Secure Positioning in Wireless Sensor Networks through Enlargement-"Miscontrol" Detection

PERICLE PERAZZO, University of Pisa
LORENZO TAPONECCO, University of Pisa
ANTONIO A. D'AMICO, University of Pisa
GIANLUCA DINI, University of Pisa

Wireless sensor networks enable a wealth of new applications in areas such as military, medical, environmental, transportation, smart city, and so on. In many of such scenarios, we need to measure in a secure way the positions of the sensors. Existing range-based techniques for secure positioning require a burdensome infrastructure, with many fixed anchors. Reducing the infrastructure would reduce the deployment cost, and foster the adoption of secure positioning solutions in wireless sensor networks. In this paper we propose SPEM, a secure positioning system based on multilateration and ultra-wideband (UWB) distance bounding protocols. The key idea behind SPEM is to leverage the low probability that an adversary has of controlling *enlargement attacks* against UWB. We estimate such a probability by a thorough study and signal-level simulations of the UWB physical layer. We test SPEM both in a simulated environment, and in a real indoor environment using real UWB transceivers. We show that SPEM needs far less infrastructure than state-of-the-art solutions (-22% to -93%, depending on the anchor deployment method), while achieving high levels of security against smart and determined adversaries.

CCS Concepts: •**Security and privacy** → **Mobile and wireless security;**

Additional Key Words and Phrases: Secure positioning, distance bounding protocols, enlargement attacks, multilateration

1. INTRODUCTION

Localizing sensor nodes is a critical function for many wireless sensor networks (WSNs) applications, such as battlefield monitoring, air pollution detection, wildlife animal habitat tracking, emergency rescue and recovery, etc. The ability to localize wireless devices and sensors is fostering new classes of location-oriented applications. However, as more location-dependent services are deployed, they will increasingly become tempting targets for malicious attacks. Therefore, it is crucial to assure the integrity of the reported locations. Equipping the sensors with GPS receivers is not enough, because GPS is vulnerable to *spoofing attacks* [Humphreys et al. 2008], which make the sensors measure false positions. A spoofing attack can be mounted by means of simple devices capable of transmitting fake GPS signals. The task of mea-

This work has been supported by the Italian Research Project TENACE (pr. no. 20103P34XC); and the research project "Analisi di dati sensoriali: dai sensori tradizionali ai sensori sociali" funded by "Progeti di Ricerca di Ateneo - PRA 2016" of the University of Pisa.
Authors' address: Department of Information Engineering, University of Pisa, Italy. Emails: [name.surname]@iet.unipi.it.

suring trustworthy positions in the presence of a spoofer adversary is called *secure positioning*. *Distance bounding protocols* [Brands and Chaum 1993] turned out to be extremely useful for this aim. The key property of these protocols is to measure a distance between two devices in a way immune to *reduction attacks*, i.e., no adversary can make the measured distance be shorter than it really is. They act by measuring the round-trip time between a request and an acknowledgment packet both carrying quantities unpredictable by the adversary. Wireless distance bounding protocols can be realized by means of the ultra-wideband (UWB) technology [Poturalski et al. 2012], and they can be fruitfully used in multilateration schemes to provide for secure positioning [Čapkun and Hubaux 2006; Chiang et al. 2012; Perazzo and Dini 2015]. These schemes leverage measurements of distances with respect to reference points with known positions (*anchors*).

Distance reduction is not the only attack that can be played against distance bounding protocols. Recently, researchers have turned their attention to the dual attacks, namely *enlargement attacks*, which aim at making the distance appear larger [Chiang et al. 2012; Taponecco et al. 2014; Dini et al. 2013]. Enlargement attacks can be mounted in a number of ways, the most promising of which is the *overshadow attack*. In this attack, the adversary replays the request or the acknowledgment packets with a certain delay and a greater power, in such a way to cause an enlargement of the measured round-trip time. Recent research has showed that, under particular conditions, the outcomes of these attacks are hardly controllable [Taponecco et al. 2014]. In practice, this depends on the fact that the adversary cannot "cancel" the ongoing communication, and can only interfere with it. The malicious signal overlaps with the legitimate one, and this can produce random outcomes, which highly depend on the particular time-of-arrival (TOA) estimation algorithm employed.

In this paper we propose SPEM, a secure positioning scheme based on UWB distance bounding and multilateration. The basic idea is to detect uncontrolled enlargement attacks by monitoring the accuracy of the position estimates, and by increasing the precision of the multilateration scheme. SPEM uses a distance bounding protocol realized on the IEEE 802.15.4a UWB standard protocol [IEEE Computer Society 2007], which is capable of sub-meter precision at low energy costs [Zhang et al. 2009]. To evaluate the security of SPEM, we first performed thorough signal-level simulations of the UWB protocol exposed to an overshadow attack. Then, we simulated SPEM to estimate its performances in terms of security, precision, and anchor saving. We also evaluated the security of SPEM in the real field, with real UWB transceivers. We show that it is possible to achieve a high level of security, while saving up to 93% of the anchors with respect to state-of-the-art solutions.

The paper brings the following novel contributions:

— We introduce the concept of *enlargement control probability*, which expresses the capability of an adversary to control the effect of an enlargement attack. We present a thorough study of the overshadow attack against the IEEE 802.15.4a UWB protocol in terms of controllability.
— We compare two classic threshold-based TOA estimation algorithms, namely jump-back search-forward and search-back [Guvenc et al. 2005; D'Amico et al. 2010]. We show that the latter is more promising for security-focused applications, because it offers more resistance against enlargement attacks.
— We introduce SPEM, a secure multilateration scheme based on enlargement detection. We study SPEM in terms of security, localization precision, and infrastructure saving.

— We develop and make publicly available the *SPEM Parametrization and Evaluation Framework*, which helps the user to configure SPEM and evaluate its security in a generic combination of deployment environment and TOA estimation algorithm.
— We parametrize and evaluate SPEM both in a simulated environment using simulated transceivers, and in a real environment using real transceivers.

The remainder of this paper is organized as follows. In Section 2 we present related work. In Section 3 we introduce our system model. In Section 4 we describe the adversary, the overshadow attack, and the concept of enlargement control probability. In Section 5 we describe the IEEE 802.15.4a UWB protocol and the classic threshold-based TOA estimation algorithms. In Section 6 we analyze the effects of overshadow attacks against IEEE 802.15.4a UWB, and evaluate the enlargement control probability of the adversary. In Section 7 we describe SPEM. In Section 8 we parametrize SPEM both in a simulated environment and in a real one, and we evaluate its security and precision. In Section 9 we evaluate SPEM in terms of infrastructure saving. Finally, we draw our conclusions in Section 10.

2. RELATED WORK

We now survey some important related work in the fields of secure positioning and security in IEEE 802.15.4a UWB ranging.

2.1. Secure positioning

Positioning systems are traditionally classified in *range-based* and *range-free*. Range-based systems leverage measurements of distances and angles with respect to some reference points with known positions (*anchors*). On the other hand, range-free systems are not based on the (direct) measurement of geometric quantities. They deduce the position from other higher-level information. A typical example is the hearing of beacon packets sent from anchors. Range-free systems are typically cheaper, as they do not require specialized hardware for distance measurements. However, they result in a worse precision in position estimation. Secure positioning systems follow the same categorization. SPEM is a range-based system, since it measures distances by means of distance bounding.

Lazos and Poovendran [2005] proposed SeRLoc, a range-free secure positioning system for wireless sensor networks, which relies on beacon packets. In SeRLoc, each sensor computes its position by hearing beacon packets from trusted anchors. The attacks are detected by checking for inconsistencies in the received beacons. SeRLoc provides for a limited localization precision, and it is not resistant to the jamming of beacon packets. Moreover, it needs directional antennas on anchors.

Park and Shin [2009] proposed a secure localization method for wireless sensor networks based on signal strength measurements and multidimensional scaling. The approach is promising principally because it requires little hardware resources. However, it defends only against naive adversaries, which try to falsify the positions to random points. Our approach permits us to defend against smarter and more determined adversaries, which choose their tactic and their objective in such a way to maximize their success probability.

Hu et al. [2003] proposed *packet leashes*, which represents the first attempt to employ distance-bounding-like techniques for secure location verification. Sastry et al. [2003] proposed the *Echo protocol* for secure location verification based on ultra-sound ranging.

Čapkun and Hubaux [2006] proposed *verifiable multilateration*. Verifiable multilateration measures the distances from a set of trusted anchors by means of distance bounding protocols. The position is computed by means of multilateration, and it is con-

sidered secure if it lies inside the polygon formed by the involved anchors (*in-polygon check*). Indeed, if the measured position has been falsified, at least one of the distance bounding protocols must have been exposed to a distance reduction attack, which is infeasible. The in-polygon check causes a significant reduction in the area covered by the positioning system. As a consequence, verifiable multilateration needs a high number of anchors. By using enlargement-detection techniques, SPEM can save up to 93% anchors with respect to verifiable multilateration.

2.2. Security in IEEE 802.15.4a UWB ranging

Poturalski et al. [2011] conducted a deep study on reduction attacks against IEEE 802.15.4a UWB distance bounding. They evaluated the impact of such attacks in terms of reduction meters, and proposed a set of countermeasures to limit their effect. In the present paper, we consider a distance bounding protocol *immune* to reduction attacks, and we focus on the detection of enlargement attacks.

Poturalski et al. [2012] studied the feasibility and the impact of *interfering attacks* against the preamble, and proposed as a countermeasure a novel TOA estimation algorithm called *PIDH* (*power independent detection with Hamming distance*). In the present paper, we consider classic TOA estimation algorithms, namely jump-back search-forward and search-back [Guvenc et al. 2005; D'Amico et al. 2010]. We leave the security analysis of SPEM with non-classic TOA estimators as future work.

Taponecco et al. [2014] showed that, in the IEEE 802.15.4a UWB ranging standard, an overshadow-based enlargement attack is poorly controllable by the adversary. We start from that result to design a multilateration scheme based on the difficulty of controlling an enlargement attack. As an additional contribution, we precisely quantify the capacity of an adversary to control an enlargement attack against IEEE 802.15.4a UWB.

3. SYSTEM MODEL

We assume a two-tier system architecture with a set of anchors $\mathcal{A}$ and a set of sensors $\mathcal{S}$. To simplify the notation, in the following the elements of $\mathcal{A}$ and $\mathcal{S}$ will indicate the devices (anchors and sensors), as well as the *positions* of such devices in the two-dimensional plane. The positions of the sensors are unknown, while the positions of the anchors are known and trusted. We assume that the anchors can communicate between each other in a secure way. We also assume that the sensors are not compromised, and thus they keep secrets and behave according to specifications.

3.1. Multilateration scheme

Our multilateration scheme measures the position of each sensor separately. Thus, we focus on the positioning of a single sensor $S \in \mathcal{S}$. We consider a multilateration scheme which determines the position S by measuring $N \geq 3$ distances $d_1, \ldots, d_N$ of the sensor from N anchors, say $A_1, \ldots, A_N$. In the absence of measurement errors, S is given by the intersection of the circumferences with centers A_i and radii d_i (*ranging circumferences*). In the presence of some imprecision, the *measured distance* $\widehat{d_i}$ will be affected by a *measurement error* e_i:

$$\widehat{d_i} = d_i + e_i. \tag{1}$$

In such a case, the ranging circumferences will not intersect in a point. The *measured position* $\widehat{S}$ will thus be the pseudo-solution in the least-squared-error sense (Figure 1):

$$\widehat{S} = \arg\min_{X} \sum_{i=1}^{N} \left(\widehat{d_i} - \|X - A_i\| \right)^2, \tag{2}$$

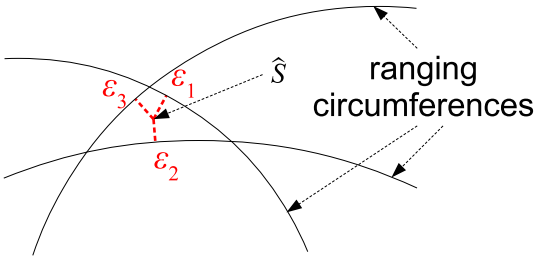

Fig. 1. Multilateration problem.

where $\|\cdot\|$ indicates the norm operation.

The output of the multilateration is the measured position $\widehat{S}$ and a set of N *residuals* $\varepsilon_1, \ldots, \varepsilon_N$, given by:

$$\varepsilon_i = \widehat{d}_i - \|\widehat{S} - A_i\|. \tag{3}$$

The residuals are an indirect estimation of the measurement errors e_i. High values of the residuals generally imply high errors.

3.2. Distance bounding protocol

Each anchor measures its distance to the sensor by means of a *distance bounding protocol* [Brands and Chaum 1993]. The anchor estimates the distance by measuring the round-trip time between the transmission of a *request packet* and the reception of an *acknowledgement packet*. Let the *processing time* T_{proc} be the period between the reception of the request at the sensor and the transmission of the acknowledgment. If the processing time is known, then the distance can be estimated by:

$$\widehat{d}_i = \frac{T_{RTT_i} - T_{proc}}{2} \cdot c, \tag{4}$$

where T_{RTT_i} is the measured round-trip time for the i-th anchor and c is the speed of light. The accuracy of distance estimation depends on the precision with which sensor and anchors estimate the time of arrival of a packet. In multipath environments, this in turn highly depends on the bandwidth of the employed radio signals. Ultra-wideband PHY protocols like IEEE 802.15.4a UWB can reach sub-meter precisions on distance estimation [Zhang et al. 2009].

A simple example of distance bounding protocol, proposed in [Poturalski et al. 2011] for IEEE 802.15.4a UWB and for external adversaries, is the following:

REQ. $A \longrightarrow S : n_A$
ACK. $S \longrightarrow A : n_S$
AUTH. $S \longrightarrow A : \mathrm{auth}(n_A, n_S)$,

where A represents the anchor and S the sensor. The request packet (REQ) and the acknowledgment packet (ACK) convey, respectively, n_A and n_S, which are two externally unpredictable sequences of bits. The *authentication packet* (AUTH) authenticates the request and the acknowledgment. The function $\mathrm{auth}(\cdot)$ represents a message authentication code (e.g., a CBC-MAC), which uses some secret shared by S and A. We assume that every sensor has a distinct secret used for authentications.

This protocol does not allow an adversary to impersonate the sensor, because the adversary cannot forge the final authentication code. Moreover, the protocol avoids reduction attacks, because an external adversary cannot predict n_A and n_S to anticipate their transmission.

4. ADVERSARY MODEL

The objective of the adversary is to spoof the position measured by the multilateration system. We denote by S' the *false position*, which is the position that the adversary wants to make the system believe the sensor is in. We define the *attack success probability* as the probability that the system accepts the false position. Different false positions give different success probabilities to the adversary. We distinguish two types of adversary, depending on the choice of the false position: the *random-objective adversary* and the *best-objective adversary*. The random-objective adversary is given a random false position among those having non-zero success probability. It models an adversary that has a predefined objective and is not free to change it. The best-objective adversary is given the false position having the greatest success probability. It models an adversary that has no predefined objective, and is free to choose the most convenient one. Of course, the best-objective adversary has more chances to succeed.

Note that, without other constraints, a convenient choice for the false position would be very close to the true position, or coincident with it in the extreme case. With these trivial false positions, the adversary has a high probability to succeed. However, such an attack would not be a true spoofing, but rather a simple degradation of the system precision. We force the adversary to cause a *minimal spoofing distance* (d_{ms}) between the false position and the true one:

$$\|S' - S\| \geq d_{ms}. \tag{5}$$

Accordingly, we assume the system to be tolerant to a precision degradation of d_{ms} meters.

In order to make the system accept the false position, the adversary has to attack a number of distance bounding protocols and make them measure *false distances* d'_i. For each anchor A_i, the adversary chooses d'_i in such a way that the multilateration gives as output S' (Figure 2(a)).

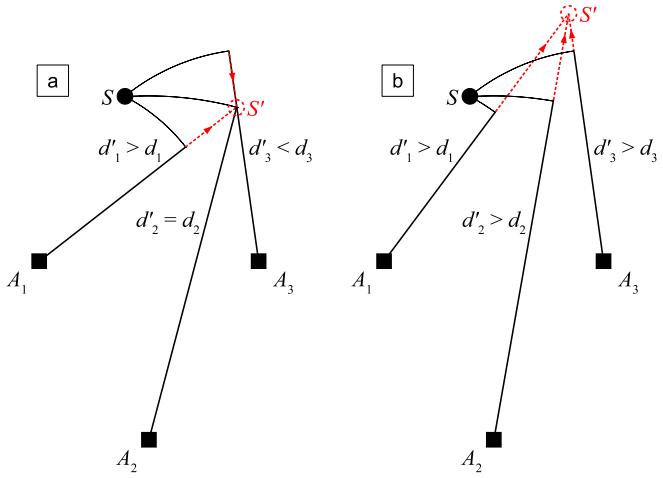

Fig. 2. Reduction and enlargement attacks needed for multilateration spoofing.

For each distance bounding protocol, three cases are possible:

(1) $d'_i < d_i$. The distance measurement has to be reduced. The adversary has to perform a *reduction attack*.
(2) $d'_i = d_i$. The distance measurement does not need to be attacked.

(3) $d'_i > d_i$. The distance measurement has to be enlarged. The adversary has to perform an *enlargement attack*.

Reduction attacks are orthogonal to enlargement attacks both in terms of objectives and related countermeasures. They essentially aim at shortening the round-trip time by predicting REQ or ACK packets, or by making the victim receive them in advance. Poturalski et al. [2011] conducted a deep study on PHY-level reduction attacks against IEEE 802.15.4a UWB distance bounding. They evaluated the impact in terms of reduction meters, and proposed a set of countermeasures to limit it. In a previously published paper [Perazzo and Dini 2015], we showed that the possibility of reduction attacks in range-based secure positioning systems causes an uncertainty on the position estimate. Basically, this is true also in enlargement-resistant solutions like SPEM. However, the security analysis of SPEM under reduction attacks falls outside the scope of the present paper, and we leave it as a future work. In this paper, we consider the distance bounding protocol to be *immune* to reduction attacks. Given that reduction attacks are impossible, only the false positions that do not require reduction attacks have a non-zero success probability (Figure 2(b)).

From now on, for the sake of exposition we will focus on a single enlargement attack against a single distance measurement, and hence we will omit the "i" subscript to ease the notation. We call *objective enlargement* ($a \geq 0$) the distance enlargement that the adversary wants to obtain:

$$a = d' - d. \tag{6}$$

Actually, with her attack the adversary causes an anomalous error in the distance measurement, which is as close as possible to her objective enlargement. The measurement error is a random variable whose characteristics change sensibly in the presence or in the absence of an attack. We discriminate between the *honest measurement error* (e) and the malicious one, which we call the *obtained enlargement* ($\widehat{a}$). We define as *enlargement control error* (e_{ctrl}) the difference between the obtained and the objective enlargement, i.e.:

$$\widehat{a} = a + e_{ctrl}. \tag{7}$$

Also, we call *enlargement control probability* (P_{ctrl}) the probability for the enlargement to be "controlled", i.e., the probability that the control error is indistinguishable to an ordinary measurement error. Assuming a *honest precision* e_{max}, such that:

$$\Pr\left[|e| \leq e_{max}\right] = 99.9\% \quad \text{(in the honest case)}, \tag{8}$$

the enlargement control probability is defined as:

$$P_{ctrl} \triangleq \Pr\left[|e_{ctrl}| \leq e_{max}\right] \quad \text{(in the malicious case)}. \tag{9}$$

Essentially, the enlargement control probability measures the capacity of the adversary to control the enlargement in a *single* ranging operation. As we will see later, this probability depends on how much the adversary wants to enlarge. In other words, some enlargements are easier to obtain than others. This is the reason why some false positions have greater success probability than others.

4.1. Enlargement attacks feasibility study

Since the distance measurement stems exclusively from the round-trip time, the aim of the adversary is to enlarge it. We note that, in the case that the sensor itself is malicious, it is quite easy for it to cause a controlled enlargement attack. The sensor can simply delay the transmission of the acknowledgment packet, thus enlarging the round-trip time. Some countermeasures exist for this, based on multiple anten-

nas [Chiang et al. 2012]. However, they can only mitigate the attack, and they are ineffective if the malicious node communicates its secret key to other malicious devices. Unfortunately, it is quite hard to detect an enlargement attack in case of untrusted sensors. In this paper we focus on external adversaries only. We are aware that in WSNs, especially in unattended ones deployed in hostile environments, the possibility of compromising sensors and keys is a concrete threat. In this cases, remote attestation techniques can be employed in order to assure the correct execution of the distance bounding protocol. In the last years, hardware-based remote attestation on WSNs has became an attractive option [Hu et al. 2010]. Off-the-shelf trusted platform modules (TPMs) come at a reasonable cost in terms of price, board space, and energy consumption. As an alternative, software-based remote attestation techniques can be used as well [Seshadri et al. 2004], though they offer weaker security in general.

For an external adversary, the only way to enlarge the round-trip time is to delay the packet TOA estimate at the sensor and/or at the anchor. We suppose the adversary is equipped with one or more devices, which are able to communicate with each other. They can eavesdrop and transmit any signal in the UWB wireless channel. The adversary does not have any limitation on the transmission power. There are several ways in which an external adversary can delay a TOA estimate. We identified three ways: (a) *jam-replay attack*, (b) *annihilation attack*, (c) *overshadow attack*. In the following, we explain and discuss the three kinds of attack. We argue that the most promising one is the overshadow attack, since it is hard to detect by the legitimate receivers.

Jam-replay attack acts by jamming a legitimate communication (for example, the REQ packet), and then replaying it afterwards. Jam-replay is maybe the simplest enlargement attack. Note however that, in order to avoid packet collisions, the adversary has to wait for the legitimate communication to end before replaying it. This, as shown by [Dini et al. 2013], forces the adversary to introduce large delays, greater than or equal to the packet transmission time. This makes the jam-replay attack suitable only for those protocols that have extremely short packet transmission times. In the specific IEEE 802.15.4a UWB protocol, the packet transmission times are in the order of milliseconds [IEEE Computer Society 2007]. Thus, a jam-replay attack would produce unrealistic enlargements, in the order of hundreds of kilometers, which are easy to detect by means of simple threshold mechanisms.

Annihilation attack acts by repeating a legitimate communication with a certain delay and a far greater power. In this way, the adversary causes an anomalous behavior of the *analog-to-digital converter* (ADC) of the receiver. Before digitizing it, the input signal is usually passed through an *automatic gain control* (AGC) stage, which levels out the peak amplitude at a constant value by reducing or increasing it. By transmitting with strong power, the adversary can force the AGC to reduce the amplitude of the signal. If the amplitude reduction is pronounced enough, the honest signal will fall below the minimal resolution of the ADC. The honest signal gets thus deleted totally, and the receiver hears exclusively the malicious one, which is delayed properly to obtain the desired enlargement. The annihilation attack can produce realistic enlargements, because it does not wait for the end of the legitimate communication, but replays it while it is still going. However, the adversary has to transmit with a strong power, which is unrealistic for a legitimate communication. In particular, the ratio of the malicious signal power to the legitimate signal power must be greater than the ratio of the ADC dynamic to the ADC resolution, which can be considerably high if the ADC has enough bits. An annihilation attack can be detected by enforcing a limit to the received power, and by providing a reasonable resolution for the ADC. We leave the detailed study of these countermeasures as future work.

Overshadow attack acts by repeating a legitimate communication with a certain delay and a (not too much) greater power. The receiver thus hears both the legitimate and

the malicious packets in a superimposed way, without being able to distinguish them. As a result, the overshadow attack can effectively introduce enlargements which are hard to detect. It does not cause unrealistically wide enlargements and does not introduce an unrealistically high power in the channel. Though the effect of an overshadow attack is not always controllable [Taponecco et al. 2014], it still results to be the most convenient strategy. We thus assume that the adversary mounts overshadow attacks in order to obtain enlargements.

4.2. Overshadow tactics

The *overshadow delay on the REQ* (T_o^{R}) and the *overshadow delay on the ACK* (T_o^{A}) are the differences between the times of arrival of the legitimate and the malicious packets (either REQ and ACK) at the victim receiver. We assume that the overshadow delays are completely controllable by the adversary. On the other hand, the *obtained delay on the REQ* (T_e^{R}) and the *obtained delay on the ACK* (T_e^{A}) are the differences between the true and the estimated TOA, respectively, on the REQ and on the ACK packet. Clearly, T_e^{R} and T_e^{A} are not zero even in the absence of an attack, due to measurement errors. In case of attack, the obtained delays differ in general from the overshadow delays. The total round-trip time enlargement obtained by the adversary equals $T_e^{\mathrm{R}}+T_e^{\mathrm{A}}$. Therefore, the obtained distance enlargement is given by:

$$\widehat{a} = (T_e^{\mathrm{R}} + T_e^{\mathrm{A}}) \cdot \frac{c}{2}, \tag{10}$$

which grows linearly with the obtained delays on the two packets.

Since T_e^{R} and T_e^{A} are, in general, different from T_o^{R} and T_o^{A}, it is convenient for the adversary to introduce the overshadow delays that, with the greatest probability, will produce obtained delays corresponding to her objective enlargement. Accordingly, we formally define an *overshadow tactic* as a couple of overshadow delays:

$$\langle T_o^{\mathrm{R}}, T_o^{\mathrm{A}} \rangle.$$

Given an objective enlargement, we also define the *best overshadow tactic* $\langle \overline{T}_o^{\mathrm{R}}, \overline{T}_o^{\mathrm{A}} \rangle$ as the one that maximizes the control probability:

$$\langle \overline{T}_o^{\mathrm{R}}, \overline{T}_o^{\mathrm{A}} \rangle = \arg \max_{\langle T_o^{\mathrm{R}}, T_o^{\mathrm{A}} \rangle} \left(P_{ctrl} \right). \tag{11}$$

Notice that, since it could be convenient for the adversary not to attack either the REQ or ACK packet, the following are valid tactics too:

$$\langle \text{no-attack}, T_o^{\mathrm{A}} \rangle,$$
$$\langle T_o^{\mathrm{R}}, \text{no-attack} \rangle,$$
$$\langle \text{no-attack}, \text{no-attack} \rangle.$$

The $\langle \text{no-attack}, \text{no-attack} \rangle$ tactic is useful when the objective enlargement is close to zero ($a \approx 0$).

5. IEEE 802.15.4a ULTRA-WIDEBAND

In order to study the best overshadow tactics against IEEE 802.15.4a UWB protocol, we now give some details on such a standard and on physical-layer procedures for threshold-based UWB ranging schemes, which are the most widely used in UWB localization applications [Guvenc et al. 2005; D'Amico et al. 2008; 2010]. In particular, we consider the *jump-back search-forward* (JBSF) and the *search-back* (SB) algorithms, which provide significantly different results from the security point of view, as we show later.

The IEEE 802.15.4a amendment [IEEE Computer Society 2007] introduces an impulse radio ultra-wideband (IR-UWB) PHY protocol capable of ranging with sub-meter precision. It has been the first standardized UWB protocol for precise ranging, and it is particularly suitable for WSNs, as it exhibits a low price and low power consumption [Zhang et al. 2009]. It is one of the most probable choices for future implementations of wireless distance bounding protocols [Poturalski et al. 2011]. From the point of view of IEEE 802.15.4a UWB, the REQ and the ACK packets constitute a *two-way ranging* operation. The REQ and the ACK packets are mapped into two 802.15.4a UWB PHY protocol data units (PPDUs).

An 802.15.4a UWB PPDU consists of three parts: a synchronization header (SHR), a PHY header (PHR), and a PHY service data unit (PSDU). The SHR part allows for the estimation of the arrival time of the packet. The PHR contains information about the modulation kind of the successive PSDU part. Finally, the PSDU part contains the information data. In our case, the unpredictable quantities n_A and n_S are conveyed by the PSDU parts. The SHR is made up of two blocks: a synchronization preamble (SYNC) and a start-of-frame delimiter (SFD). In particular, here we are interested in the synchronization preamble. The mathematical model of the signal transmitted during the SYNC is:

$$s\left(t\right) = \sum_{i=0}^{N_{SYNC}-1} \psi\left(t - iT_{sym}\right), \tag{12}$$

where $N_{SYNC} = 1024$ is the number of symbols belonging to the SYNC, $T_{sym} = 3968\,\mathrm{ns}$ is the symbol duration, and $\psi(t)$ has the following expression:

$$\psi\left(t\right) \triangleq \sum_{k=0}^{K_{pbs}-1} d_k p\left(t - kT_{pr}\right). \tag{13}$$

In (13), $\{d_k\}_{k=0}^{K_{pbs}-1}$ is a *perfectly balanced sequence* of $K_{pbs} = 31$ elements with values $\{-1, 0, +1\}$, $p(t)$ is an ultra-short causal pulse (*monocycle*) and $T_{pr} \triangleq T_{sym}/K_{pbs} = 128\,\mathrm{ns}$ is the *pulse repetition period*. As shown in (12), the signal transmitted during the preamble is the periodic repetition of the waveform $\psi(t)$ with period T_{sym}.

Propagation occurs on a *multipath channel*, in which each propagation path is characterized by different attenuations and delays. Denoting by $h(t)$ the *channel response* to $p(t)$[1], the received signal can be written as:

$$r\left(t\right) = \sum_{i=0}^{N_{SYNC}-1} \sum_{k=0}^{K_{pbs}-1} d_k h\left(t - kT_{pr} - iT_{sym} - t_{TOA}\right) + w\left(t\right), \tag{14}$$

where $w(t)$ is thermal noise with a flat two-sided power spectral density. In the above equation, t_{TOA} is the time of arrival of the signal at the receiver and represents the parameter to be measured. It can be the time of arrival of either the REQ packet at the sensor or the ACK packet at the anchor.

We consider a simple non-coherent energy-based receiver, which guarantees high ranging precision with low cost and low power consumption. Here, $r(t)$ is first passed through a band-pass filter (BPF), to remove the extra-band noise, and then is demodulated in a square-law device followed by a low-pass filter (LPF). Assuming that the $h(t)$-pulses in (14) do not overlap, it is readily shown that the LPF output, $y(t)$, has the

[1]It is assumed that $h(t) = 0\ \forall t < 0$.

following form:

$$y(t) = \sum_{i=0}^{N_{SYNC}-1} \sum_{k=0}^{K_{pbs}-1} d_k^2 q\left(t - kT_{pr} - iT_{sym} - t_{TOA}\right) + n_y(t). \tag{15}$$

In this equation $n_y(t)$ is a noise term originating from signal×noise and noise×noise interactions in the squarer, and $q(t) \triangleq h^2(t) \otimes h_{LP}(t)$, where $h_{LP}(t)$ is the impulse response of the LPF and $\otimes$ denotes the convolution operation.

In general, $q(t)$ shows different peaks each of which corresponds to the arrival of a signal echo through a propagation path. The first peak indicates the arrival through the shortest path. The ranging operation is concerned with the estimation of t_{TOA}, the time of arrival of the first peak of the first $q(t)$-pulse of the preamble.

The TOA estimation algorithm considered in the present paper is that described and analyzed in [D'Amico et al. 2010]. Specifically, with both JBSF and SB algorithms the TOA estimation is performed in three phases (Figure 3):

(1) *Frame detection*: decides through energy measurements whether a packet is present or not.
(2) *Fine time acquisition*: produces a fine estimate of the arrival time t_{TOA} with an ambiguity of multiples of T_{sym}.
(3) *Start-of-Frame Delimiter (SFD) detection*: disambiguates the estimate of t_{TOA} through a correlation mechanism.

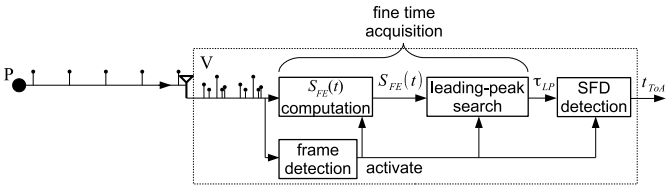

Fig. 3. TOA estimation block diagram.

The *fine time acquisition* phase provides a measure of a timing parameter, say $\tau_{LP} \in [0, T_{sym})$, which is related to t_{TOA} by $t_{TOA} = t_{fd} + \tau_{LP} - N_{fd}T_{sym}$, where t_{fd} is the time at which the *frame detection* phase declares the presence of the packet. The *SFD detection* phase resolves the T_{sym}-ambiguity by estimating N_{fd}.

We now focus on the *fine time acquisition* phase. Indeed, as shown in [Taponecco et al. 2014], this is the only phase of the ranging operation that the adversary attacks. The *fine time acquisition* phase consists in the correlation of the signal $y(t)$ at the output of LPF with K_{pbs} cyclic-shifted versions of the sequence $\{d_k^2\}_{k=0}^{K_{pbs}-1}$. This produces a T_{sym}-long signal, say $S_{FE}(t)$, whose support is the interval $[0, T_{sym})$, which is used for the estimation of τ_{LP}. More precisely, for $t \in [mT_{pr}, (m+1)T_{pr})$, with $m = 0, 1, \ldots, K_{pbs} - 1$, $S_{FE}(t)$ is given by:

$$S_{FE}(t) = \frac{1}{M} \sum_{i=0}^{M-1} \left[\sum_{k=0}^{K_{pbs}-1} d_{|k-m|_{K_{pbs}}}^2 \, y\left(t + t_{fd} + (k-m)T_{pr} + iT_{sym}\right) \right], \tag{16}$$

where $M < N_{SYNC}$ is the number of preamble symbols exploited by the *fine time acquisition* phase, and $|u|_U \triangleq u \bmod U$.

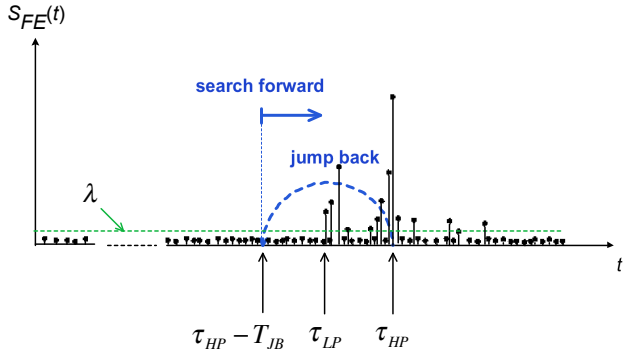

Fig. 4. Jump-back search-forward algorithm.

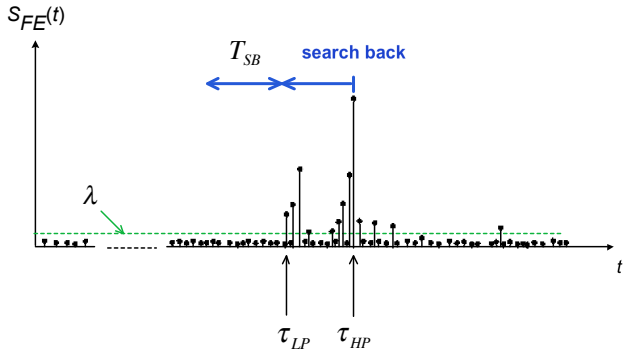

Fig. 5. Search-back algorithm.

Mathematical details apart, the computation of $S_{FE}(t)$ essentially leverages the periodicity of the preamble signal and the correlation properties of the sequence $\{d_k^2\}_{k=0}^{K_{pbs}-1}$ to improve the signal-to-noise ratio. The estimation of τ_{LP} is performed in two steps. In the first step (*highest-peak search*) the position τ_{HP} of the maximum of $S_{FE}(t)$ is sought for. In the second step (*leading-peak search*), starting from τ_{HP} the parameter τ_{LP} is determined through a threshold-based mechanism.

The JBSF and SB algorithms differ only for the leading-peak search step. In particular, the JBSF algorithm (Figure 4) starts from the maximum of $S_{FE}(t)$, jumps back by T_{JB} seconds and proceeds forward looking for the first time $S_{FE}(t)$ goes beyond a given *noise threshold* (λ). The distance of such a crossing time from the beginning of $S_{FE}(t)$ provides an estimate of τ_{LP}.

On the other hand, the SB algorithm (Figure 5) starts from τ_{HP}, and searches backward until $S_{FE}(t)$ goes below the noise threshold and continues to be under the threshold level for T_{SB} seconds (T_{SB}*-long noise-only region*). Just as for the JBSF algorithm, the distance of such a crossing time from the beginning of $S_{FE}(t)$ provides an estimate of τ_{LP}.

The noise threshold is fixed on the basis of the thermal noise statistics. We set it in such a way that a noise-only sample of $S_{FE}(t)$ has a probability of 10^{-5} of being above λ, i.e., of being wrongly interpreted as signal (*false alarm*). In addition, we set $T_{JB} = 60\,\text{ns}$, as recommended by [D'Amico et al. 2010], and $T_{SB} = 30\,\text{ns}$, as determined experimentally through computer simulations to guarantee the optimal performance (in terms of mean squared error) of the SB algorithm in indoor environments.

The distance measurement error mainly depends on the TOA estimation error at the prover and at the verifier, which in turn depends on the signal-to-noise ratio. The TOA estimation error of a threshold-based algorithm like those considered in this paper follows a distribution studied by Sharp and Yu [2014]. Figure 6 shows the histograms of the measurement error of our simulated receivers, with different values of the signal-to-noise ratio.

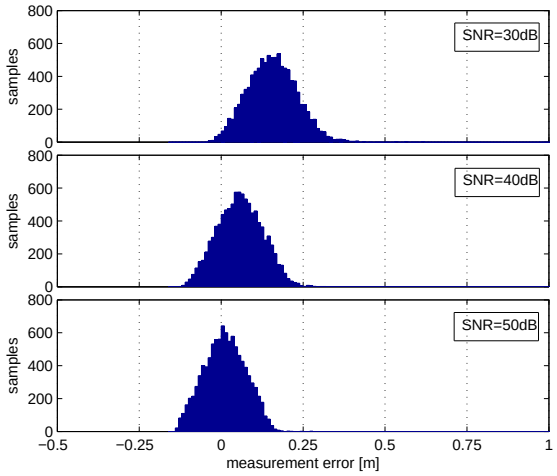

Fig. 6. Histograms of the distance estimation error (10,000 samples).

The histograms have been obtained by means of 10,000 signal-level simulations of the search-back TOA estimation algorithm with the standard channel model CM1 [Molisch et al. 2006]. The jump-back search-forward algorithm has a similar performance.

6. OVERSHADOW ATTACK AGAINST IEEE 802.15.4a UWB

Now we analyze in detail the effects of an overshadow attack against the TOA estimation algorithm. We suppose that the adversary knows the TOA estimation algorithm implemented at the receivers and the statistical characteristics of the channel. In the presence of a legitimate transmission of a REQ or an ACK packet, the adversary first synchronizes with the ongoing communication. It takes some of the initial symbols of the synchronization preamble to do that. Then, she starts transmitting the replayed copy with a greater power, skipping those initial preamble symbols. The replayed signal is timed to arrive at the receiver shifted of a certain delay (the overshadow delay) with respect to the legitimate one. In so doing, the adversary delays the estimate of the packet TOA. We assume that she knows exactly the distance between the transmitter and the receiver, and her distance from the receiver. This is a necessary condition for the malicious signal to arrive at the victim receiver with the desired delay. During the successive payload transmission, a copy of the payload is replayed in the same way.

An adversary enjoying a single-path channel toward the victim is generally more powerful, since she can control more precisely her attack. She can obtain this either by deploying a transmitter very close to the victim, or by using a highly directive antenna toward it. We assume the worst-case scenario, in which the adversary is capable of establishing a single-path channel toward the sensor and toward each anchor.

As a consequence of the overshadow attack, the waveform $S_{FE}(t)$ has a component due to the legitimate signal, and an additional component (the strongest one) associ-

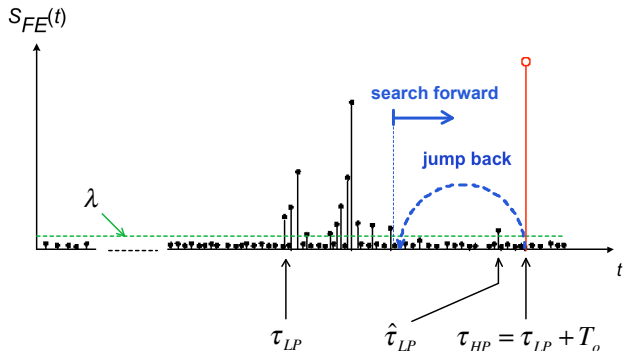

Fig. 7. Overshadow attack against JBSF algorithm. The black-headed pulses correspond to the legitimate signal while the white-headed one is associated to the malicious one. We indicated the overshadow delay with a generic symbol T_o, abstracting from which packet (ACK or REQ) is being attacked.

ated to the malicious transmission. The latter arrives T_o seconds after τ_{LP}. Figure 7 shows an example of overshadow attack against the JBSF algorithm. As it can be seen, the estimated $\widehat{\tau}_{LP}$ is different from the true τ_{LP}. In particular, one of the secondary peaks of the channel response is wrongly identified as the leading peak. Depending on which peak is identified as the leading one, the attack can fall into three cases:

— *Case 1.* The first peak of the legitimate signal is identified as the leading peak. Case 1 captures the case in which the attack has no effect.
— *Case 2.* A secondary peak of the legitimate signal is identified as the leading peak. This is the case of Figure 7.
— *Case 3.* The malicious peak is identified as the leading peak.

This categorization can be applied to an attacked JBSF algorithm, as well as to an attacked SB algorithm. With her overshadow attack against a packet (REQ or ACK), the adversary aims at falling into Case 1, Case 2, or Case 3 in the TOA estimation of that packet.

6.1. Evaluation of the best overshadow tactics

In order to determine the best overshadow tactics, we simulated the TOA estimation algorithms described in Section 6 under an overshadow attack. We introduced different overshadow delays (from $T_o = 0.0\,\text{ns}$ to $T_o = 270.0\,\text{ns}$, with steps of $0.1\,\text{ns}$) over 100 randomly generated UWB channels, and we measured their effects. The UWB channels follow the standard statistical model CM1 [Molisch et al. 2006]. Finally, for each objective enlargement, we tested all the combinations of overshadow delays on the REQ and the ACK packets, selecting the one giving the highest control probability[2]. Due to the symmetry between the REQ and the ACK transmissions, symmetric tactics (e.g. $\langle 50\,\text{ns}, 100\,\text{ns}\rangle$ and $\langle 100\,\text{ns}, 50\,\text{ns}\rangle$) result in exactly the same control probability. Without loss of generality, we thus impose $\overline{T}_o^{\text{R}} \geq \overline{T}_o^{\text{A}}$.

Figure 8(a) shows the best tactics against the JBSF algorithm, found by means of the method described above. For example, in order to cause an enlargement of 10 meters against a JBSF algorithm, the best tactic is to overshadow the REQ packet with a 98.3-nanosecond delay and the ACK packet with a 82.4-nanosecond delay. The corresponding control probability, indicated in Figure 8(b), is 91.88%. Note that the trend

[2]To compute the control probability, we used a honest precision of $e_{max} = 39.7\,\text{cm}$, as we found during the parametrization of SPEM in the simulated environment (see Section 8.1).

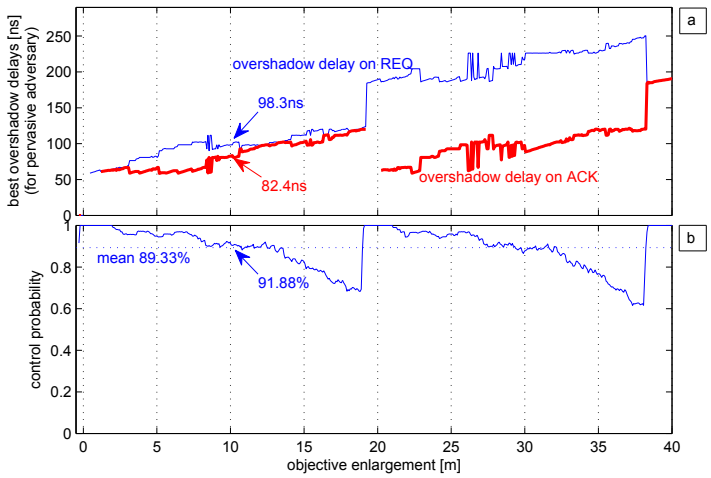

Fig. 8. Best tactics for an overshadow attack against JBSF: (a) best overshadow delays on ACK and REQ; (b) resulting enlargement control probability.

of the control probability is roughly periodic, with a period of $T_{pr} \cdot c/2 = 19.2\,\mathrm{m}$. This is due to the periodic structure of $S_{FE}(t)$.

To explain the results of Figures 8(a) and 8(b) it is necessary to consider the channel model of a typical indoor environment. Due to the presence of many reflecting objects, the multipath echoes arrive at the receiver grouped into short trains called *clusters* [Molisch et al. 2006]. If the overshadow delay is such that the leftward jump falls within a cluster, then the noise threshold will be crossed at the very beginning of the rightward search (Figure 9). In this situation, the attack falls into Case 2, and the leading-peak search provides an estimate $\widehat{\tau}_{LP} \approx \tau_{HP} - T_{JB}$. Thus, the adversary obtains two controlled delays of about $T_e^{\mathrm{R}} \approx T_o^{\mathrm{R}} - T_{JB}$ on the REQ, and $T_e^{\mathrm{A}} \approx T_o^{\mathrm{A}} - T_{JB}$ on the ACK. Moreover, from Figure 8(a) we note that the best tactic is (almost always) to attack both the REQ and the ACK. This is because the clusters are more frequent at the beginning of the channel response. Thus, two small delays are more controllable than a single large one. This explains also the decreasing trend of the control probability (apart from the periodicity) shown in Figure 8(b). As an exception, with an objective enlargement between 0.0 and 1.2 meters, and between 19.2 and 20.2 meters, the adversary attacks only the REQ packet, aiming at Case 2. In these zones, because of the periodicity of $S_{FE}(t)$, the probability of controlling a single long delay is greater than that of controlling two short ones.

A more complete motivation of the best overshadow tactics shown in Figure 8(a) requires a deeper analysis of the JBSF algorithm under attack. This analysis includes studying the probabilities of Cases 1, 2, 3 with respect to the overshadow delay, and the probability density of the TOA error in the three cases. Moreover, we have to take into account several low-level factors, as the thermal noise characteristics, the shape of the UWB pulse $p(t)$, etc. This analysis deserves a dedicated study, which falls outside the scope of the present paper.

Figures 10(a) and 10(b) show, respectively, the best tactics and the corresponding control probability for the SB algorithm. As we can see from the plots, the average control probability against the SB algorithm is lower than that against the JBSF algorithm (15.90% versus 89.33%). This is due to the fact that JBSF is intrinsically weaker than SB. Indeed, with JBSF the adversary can control where the jump-back falls, which is (with high probability) where the search forward stops. On the other

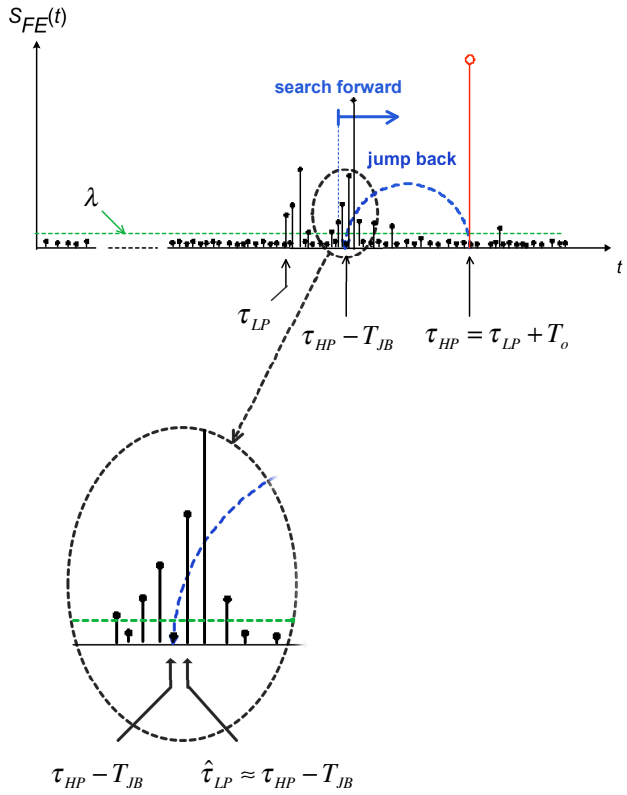

Fig. 9. Controlled overshadow attack against JBSF algorithm.

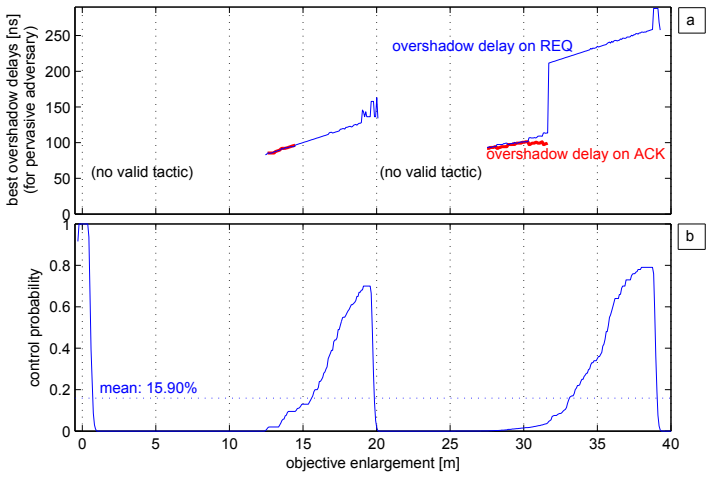

Fig. 10. Best tactics for an overshadow attack against SB: (a) best overshadow delays on ACK and REQ; (b) resulting enlargement control probability.

hand, with SB the adversary cannot control where the search-back stops, because this occurs at the first T_{SB}-long noise-only region. In addition, the adversary cannot cre-

ate an "artificial" noise-only region, since she cannot delete echoes from the received signal. As a consequence, the same overshadow attack may succeed against JBSF and have no effect against SB. For example, the attack shown in Figure 9 introduces a delay approximately equal to $T_o - T_{JB}$ with JBSF algorithm, but it is ineffective against SB, as shown in Figure 11.

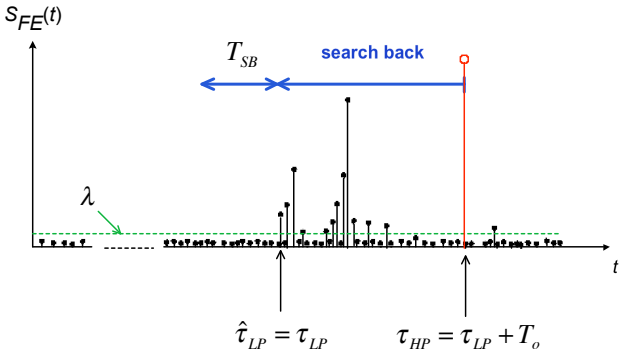

Fig. 11. Ineffective overshadow attack against SB algorithm.

With the SB algorithm, the enlargement attack is successful only if the legitimate echoes are sufficiently sparse. This condition occurs with higher probability at the end of the channel response (Figure 12).

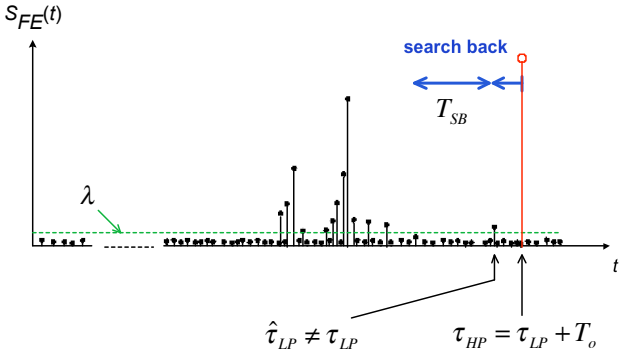

Fig. 12. Effective overshadow attack against SB algorithm.

Accordingly, as shown in Figure 10(b), the control probability has an increasing trend (apart from the periodicity stemming from the periodic structure of $S_{FE}(t)$). Moreover, it takes significant values only for objective enlargements greater than 12.6 meters. For smaller enlargements there is not any valid tactic, meaning that the control probability is zero or negligible. From Figure 10(a), we note that the most convenient tactic changes depending on the objective enlargement. With an objective enlargement between 14.5 and 19.2 meters and between 31.7 and 38.4 meters, the best tactic is to attack the REQ with an overshadow delay corresponding to the objective enlargement, and leave the ACK unattacked. This is because the adversary aims at falling into Case 3 for the REQ packet. Indeed, in SB the adversary cannot leverage Case 2 as in JBSF. Thus, Case 3 turns out to be the most convenient because it is the one giving most control to the adversary. On the other hand, with an objective enlargement between

12.6 and 14.5 meters, the adversary attacks both the REQ and the ACK with the same delay. She aims at falling into either Case 3 for the REQ and Case 1 for the ACK, or vice versa. In this zone, the probability of falling into Case 1 is quite high, so this tactic turns out to be more convenient. Finally, with an objective enlargement between 27.5 and 31.7 meters, the adversary attacks both the REQ and the ACK with two different delays, whose sum corresponds to the objective enlargement. She aims at falling into Case 3 both for the REQ and the ACK. Indeed, in this zone the probability of controlling two short delays is greater than that of controlling a single long one. Even in SB, a more complete motivation of the best overshadow tactics requires a deeper analysis of the TOA estimation algorithm under attack, which falls outside the present scope.

To sum up, the outcome of an overshadow attack is not always controllable, even with the best tactics. Moreover, the SB TOA estimation algorithm provides the adversary with less controllability than the JBSF one (15.90% versus 89.33%). This makes SB more promising for security-focused applications.

7. SPEM

In this section we introduce SPEM (*Secure Positioning through Enlargement Miscontrol detection*), a range-based secure positioning system leveraging the difficulty of the adversary to control the effect of the enlargement attacks. SPEM is based on distance bounding performed on the IEEE 802.15.4a UWB protocol. The basic idea is to detect the presence of an attacker by means of the residuals $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_N$ (see Figure 1), which have larger values in case of lower precision. In the honest case, a distance measured by a ranging operation will be the real distance plus a honest measurement error. In the presence of an attacker, it will be the false distance plus an enlargement control error. However, due to the difficulty for the adversary to control precisely the outcome of her attack, the expected value of the enlargement control error is greater than that of the honest measurement error. As a consequence, multilateration produces higher residuals in case of attack.

According to the Neyman-Pearson lemma [Lehmann and Romano 2006], the most sensitive test to detect the presence of an attack is the *likelihood-ratio test*. It is based on the probability density function of the residuals in the absence and in the presence of an attack. If the ratio of the two densities is lower than a given threshold, then the system detects the presence of an attack. Unfortunately, the probability density of the residuals in case of attack highly depends on the adversary's behavior, which is hard to predict. Designing a system assuming a precise behavior of the adversary is not a good security practice in general. We preferred a simpler criterion, based on a threshold on the value of the residuals. Such a threshold is tailored on the value of the residuals in the honest case, so no assumption is made on the behavior of the adversary to design the system.

More precisely, SPEM detects an attack when at least one of the residuals exceeds (in absolute value) a *residual threshold* (ε_{max}). When an attack is detected, the position measurement is discarded as untrusted. The value of the residual threshold is a design parameter. It constitutes a trade-off between sensitivity (i.e., probability of true positive in the presence of an attack), and specificity (i.e., probability of true negative in the absence of an attack). We suppose that the system has a table of pre-computed residual thresholds, depending on the number of involved anchors.

Note that also in verifiable multilateration the residuals are compared to a threshold, and the measured position is discarded if the threshold is exceeded [Čapkun and Hubaux 2006]. However this is done for a different objective than in SPEM. Indeed, in verifiable multilateration the residuals are large because the adversary *cannot perform reduction attacks*, while in SPEM the residuals are large because the adversary *cannot*

control enlargement attacks. This allows us to accept honest positions also outside the polygon formed by the anchors, which would be rejected by verifiable multilateration. As a result, SPEM can cover the same area with far less anchors, and avoids the need of deploying additional anchors at the borders of the deployment area, as it will be shown later.

The number of anchors involved in the multilateration affects the security of the system. Indeed, the more anchors are involved, the more enlargement attacks the adversary has to control, and thus the lower probability she has to succeed. SPEM establishes a *safe number of anchors* (N_{safe}) to assure a given level of security.

7.1. Improving SPEM by distance bounding repetition

It is possible to improve the security of SPEM by increasing the precision of the distance estimates. Indeed, more precise measurements produce smaller residuals, which in turn permit us to narrow the residual threshold. A simple way to achieve this is to perform more than one distance measurement at each anchor, and then to compute an estimate of the distance as the arithmetic mean of all the measurements (*distance bounding repetition*). Note that the system increases its measurement precision, but the adversary does not increase her control precision in the same way. Indeed, the control error does not tend to zero as the number of repetitions grows, as it depends mainly on the channel response, which does not change sensibly from a repetition to another. This is a reasonable hypothesis, since the whole series of distance bounding protocols takes less than 80 milliseconds to run. The synchronization header (SHR) takes the biggest part of the time: 4.1 milliseconds for each UWB packet with typical values of the IEEE 802.15.4a UWB parameters [IEEE Computer Society 2007]. The following PHR and PSDU parts are much shorter (tens of microseconds at a rate of $0.85\,\mathrm{Mb/s}$) and can be neglected. The $\mathrm{auth}(\cdot)$ function can be computed efficiently via hardware. For example, the IEEE 802.15.4-compliant CC2420 chip by Texas Instruments is capable of computing a 64-bit CBC-MAC in about 0.22 milliseconds [Roman et al. 2007]. Note that the AUTH packet can be sent in piggyback to the payload of the ACK packet, so we need only two UWB packets to implement the distance bounding protocol. Thus, a single distance bounding execution stays within 10 milliseconds. Supposing 8 repetitions, the whole series of distance bounding protocols takes less than 80 milliseconds. It is safe to suppose that the UWB channel does not change sensibly in such a short period. By repeating the distance bounding we consume more energy, but we increase the resistance against spoofing attacks.

As an additional security precaution, we check for anomalous variations on the measured distance from a repetition to another. In particular, we detect an attack if the dynamic of the distance estimates is beyond $2e_{max}$. This is to avoid that the adversary introduces different enlargements on different repetitions, in such a way to make the average distance coincide with the desired false distance.

7.2. Final algorithm description

Algorithm 1 shows a pseudo-code description of SPEM. First of all, we discover the set of reachable anchors ($\mathcal{A}_{prox}$) with a simple beacon-based method (Lines 2–3). We suppose that the anchor discovery is initiated by the sensor. However, the algorithm can be extended to let the infrastructure start it. Note that the anchor discovery does not need to be performed in a secure way. In other words, $\mathcal{A}_{prox}$ does not need to be the set of *truly* near anchors. Then, we perform the distance bounding protocols (Lines 6–8). We indicate with k the number of distance bounding repetitions for each anchor, and by $\widehat{d}_i^{(1)}, \widehat{d}_i^{(2)}, \ldots, \widehat{d}_i^{(k)}$ the distance estimates at the anchor A_i. We check for anomalous variations on the measured distances (Line 9). If no anomalous variation is detected, we av-

Algorithm 1 SPEM algorithm

1: **procedure** SPEM(S, N_{safe}, k)
2: S sends a broadcast beacon
3: $\mathcal{A}_{prox} \leftarrow \{A_i$ that answered the beacon$\}$
4: $N \leftarrow |\mathcal{A}_{prox}|$
5: **for all** $A_i \in \mathcal{A}_{prox}$ **do**
6: **for** $j \leftarrow 1 \dots k$ **do**
7: $\widehat{d}_i^{(j)} \leftarrow$ distance-bounding(A_i)
8: **end for**
9: **if** $\max_j(\widehat{d}_i^{(j)}) - \min_j(\widehat{d}_i^{(j)}) > 2e_{max}$ **then**
10: **return** "reject position measurement"
11: **end if**
12: $\widehat{d}_i \leftarrow \frac{1}{k}\sum_j \widehat{d}_i^{(j)}$
13: **end for**
14: $\langle \widehat{S}, \varepsilon_i \rangle \leftarrow$ multilaterate$(\mathcal{A}_{prox}, \widehat{d}_1, \widehat{d}_2, \dots)$
15: **if** $\widehat{S} \in$ polygon$(\mathcal{A}_{prox})$ **and** $\forall i\ |\varepsilon_i| \leq \varepsilon_{max}(N,k)$ **then**
16: **return** $\widehat{S}$
17: **end if**
18: **if** $N \geq N_{\text{safe}}$ **and** $\forall i\ |\varepsilon_i| \leq \varepsilon_{max}(N,k)$ **then**
19: **return** $\widehat{S}$
20: **end if**
21: **return** "reject position measurement"
22: **end procedure**

erage the measured distances (Line 12). Successively, we solve the least-squared-error multilateration problem (Line 14). We accept the positions that meet the conditions of verifiable multilateration (Line 15). In addition, we also accept positions measured by means of at least N_{safe} anchors and whose residuals are lower than ε_{max} (Line 18). In this way we make sure to cover a larger area than verifiable multilateration. We use the notation $\varepsilon_{max}(N,k)$ because the value of the residual threshold depends on the number of involved anchors and on the number of distance bounding repetitions.

8. PARAMETRIZATION AND EVALUATION

Choosing a good set of parameters for SPEM (residual threshold, safe number of anchors, and number of repetitions) is important to guarantee the desired level of security. Nevertheless, this parametrization is not a trivial task, since it depends on many factors: the deployment environment, the channel statistics, and the TOA estimation algorithm influence the security. To solve the problem, we developed and made publicly available the *SPEM Parametrization and Evaluation Framework*[3], which helps the user to configure SPEM and evaluate its security. The framework is written in the Matlab language. It is highly configurable, and can parametrize SPEM in a generic combination of deployment environment, channel statistics, and TOA estimation algorithm. It consists of a pair of tools: (1) the *residual threshold estimator*, and (2) the *security estimator*. The residual threshold estimator determines the value of ε_{max} that offers a given probability of true negative in the absence of an attack (*specificity*). It works by simulating a number of non-attacked multilateration scenarios. A multilateration scenario consists of the position of a sensor plus the positions of N anchors

[3]www.iet.unipi.it/g.dini/download/pubs/SpemParametrization.zip.

reachable by that sensor. Note that the total number of sensors in the WSN is not a parameter of the residual threshold estimator. This is because SPEM measures the position of each sensor separately, so the presence of the other sensors does not influence the security of the single execution of SPEM. For each scenario, the tool simulates the execution of SPEM *without any security check*, i.e., with $e_{max} = +\infty$, $\varepsilon_{max} = +\infty$, and $N_{\text{safe}} = 3$. Then, it saves the residuals of the multilateration problem. Once the residuals of all the scenarios have been collected, the tool determines the residual threshold that accepts a percentage of them, equal to the desired specificity. The user must run the tool several times, varying the number of anchors involved in SPEM and the number of repetitions, in such a way to obtain the table of residual thresholds that will be used by SPEM. As a side result, the residual threshold estimator determines also the honest precision e_{max}. To do this, it saves all the distance measurement errors generated during the simulations, and takes the value of e_{max} that comprises 99.9% of them (as specified in (8)). Such a honest precision will be used by SPEM to check for anomalous variations on the measured distances throughout the distance bounding repetitions (see Algorithm 1, Line 9).

The security estimator determines the adversary's success probability against SPEM. It works by simulating a number of attacked multilateration scenarios. For each scenario, the tool simulates the execution of SPEM, now with the residual threshold and the honest precision determined by the residual threshold estimator. The security check on the safe number of anchors is still disabled, i.e., $N_{\text{safe}} = 3$. The adversary can be either random-objective or best-objective. The best-objective adversary chooses her false position based on precomputed values of the control enlargement probability. For each distance bounding performed by SPEM, the adversary mounts a pair of overshadow attacks against the REQ and the ACK, following some precomputed best overshadow tactics. After the execution of SPEM, the tool saves the outcome of the attack: success or failure. Once the outcomes of all the scenarios have been collected, it determines the adversary's success probability. The user must run the tool several times, varying the number of anchors involved in SPEM and the number of repetitions, in such a way to obtain a table of success probabilities. From this table, the user can extract a pair of parameters $\langle N_{\text{safe}}, k \rangle$ that gives the desired security level in terms of success probability. If more than one pair gives the desired security level, the user can choose among them based on other performance objectives.

The framework is highly configurable, and can parametrize SPEM in a generic combination of deployment environment, channel statistics, and TOA estimation algorithm. The user must provide the framework for the best overshadow tactics, the trend of the enlargement control probability, and four routines: (i) the *scenario generator*, (ii) the *false position generator*, (iii) the *measurement error generator*, (iv) the *enlargement control error generator*. These routines are Matlab functions coded by the user and passed to the framework via their function handles. They are invoked as callbacks by the tools during their simulation cycles. Note that this gives an extreme flexibility to the framework. The scenario generator allows the user to model scenarios of any shape, in which the anchors have non-circular or time-varying coverage areas as well. The measurement error generator and the enlargement control error generator allow the user to generate synthetical errors, or to extract them from datasets stemming from real experiments. The user can also take into account the effect of the signal-to-noise ratio on the precision of the distance estimation, both in a non-attacked and attacked distance bounding.

We used the SPEM Parametrization and Evaluation Framework to configure SPEM (1) in a simulated environment with simulated UWB transceivers, and (2) in a real environment with real UWB transceivers.

8.1. SPEM in a simulated environment

We parametrized and evaluated the security of SPEM in a simulated environment, using simulated UWB transceivers implementing the search-back TOA estimation algorithm described in Section 5. For the generation of the multilateration scenarios, we did not suppose a specific area shape with a given size and a given number of anchors inside. Rather, we first placed the sensor in a fixed position and then we placed N anchors in random positions around it, within a communication range of 40 meters. This allows us to parametrize SPEM independently of the shape and the size of the deployment area. To generate the measurement errors of the non-attacked distance bounding protocols, we simulated the two-way ranging mechanism at the signal level, including the TOA estimation algorithm for both receivers. We simulated the $S_{FE}(t)$ computation and the fine time acquisition. The TOA estimation error is affected by the signal-to-noise ratio (as shown in Figure 6). To determine the signal-to-noise ratio, we used the standard path-loss model of CM1 [Molisch et al. 2006]. The resulting signal-to-noise ratio $\mathrm{SNR}(d)$ at a distance d follows the law:

$$\mathrm{SNR}(d) = \mathrm{SNR}_0 - 10n \log_{10}(d/d_0), \tag{17}$$

where d_0 is a reference distance, SNR_0 is the signal-to-noise ratio at the reference distance, and $n = 1.79$ is the path-loss exponent. The transmission power is supposed to be such that the signal arrives with a reasonable SNR (30 decibels) even at the maximum distance (40 meters). So we put $d_0 = 40\,\mathrm{m}$ and $\mathrm{SNR}_0 = 30\,\mathrm{dB}$. We used the residual threshold estimator with a specificity of 99.9%, which means that we want SPEM to accept 99.9% of the honest position measurements. Table I shows the resulting residual thresholds as a function of the number of anchors and the number of distance bounding repetitions. Every threshold stems from 100,000 randomly generated multilateration scenarios. All 95%-confidence intervals are within $-1.25\,\mathrm{cm}$ and $+1.33\,\mathrm{cm}$.

Table I. SPEM residual thresholds (ε_{max}) for the simulated environment

dist. bound. repetitions:	number of anchors: 3	4	5	6
1	32.58cm	37.28cm	42.22cm	44.16cm
2	27.71cm	31.65cm	34.20cm	35.87cm
4	25.50cm	28.83cm	30.93cm	32.24cm
8	24.50cm	27.91cm	29.61cm	31.19cm

It can be seen that the threshold value decreases as the number of distance bounding repetitions increases, because the distance measurements become more and more precise. The honest precision computed by the residual threshold estimator is $e_{max} = 39.7\,\mathrm{cm}$.

We used the security estimator to quantify the attack success probability. For each generated multilateration scenario, we simulated a random-objective adversary and a best-objective adversary. We supposed the minimal spoofing distance (as specified in (5)) to be $d_{ms} = 1\,\mathrm{m}$. For each distance bounding, we simulated an overshadow attack following the best overshadow tactics against SB shown in Section 6.1 (see Figure 10(a)). We considered four different values of the number of anchors covering the sensor position, namely $N = 3, 4, 5, 6$, and six different values of the number of distance bounding repetitions, namely $k = 1, 2, 4, 8$. Figures 13 and 14 show the success probability of a random-objective and a best-objective adversary, respectively, as a function of the number of distance bounding repetitions at each anchor. It can be seen that the success probability of an attack decreases with the growing of the number of anchors

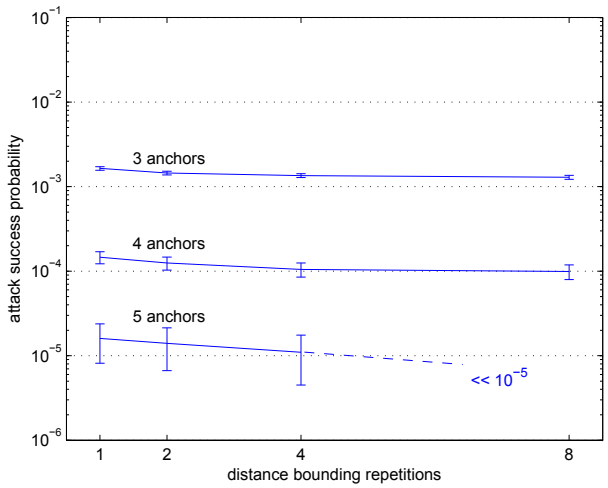

Fig. 13. Mean success probability of random-objective adversary in the simulated environment. Each value comes from 1,000,000 Monte Carlo runs. 99%-confidence intervals are displayed in error bars.

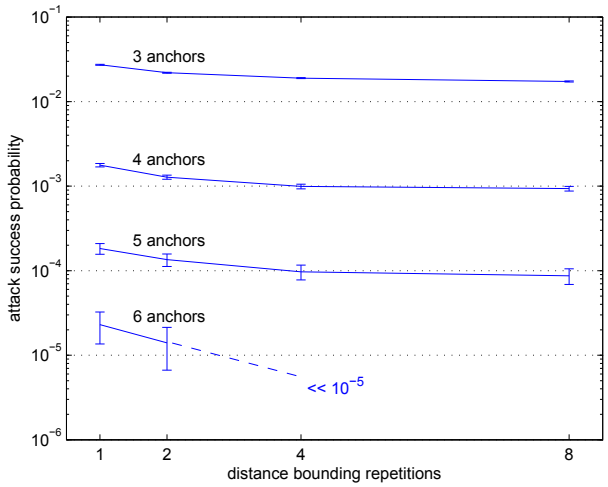

Fig. 14. Mean success probability of best-objective adversary in the simulated environment.

involved in the multilateration, and with the growing of the number of distance bounding repetitions. Based on the results of Figures 13 and 14, it is possible to tailor the parameters of SPEM (N_{safe} and k) in order to assure a given level of security. As an example, Table II shows possible configurations to offer a given level of security against a best-objective adversary. For each security level ("attack success probability" column), Table II gives two possible SPEM configurations. The first one ("anchor saving configuration" column) aims at minimizing the number of needed anchors (i.e., N_{safe}). The second one ("energy saving configuration" column) aims at minimizing the number of needed distance bounding operations (i.e., $N_{\text{safe}} \cdot k$). For example, if we want to minimize the number of anchors while limiting the attack success probability to 10^{-4}, we

Table II. SPEM configuration parameters

attack success probability:	anchor saving configuration:		energy saving configuration:	
$< 10^{-2}$	$N_{\text{safe}} = 4$	$k = 1$	$N_{\text{safe}} = 4$	$k = 1$
$< 10^{-3}$	$N_{\text{safe}} = 4$	$k = 8$	$N_{\text{safe}} = 5$	$k = 1$
$< 10^{-4}$	$N_{\text{safe}} = 5$	$k = 8$	$N_{\text{safe}} = 6$	$k = 1$
$< 10^{-5}$	$N_{\text{safe}} = 6$	$k = 4$	$N_{\text{safe}} = 6$	$k = 4$

can set $N_{\text{safe}} = 5$ and $k = 8$. In this way, a secure positioning operation will require 40 distance bounding operations. On the other hand, if we want to minimize the energy consumption at the sensor (and thus prolong the WSN lifetime) at the same level of security, we can set $N_{\text{safe}} = 6$ and $k = 1$. In this way, a secure positioning operation will require only 6 distance bounding operations.

Figure 15 shows the mean error on the position estimation as a function of the distance bounding repetitions, for different values of the number of anchors.

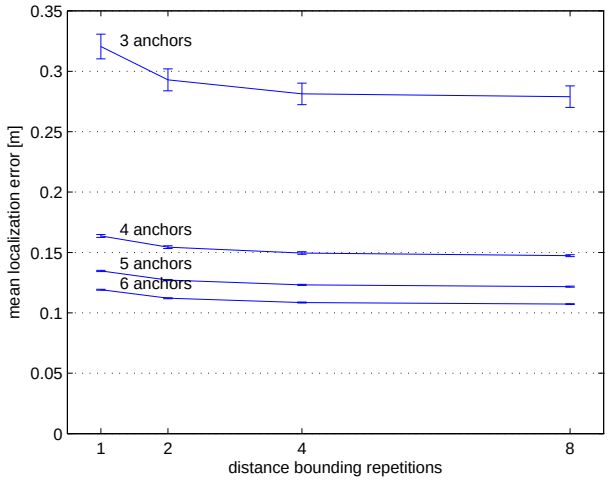

Fig. 15. Mean localization error in the simulated environment.

As expected, localization precision increases with the anchors and the repetitions. The results in Figures 13, 14 and 15 suggest that the security level is strictly related to the localization precision. In general, a more precise SPEM system is also more secure. This is evident from Figure 16, which shows the trend of the localization error with respect to the security level, under the anchor saving configuration and the energy saving configuration. We note that the energy saving configuration is better in terms of localization precision with the same security. This is because it involves more anchors in the multilateration, and this is advantageous for the precision.

8.2. SPEM in a real environment

We configured and evaluated SPEM in a real indoor environment, using real UWB transceivers. The environment is a room of the Department of Information Engineering (University of Pisa), shown in Figure 17. Such an environment is characterized by non-ideal non-line-of-sight (NLOS) propagation conditions due to the internal walls, that reduce the speed of the electromagnetic wave and hence introduce an extra delay in the direct path propagation time. Also, we experimentally noticed that the signal

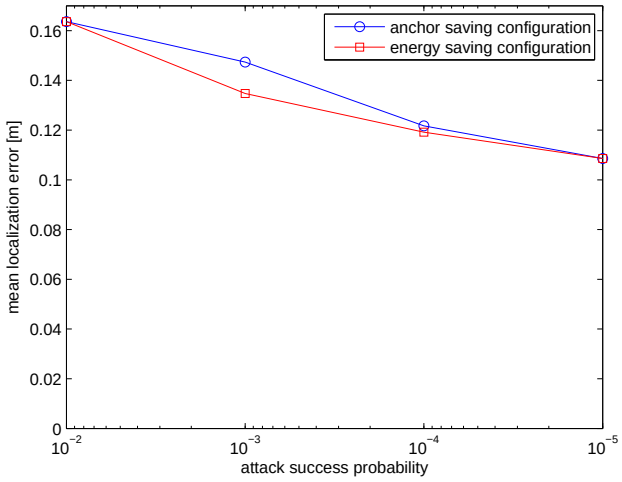

Fig. 16. Localization error with respect to the security level.

Fig. 17. Picture of the environment used for SPEM parametrization.

along the direct path is not totally blocked, but only attenuated by the internal walls. As a consequence, every anchor is reachable from every point in the room. Finally, the adversarial channel is not perfectly Gaussian due to spurious reflections. Figure 18 shows a map of the environment. The black squares indicated with letters represent the anchors (6 anchors in total). Note that, with this anchor deployment, the light gray areas could not be covered by verifiable multilateration because outside the convex hull of the anchors. In contrast, SPEM covers these areas as well.

As UWB transceivers, we used a pair of EVB1000 evaluation boards by DecaWave Ltd. [DecaWave Ltd. 2016], which are capable of implementing the IEEE 802.15.4a UWB protocol. These boards are shown in Figure 19. Note that a pair of EVB1000 boards are sufficient to parametrize and evaluate a SPEM system in a real indoor environment. On the other hand, a complete *implementation* of a 6-anchor SPEM system would have required at least 6 EVB1000 boards for the anchors plus 1 for the sensor. The EVB1000 boards do not offer the necessary functionalities to implement a real overshadow adversary. In particular, they cannot synchronize with an ongoing com-

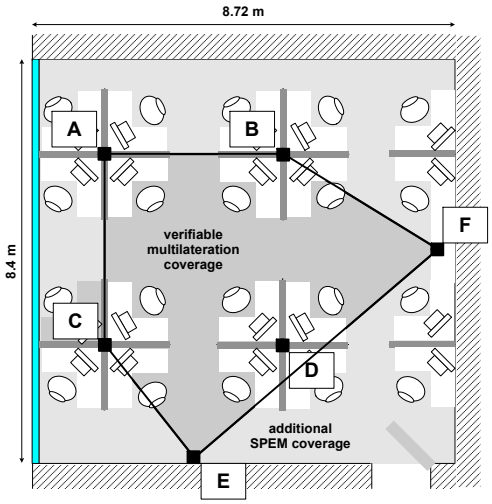

Fig. 18. Map of the environment used for SPEM parametrization.

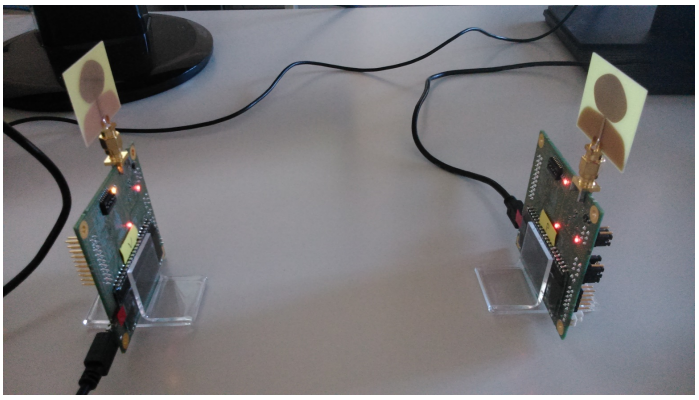

Fig. 19. DecaWave EVB1000 UWB boards used in the real environment.

munication and then replay it with a delay while it is still going. However, they permit us to simulate the presence of an overshadow attack on the real UWB channels. This is done by means of the diagnostic features of the EVB1000 boards, which can return the channel response recorded during the execution of a distance bounding protocol. Such a channel response can be considered the analogous of the $S_{FE}(t)$ function of our simulated receivers, except that it does not have a periodic structure. To simulate the attack, we proceeded in the following way. First, we performed a distance bounding protocol, with an EVB1000 board representing the anchor and the other one the sensor. We extracted the channel response, which represents the honest channel. Then, we performed another distance bounding protocol, with an EVB1000 board representing the sensor (or, equivalently, the anchor) and the other one the adversary. The boards have been positioned at 20cm from each other, in order to realize the (quasi-)Gaussian channel of the adversary. We extracted the channel response, which represents the adversarial channel. Finally, we superimposed the two channels with a delay of T_o between each other (Figure 20). This represents the attacked channel. On this channel we executed the fine time acquisition phase to estimate the TOA in the presence of an

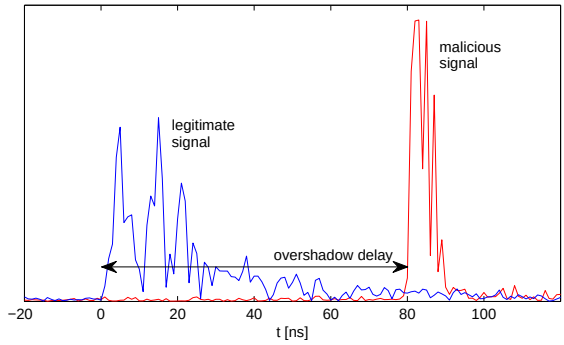

Fig. 20. Example of overshadow attack on real UWB channels with $T_o = 80\,\mathrm{ns}$. The ordinate is unitless, because DecaWave does not release documentation about the unit of measurement of its $S_{FE}(t)$.

overshadow attack. DecaWave does not release the details of the fine time acquisition performed by its EVB1000 boards[4]. We thus implemented the simple search-back algorithm described in Section 5. The threshold has been set by measuring the noise level on a noise-only region of the channel response, in such a way to have 10^{-5} false alarm probability (similarly to what we supposed for the simulated receivers, see Section 5). By realizing the overshadow attack in this way, the only simulated aspects are the superimposition of the legitimate and the malicious signals, and the fine time acquisition phase. All the other aspects, including the UWB channel in the real environment, the adversarial channel, and the $S_{FE}(t)$ computation are based on the physical experiments. We mounted a total of 240 attacks: 30 honest channels (given by 6 anchor positions, each of which with 5 random sensor positions) times 8 distance bounding repetitions for each channel. Using the outcomes of these attacks, we computed the best overshadow tactics, shown in Figure 21(a), and the enlargement control probability, shown in Figure 21(b). Note that the trends are quite similar to the tactics against the simulated search-back receivers (see Figure 10), except that they are not periodic after $19.2\,\mathrm{m}$. This is due to the non-periodicity of the $S_{FE}(t)$ provided by the EVB1000 boards.

We have employed the SPEM Parametrization and Evaluation Framework to configure SPEM in this environment. Table III shows the resulting residual threshold table. Every threshold stems from 100,000 randomly generated multilateration scenarios. All 95%-confidence intervals are within $-1.29\,\mathrm{cm}$ and $+2.10\,\mathrm{cm}$.

Table III. SPEM residual thresholds (ε_{max}) in the real environment

dist. bound. repetitions:	number of anchors:			
	3	4	5	6
1	59.21cm	67.88cm	75.23cm	79.86cm
2	38.92cm	44.02cm	48.22cm	51.44cm
4	31.08cm	35.98cm	41.32cm	44.36cm
8	28.90cm	34.39cm	40.38cm	43.40cm

It can be noted that the thresholds are generally larger than those found in the simulated environment, as the single ranging operations have a worse precision. This is probably due to the sources of non-ideality of the real environment. In particular, we observed that the precision gets particularly bad in NLOS conditions, when the direct

[4]From private communication. However, we inferred from the name of some configuration registers that it is a threshold-based algorithm, like that considered in this paper.

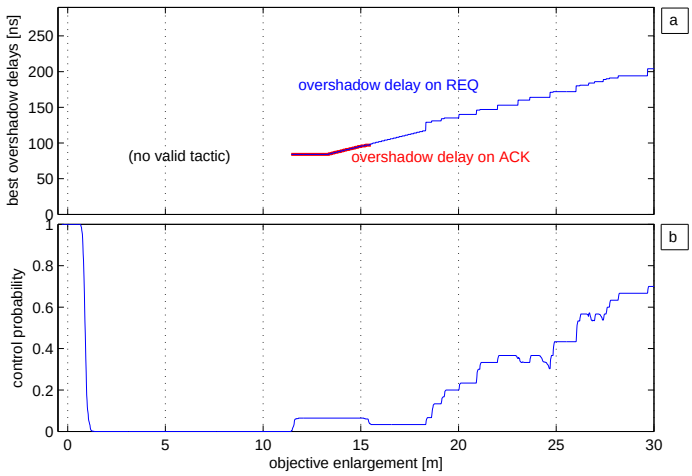

Fig. 21. Best tactics for an overshadow attack against EVB1000 transceivers with SB: (a) best overshadow delays on ACK and REQ; (b) resulting enlargement control probability.

path is attenuated by more than one thin obstacle or by a single thick obstacle. The honest precision computed by the residual threshold estimator is $e_{max} = 96.8\,\mathrm{cm}$.

We used the security estimator to quantify the attack success probability. For each generated multilateration scenario, we simulated a best-objective adversary. We supposed the minimal spoofing distance (as specified in (5)) to be $d_{ms} = 1\mathrm{m}$. Figure 22 shows the success probability of a best-objective adversary in the real environment.

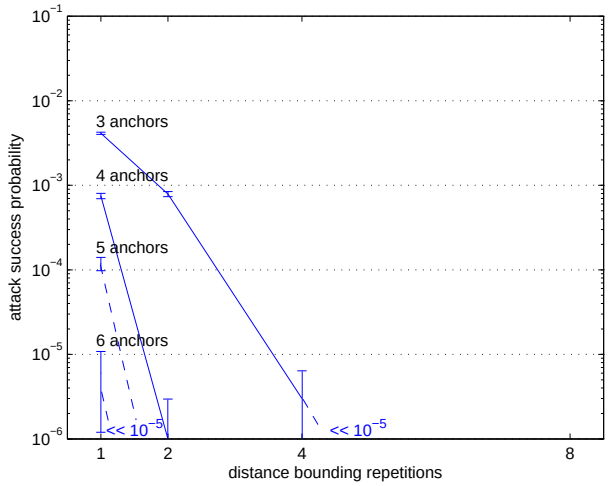

Fig. 22. Mean success probability of best-objective adversary in the real environment. Each value comes from 1,000,000 Monte Carlo runs. 95%-confidence intervals are displayed in error bars.

It can be noted that, despite the residual thresholds are larger, the success probability in the real environment is smaller. This is due to the size of the room, which does not leave much freedom to the adversary in the choice of the false position. On the other

hand, in the simulated environment, the adversary had more freedom to choose a more convenient false position, which gave her a greater success probability.

Table IV shows the final SPEM parametrization (N_{safe} and k) for the real environment with respect to the desired security level, in both the anchor-saving configuration and the energy-saving configuration.

Table IV. SPEM configuration parameters for the real environment

attack success probability:	anchor saving configuration:		energy saving configuration:	
$< 10^{-2}$	$N_{\text{safe}} = 3$	$k = 1$	$N_{\text{safe}} = 3$	$k = 1$
$< 10^{-3}$	$N_{\text{safe}} = 3$	$k = 2$	$N_{\text{safe}} = 4$	$k = 1$
$< 10^{-4}$	$N_{\text{safe}} = 3$	$k = 4$	$N_{\text{safe}} = 6$	$k = 1$
$< 10^{-5}$	$N_{\text{safe}} = 3$	$k = 4$	$N_{\text{safe}} = 6$	$k = 1$

9. INFRASTRUCTURE SCALABILITY

SPEM permits us to cover the same area with less anchors with respect to verifiable multilateration [Čapkun and Hubaux 2006], while maintaining a high level of security. From Algorithm 1 we can see that, given a set of anchors $\{A_i\}$, a sensor S is covered by SPEM iff one of the following conditions is true: (1) there exist at least three anchors within the communication range of the sensor and the polygon formed by them contains S; or (2) there exist at least N_{safe} anchors within the communication range. The first condition is present also in verifiable multilateration. The second condition gives us additional coverage. An example of this is given in Figure 23.

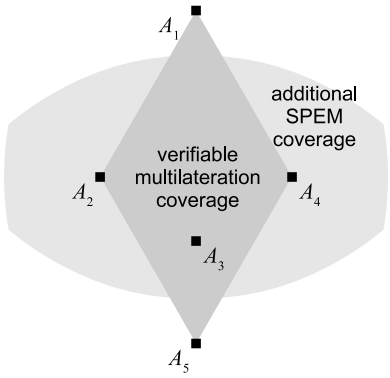

Fig. 23. Coverage area comparison with $N_{\text{safe}} = 5$. The dark and the light grey regions represent respectively the area covered with verifiable multilateration and the additional coverage provided by SPEM.

The additional coverage lowers the number of anchors that must be deployed in order to cover a given area. Figure 24 shows the mean number of anchors needed to cover 90% of a square area. The anchors are positioned randomly. The communication range is 40 meters. The two "SPEM" curves represent the number of anchors needed by SPEM configured as in the simulated environment (see Table II) with an attack success probability $< 10^{-4}$, respectively under the energy saving configuration ($N_{\text{safe}} = 6$, $k = 1$) and the anchor saving configuration ($N_{\text{safe}} = 5$, $k = 8$). It can be seen that SPEM greatly improves the anchor scalability, and it gets close to the theoretical limit of the classic (insecure) multilateration with $N = 3$ anchors. This big difference (-93% to -71% of needed anchors under the anchor saving configuration) is also due to the fact that verifiable multilateration cannot provide coverage outside the polygon formed

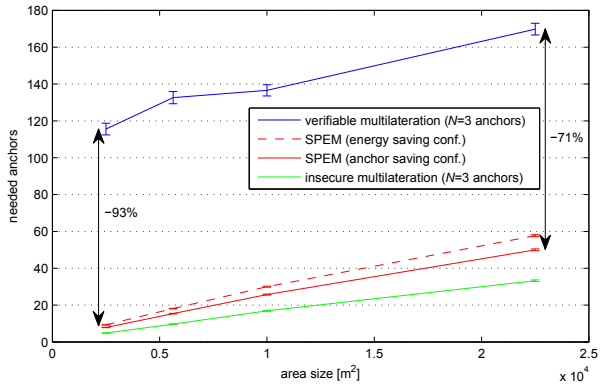

Fig. 24. Mean number of anchors needed to cover 90% of the area. Each value comes from 1,000 Monte Carlo runs. 99%-confidence intervals are displayed in error bars.

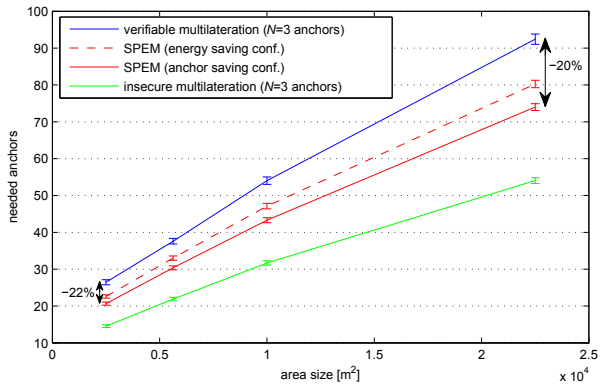

Fig. 25. Mean number of anchors needed to cover 90% of the area with external deployment.

by the anchors. Thus covering the zones at the border of the deployment area is quite hard. To avoid the problem, for the evaluation of verifiable multilateration Čapkun and Hubaux [2006] used a special deployment scheme. In this scheme, the anchors are randomly deployed in the area and also in external bands with width equal to the communication range. Deploying anchors also in the external bands can significantly increase the total number of anchors, their density being equal. Moreover, the anchors could have been already deployed for other purposes, and the user wanting to perform secure positioning could not add new anchors or change the positions of the existing ones. Figure 25 shows the mean number of anchors needed to cover 90% of a square area with such a deployment scheme. It can be seen that SPEM improves the anchor scalability (-22% to -20% of needed anchors under the anchor saving configuration) also with this deployment scheme, which is ad-hoc for verifiable multilateration. It is worth observing that the total number of anchors can be considerably reduced also with the energy saving configuration of SPEM, as shown by the results of Figures 24 and 25. Accordingly, we can use this configuration to save energy at the sensors (end hence prolonging the lifetime of the WSN) while at the same time reducing the number of anchors with respect to state-of-the-art solutions.

10. CONCLUSIONS

In this paper we proposed SPEM, a secure positioning scheme based on UWB distance bounding and multilateration. The basic idea is to detect uncontrolled enlargement attacks by monitoring the accuracy of the position estimates, and by increasing the precision of the multilateration scheme. SPEM uses a distance bounding realized on IEEE 802.15.4a UWB [IEEE Computer Society 2007], which is capable of sub-meter precision at low energy costs [Zhang et al. 2009]. To evaluate the security of SPEM, we first performed thorough signal-level simulations of the UWB protocol exposed to an overshadow attack. Then, we simulated SPEM to estimate its performances in terms of security, precision, and anchor saving. We also evaluated the security of SPEM in the real field, with real UWB transceivers. We showed that it is possible to achieve a high level of security, while saving up to 93% of the anchors with respect to state-of-the-art solutions.

REFERENCES

Stefan Brands and David Chaum. 1993. Distance bounding protocols. In *Proceedings of the Workshop on the Theory and Application of Cryptographic Techniques (EUROCRYPT'93)*. 344–359.

Jerry T. Chiang, Jason J. Haas, Jihyuk Choi, and Yih-Chun Hu. 2012. Secure Location Verification Using Simultaneous Multilateration. *IEEE Transactions on Wireless Communications* 11, 2 (February 2012), 584–591.

Antonio A. D'Amico, Umberto Mengali, and Lorenzo Taponecco. 2008. Energy-Based TOA Estimation. *IEEE Transactions on Wireless Communications* 7, 3 (March 2008), 838–847.

Antonio A. D'Amico, Umberto Mengali, and Lorenzo Taponecco. 2010. TOA estimation with the IEEE 802.15.4a standard. *IEEE Transactions on Wireless Communications* 9, 7 (2010), 2238–2247.

DecaWave Ltd. 2016. ScenSor EVK1000 Evaluation Kit. www.decawave.com/products/evk1000-evaluation-kit. (2016).

Gianluca Dini, Francesco Giurlanda, and Pericle Perazzo. 2013. SecDEv: Secure Distance Evaluation in Wireless Networks. In *Proceedings of the 9th International Conference on Networking and Services (ICNS'13)*. 207–212.

Ismail Guvenc, Zafer Sahinoglu, Andreas F Molisch, and Philip Orlik. 2005. Non-coherent TOA estimation in IR-UWB systems with different signal waveforms. In *Proceedings of the 1st IEEE/CreateNet International Workshop on Ultrawideband Wireless Networking (UWBNETS'05)*. 245–251.

Wen Hu, Hailun Tan, Peter Corke, Wen Chan Shih, and Sanjay Jha. 2010. Toward trusted wireless sensor networks. *ACM Transactions on Sensor Networks* 7, 1 (August 2010), 1–25.

Yih-Chun Hu, Adrian Perrig, and David B. Johnson. 2003. Packet leashes: a defense against wormhole attacks in wireless networks. In *Proceeding of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'03)*, Vol. 3. 1976–1986.

Todd E. Humphreys, Brent M. Ledvina, Mark L. Psiaki, Brady W. O'Hanlon, and Paul M. Kintner Jr. 2008. Assessing the spoofing threat: Development of a portable GPS civilian spoofer. In *Proceedings of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS'08)*, Vol. 55. 2314–2325.

IEEE Computer Society. 2007. IEEE Std 802.15.4a-2007 (Amendment 1: Add Alternate PHYs). (2007).

Loukas Lazos and Radha Poovendran. 2005. SeRLoc: Robust localization for wireless sensor networks. *ACM Transactions on Sensor Networks* 1, 1 (2005), 73–100.

Erich L. Lehmann and Joseph P. Romano. 2006. *Testing statistical hypotheses*. Springer Science & Business Media.

Andreas F. Molisch, Dajana Cassioli, Chia-Chin Chong, Shahriar Emami, Andrew Fort, Balakrishnan Kannan, Johan Karedal, Juergen Kunisch, Hans Gregory Schantz, Kazimierz Siwiak, and others. 2006. A comprehensive standardized model for ultrawideband propagation channels. *IEEE Transactions on Antennas and Propagation* 54, 11 (2006), 3151–3166.

Taejoon Park and Kang G. Shin. 2009. Attack-tolerant Localization via Iterative Verification of Locations in Sensor Networks. *ACM Transactions on Embedded Computing Systems* 8, 1, Article 2 (2009), 24 pages.

Pericle Perazzo and Gianluca Dini. 2015. Secure Positioning with Non-Ideal Distance Bounding Protocols. In *Proceedings of the 20th IEEE Symposium on Computers and Communications (ISCC'15), to appear, Larnaca (Cyprus), 6-9 July*. 1–8.

Marcin Poturalski, Manuel Flury, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. 2012. On Secure and Precise IR-UWB Ranging. *IEEE Transactions on Wireless Communications* 11, 3 (March 2012), 1087–1099.

Marcin Poturalski, Manuel Flury, Panos Papadimitrios, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. 2011. Distance bounding with IEEE 802.15.4a: Attacks and countermeasures. *IEEE Transactions on Wireless Communications* 10, 4 (2011), 1334–1344.

Rodrigo Roman, Cristina Alcaraz, and Javier Lopez. 2007. A survey of cryptographic primitives and implementations for hardware-constrained sensor network nodes. *Mobile Networks and Applications* 12, 4 (2007), 231–244.

Naveen Sastry, Umesh Shankar, and David Wagner. 2003. Secure verification of location claims. In *Proceedings of the 2nd ACM Workshop on Wireless Security (WiSE'03)*. 1–10.

Arvind Seshadri, Adrian Perrig, Leendert Van Doorn, and Pradeep Khosla. 2004. SWATT: Software-based attestation for embedded devices. In *Proceedings of the 2004 IEEE Symposium on Security and Privacy (S&P'04)*. 272–282.

Ian Sharp and Kegen Yu. 2014. Indoor TOA Error Measurement, Modeling, and Analysis. *IEEE Transactions on Instrumentation and Measurement* 63, 9 (September 2014), 2129–2144.

Lorenzo Taponecco, Pericle Perazzo, Antonio D'Amico, Gianluca Dini, and others. 2014. On the Feasibility of Overshadow Enlargement Attack on IEEE 802.15.4a Distance Bounding. *IEEE Communications Letters* 18, 2 (2014), 257–260.

Srdjan Čapkun and Jean-Pierre Hubaux. 2006. Secure Positioning in Wireless Networks. *IEEE Journal on Selected Areas in Communications* 24, 2 (February 2006), 221–232.

Jinyun Zhang, Philip V Orlik, Zafer Sahinoglu, Andreas F Molisch, and Patrick Kinney. 2009. UWB systems for wireless sensor networks. *Proc. IEEE* 97, 2 (2009), 313–331.