attacks possible. The level of physical security in an installation dictates whether these attacks can be mounted: compared to a cafeteria, it is harder for an attacker in an enterprise building to get close enough to his victim to mount an undetected masquerading attack.

Our mechanism may also not be able to detect DoS attacks composed of few packets. The more packets are involved in an attack, the more signalprints are available for processing and the higher the probability of detection. A single-packet deauthentication attack in a 802.11 network may go unnoticed – for example if APs are sensing other channels – or not provide enough confidence as to raise an alarm. In most situations, however, attacks require high packet rates to be effective, increasing chances of detection.

An attacker may be able to avoid detection if provided with multiple antennas ($A > 1$). Suppose that an attacker configures its antennas so that each sensor can only listen to transmissions from a single antenna (e.g. using directional antennas with narrow beamwidth values). To successfully mount a resource depletion attack, the attacker can simultaneously transmit a different packet through each antenna. As a single sensor detects each transmission, the signalprints produced are too short to satisfy the rules presented in section 5.1. To mount a masquerading attack, the attacker simultaneously transmits the same packet using all antennas. By choosing the proper transmission power level for each of them, he is able to “compose” any arbitrary signalprint with A values. In both scenarios, attacks would be detected if some of the packets – even a small fraction – were detected by multiple access points.

8. RELATED WORK

Bellardo and Savage have shown that effective DoS attacks in 802.11 networks can be mounted with standard hardware [5]. They measured the impact of several identity-based attacks, including the ones targeting authentication and association services, and presented practical solutions

that can be realized with low overhead and without modifying clients. For example, the authors suggest that access points buffer deauthentication and disassociation requests for brief periods of time (5-10 seconds) before processing them. In this case, conflicting requests would be taken as indications of an attack.

Concurrently to our work, Demirbas *et al.* have proposed the use of RSSI measurements from multiple sensors to detect sybil attacks in wireless sensor networks, where a node uses multiple identities [8]. As testbed, the authors use up to four Mica2 motes operating as sensors at 433 MHz, with motes always located in close proximity to each other (30 cm to 10 m). Our research demonstrates that reliable attack detection is possible for larger 802.11 installations, where clients can be more than 40 meters from access points.

A technique called RF fingerprinting (RFF) has been developed to identify distinct transceivers across multiple wireless systems [22, 9]. The fingerprint for a transmitter is created from several features (such as phase, and amplitude) extracted from a period of transient behavior that occurs as the device powers up before a transmission. These turn-on transients are different for each transceiver, allowing even units build on the same factory to be distinguished. RFF systems have been used to detect cloned phones in cellular systems [19], and several researchers have proposed their use in wireless LANs [12, 23]. One disadvantage of RFF is that it requires specialized hardware to measure the signal properties needed with enough precision.

Gruteser *et al.* have proposed the use of temporary interface identifiers to improve privacy in WLANs: clients change their MAC addresses whenever they associate with an access point, reducing the chances of being tracked [10]. The authors evaluate this mechanism against an attacker that uses signal strength information to identify MAC addresses used by the same client. Our research extends this analysis to show that with higher number of access points, attackers may be able to track clients even after address changes, unless the number of active devices in the network is large enough as to create multiple similar signalprints.

Mechanisms such as client puzzles have been designed to *slow down* attack sources, reducing the damages caused by resource depletion attacks [15, 7, 3]. Before any resource is committed to an incoming request, computational puzzles are sent back to clients that require CPU- or memory-intensive operations. Despite being protocol-agnostic, puzzles demand that both clients and servers be modified, increasing deployment overhead when compared to a signalprint-based mechanism, implemented solely at the WA.

Our work also relates to localization algorithms, from pioneer systems such as RADAR [4] and SpotON [14], to more recent approaches that use probabilistic techniques, including the work of Roos *et al.* [20], Ladd *et al.* [17], Haeberlen *et al.* [11], and the Horus system [24]. By achieving average localization errors below 3 meters these systems have demonstrated that signalprints are strongly correlated with the location of a wireless client. Moreover, they can be used to complement signalprint-based mechanisms with localization services: when an attack is detected, the corresponding signalprint can be used as input to such systems so the location of the offending device can be determined. Tao *et al.* have in fact used differential signal strength values to make localization services more robust against variations in transmission power [21].

9. CONCLUSION

In this paper we showed that reliable client identifiers, which we call signalprints, can be created using signal strength measurements reported by access points acting as sensors. We showed that while malicious clients can lie about their MAC addresses, the signalprints they produce are strongly correlated with their physical location. We demonstrated that by tagging packets with their signalprints and crafting proper matching rules, a wireless network is able to detect a large class of effective denial-of-service attacks based on MAC address spoofing. We presented several examples of attacks that can be easily mounted in IEEE 802.11 networks and that can be detected by our proposed mechanism with high probability.

Measurements in our network testbed demonstrate that multiple packets transmitted by a stationary device produce similar signalprints with high probability. In our test dataset, most RSSI variations for a stationary client with respect to a single access point are small, within 5 dB from the median signal strength level. This allows the network to detect resource depletion attacks, in which a malicious device transmits high rates of packets (e.g. DHCP or authentication requests) containing random, forged MAC addresses. We presented matching rules able to detect that a large percentage of these packets were indeed generated by a single device, despite the different MAC addresses.

We also showed that similar signalprints are mostly found in close proximity. First, using 6 of our deployed access points, we showed that locations that produce signalprints with multiple similar RSSI values tend to be within 5 meters from each other. Then we showed that large RSSI differences provide strong evidence that packets were generated by distinct devices. Consequently, an attacker needs to be *physically* close to his intended victim in order to mount undetected masquerading attacks.

Overall, we showed that signalprints are tags that allow a wireless network to identify mobile devices according their physical location, improving security in a cost-effective manner. Although signalprints can be defeated, such as by the use of multiple synchronized direction antennas, these situations present a challenge for an intruder and increase the likelihood of detection by physical security measurements. Thus, like the use of fingerprints to identify humans, the mechanism is not infallible but a significant improvement over just believing the identity that the individual claims.

10. REFERENCES

- [1] LAN MAN Standards Committee of the IEEE Computer Society. Standard for Port based Network Access Control. Technical Report Draft P802.1X/D11, IEEE Computer Society, Mar. 2001.
- [2] LAN MAN Standards Committee of the IEEE Computer Society. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 6: Medium Access Control (MAC) Security Enhancements. Technical Report 2004 Edition, IEEE Std 802.11i, July 2004.
- [3] M. Abadi, M. Burrows, and T. Wobber. Moderately Hard, Memory-Bound Functions. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, San Diego, USA, Feb. 2003.
- [4] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-Based User Location and Tracking System. In *Proc. of IEEE INFOCOM*, Tel-Aviv, Israel, Mar. 2000.
- [5] J. Bellardo and S. Savage. 802.11 Denial-of-Service Attacks: Real Vulnerabilities and Practical Solutions. In *Proceedings of the USENIX Security Symposium*, Washington, DC, USA, Aug. 2003.
- [6] P. Calhoun, M. Montemurro, and D. Stanley. CAPWAP Protocol Specification. IETF Internet Draft, `draft-ietf-capwap-protocol-specification-01`, May 2006.
- [7] D. Dean and A. Stubblefield. Using Client Puzzles to Protect TLS. In *Proceedings of the Tenth USENIX Security Symposium*, Washington, DC, USA, Aug. 2001.
- [8] M. Demirbas and Y. Song. An RSSI-based Scheme for Sybil Attack Detection in Wireless Sensor Networks. In *Proc. of International Workshop on Advanced Experimental Activities on Wireless Networks and Systems*, June 2006.
- [9] K. J. Ellis and N. Serinken. Characteristics of Radio Transmitter Fingerprints. *Radio Science*, 36:585-598, 2001.
- [10] M. Gruteser and D. Grunwald. Enhancing Location Privacy in Wireless LAN Through Disposable Interface Identifiers: A Quantitative Analysis. *Mobile Networks and Applications*, 10(3):315-325, June 2005.
- [11] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. Kavraki. Practical Robust Localization over Large-Scale 802.11 Wireless Networks. In *Proc. of ACM MobiCom*, Philadelphia, PA, Sept. 2004.
- [12] J. Hall, M. Barbeau, and E. Kranakis. Enhancing Intrusion Detection in Wireless Networks Using Radio Frequency Fingerprinting. In *Proc. of The IASTED Conference on Communications, Internet and Information Technology*, Nov. 2004.
- [13] H. Hashemi. The Indoor Radio Propagation Channel. *Proceedings of IEE*, 81(7):943-968, July 1993.
- [14] J. Hightower, R. Want, and G. Borriello. SpotON: An Indoor 3D Location Sensing Technology Based on RF Signal Strength. Technical Report UW CSE 2000-02-02, University of Washington, Feb. 2000.
- [15] A. Juels and J. Brainard. Client Puzzles: A Cryptographic Defense Against Connection Depletion Attacks. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, pages 151-165, San Diego, USA, Feb. 1999.
- [16] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *Proc. of ACM MobiCom*, pages 107-118, Atlanta, GA, Sept. 2002.
- [17] A. M. Ladd, K. E. Bekris, A. Rudys, G. Marceau, L. E. Kavraki, and D. S. Wallach. Robotics-Based Location Sensing using Wireless Ethernet. In *Proc. of ACM MobiCom*, Atlanta, GA, USA, Sept. 2002.
- [18] T. S. Rappaport. *Wireless Communications - Principles and Practice*. Prentice Hall PTR, 2nd edition, Jan. 2002.
- [19] M. J. Riezenman. Cellular security: better, but foes still lurk. *IEEE Spectrum*, 37(6):39-42, June 2000.
- [20] T. Roos, P. Myllymäki, H. Tirri, P. Misikangas, and J. Sievänen. A Probabilistic Approach to WLAN User Location Estimation. *International Journal of Wireless Information Networks*, 9(3):155-164, July 2002.
- [21] P. Tao, A. Rudys, A. Ladd, and D. S. Wallach. Wireless LAN Location-Sensing for Security Applications. In *Proc. of the Second ACM Workshop on Wireless Security (WiSe'03)*, pages 11-20, Sept. 2003.
- [22] O. Ureten and N. Serinken. Detection of Radio Transmitter Turn-On Transients. *Electronic Letters*, 35(23):1996-1997, Nov. 1999.
- [23] O. Ureten and N. Serinken. Bayesian Detection of Wi-Fi Transmitter RF Fingerprints. *Electronic Letters*, 41(6):373-374, Mar. 2006.
- [24] M. Youssef and A. Agrawala. The Horus WLAN Location Determination System. In *Proc. of ACM/USENIX Mobisys*, Seattle, WA, June 2005.