Attacks on Public WLAN-based Positioning Systems

Nils Ole Tippenhauer, Kasper Bonne Rasmussen, Christina Pöpper, and Srdjan Čapkun
Department of Computer Science, ETH Zurich
8092 Zurich, Switzerland
{tinils, kasperr, poepperc, capkuns}@inf.ethz.ch

ABSTRACT

In this work, we study the security of public WLAN-based positioning systems. Specifically, we investigate the Skyhook positioning system, available on PCs and used on a number of mobile platforms, including Apple's iPod touch and iPhone. By implementing and analyzing several kinds of attacks, we demonstrate that this system is vulnerable to location spoofing and location database manipulation. In both, the attacker can arbitrarily change the result of the localization at the victim device, by either impersonating remote infrastructure or by tampering with the service database. Our attacks can easily be replicated and we conjecture that—without appropriate countermeasures—public WLAN-based positioning should therefore be used with caution in safety-critical contexts. We further discuss several approaches for securing WLAN-based positioning systems.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design —*Distributed networks, Wireless communication*; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms

Security

Keywords

Public WLAN localization, Localization Attacks

1. INTRODUCTION

In the last decade, researchers have proposed a number of WLAN positioning techniques for (local area) wireless networks [7,12,23,54]. The applications of these techniques are broad and range from improving networking functions (i.e., position-based routing) to enabling location-related applications (e.g., access control and data harvesting).

WLAN positioning systems are now being commercialized and are being used as a substitution and/or complement to the Global Positioning System [19]. One such system is the Wi-Fi positioning system (WPS) from Skyhook [6], available for PCs (as a plug-in) and on a number of mobile platforms, including the Apple iPod touch and iPhone [1] as well as Nokia mobile phones based on Symbian [5]. The resulting position can also be used by other services, such as the CyberAngel Security and Recovery System [2]. The Skyhook WPS relies on existing WLAN access points for localization of devices that have 802.11a/b/g wireless interfaces. In WPS, a mobile device collects information about all visible WLAN access points in its vicinity, sends this information to the Skyhook location database which replies with a position estimate based on the aggregated information. The position estimate can then be directly used by a mapping application like Google maps or can be combined with other sources of location information, such as those from GSM stations or GPS. Positioning systems by Mexens [31] and the Fraunhofer institute [17] have a similar mode of operation. We call these systems *public WLAN-based positioning systems*, since they rely on public WLAN access points which are not under control of the service operator that provides the positioning service.

In this work, we analyze the security of public WLAN-based positioning systems. Using the example of the Skyhook WPS, we demonstrate that such positioning systems are vulnerable to location-spoofing attacks: by jamming and replaying localization signals, an attacker can convince a device that it is at a position which is different from its actual physical position. Public WLAN-based positioning systems also rely on large databases that contain information about the position of the infrastructure. This information is often gathered by using the data reported by the users—either manually or automated during the positioning process. We show that this basic principle makes the Skyhook WPS vulnerable to database manipulation attacks, which can equally be used for location spoofing. We further discuss possible approaches for securing public positioning systems and show their potential advantages and drawbacks, given the constraints of the application scenarios in which they are used.

By performing these attacks, we demonstrate the limitations of Skyhook and similar positioning systems, in terms of the guarantees that they provide and the applications that they can be used for. Given the relative simplicity of the attacks and the availability of the equipment used to perform the attacks, we conclude that, without appropriate modifications, these positioning systems cannot be used in security- and safety-critical applications.

To the best of our knowledge, this work is the first that an-

MobiSys'09, June 22–25, 2009, Kraków, Poland.

alyzes the security of public WLAN positioning systems and the first that demonstrates the implementation of location-spoofing attacks in WLAN networks. Equally, we are unaware of any prior work that discusses location database manipulation attacks.

The structure of the paper is as follows. We give background information on public WLAN-based positioning in Section 2. We describe location-spoofing attacks in Section 3 and database manipulation attacks in Section 4. Solutions for securing public WLAN-based positioning systems are discussed in Section 5. In Section 6, we describe related work and we conclude the paper in Section 7.

2. BACKGROUND: WLAN-BASED POSITIONING SYSTEMS

WLAN positioning systems include range-based systems, which rely on RSS (Received Signal Strength) measurements, and range-free systems, which rely on the presence of localization beacons. Both types of systems make use of WLAN access points (AP) as localization stations which typically broadcast service announcement beacons from fixed known locations. Based on the reception of these beacons, devices compute their positions. In range-based systems [7,18,20,36, 39], the localized node (LN) records the signals received from access points, measures their RSS values, converts them into ranges, and estimates its own positions using the measured ranges. In range-free systems, the LN registers which APs are in its reception range and, based on this information, estimates its position. In most proposed WLAN-based positioning systems, the APs are controlled by the authority that operates the system. Typically, range-free and range-based WLAN positioning systems achieve a positioning accuracy in the order of meters, and if they rely on location fingerprinting, they can achieve accuracy in the order of tens of centimeters [7, 36, 39].

Skyhook's Wi-Fi Positioning System (WPS) is a metropolitan area public positioning system; it is a software-only system and requires a LN solely to have a WLAN-capable card and an Internet connection. Skyhook's WPS differs from existing WLAN positioning systems in that it does not maintain its own AP infrastructure; instead, it relies on the existing commercial, public, and private access points. In WPS, the operator (Skyhook) creates a location lookup table (LLT), which contains data samples taken from different locations. To develop and extend its LLT, the operator is sending vehicles with GPS and roof-mounted antennas through urban and suburban areas to scan the present APs. For each location, the Medium Access Control (MAC) addresses of all visible access points are stored. This lookup table can then be queried by the software on the LN. Since obtaining information about available access points in an area can be a work-intensive process (that needs to be constantly updated due to the dynamics of the WLAN networks), WPS also allows users to enter (on-line) locations and MAC addresses of their access points and of access points that they observe in their vicinity. As we will discuss later, Skyhook also leverages information obtained from location requests to update its WLAN location database (LLT).

WPS localization can be divided into five phases, as shown in Figure 1. In phase 1, the LN scans all (802.11a/b/g) WLAN channels for access points by broadcasting a probe request frame on all channels.[1] In phase 2, the APs in range reply to the LN with network announcement beacon frames containing, among other parameters, their MAC addresses. After having detected these beacons and recorded their corresponding signal strengths, the LN sends the identified MACs over an encrypted channel to the Skyhook LLT (phase 3); this step requires that the LN has an (internet) connection to the service provider. In phase 4, the server compares the reported MACs to the data stored in the LLT and returns the locations of the access points to the LN, again in encrypted form. In phase 5, the LN computes its position based on the received access point location information using a non-disclosed algorithm[2].

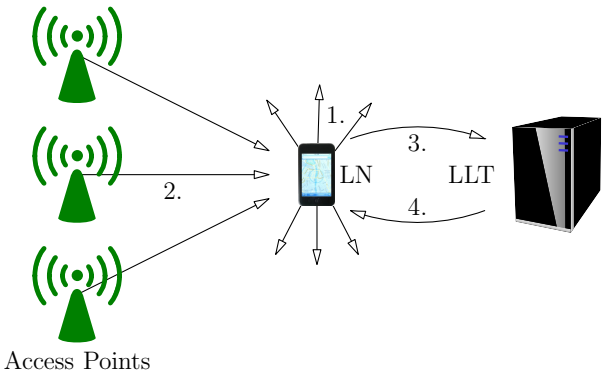

Figure 1: The Skyhook localization process. 1. The LN broadcasts a probe request frame. 2. APs reply with a response beacon frame. 3. The LN queries the LLT server. 4. The server returns location data about the observed APs. 5. The LN computes its location.

Note that by sending information about its neighboring access points to the WPS database, the LN also allows Skyhook to update its database.

This description of WPS is based on our experiments. According to the tests we performed, other factors such as the received signal strength of the individual AP beacons do not seem to influence the localization result.

3. LOCATION SPOOFING

In this section, we analyze the security of the Skyhook WPS and we show attacks on its positioning service. We demonstrate that the Skyhook WPS is vulnerable to attacks in which signal insertions, replays, and/or jamming allow an attacker to either prevent the localization or to convince a device that it is at a position which is different from its actual physical position (location spoofing). Our attacks are composed of two actions: (1) impersonation of access points (from one location to another) and (2) elimination of signals sent by legitimate access points. Since rogue access points can forge their MAC addresses and can transmit at arbitrary power levels within their physical capabilities, access point impersonation can be easily done in WPS (see Section 3.2). Equally, since WLAN signals are easy to jam (see

[1]In our experiments, we found that some devices performed active scanning while others only collected beacons passively. If only passive scanning is used, phase 1 is redundant.

[2]We were not able to find the description of the Skyhook's position computation algorithm in the open literature.

Figure 2: Equipment used in our experiment. The iPod and iPhone devices in front are being located, the laptops are used to impersonate access points (APs), and the software radios on the left are used to jam legitimate APs.

Section 3.3), signals from legitimate access points can be eliminated, thus enabling location spoofing.

In what follows we will demonstrate location-spoofing attacks in three scenarios: (i) the LN is not in the range of legitimate APs (AP impersonation), (ii) the LN is in the range of legitimate APs and uses only WLAN-based localization (AP replacement) and (iii) the LN is in the range of legitimate APs and uses a hybrid WLAN/GSM-based localization system. Further we will show that the same attacks can be performed on Skyhook's Loki browser plug-in on a standard PC (see Section 3.4).

3.1 Equipment

In our experiments, we used the following equipment. As positioning devices, we used Apple's iPhone and iPod touch [1] devices with the WPS-enabled Google maps application as well as a laptop with an installed Skyhook's Loki browser plug-in [6]. The iPhone was a first generation, original model without GPS support (OS version 1.1.4, model MB384LL, firmware 04.04.05_G), the iPod model was MA632ZD. To perform the attacks, the attacker both needs devices that impersonate legitimate access points and devices that eliminate legitimate access point signals. For access point impersonation, we used two laptops running Ubuntu Linux configured as wireless access points (using the Scapy packet manipulation program, v. 1.2); the laptops were transmitting on channels which were not occupied by existing access points, and they were configured such that they could modify the MAC addresses of their Wi-Fi interfaces, their network names (SSID), and signal strengths. To eliminate signals from legitimate access points we jammed these signals using a software radio platform (USRP Rev. 4.2 [15] with daughterboards for the 2.4 GHz band (FLEX2400 2-6-2006), operated by Gnuradio 3.0). Our equipment is displayed in Figure 2.

3.2 AP Impersonation

First, we performed an attack which we call *access point impersonation* attack. The idea of an AP impersonation attack is to report remote access points to the attacked device, which will then compute a location that is in the proximity

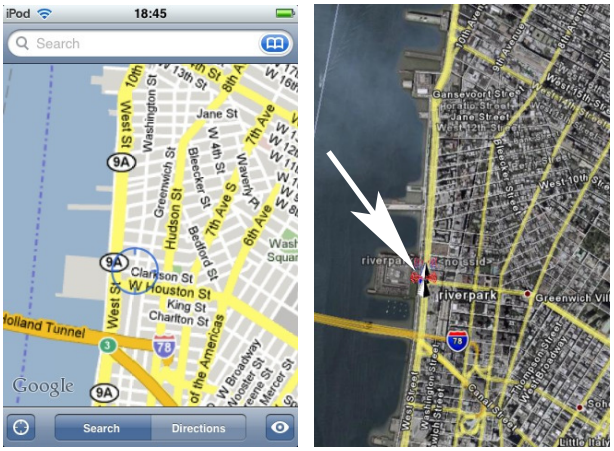

(a) (b)

Figure 3: Location-spoofing attack. (a) Location in New York City (circle in the center) displayed by an iPod physically located in Europe (caused by an AP impersonation/replacement attack) (b) Physical location of the impersonated APs (indicated by the arrow and displayed using Google Earth [4]).

of the remote APs. This attack exploits the fact that WPS localization relies on MAC addresses for the identification of APs. AP MAC addresses are public since they are contained in the network announcement beacons and can thus be replayed.

To execute the AP impersonation attack, we used a Laptop with a WLAN card, running a purpose-written program to impersonate APs. Our program waits for probe requests sent by the LN (iPhone or iPod) and replies to these requests with custom-made beacon responses that correspond to the beacon responses from the impersonated remote APs. Each beacon response $\hat{r}_i$ contains a MAC_i address that equals the MAC address of the spoofed network i; the beacon also contains an SSID_i which is not necessarily equal to the one of the spoofed network. We note that network SSIDs are not used by the WPS to identify the APs, but they helped us to distinguish the impersonated from legitimate APs. All other parameters in the beacon responses were set to their default values (e.g., signal strength was set to 17 dBm). This setup enabled us to impersonate an almost arbitrary number of access points in the vicinity of the LN.

In our experiment, we chose to impersonate four geographically mutually close access points located in New York City, and we set their SSIDs to: NY_1, NY_2, NY_3, NY_4. In order to find the MAC addresses corresponding to these access points, we used the WiGLE database [57], which provides information about worldwide wireless networks.

We first performed the experiment in an environment without WLAN coverage, i.e., no legitimate APs were visible to the iPod at its physical location at the time of localization. By impersonating the APs as described, the localization process on the iPod returned a location in New York City, while the device was physically located in Europe; this is shown in Figure 3. The displayed location was close to the position of the spoofed access points. We then successfully performed a more fine-grained spoofing attack and modified the displayed location of an iPod such that it displayed a position in the city center of our city, approx. 1 km from its

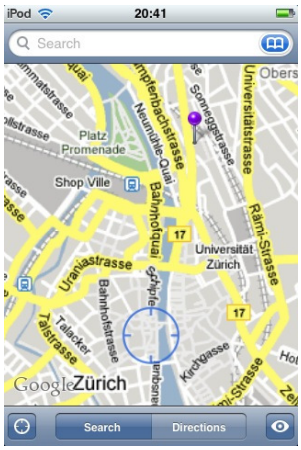

Figure 4: Location-spoofing attack. Location displayed by the iPod in city downtown (marked by a circle) at about 1 km distance from the iPod's actual position (marked by a pin) at university campus (caused by an AP impersonation/replacement attack).

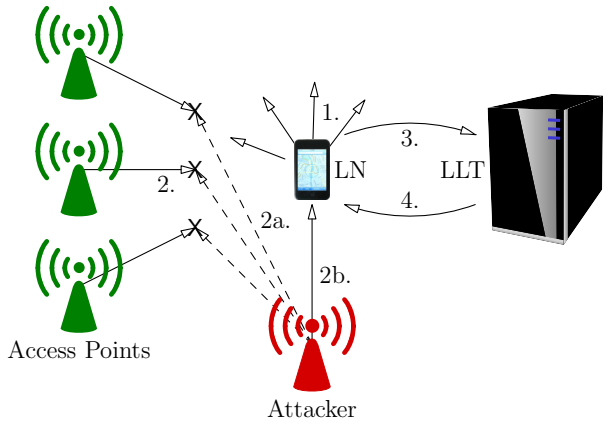

Figure 5: AP replacement attack. The beacons of legitimate APs are jammed (2a) and the attacker sends spoofed beacon responses to the LN (2b). The LN processes the spoofed beacons instead (3), as they pretend to be legitimate APs from a different location. As a result, the LLT will return location data for the remote position (4).

actual position (at the university campus). This is shown on Figure 4. Beyond succeeding in performing the AP impersonation attack, we learned that at least some of the access points that we impersonated are registered in the Skyhook database (LLT).

We then performed the same attack in an area covered by public APs. We impersonated wireless networks with MAC addresses of access points that are contained in the Skyhook database, but are located far (New York) from the actual physical location of the device. As a result, the WPS algorithm failed since the LN (iPod) registered access points at locations which are physically too far apart and was thus not able to resolve its own position. Although, in this scenario, the location-spoofing attack failed, it unveiled a simple denial-of-service (DoS) attack on WPS localization. To perform this DoS attack, the attacker only needs to impersonate an AP which is contained in the Skyhook database and is located far from the actual physical location of the localized device.

In the following section, we show how to successfully spoof a location of a device even if it resides in an area covered by public (legitimate) APs.

3.3 AP Replacement

Nowadays, most urban and suburban regions as well as other popular areas are covered by a large number of legitimate APs and will—with increasing probability—be categorized by Skyhook. In order for the AP impersonation attack to succeed despite the presence of known APs, we need to eliminate the announcement beacons sent by the legitimate APs and replace them with our impersonated AP beacons. We call this an *AP replacement* attack and consider it as a more comprehensive form of the AP impersonation attack.

The idea behind the AP replacement attack is shown in Figure 5. In this attack, we use standard wireless tools to detect the channels on which the legitimate APs are transmitting beacons and then launch a physical-layer jamming attack to disable the reception of those beacons on the identified channels. The jamming is not noticed by the user because the simple user interface on the iPod does not provide enough information to detect ongoing jamming. Simultaneously, we insert signals from impersonated APs on non-jammed channels.

In our attack setup, the legitimate access points were transmitting on WLAN channels 6, 10, and 11 (up to 13 channels are available in 802.11b/g). We used software-defined radios (USRPs [3], Figure 2) to emit uniform noise on those channels, which blocked the communication between legitimate APs and the iPod. By using physical-layer jamming, we had full control over the transmission power and bandwidth of the jamming signals and could easily elude the 802.11 protocol standard. We then announced the impersonated networks using channel 2. The localization results on the iPhone resembled the ones of the AP impersonation attack in Section 3.2. Equally the AP replacement caused the iPod to report an incorrect, attacker-chosen location (in New York or in city downtown, as shown in Figures 3(a) and 4).

If an attacker targets devices that use passive network scanning, a more stealthy attack can be performed by jamming according to the AP beacon schedules (this was however not reasonable for our attacks on the iPhone due to the active network scanning mode). Given sufficiently fast hardware, the attacker can also jam the beacons reactively (for passive and active network scanning); in this case the jamming device first senses for ongoing beacon transmissions (typically a beacon frame has a length of approx. 100 byte, taking about 10 to 100 μs) and then jams them in a targeted (or selective) way after their detection.

Instead of physical-layer jamming, an attacker could also use MAC layer jamming [58] or signal overshadowing, which would likewise eliminate the legitimate AP signals, but would still allow the insertion of fake beacons even if all non-overlapping frequency ranges in 802.11b/g are used for legitimate beacon transmissions. Using MAC layer jamming, the attacker can prevent the APs from sending beacons by keeping the channel busy all the time.

(a) (b)

Figure 6: Attack on iPhone WLAN/GSM localization. (a) Location displayed by the iPhone in the city center (marked by a circle) at about 1 km distance from the iPhone's actual position (marked by a pin) at university campus (caused by an AP impersonation/replacement attack). (b) Location displayed by the iPhone when spoofing failed. The attacker's target location was in New York City, far from the location that the iPhone's GSM localization computed (in Europe). The attack thus failed and the iPhone displayed the location computed using GSM localization (marked by a circle).

3.4 Location Spoofing Attacks on Hybrid WLAN/GSM Localization Systems

The attacks described in Sections 3.2 and 3.3 were performed on an iPod touch device. Regarding localization, the iPhone differs from the iPod in the sense that it applies a hybrid localization technique that combines WLAN and GSM base station localization. On the iPhone, GSM-based localization provides a rough position estimate, while WLAN localization provides the device with a fine-grained position estimate.

GSM, the second source of location information, can be used to detect displaced (impersonated) locations in the WPS process and enables the devices to fall back to GSM localization. As our experiments showed, if the position that the device computes using WPS is too far away from the position that it obtains using GSM information, the iPhone will only display the position computed using GSM (i.e., it ignores the WPS position). This is shown on Figure 6b.

However, since the GSM localization is significantly less accurate than WPS localization, using the attack described in Section 3.3, we were still able to displace the iPhone within its GSM localization accuracy (in our test, within 1 km distance). The result of this attack can be seen in Figure 6a; the figure shows the real physical location of the device (marked by a pin) and the location displayed by the iPhone (marked by a circle). This result is similar to the one obtained for iPod location spoofing (Figure 4).

In order to spoof the position of an iPhone to different cities or countries, the attacker either needs to disable the GSM signal reception (e.g., using widely available GSM jammers) or spoof GSM base stations, which has been shown to be feasible for GSM [33] and even for UMTS networks [32]. Spoofing a GSM base station is more complex than impersonating a WLAN AP because it involves a functional protocol interaction with the mobile device. Nevertheless, in GSM, base stations do not authenticate to the user and are therefore inherently susceptible to impersonation attacks. Attacks using GSM base station spoofing or jamming were not part of our experiments.

3.5 PC Location Spoofing

We further tested the same attacks on a Laptop with an installation of Skyhook's Loki browser plug-in [6]. This plug-in is installed into the browser like a toolbar and is able to provide web sites with location information. We repeated the above described location-spoofing attack (AP impersonation and replacement) and the results we got on the Laptop were identical to the ones reported by iPod touch.

Consequently, if users rely on Loki to provide their web applications with location information, this can be misused by the attacker and possibly lead to a wider system or data compromise, since the attacker can fully control the location that Loki provides to the application.

One can further imagine a web service that provides information to the user only if the user is at a given location. Given the described attacks, Loki cannot be used for location verification, since even the user could spoof his own location to get access to the service (e.g., pretending to be in New York, while being in Europe). Loki results could equally be modified in a number of other ways, including the manipulation of input that the plug-in gets from the networking interfaces.

4. LOCATION DATABASE MANIPULATION

AP location/MAC data enters the Skyhook database in one of three ways: (1) The database is extended and updated by vehicle-based signal scanning and data collection performed by the company, (2) new access points can be inserted manually online (by users and nonusers of the service), and (3) Skyhook incorporates data that was submitted by users in localization queries in order to improve the accuracy of its reference database. As we show in what follows, Skyhook's WPS database (LLT) is not resistant to targeted *location database manipulation attacks*, although Skyhook tries to counteract this threat by applying error-detection and error-correction methods (surveying the age and consistency of data and executing periodic rescans of outdated areas [6]).

In database manipulation attacks, the attacker tries to actively interfere with the database underlying the localization process by inserting wrong data and/or by modifying existing entries, e.g., by changing the recorded positions of access points to remote positions. Consequently, the attacker may not only change the result of an individual localization, but influence many localizations that all use the common database (in this case, the LLT).

4.1 Injection of False Data

Here we show how we successfully inserted false location information for an access point into the Skyhook LLT. During the location-spoofing experiments described in Section 3, access point AP_k was used to provide internet access over

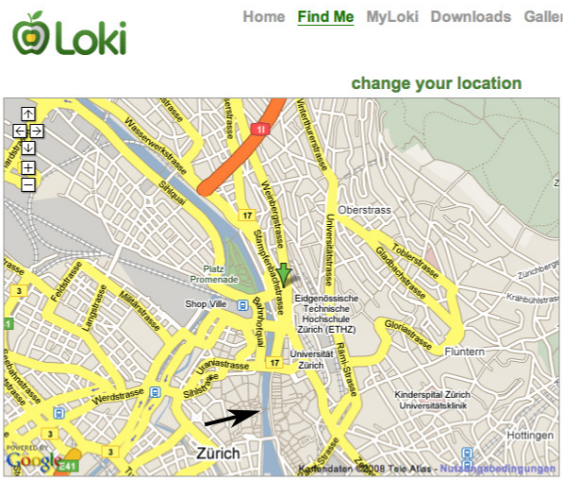

Figure 7: Data corruption in the LLT (using Loki). The lower (black) arrow displays the original location D of AP_k in the LLT. The upper (green) arrow shows the new location U, after impersonating an access point with AP_k's MAC.

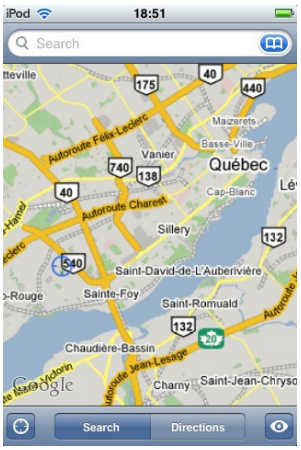

Figure 8: Reverse AP location lookup in the LLT. Using a spoofed access point and modifying its MAC address resulted in a localization in Québec, Canada, though being physically located in Europe.

WLAN to the iPod in order to enable the communication between the iPod and Skyhook's database. The MAC of AP_k was, before the experiments, unknown to Skyhook.

When we spoofed the location of the iPod (located at the university campus at location U), the iPod reported to Skyhook not only the impersonated APs (from location D at city downtown), but also the MAC of the AP_k (located at U) that it used for its internet connection. As a consequence, the MAC address of AP_k got added by WPS to the LLT with location D although being physically located approx. 1 km away (at U). Subsequent tests in which the iPod was connected only to AP_k showed that the iPod displayed a location at D, confirming that, in the LLT database, AP_k had a position in the city downtown.

This attack is a direct consequence of Skyhook's WPS mode of operation, in which the LLT database is not being updated only by company employees and on-line users, but also through the localization requests that users send when they want to determine their positions.

4.2 Corruption of Existing Data

A more severe form of the data manipulation attack is the virtual relocation of access points that have already been categorized in the database. To demonstrate this, we used the access point AP_k that was previously injected at the location in city downtown (position D). To change the location of AP_k in LLT, we set up a second access point using AP_k's MAC at our university campus (position U) and kept it active over a longer time (several days). All localizations executed by users in this area were now also submitting AP_k's MAC when localizing themselves—however this time for the new position U. As a consequence, Skyhook started to resolve the old location of AP_k to the new position U, assuming that the majority of the reported MAC addresses define the correct location. Consequently, the access point AP_k was moved in the LLT database from location D to its new position U, as shown in Figure 7. As before, we verified this result by localizing a device using only AP_k.

Because this database manipulation attack is conducted involuntarily by all users using the WPS service in range of AP_k, this attack is hard to detect by the service provider (Skyhook). Although we acknowledge that this manipulation attack only succeeds in settings where the localization service is triggered (much) more often at the new location than at the old one, this attack represents a powerful threat to the correctness of the LLT database and may affect the results of the localization service for many users and at many locations.

4.3 Reverse Location Lookups

Using our attack equipment from Section 3, we were able to find the exact positions of APs recorded in the LLT with known MAC addresses but unknown positions. Although this does not manipulate the database contents, it still represents an undesired function of WPS revealing (confidential) positions of access points; here we assume that Skyhook wants to maintain the confidentiality of the AP MAC-AP location bindings.

We performed a *reverse AP location lookup* by changing the MAC address (switching two bits) of one of the spoofed access points in the downtown area which we used in the prior attacks. The position resulting from an iPod localization using this access point was Québec, Canada, as shown in Figure 8. Since manufacturers often assign MAC addresses linearly to their products, this also allows us, e.g., to look up the locations of access points from one production charge.

4.4 Conclusion

From the above analysis and attacks we conclude that WLAN localization systems that use the data originating from clients to update their database (LLT)—either explicitly by manual insertions or implicitly during the localization process—are susceptible to *database manipulation* attacks. Both attacks can be easily performed by the attacker, with the additional advantage that, in the latter case, the update requests are actually reported by legitimate, honest users. In Section 5, we will outline mechanisms that protect public (cooperative) WLAN localization systems from such database manipulations and/or that mitigate the attacks.

5. SECURING PUBLIC WLAN-BASED LOCALIZATION

In this section, we discuss solutions to protect the LN against location-spoofing attacks. Our design space includes solutions that are based on client, transmission, and service-provider mechanisms. The first countermeasure we present is based on client-side integrity checks. We then discuss secure data acquisition techniques; achieving authentication for the received signals is technically challenging. Finally, we propose techniques for thwarting database manipulation attacks by detecting and eliminating false data in the database.

5.1 Client-Based Integrity Checks

One approach to detect attacks is based on location history recordings on the LN. For each localization request by the user, the current position is computed as in the original WPS system. To detect displacement attacks, the resulting position is then compared to the latest stored position in the history record. This allows the detection of attacks that try to displace the LN over a distance that the LN is unlikely to cover within the given time. If the new location would require a speed exceeding a maximal average speed, an attack can be suspected. Additionally, a trace of multiple past locations can be examined to prevent the attacker from starting the attack before the real measurements are taken. Unfortunately, due to greatly varying speed of travel in typical urban areas (e.g., public transportation systems), the tolerated maximum average speed will have to be high. Therefore, the attacker can still displace the LN within a large area in such scenarios.

Using location history records does not require any hardware or software modifications in the APs or in the (Skyhook) WLAN-positioning system but can be implemented on the LN. However, in order to cover a wide range of attacks, automatically triggered (background) localization would be required so that the latest recordings are still recent enough.

5.2 Secure Data Acquisition

The first observation we make regarding secure data acquisition is that the localization beacons of WLAN access points can be easily forged and replayed. Traditional authentication mechanisms would not help much here because they would require software modifications at the access points, which are not under the control of the service provider (e.g., Skyhook). Furthermore, even if appropriate modifications would be made to the access points and AP beacons would be properly authenticated, public WLAN based localization systems would still be vulnerable to jamming and wormhole-based [24] signal relay attacks [28, 51]. To prevent relay (wormhole) attacks, the access points and the localized nodes (LNs) would therefore need to mutually authenticate their communication and would either need to be tightly time synchronized, or use challenge-response protocols with accurate (i.e., *ns*) time measurements [8]; both would require significant hardware and software modifications to both the access points and the LNs.

Given that such modifications of access points and LNs are not feasible in public WLAN positioning systems, authentication of access point beacons needs to be done in a manner that does not require any pre-shared cryptographic material between the APs and the LNs. For this, we propose to use unique AP characteristics such as their *traffic* or *signal fingerprints*. These fingerprints should be difficult to forge and easy to measure (by the LNs); if they are not, APs could be impersonated in the same manner as in the WPS system that uses (easily forgeable) MACs as device identifiers. Equally, traffic and signal fingerprints of APs need to be chosen such that they are unique or mutually distinguishable with high probability. Since access points do not cooperate in fingerprint extraction with the LNs or with the system provider, the fingerprints also need to be easily measurable by the provider to build the LLT database and by the LN for AP identification.

Assuming that such AP fingerprints can be measured by the LNs, our fingerprint-enhanced WPS would then work in the following manner. The service provider measures the fingerprint data of the APs using appropriate software and hardware, and stores it in the LLT along with the MAC addresses and RSS values for each access point; as a side effect, manual user input, which represents a source of false information, would be precluded and reverse location lookups could be prevented. During the localization, the fingerprint-enhanced LNs measure the fingerprint data from the surrounding APs and report this data along with the MAC addresses and signal strengths to the service provider that, in turn, compares it to the data in the LLT. Based on a probabilistic analysis, the service provider returns the location information which is then used by the LN to compute its position. If the analysis fails, i.e., if the fingerprint data does not correspond to the stored data up to a pre-defined degree, possibly indicating an attack, no location information is returned. Alternatively the most likely position could be returned along with a warning that the location could not be verified. This could then be presented to the user, e.g., as a red circle instead of a blue marking the found location.

Recently, a number of results have emerged that show how unique characteristics of WLAN access points and of other wireless devices can be measured. One way of identifying access points is by collecting data specific to their configuration or model. The feasibility of this approach was discussed in [9]. This approach does neither require hardware modification of the LN device nor changes on the scanned access points but instead relies on characteristic behavior of different AP models on malformed 802.11 frames. Although this does not completely prevent location-spoofing attacks, it makes them more difficult since the attacker has to extract AP device specific data in order to compose adequately forged response frames. This would require his prior physical presence at the access point whose location is to be spoofed. In [25], the authors use intra-device clock skews to differentiate individual devices on the internet. In [10, 13, 41, 48], the respective authors discuss signal fingerprints based on physical characteristics of individual device radios. While [41] and [13] focus on the fingerprinting of CC2420 and Chipcon 1000 (433MHz) wireless sensor motes, [48] and [10] demonstrate the successful fingerprinting of 802.11b WLAN network interface cards. Different distinction features may be extracted: e.g., unique transient characteristics [48] or unique timing behavior in the modulation (frequency magnitude, phase errors, I/Q origin offset, etc.) [10].

The process of collecting these fingerprints requires specialized hardware which would have to be added to the LN devices. Signal fingerprints would prevent attacks by attackers using off-the-shelf hardware, but they would not prevent

a sophisticated attack by an attacker that samples and replays the signals on the physical layer. Spoofing fingerprints based on physical characteristics of the transmitted signal requires a high frequency sampling oscilloscope with a sampling rate at least as high as the fingerprinting hardware in the LNs. In addition, an attacker would need an arbitrary waveform generator capable of reconstructing the sampled signal without adding any distortion or noise. This is very hard because the oscilloscope, waveform generator, and controlling computer all have finite dynamic ranges, i.e., they can only represent the captured signal in steps. These hardware requirements make signal fingerprints much harder to forge than behavioral fingerprints. The usability of the fingerprint-enhanced WPS stimulates further research on identifying non-forgeable and easily measurable AP fingerprints, in particular regarding fingerprinting stability with respect to mobility of the capturing device and to environmental effects such as multipath propagation and interference.

Another technique for detecting and preventing location spoofing in WPS is by geo-locating the IP address of the AP that is used to query the location database. Although this information is relatively coarse and can be spoofed using IP tunneling, it can be used to make simple location spoofing attacks at larger distances (e.g., to different countries) more difficult.

5.3 Mitigating Database Poisoning

We now discuss techniques for the mitigation of database poisoning and their implications on the system behavior. The risk of database poisoning based on user-supplied data (Section 4) can be mitigated in several ways. We distinguish temporal rules and update rules.

Temporal rules for system updates determine the reaction time of the LLT towards new or changed data. A system which updates its LLT immediately with new user data will be closest to mapping the real situation, but it also enables an attacker to more easily influence localization results in a targeted manner. In contrast, a system which only introduces a new AP once and will never update the respective location will be most secure against database attacks, but cannot reflect real-world changes. User-based database updates may follow different temporal rules based on their respective confidence levels. In other words, the more likely the correctness of a database update is, the quicker it should be represented in the database. Nevertheless, there remains a trade-off between database freshness and resistance against attacks.

Database update rules determine if a newly reported AP enters the LLT and/or if the information stored for it is modified (e.g., a new location is assigned); update rules may incorporate information from multiple users or nearby APs. New entries as reported by one or multiple users (either automatically during the localization or by manual insertion) may contradict existing entries—in such a case update rules are required that determine how the system reacts to them. We define a *contradicting location report* as a report that claims a location x for an access point AP_i while the active location y stored for AP_i in the database is further away from x than the transmission range allows. Without attackers and without AP repositioning, contradicting reports would never occur; instead, the reported APs would either confirm the data in the LLT or could be used for new

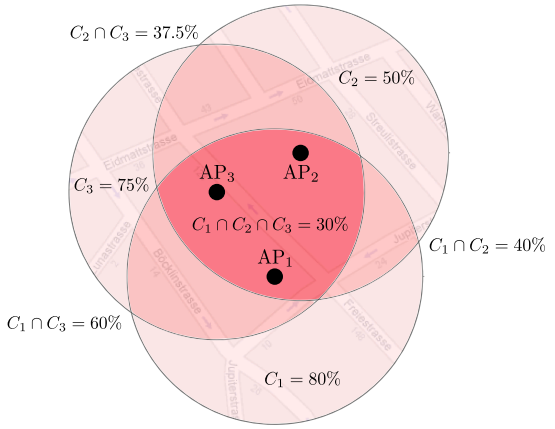

Figure 9: Intersection of transmission ranges for three access points. Each access point AP_i is associated with a confidence value C_i between zero and one. We propose to return a confidence indication together with the localized region. A darker color indicates a more precise location estimate but with lower confidence. In our example, the probability that the LN is located outside of the entire region is $\prod_i (1 - C_i) = 2.5\%$ (assuming the confidence values are independent).

LLT entries. However, given the possibility of legitimate AP repositioning, the LLT should remain consistent in the presence of attackers. As a solution, we propose to store confidence values together with the location data in the LLT; if a location x has the highest confidence value for AP_i then location x is *active*. An AP is relocated in the database if the confidence in location y exceeds the confidence for any other location; then y becomes active. The confidence values are updated based on majority updates and consistency checks; the former (*Maj*) incorporates majority user reports, the latter (*Con*) incorporates consistency with the neighboring APs. More precisely, typical database update rules include:

- *New*: Locations of APs collected by the service provider enter the database with the maximum confidence value, user-reported APs with the minimum confidence value.
- *Maj*: If, over a period of time (hours to few days), more user reports on access point AP_i occur for location y than for the currently active location x and, at the same time, the number of reports for x significantly drop, the confidence in y increases and the confidence in x decreases (majority update rule).
- *Con*: The confidence in a position y of access point AP_i increases if the frequency of localization and the localization pattern for AP_i are comparable to those of physically close APs with maximum confidence value, i.e., the reported environment close to y matches the system model in the database (consistency check rule).

This approach leads to consistent data in the LLT but does not entirely rule out database attacks (e.g., if an attacker physically moves an AP or if (sets of) isolated APs are relocated). Therefore, we propose that the confidence values should be part of the localization result (for all LLTs with user-supplied data).

In Figure 9, we give an example of the data returned to the client. This data could consist of the location of each observed access point along with confidence values. This will

enable the client to make a probabilistic map of the area and present it to the user, using an aggregation function of its choice. The example in Figure 9 uses a product-based aggregation function. The three observed access points have confidence values $C_1 = 0.8$, $C_2 = 0.5$, and $C_3 = 0.75$ respectively. That means that if the client chooses to only rely on the information from one AP, say AP_1, it can do so with a confidence of 80%. If the client requires a more precise location it will have to rely on more access points. Each access point has a confidence level for the correctness of its location in the LLT (i.e., that it was neither spoofed nor moved by the attacker), so the more access points the client utilizes the more precise its location estimate can be (e.g., within the center region) but the less confident it can be that the result is correct (30% in the example). We note that, in general settings and despite the countermeasures described above, the transmission ranges of the sensed APs may not necessarily be overlapping and the same AP might have multiple location entries in the LLT (if attacks are taking place). Hence, the locations of the sensed AP's may be remote or even ambiguous. The best the LN can do is to provide this conflicting information to the user in form of graphically (non-overlapping) localization regions including the varying confidence levels of the entries in the LLT.

During our experiments, we observed that APs are often abundant in urban environments. In addition, both commercial and private AP will typically rarely change their positions. Both assumptions suggest that the service provider can choose more stringent temporal rules without impacting the performance too much. Indeed, we argue that even if the provider chooses to never relocate an AP, the performance of the system in terms of localization precision will only be affected marginally. This would prevent database poisoning attacks on existing records. Nevertheless, an attacker could mount a denial of service (DoS)-like attack by registering MAC addresses of APs at arbitrary locations before they are detected by the service provider or reported by other users. But since it is rather hard for an attacker to predict the MAC address of a specific AP, this will not allow targeted DoS attacks.

6. RELATED WORK

In the last decade, a number of outdoor localization systems for mobile devices were proposed and implemented, based on satellite communication (GPS [19]), cellular networks (GSM), Wi-Fi networks or specialized platforms [7, 12, 16, 23, 40, 54, 55]. These techniques differ in terms of accuracy, reliability, and hardware requirement. Positioning techniques were also extended and used for positioning in wireless ad-hoc networks [11, 14, 34, 35, 44, 50].

Later security analysis has shown many of these systems or underlying technologies to be vulnerable to attacks [32,33, 37,38,56]. Proposals followed that aim at securing GPS [26]. Furthermore, several proposals have been made to improve the security in WLAN-based localization. However, these solutions require the cooperation of the APs [37–39, 46].

To secure systems not based on WLAN, several secure ranging and secure localization systems were proposed in the open literature. The first secure ranging protocol was described in [8]; this protocol was later applied to a wireless scenario and extended to provide mutual authentication in [49]. To allow more resource constrained devices to perform secure ranging in noisy environments, Hancke and Kuhn proposed an alternative protocol in [21]. This paper also discussed possible implementations of secure ranging in hardware. An authenticated ranging protocol for wireless devices was proposed in [53]. Attacks on possible implementations of secure ranging protocols were discussed in [22].

A system for secure localization was proposed in [43], based on ultrasonic and radio wireless communications; this system is limited by the use of ultrasonic signals, which requires that no attackers are present in the area of interest as demonstrated in [45]. Kuhn [26] proposes an asymmetric security mechanism for navigation signals, based on hidden message spreading codes. Čapkun and Hubaux [51, 52] propose a secure localization technique called verifiable multilateration, based on secure ranging, which further enables a local infrastructure to verify positions of the localized devices. Authenticated ranging and secure localization (verifiable multilateration) were implemented in [47]. In [53], Čapkun et al. propose a location verification scheme based on hidden and mobile base stations.

Lazos et al. [27] propose a technique for secure positioning of a network of sensors based on directional antennas. Lazos et al. [28] propose an extension of this technique that copes with jamming and replays of localization signals. In [42], the authors propose and implement a system for broadcast localization and time-synchronization, based on navigation signal encoding, that prevents signal replay and time-shift attacks. Li et al. [29] and Liu et al. [30] propose statistical methods for securing localization in wireless sensor networks.

To the best of our knowledge, this work is the first that analyzes the security of public WLAN positioning systems and the first that demonstrates the implementation of location-spoofing attacks in WLAN networks. Equally, we are unaware of any prior work that discusses location database manipulation attacks.

7. CONCLUSION

In this work, we studied the security of public WLAN-based positioning systems. Specifically, we investigated the Skyhook positioning system [6], available for PCs and used on a number of mobile platforms, including Apple's iPod touch and iPhone. We demonstrated that this system is vulnerable to location spoofing and location database manipulation attacks. By demonstrating these attacks, we showed the limitations of Skyhook and similar public WLAN-based positioning systems, in terms of the guarantees that they provide and the applications that they can be used for. Given the relative simplicity of the described attacks, we conclude that, without appropriate modifications, these positioning system cannot be used in security- and safety-critical applications. We further discussed approaches for securing public WLAN positioning systems based on client-side integrity checks, secure data acquisition, and the mitigation of database poisoning. We call for more research on the development of usable and secure public positioning solutions, based on, e.g., access point signal fingerprints and/or calibration of location databases.

8. ACKNOWLEDGMENTS

The work presented in this paper was in part supported by the Swiss National Science Foundation under Grant 200021-116444 as well as by the Zurich Information Security Center.

9. REFERENCES

[1] Apple Inc. http://www.apple.com.

[2] Cyberangel security and recovery system. http://www.skyhookwireless.com/press/skyhookcyberangel.php.

[3] GNU Radio: The gnu software radio. http://gnuradio.org/trac.

[4] Google earth. http://earth.google.com.

[5] Loki Mobile applet for Nokia phones using Symbian. http://loki.com/download/mobile.

[6] Skyhook, Inc. http://www.skyhookwireless.com.

[7] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-Based User Location and Tracking System. In *Proceedings of the IEEE Conference on Computer Communications (InfoCom)*, volume 2, 2000.

[8] S. Brands and D. Chaum. Distance-bounding protocols. In *Workshop on the Theory and Application of Cryptographic Techniques (EUROCRYPT)*. Springer, 1994.

[9] S. Bratus, C. Cornelius, D. Kotz, and D. Peebles. Active behavioral fingerprinting of wireless devices. In *Proceedings of the ACM Conference on Wireless Network Security (WiSec)*, 2008.

[10] V. Brik, S. Banerjee, M. Gruteser, and S. Oh. Wireless device identification with radiometric signatures. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, 2008.

[11] N. Bulusu, J. Heidemann, and D. Estrin. GPS-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5), October 2000.

[12] P. Castro, P. Chiu, T. Kremenek, and R. Muntz. A Probabilistic Room Location Service for Wireless Networked Environments. In *Proceedings of the International Conference on Ubiquitous Computing (Ubicomp)*, volume 2201, 2001.

[13] B. Danev and S. Čapkun. Transient-based identification of wireless sensor nodes. In *Proceedings of the ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2009.

[14] L. Doherty, K. Pister, and L. El Ghaoui. Convex position estimation in wireless sensor networks. In *Proceedings of the IEEE Conference on Computer Communications (InfoCom)*, April 2001.

[15] Ettus. Universal software radio peripheral (USRP). http://www.ettus.com.

[16] R. J. Fontana, E. Richley, and J. Barney. Commercialization of an ultra wideband precision asset location system. *IEEE Conference on Ultra Wideband Systems and Technologies*, 2003.

[17] Fraunhofer IIS. Autonomous WLAN positioning system. press release. http://www.fraunhofer.de/EN/press/pi/2008/01/Presseinformation14012008.jsp, 2008.

[18] S. Ganu, A. Krishnakumar, and P. Krishnan. Infrastructure-based location estimation in WLAN networks. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, March 2004.

[19] I. Getting. The Global Positioning System. *IEEE Spectrum*, December 1993.

[20] Y. Gwon, R. Jain, and T. Kawahara. Robust indoor location estimation of stationary and mobile users. In *Proceedings of the IEEE Conference on Computer Communications (InfoCom)*, March 2004.

[21] G. Hancke and M. Kuhn. An RFID Distance Bounding Protocol. In *Proceedings of the International Conference on Security and Privacy for Emerging Areas in Communications Networks (SecureComm)*. IEEE Computer Society, 2005.

[22] G. Hancke and M. Kuhn. Attacks on 'Time-of-Flight' Distance Bounding Channels. In *Proceedings of the ACM Conference on Wireless Network Security (WiSec)*. ACM, 2008.

[23] J. Hightower, G. Boriello, and R. Want. SpotON: An indoor 3D location sensing technology based on RF signal strength. Technical Report 2000-02-02, University of Washington, 2000.

[24] Y.-C. Hu, A. Perrig, and D. B. Johnson. Packet Leashes: A Defense against Wormhole Attacks in Wireless Networks. In *Proceedings of the IEEE Conference on Computer Communications (InfoCom)*, San Francisco, USA, April 2003.

[25] T. Kohno, A. Broido, and K. C. Claffy. Remote physical device fingerprinting. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2005.

[26] M. Kuhn. An asymmetric security mechanism for navigation signals. In *Proceedings of the Information Hiding Workshop*, 2004.

[27] L. Lazos and R. Poovendran. SeRLoc: secure range-independent localization for wireless sensor networks. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, 2004.

[28] L. Lazos, R. Poovendran, and S. Čapkun. ROPE: robust position estimation in wireless sensor networks. In *Proceedings of the symposium on Information processing in sensor networks (IPSN)*. IEEE Press, 2005.

[29] Z. Li, W. Trappe, Y. Zhang, and B. Nath. Robust Statistical Methods for Securing Wireless Localization in Sensor Networks. In *Proceedings of the symposium on Information processing in sensor networks (IPSN)*, 2005.

[30] D. Liu, P. Ning, and W. Du. Attack-Resistant Location Estimation in Sensor Networks. In *Proceedings of the symposium on Information processing in sensor networks (IPSN)*, 2005.

[31] Mexens LLC. Navizon virtual GPS service. http://www.navizon.com.

[32] U. Meyer and S. Wetzel. A man-in-the-middle attack on UMTS. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, 2004.

[33] C. Mitchell. The security of the GSM air interface protocol. Technical report, RHUL-MA-2001-3, Royal Holloway University of London, 2001.

[34] D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with noisy range measurements. In *Proceedings of the ACM Conference on Networked Sensor Systems (SenSys)*, 2004.

[35] D. Niculescu and B. Nath. Ad hoc positioning system (APS) using AoA. In *Proceedings of the IEEE*

Conference on Computer Communications (InfoCom), San Francisco, USA, April 2003.

[36] S. Pandey and P. Agrawal. A survey on localization techniques for wireless networks. *Journal of the Chinese Institute of Engineers*, 29(7), 2006.

[37] S. Pandey, F. Anjum, and P. Agrawal. *TRaVarSeL–Transmission Range Variation based Secure Localization*, pages 215–236. 2007.

[38] S. Pandey, F. Anjum, B. Kim, and P. Agrawal. A low-cost robust localization scheme for WLAN. In *Proceedings of the International Workshop on Wireless Internet*, New York, NY, USA, 2006. ACM.

[39] S. Pandey, B. Kim, F. Anjum, and P. Agrawal. Client assisted location data acquisition scheme for secure enterprise wireless networks. *IEEE Wireless Communications and Networking Conference (WCNC)*, 2, March 2005.

[40] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket location-support system. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, 2000.

[41] K. B. Rasmussen and S. Čapkun. Implications of radio fingerprinting on the security of sensor networks. In *Proceedings of the International Conference on Security and Privacy for Emerging Areas in Communications Networks (SecureComm)*, 2007.

[42] K. B. Rasmussen, S. Čapkun, and M. Čagalj. SecNav: secure broadcast localization and time synchronization in wireless networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, 2007.

[43] N. Sastry, U. Shankar, and D. Wagner. Secure verification of location claims. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, 2003.

[44] A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in Ad-Hoc networks of sensors. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, 2001.

[45] S. Sedihpour, S. Čapkun, S. Ganeriwal, and M. Srivastava. Implementation of Attacks on Ultrasonic Ranging Systems.Demo at the ACM Conference on Networked Sensor Systems (SenSys), 2005.

[46] P. Tao, A. Rudys, A. M. Ladd, and D. S. Wallach. Wireless LAN location-sensing for security applications. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, 2003.

[47] N. O. Tippenhauer and S. Čapkun. UWB-based Secure Ranging and Localization. Technical Report 586, ETH Zurich, January 2008.

[48] O. Ureten and N. Serinken. Wireless security through RF fingerprinting. *Canadian Journal of Electrical and Computer Engineering*, 32, 2007.

[49] S. Čapkun, L. Buttyan, and J.-P. Hubaux. Sector: Secure tracking of node encounters in multi-hop wireless networks. In *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, October 2003.

[50] S. Čapkun, M. Hamdi, and J.-P. Hubaux. GPS-free Positioning in Mobile Ad-Hoc Networks. *Cluster Computing*, 5(2), April 2002.

[51] S. Čapkun and J.-P. Hubaux. Secure positioning of wireless devices with application to sensor networks. In *Proceedings of the IEEE Conference on Computer Communications (InfoCom)*, volume 3, 2005.

[52] S. Čapkun and J.-P. Hubaux. Secure positioning in wireless networks. *IEEE Journal on Selected Areas in Communications*, 24(2), February 2006.

[53] S. Čapkun, M. Čagalj, and M. Srivastava. Secure localization with hidden and mobile base stations. In *Proceedings of the IEEE Conference on Computer Communications (InfoCom)*, April 2006.

[54] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The Active Badge Location system. *ACM Transactions on Information Systems*, 10(1), 1992.

[55] A. Ward, A. Jones, and A. Hopper. A New Location Technique for the Active Office. *IEEE Personal Communications*, 4(5), October 1997.

[56] J. S. Warner and R. G. Johnston. Think GPS Cargo Tracking = High Security? Think Again. *Technical report, Los Alamos National Laboratory*, 2003.

[57] WiGLE. Wireless Geographic Logging Engine. http://wigle.net/.

[58] W. Xu, W. Trappe, Y. Zhang, and T. Wood. The feasibility of launching and detecting jamming attacks in wireless networks. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2005.