Modeling Enlargement Attacks Against UWB Distance Bounding Protocols

Alberto Compagno, Mauro Conti, Antonio A. D'Amico, Gianluca Dini, Pericle Perazzo, Lorenzo Taponecco

***Abstract*—Distance bounding protocols make it possible to determine a trusted upper bound on the distance between two devices. Their key property is to resist *reduction attacks*, i.e., attacks aimed at reducing the distance measured by the protocol. Recently, researchers have focused also on *enlargement attacks*, aimed at enlarging the measured distance. Providing security against such attacks is important for secure positioning techniques. The contribution of this paper is to provide a probabilistic model for the success of an enlargement attack against a distance bounding protocol realized with the IEEE 802.15.4a UWB standard. The model captures several variables, like the propagation environment, the signal-to-noise ratio, and the time-of-arrival (TOA) estimation algorithm. We focus on non-coherent receivers, which can be used in low-cost low-power applications. We validate our model by comparison with physical-layer simulations and goodness-of-fit tests. The results show that our probabilistic model is sufficiently realistic to replace physical-layer simulations. Our model can be used to evaluate the security of the ranging/positioning solutions that can be subject to enlargement attacks. We expect that it will significantly facilitate future research on secure ranging and secure positioning.**

I. Introduction

Distance bounding protocols [1] are security protocols that make it possible to determine a trusted upper bound on the distance between two devices. They do this by measuring the round-trip time between two messages that are unpredictable for an adversary. The basic property of a distance bounding protocol is to resist *reduction attacks*, which aim at reducing the measured distance with respect to the real one. In short, the adversary cannot reduce the distance, because she should anticipate the messages which are instead unpredictable.

Recently, researchers have turned their attention also to the opposite kind of attacks: *enlargement attacks* [2], [3], [4], [5]. In these attacks, the adversary aims at enlarging the measured distance. The interest in solutions against enlargement attacks is growing because they open the door to more scalable secure positioning techniques [3], [4]. In this paper, we focus on distance bounding protocols performed with the IEEE 802.15.4a ultra-wideband (UWB) physical protocol [6]. IEEE 802.15.4a UWB has been the first standardized ultra-wideband protocol for precision ranging, capable of reaching sub-meter precision in distance estimations. It is one of the most convenient choices for future implementations of wireless distance bounding protocols [7]. We consider non-coherent architectures, which are designed to be used in low-cost low-power receivers, as indicated by the IEEE 802.15.4a UWB standard [6]. Non-coherent schemes are the most interesting ones for practical applications at the moment. Moreover, we focus on *external adversaries* only, i.e., we suppose that the devices that execute the distance bounding protocol are trusted.

To cause an enlargement on the measured distance, an adversary must introduce a delay in the round-trip time. One way in which an external adversary can do this is to mount an *overshadowing attack*. This attack is hard to detect because, unlike other attacks, it produces realistic enlargements without introducing an unrealistic quantity of energy in the channel. In the overshadowing attack, the adversary replays the messages sent by the legitimate devices with a certain delay and a greater power. In this way, she tries to make the victim receivers "hook" to the malicious signals instead of the legitimate ones.

Contribution This paper brings the following contributions:

- We provide a probabilistic model of the outcome of an overshadowing attack against a distance bounding protocol realized with IEEE 802.15.4a UWB. Our model takes into consideration several variables, like the propagation environment, the signal-to-noise ratio, and the time-of-arrival (TOA) estimation algorithm.
- We evaluate the soundness of our model by comparing it to attack outcomes generated by physical-layer simulations, and by performing goodness-of-fit tests. The results show that our model is sufficiently realistic to replace physical-layer simulations.
- We finally develop a Matlab tool based on our model, capable of simulating attacked and non-attacked TOA estimations. The tool allows researchers to evaluate the security of the ranging/positioning solutions that can be subject to enlargement attacks. We make such a tool available to the research community.

Organization The rest of the paper is organized as follows. Section II introduces the state of the art. Section III motivates the importance of overshadowing attacks. Section IV introduces the reference distance bounding protocol, the IEEE 802.15.4a UWB signal format, and the TOA estimation algorithm of the receiver. Section V analyzes the effects of an overshadowing attack. Basing on this analysis, Section VI introduces the probabilistic model of the attack outcome. Section VII studies the parameters of the model and evaluates its soundness. Final conclusions are drawn in Section VIII.

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

A. Compagno is with the Department of Computer Science, Sapienza University of Rome, Italy (email: compagno@di.uniroma1.it). M. Conti is with the Department of Mathematics, University of Padova, Italy (email: conti@math.unipd.it). A.A. D'Amico, G. Dini, P. Perazzo, and L. Taponecco are with the Department of Information Engineering, University of Pisa, Italy (email: {firstname.lastname}@iet.unipi.it).

II. Related Work

Distance bounding protocols were first proposed by Brands and Chaum [8]. Such protocols leverage single-bit challenge-response rounds to establish a secure upper bound on the distance between two devices. Several variants of distance bounding protocols have been proposed in the literature [1], [9], [10], [11], having different properties in terms of resistance to different threats, adversary's success probability, bit error tolerance, memory requirements, and so on. Three classic threats are addressed in distance bounding: an internal adversary (distance fraud), an external one (mafia fraud), or a collusion between the two (terrorist fraud). While the original protocols in [8] resisted only to the first two frauds, more recent proposals (e.g., [10], [11]) offer provable resistance also to the latter one. Distance bounding protocols have been adapted also for packet-based communications [7].

Clulow et al. [12] showed that distance bounding protocols are vulnerable to low-level attacks, in which an adversary attacks directly the physical-layer procedures in order to obtain a distance reduction. These kinds of reduction attacks was analyzed by Poturalski et al. [7] within the IEEE 802.15.4a UWB protocol. The authors of [7] proposed a set of countermeasures as well. In this paper we focus on enlargement attacks, not on reduction ones. Notably, all the countermeasures in [7] involve the format and the decoding method of the payload part of the packet. Our probabilistic model is instead based on orthogonal aspects, namely the format of the preamble part and the TOA estimation algorithm. As a consequence, our model remains valid when the countermeasures in [7] are applied.

Poturalski et al. [13] presented the Cicada attack, a simple physical-layer reduction attack against IEEE 802.15.4a UWB. The authors proposed also a set of countermeasures, some of which introduce new security-driven TOA estimation algorithms. They tested these algorithms only against reduction attacks, not against enlargement attacks. In this paper, we refer only to "classic" TOA estimation algorithms [14]. We leave the study of enlargement attacks against non-classic ones as future work.

On the other hand, the interest in solutions against enlargement attacks is growing. Chiang et al. [2] proposed a way to detect enlargement attempts by measuring the difference of the received power on two antennas. Given the way a signal propagates in a medium, if such difference is low, then the signal will come from far away, and vice versa. Wang et al. [5] proposed a similar countermeasure, measuring the difference on received power of two multipath components on the same antenna. Both these countermeasures ([2] and [5]) defend against a (single) internal adversary, i.e., a participant to the protocol which tries to make its distance from the other participant appear larger. In the case of an external adversary, like the one we consider in this paper, these countermeasures are ineffective. Indeed, an external adversary can deploy a malicious device farther than the honest one, and then replay the legitimate signal from there. In so doing, the replayed signal *actually* comes from a far source, so the countermeasures are deceived.

A solution to enlargement attacks mounted by external adversaries was proposed by Dini et al. [3]. The authors focused on jam-and-replay attacks, which consist in jamming the reception of a legitimate packet, and then replaying it afterwards. The solution is based on the observation that, to avoid packet collisions, the adversary must wait for the legitimate transmission to end. As a consequence, the resulting enlargement is quite big, because proportional to the packet transmission time. Thus, jam-and-replay attacks can be detected by a threshold on the maximum measured distance. In this paper, we focus on overshadowing enlargement attacks, which are not detectable by simple threshold mechanisms.

Taponecco et al. [4] showed that, in the IEEE 802.15.4a UWB ranging standard, an overshadowing-based enlargement attack has a random outcome and is poorly controllable by the adversary. In this paper, we analyze the different sources of randomness, and we provide a probabilistic model of the outcome of an overshadowing attack.

To develop our model, we used the results of Sharp and Yu [15], which introduced a probabilistic model for the error of threshold-based TOA estimation algorithms. We focus on *attacked* TOA estimations, while [15] focused on non-attacked ones. Nevertheless, we will see later that an attacked TOA estimation sometimes behaves like a non-attacked one. To represent these cases, we used the Sharp and Yu's model inside ours as a building block.

III. Motivation

In the original applications of distance bounding protocols, the resistance against enlargement attacks was not necessary, as the only objective was to assure the proximity of two devices. More recently, with the emergence of secure positioning, the enlargement resistance has become an attractive feature. Basically, this is because a distance bounding protocol which resists also enlargement attacks is de facto a *secure distance estimation technique*, which can be fruitfully applied to trilateration in order to estimate a position in a secure manner. Though some secure positioning techniques based on "classic" distance bounding exist (e.g., [16], [17]), they generally offer less coverage compared to ordinary trilateration, as they must tolerate the possibility of distance enlargements [3], [4]. As shown by [16], if the anchors use distance bounding protocols to measure the distances from the device being localized, only the positions inside the polygon formed by the anchors are trusted (and thus covered). With reference to Fig. 1, position A (inside the polygon) cannot be falsified without reducing at least one of the three distances from the anchors.

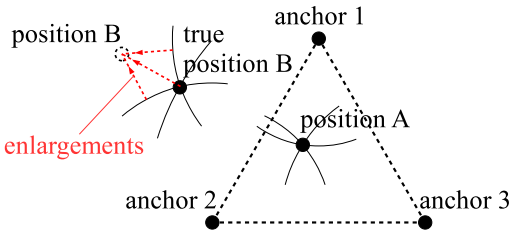

Fig. 1. Trusted and untrusted measured positions.

Since reducing the measured distance is infeasible in a distance

bounding protocol, position A is trusted. On the other hand, position B (outside the polygon) is not trusted, because it could have been falsified by an adversary performing three distance enlargements. If the distance bounding protocols resist enlargement too, these positions are covered as well, and the localization system generally needs less anchors to cover the same area. To sum up, solutions against enlargement attacks are relevant because they open the door to more scalable secure positioning techniques.

For an external adversary, the only way to obtain an enlargement on the distance estimate is to delay the packet TOA estimate. We identified three ways to do that: (a) *jam-and-replay attack*, (b) *ADC resolution attack*, (c) *overshadowing attack*. In the following, we explain and discuss these three ways, and we argue that the most promising one is overshadowing, since it produces realistic enlargements and, at the same time, it does not introduce an unrealistic amount of energy into the channel.

Jam-and-replay attack acts by jamming a legitimate packet, and then replaying it afterwards. Note that, in order to avoid packet collisions, the adversary has to wait for the legitimate communication to end before replaying it. This, as shown in [3], forces her to introduce large delays, greater than or equal to the packet transmission time. In the specific IEEE 802.15.4a UWB protocol, the packet transmission times are in the order of milliseconds [6]. Thus, jam-and-replay attack would produce unrealistic enlargements, in the order of hundreds of kilometers, which are easy to detect by means of simple threshold mechanisms.

On the other hand, ADC resolution 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 forces the AGC to reduce the signal's amplitude so much that the honest signal falls below the minimal resolution of the ADC, and thus gets deleted. ADC resolution attack can produce realistic enlargements, because it has not to wait for the end of the legitimate communication. However, the adversary has to transmit with a very strong power, which is unrealistic for any legitimate communication. ADC resolution attack can be detected by enforcing a limit on the received power.

Finally, overshadowing 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 overshadowing attack can effectively introduce enlargements which are hard to detect. It does not cause unrealistically wide enlargements and does not introduce an unrealistically high energy into the channel. To sum up, though the effect of an overshadowing is not always controllable [4], it however results to be the most convenient strategy by which an external adversary can obtain enlargements.

A. Alternatives to IEEE 802.15.4a UWB

IEEE 802.15.4a UWB is not the only PHY protocol which has been proposed to implement distance bounding. Some researchers proposed that ad-hoc UWB protocols should be designed from scratch with distance bounding in mind [12]. To the best of our knowledge, no such ad-hoc protocol underwent a standardization process, nor found any commercially available implementation. The non-security properties of such protocols (e.g., robustness, ease of use, power efficiency) have not been studied in deep yet. On the other hand, transceivers implementing the IEEE 802.15.4a UWB standard are already present in the market [18]. Poturalski et al. [7] showed how to adapt IEEE 802.15.4a UWB for distance bounding with a limited number of changes, in such a way to preserve its properties. Another alternative is the IEEE 802.15.4a chirp spread spectrum (CSS) physical protocol [6], standardized by the same amendment of UWB. IEEE 802.15.4a CSS allows for TOA-based ranging as well, and transceivers implementing it are already present in the market [19]. However, the UWB standard is more suitable than the CSS one for distance bounding applications, because it can have shorter symbol durations. As shown by [12], shorter symbols make the protocol less vulnerable to low-level reduction attacks. We therefore conclude that IEEE 802.15.4a UWB is definitely a convenient choices for future implementations of wireless distance bounding protocols.

IV. Preliminaries

In the following, we introduce the reference distance bounding protocol, the IEEE 802.15.4a UWB signal format, and the TOA estimation algorithms of the receiver.

A. Threat Model and Reference Distance Bounding Protocol

A distance bounding protocol determines a secure upper bound on the distance between two devices, namely a *verifier* (V) and a *prover* (P), by measuring the round-trip time between unpredictable messages. Originally, distance bounding protocols involved the transmission of single bits (rapid bit exchange) [8]. In recent years, they have been adapted for packet-based communications [7], which are more efficient for wireless protocols. Distance bounding protocols can defend against reduction attacks under three adversary models: a *dishonest prover* wanting to cheat about its distance from the verifier (distance fraud), an *external adversary* wanting to make the prover-verifier distance appear different (mafia fraud), and a collusion of the two (terrorist fraud). In contrast, it is hard to defend against a dishonest prover playing an enlargement attack, since she can introduce an arbitrary delay in the round-trip time. In this paper, we consider external adversaries only, and thus assume the prover to be honest. Under this assumption, to ease the presentation, we consider a simpler class of distance bounding protocols, which resist external adversaries only. An example of such protocols (taken from [7]) is the following:

REQ $\quad V \longrightarrow P : n_V$
ACK $\quad P \longrightarrow V : n_P$
AUTH $\quad P \longrightarrow V : \mathrm{auth}_K(n_V, n_P)$.

The *request packet* (REQ) and the *acknowledgment packet* (ACK) convey respectively n_V and n_P, which are random sequences of bits, unpredictable by the adversary. The *authentication packet* (AUTH) authenticates the whole communication by means of a secret key K shared by prover and verifier. The function $\mathrm{auth}_K(\cdot)$ represents a message authentication code, for example an HMAC. This protocol does not allow an external adversary to impersonate the prover, because she cannot forge the final authentication code. Moreover, the protocol avoids reduction attacks, because the adversary cannot predict n_V or n_P and transmit them in advance. The probabilistic model presented herein is valid as well for more complex distance bounding protocols, as long as IEEE 802.15.4a UWB is used.

The verifier measures the *round-trip time* (T_{RTT}) between the request packet and the acknowledgment packet. To do this, both packets are transmitted by means of the IEEE 802.15.4a UWB protocol, which permits us to measure the time of arrival of a packet with nanosecond precision, corresponding to centimeters in terms of distance. We call *processing time* (T_{proc}) the time interval between the reception of the REQ at the prover and the transmission of the ACK. The processing time is assumed to be known by the verifier. The distance (d) can thus be estimated as:

$$d = \frac{T_{RTT} - T_{proc}}{2} \cdot c, \tag{1}$$

where c is the speed of light.

B. IEEE 802.15.4a UWB Signal Format

IEEE 802.15.4a UWB [6] has been the first standardized Impulse-Radio Ultra-Wideband (IR-UWB) protocol for precision ranging, and it is a convenient choice for future implementations of wireless distance bounding protocols [7].

An IEEE 802.15.4a UWB packet is called *PHY protocol data unit* (PPDU), and consists of three parts: a *synchronization header* (SHR), a *PHY header* (PHR), and a *PHY service data unit* (PSDU). The SHR part is the one allowing for the estimation of the time of arrival of the packet. The PHR contains information about the modulation kind of the successive PSDU part. Finally, the PSDU part contains the payload. In our case, the unpredictable quantities n_V and n_P are conveyed by the PSDU. The SHR is made up of two blocks: a *synchronization preamble* (SYNC) and a *start-of-frame delimiter* (SFD).

As shown in [4], the overshadowing attack influences the outcome of the SYNC processing at the receiver. The mathematical model of the signal transmitted during the SYNC is [6]:

$$s\left(t\right) = \sum_{i=0}^{N_{\mathrm{SYNC}}-1} \psi\left(t - iT_{sym}\right), \tag{2}$$

where $N_{\mathrm{SYNC}} = 1024$ is the number of symbols belonging to the SYNC, and $T_{sym} = 3968\mathrm{ns}$ is the symbol duration. The signal $\psi(t)$ is described as follows:

$$\psi\left(t\right) = \sum_{k=0}^{K_{pbs}-1} d_k p\left(t - kT_{pr}\right), \tag{3}$$

where $\{d_k\}_{k=0}^{K_{pbs}-1}$ is a *perfectly balanced sequence* of $K_{pbs} = 31$ elements with values $\{-1, 0, +1\}$, and $T_{pr} = T_{sym}/K_{pbs} = 128\mathrm{ns}$ is the *pulse-repetition period*. The signal $p(t)$ is an ultra-short causal pulse having a band-pass spectrum, with a bandwidth $B = 500\,\mathrm{MHz}$ centered around the frequency $f_0 = 4.5\,\mathrm{GHz}$. More precisely, $p(t) = c(t)\cos(2\pi f_0 t)$, where $c(t)$ has the shape shown in Fig. 2. The *rise time* (T_r) of $c(t)$ is the time period from the beginning of the pulse to its peak.

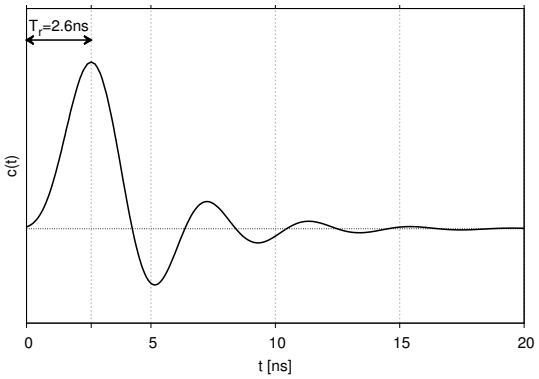

Fig. 2. Monocycle pulse shape.

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

$$r(t) = \sum_{i=0}^{N_{\mathrm{SYNC}}-1} \sum_{k=0}^{K_{pbs}-1} d_k h(t - kT_{pr} - iT_{sym} - t_{\mathrm{TOA}}) + w(t), \tag{4}$$

where $w(t)$ is white thermal noise. In the above equation, t_{TOA} is the *time of arrival* of the signal at the receiver, i.e., the parameter we want to measure. It can represent the time of arrival of the REQ at the prover, as well as that of the ACK at the verifier.

C. Receiver Architecture and TOA Estimation Algorithm

To better understand the effects of an overshadowing attack against an IEEE 802.15.4a UWB receiver, it is necessary to give some details on the physical-layer procedures. We focus on threshold-based UWB ranging schemes, which are the most widely used in UWB localization applications [20], [21], [22]. Moreover, we consider a simple non-coherent energy-based receiver, which guarantees high ranging precision with low cost and low power consumption. Coherent receivers are capable of more precision at the same signal-to-noise ratio, but they are more expensive and power consuming. We leave the study of overshadowing attacks against coherent receivers for future work.

The received signal $r(t)$ 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 Equation 4 do not overlap, it is readily shown that the LPF output, $y(t)$, has the following form:

$$y(t) = \sum_{i=0}^{N_{\mathrm{SYNC}}-1} \sum_{k=0}^{K_{pbs}-1} d_k^2 \times q(t - kT_{pr} - iT_{sym} - t_{\mathrm{TOA}}) + n_y(t). \tag{5}$$

In this equation, $n_y(t)$ is a noise term originating from the signal×noise and the noise×noise interactions in the square-law device, and $q(t) \triangleq h^2(t) \otimes h_{LPF}(t)$, where $h_{LPF}(t)$ is the impulse response of the LPF, and $\otimes$ denotes the convolution operation. In general, $q(t)$ shows a number of peaks, each of which corresponds to the arrival of a signal echo through a different propagation path. The first peak indicates the arrival through the shortest path. The TOA estimation algorithm is concerned with the estimation of t_{TOA}, which is the time of arrival of the first peak of the first $q(t)$-pulse of the preamble.

The signal $y(t)$ is passed through an analog-to-digital converter (ADC) before being processed by the TOA estimation algorithm. The ADC takes samples of the signal with a *sampling period* T_s. We fixed $T_s = 1\mathrm{ns}$, as it is the minimum frequency to correctly sample the 500MHz-bandwidth UWB signal according to the Nyquist-Shannon theorem.

In the present paper, we consider two classic TOA estimation strategies, namely *jump-back search-forward* (JBSF) and *search-back* (SB), which provide significantly different outcomes when attacked by overshadowing. In particular, we refer to the algorithms described and analyzed in [21]. With both JBSF and SB, the TOA estimation is performed in three phases, as shown in Fig. 3:

1) *Frame detection*. This phase decides through power measurements whether a packet is present or not.
2) *Fine time acquisition*. This phase produces an estimate of the arrival time t_{TOA} with an ambiguity of multiples of T_{sym}.
3) *SFD detection*. This phase disambiguates the estimate of t_{TOA} by detecting the position of the SFD through correlation.

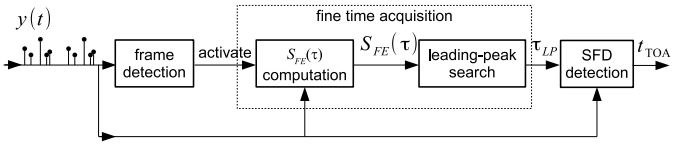

Fig. 3. TOA estimation block diagram.

The fine time acquisition phase provides a measure of a time parameter, say $\tau_{LP} \in [0,\ T_{sym})$, which is related to t_{TOA} by $t_{\mathrm{TOA}} = t_{fd} + \tau_{LP} - N_{fd} T_{sym}$, where t_{fd} is the time at which the frame detection phase declares the presence of the packet, and $N_{fd} \in \mathbb{N}$ is the number of preamble symbols used for frame detection. The successive 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 [4], this is the only phase of the TOA estimation whose result is influenced by the overshadowing attack. The fine time acquisition phase is split in two sub-phases, namely the $S_{FE}(\tau)$ *computation* and the *leading-peak search*. The $S_{FE}(\tau)$ computation consists in combining the signal $y(t)$ at the output of LPF with cyclic-shifted versions of the sequence $\{d_k^2\}_{k=0}^{K_{pbs}-1}$. This produces a T_{sym}-long signal, $S_{FE}(\tau)$, whose support is the interval $\tau \in [0,\ T_{sym})$, which is used for the estimation of τ_{LP}. More precisely, in each interval $\tau \in [mT_{pr},\ (m+1)T_{pr})$, with $m = 0, 1, \ldots, K_{pbs} - 1$, $S_{FE}(\tau)$ is given by:

$$S_{FE}(\tau) = \frac{1}{M} \sum_{i=0}^{M-1} \sum_{k \in \mathcal{I}(m)} d^2_{|k-m|_{K_{pbs}}}\, y(\tau + t_{fd} + (k-m)T_{pr} + iT_{sym}), \qquad (6)$$

where $M < N_{\mathrm{SYNC}}$ is the number of preamble symbols exploited for the fine time acquisition, and $|u|_U$ means "u modulo U." The set $\mathcal{I}(m)$ contains the indices k such that $d^2_{|k-m|_{K_{pbs}}} = 1$ and $d^2_{|k-m-1|_{K_{pbs}}} = 0$. Mathematical details apart, the computation of $S_{FE}(\tau)$ essentially leverages the periodicity of the preamble signal and the autocorrelation properties of the sequence $\{d_k^2\}_{k=0}^{K_{pbs}-1}$ to improve the signal-to-noise ratio.

Fig. 4 shows an example of $S_{FE}(\tau)$ function.

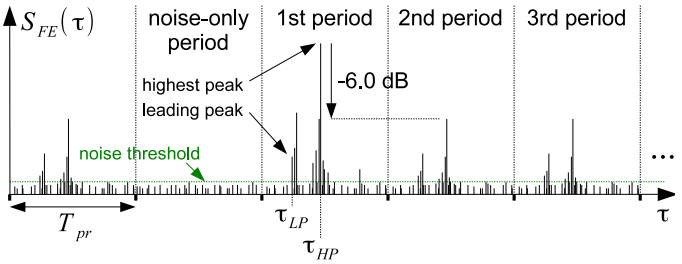

Fig. 4. Example of $S_{FE}(\tau)$.

From now on, the pulses will be represented in the figures as needle-shaped. Actually, the real pulses have a non-negligible time duration, so they sometimes overlap to each other in the $S_{FE}(\tau)$ function. We neglect this when possible, in order to ease the explanation. As seen in Fig. 4, $S_{FE}(\tau)$ is T_{pr}-periodic, except that the periods repeat with different amplitudes, and some of them contain only thermal noise (*noise-only periods*). Conventionally, we refer to the period with the greatest amplitude as the *first period*. The successive one is the second period, and so on. If no overshadowing takes place, the first period contains the maximum of $S_{FE}(\tau)$ (*highest peak*), and, shortly before, a pulse corresponding to the signal echo through the shortest path (*leading peak*). The highest peak is in position τ_{HP}, while the leading peak is in position τ_{LP}. The second and the third periods are attenuated by 6 decibels with respect to the first one. The period preceding the first one is a noise-only period.

The leading-peak search is the sub-phase by which JBSF and SB differ. In particular, the JBSF algorithm (Fig. 5) starts from the highest peak position τ_{HP}, jumps back by T_{JB} seconds, and then proceeds forward looking for the first time $S_{FE}(\tau)$ goes beyond a given *noise threshold*. This gives the estimate of τ_{LP}. On the other hand, the SB algorithm (Fig. 6) starts from τ_{HP}, and searches backward until $S_{FE}(\tau)$ goes below the noise threshold and continues to be under for a T_{SB}-wide window (*noise-only region*). This gives the estimate of τ_{LP}. The JBSF and the SB algorithms treat $S_{FE}(\tau)$ as a "circular" function. That is, if they jump or search below the $\tau = 0$ limit, they continue from $\tau = T_{sym}$, and vice versa.

Note that the SB algorithm could search backward indefinitely if it does not find a T_{SB}-long noise-only region. However, in those environments where the IEEE 802.15.4a

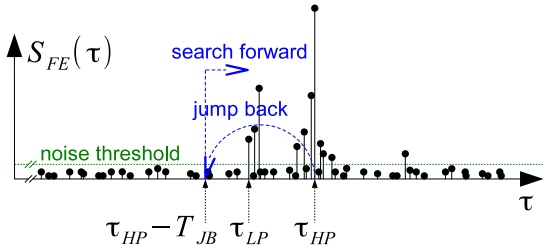

Fig. 5. Jump-back search-forward algorithm.

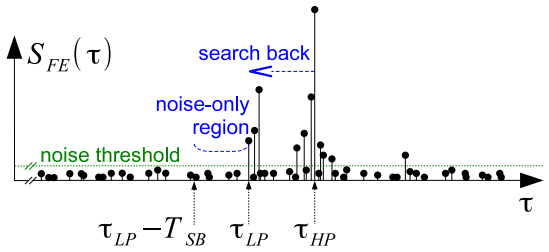

Fig. 6. Search-back algorithm.

UWB has been designed to work, it is statistically impossible to find the leading peak more than T_{pr} seconds earlier than the highest peak. If a T_{SB}-long noise-only region is not found after a backward search of T_{pr}, it is highly probable that an interfering signal (not necessary malicious) is present. We assume that, after a backward search of T_{pr}, the TOA estimation fails and produces an error (*long-search error*).

The noise threshold is tailored on the basis of the thermal noise statistics. We fixed it in such a way that a noise-only sample of $S_{FE}(\tau)$ has a probability of 10^{-5} of being above it, i.e., of being wrongly interpreted as a signal (false alarm). In addition, we set $T_{JB} = 60\text{ns}$, as recommended by [21], 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. Finally, we denote by SNR_h the signal-to-noise ratio of the (legitimate) signal at the receiver.

V. Overshadowing Attack

Since the distance measurement comes from the round-trip time, the adversary's aim is to enlarge it. The only way for an external adversary to do that is to delay the packet TOA estimate at the verifier and/or at the prover. The *overshadowing attack* is one of the most promising ways in which an external adversary can cause an enlargement against a distance bounding protocol without being detected. This is because it produces realistic enlargements and at the same time it does not introduce an unrealistic amount of energy into the channel.

In the overshadowing attack, the adversary repeats a legitimate packet with a certain delay and a higher power. In this way, she tries to "overshadow" the legitimate communication with a delayed copy of it. At the signal level, the adversary tries to deceive the receiver into thinking that the malicious signal is the legitimate one that took the shortest path, while the (true) legitimate signal is an echo of it. The attack starts in the presence of a legitimate transmission of a request or an acknowledgment packet. The adversary has first to synchronize with the ongoing communication. It takes some of the SYNC initial symbols to do that. Then, she starts transmitting the replayed copy (skipping those initial symbols). The replayed signal is timed in such a way to arrive at the receiver shifted of a certain delay (*overshadow delay*, Δ_{ovrs}) with respect to the legitimate one. During the transmission of the successive PSDU part, a copy of it is replayed in the same way. We denote by SNR_m the *adversarial signal-to-noise ratio*, i.e. the signal-to-noise ratio of the malicious signal at the receiver. The adversarial signal-to-noise ratio must be greater than the legitimate one ($\text{SNR}_m > \text{SNR}_h$), but not too much greater, otherwise the attack will be easily detectable.

We assume that the adversary knows exactly the distance between transmitter and receiver, and her distance from the receiver. This is a necessary condition for the malicious signal to arrive at the receiver with the desired delay. Note that if the adversary is too much far from both transmitter and receiver, she could not be able to replay the unpredictable payload bits timely. With reference to Fig. 7, the replayed payload bits propagate through the prover-adversary-verifier (PMV) path, while the legitimate ones through the direct prover-verifier (PV) path.

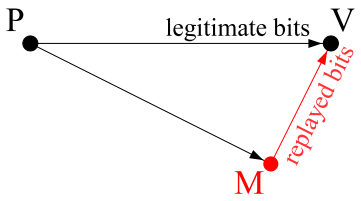

Fig. 7. Possible positions of prover (P), verifier (V), and adversary (M) in an overshadowing attack (against ACK).

The difference between the two propagation times cannot be greater than the overshadow delay that the adversary wants to introduce. More precisely:

$$\left(\overline{\text{PM}} + \overline{\text{MV}} - \overline{\text{PV}}\right)/c + T_{\text{M},proc} \leq \Delta_{ovrs}, \qquad (7)$$

where $\overline{\text{PM}}$, $\overline{\text{MV}}$, $\overline{\text{PV}}$ are respectively the prover-adversary, adversary-verifier, and prover-verifier distances, and $T_{\text{M},proc}$ is the processing time of the adversary. Such a processing time could be zero or even negative, if the adversary employs time-gaining replay techniques like early bit detection or deferred bit signalling [12].

In addition, we assume that the adversary enjoys a *Gaussian channel* towards the victim receiver. A Gaussian channel is characterized by an impulse response showing a single pulse, without secondary echoes. In practice, a Gaussian channel can be obtained with a transmitter very close to the victim receiver, or by employing a highly directional antenna towards it. This gives more power to the adversary, because she can control precisely how the malicious signal is received by the victim. As a consequence, the attack is more controllable [4].

Fig. 8 shows the effect of an overshadowing attack against the JBSF algorithm. The waveform $S_{FE}(\tau)$ has a component due to the legitimate signal, and a component (the strongest

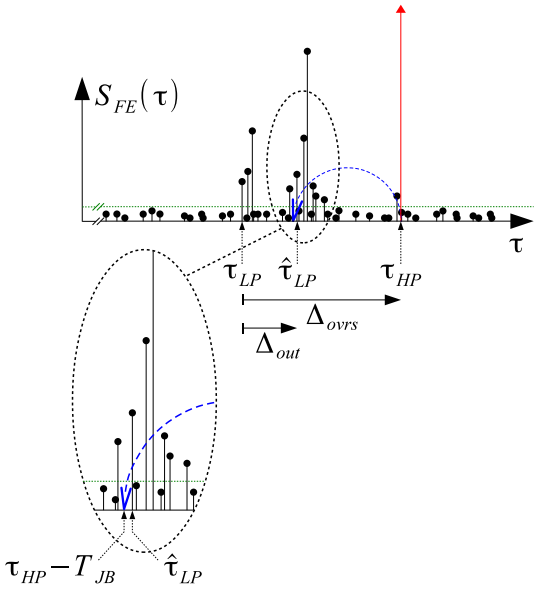

Fig. 8. Overshadowing attack against JBSF. The round-headed pulses correspond to the legitimate signal, the triangle-headed pulse to the malicious one.

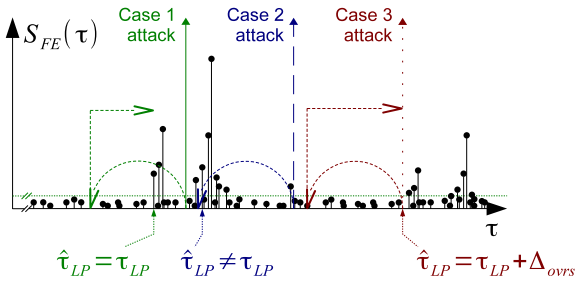

Fig. 9. Three example attacks against JBSF. The solid-line attack falls into Case 1, the dashed-line attack falls into Case 2, and the dotted-line attack falls into Case 3.

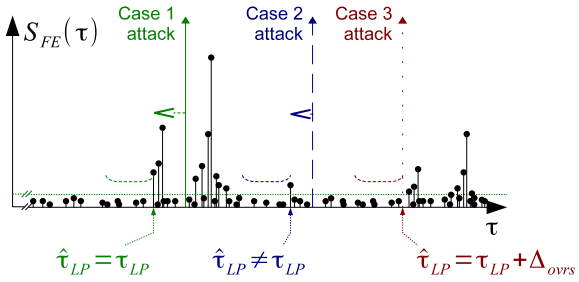

Fig. 10. Three example attacks against SB. The solid-line attack falls into Case 1, the dashed-line attack falls into Case 2, and the dotted-line attack falls into Case 3.

one) due to the malicious signal, which arrives at time $\tau_{LP} + \Delta_{ovrs}$. From Fig. 8 it is clear that, if the overshadow delay is such that the leftward jump falls after the true leading peak, the leading-peak search will provide for a wrong estimate of it ($\hat{\tau}_{LP} \neq \tau_{LP}$). A similar effect can happen on the SB algorithm as well. We define the *outcome delay* ($\Delta_{out} = \hat{\tau}_{LP} - \tau_{LP}$) as the timing delay actually obtained by the adversary in the estimate of the leading peak. The outcome delay directly translates into an equal delay in the packet TOA estimate, which in turn translates into an enlargement in the estimated distance ($\hat{d}$):

$$\hat{d} = d + \frac{\Delta_{out}}{2} \cdot c. \tag{8}$$

Depending on the legitimate propagation channel and the overshadow delay, the overshadowing attack can fall into three cases:

- *Case 1*. The first pulse of the legitimate signal is identified as the leading peak. Case 1 captures the case in which the attack has no effect. However, due to the (quasi-)periodicity of $S_{FE}(\tau)$, replicas of the leading peak repeat with a period of T_{pr}. Accordingly, Case 1 captures also the cases in which a *replica* of the first pulse is identified as the leading peak, causing an enlargement of multiples of the pulse-repetition period.
- *Case 2*. A non-first pulse of the legitimate signal, or a replica of it, is identified as the leading peak. In this case, the attack produces an enlargement which may be not controllable by the adversary, since it depends on the propagation channel between the verifier and prover.
- *Case 3*. The malicious pulse is identified as the leading peak. In this case, the adversary is always able to control the outcome delay.

Figs. 9 and 10 show examples of overshadowing attacks falling into the three cases, against JBSF and SB, respectively. This categorization is useful to understand the behavior of an attacked TOA estimation algorithm, and to develop the probabilistic model of the attack. Which case the attack falls into strictly depends on the propagation channel between prover and verifier, which is unknown by the adversary and hard to estimate in practice. In addition, in Case 2 the adversary cannot predict which secondary pulse will be identified as the leading peak. As a consequence, the outcome of an overshadowing attack is essentially random.

VI. Our Proposed Probabilistic Model

To model the outcome of the attack, we find an approximation of the *probability distribution function* (*pdf*) of the outcome delay: $f(\Delta_{out})$. Such a *pdf* is a composite of three *basic pdf's*, one for each case (Case 1, Case 2, Case 3). The basic *pdf* for Case 1 represents the distribution of the outcome delay conditioned to the event that the attack falls into Case 1, and so on. The resulting outcome delay will be described by a composite *pdf*, obtained by summing scaled (and shifted, see after) versions of the basic *pdf*'s.

We designed heuristically the three basic *pdf*'s. In particular, we modeled them on the basis of histograms of the outcome delay, obtained by signal-level simulations of the TOA estimation algorithms. We justify the shape of such *pdf*'s with statistical considerations on the propagation channel and the TOA estimation algorithm. The probabilistic model we present here has been tested to be valid for $\mathrm{SNR}_h \in [20\mathrm{dB}, 50\mathrm{dB}]$, $\mathrm{SNR}_m \in (\mathrm{SNR}_h, 51\mathrm{dB}]$, and $\Delta_{ovrs} \in (0\mathrm{ns}, 384\mathrm{ns} = 3T_{pr}]$. Such ranges permit us to capture the majority of practical attack scenarios.

A. *Basic* pdf *for Case 1*

In Case 1, the outcome delay is equivalent to an ordinary TOA estimation error of a non-attacked receiver. Sharp and Yu [15] introduced a *pdf* which models the error of a generic threshold-based TOA estimation algorithm:

$$\mathrm{SharpYu}(t) \triangleq \begin{cases} (1/T_r) \cdot \frac{3\alpha^2 (t/T_r)^2}{(\alpha^2 + (t/T_r)^3)^2} & \text{if } t \geq 0 \\ 0 & \text{else.} \end{cases} \quad (9)$$

where T_r is the rise time of the the signal pulse (the monocycle in our case), and α is a shape factor which depends on the signal-to-noise ratio (the legitimate one in our case). The SharpYu distribution describes the random instant in which the leading edge of the observed signal, which has a random amplitude, passes the noise threshold and is thus detected by the algorithm. It is a very general model which can be applied to a large class of ranging systems, from relatively narrow bandwidth Wi-Fi-based systems to UWB systems. We found that the SharpYu *pdf* fits our simulated outcome delays better[1] than other common error models (e.g., the Gaussian model). It has such a simple analytical form because it simplifies out a number of aspects, among which the error component due to the time sampling of the signal. The time-sampling error is distributed uniformly in $[0, T_s)$, and it is added to the other error components captured by the SharpYu distribution. We therefore use a "sampled" version of the SharpYu *pdf*, given by:

$$\mathrm{SampledSharpYu}(t) \triangleq \mathrm{SharpYu}(t) \otimes \mathrm{rect}(t), \quad (10)$$

where $\otimes$ is the convolution operation, and $\mathrm{rect}(t)$ is a T_s-wide rectangular function:

$$\mathrm{rect}(t) \triangleq \begin{cases} 1/T_s & \text{if } t \in [0,\ T_s) \\ 0 & \text{else.} \end{cases} \quad (11)$$

We found that the SampledSharpYu distribution fits our simulated outcome delays better than the SharpYu distribution.

The TOA estimation error of a non-attacked receiver follows a distribution exemplified by the histogram in Fig. 11. The figure shows also the SampledSharpYu *pdf*, with three different values of the parameter α. Lower signal-to-noise ratios correspond to higher α's. Note that, as α increases (and thus SNR_h decreases), the *pdf* spreads towards the right. This catches the fact that, with a lower signal-to-noise ratio, the first pulse passes the noise threshold later in time. The SampledSharpYu distribution fits for Case 1 with both JBSF and SB algorithms. Indeed, the errors of these algorithms (when not attacked) differ to a negligible extent. The trend of the α parameter with respect to the legitimate signal-to-noise ratio will be studied in the parametrization section (Section VII).

Note that, for the purpose of evaluating the security of ranging/positioning solutions, it is often important to simulate also a *honest scenario*, in which the adversary is absent. The SampledSharpYu *pdf* can be used for this aim, as it models the TOA estimation error of a non-attacked receiver.

[1]In terms of mean log-likelihood (see Section VII).

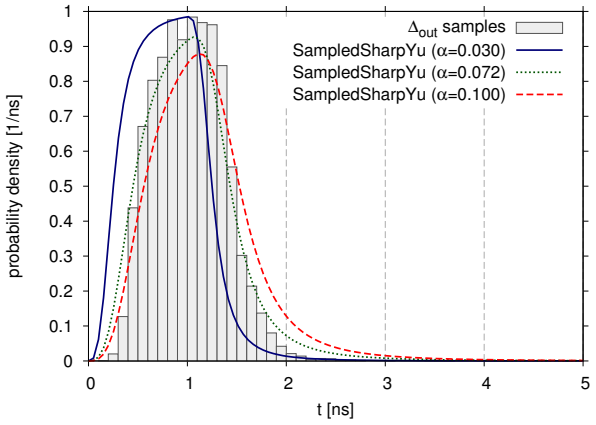

Fig. 11. Histogram of the TOA estimation error of a non-attacked receiver, compared to SampledSharpYu *pdf*'s with different values of α. The histogram comprises 10,000 outcome delays, with $\mathrm{SNR}_h = 30\mathrm{dB}$. The SampledSharpYu with $\alpha = 0.072$ is the one that best fits the histogram, according to the maximum log-likelihood criterium (see Section VII).

B. *Basic* pdf *for Case 2*

In Case 2, the outcome delay with JBSF follows a distribution exemplified by the histogram in Fig. 12.

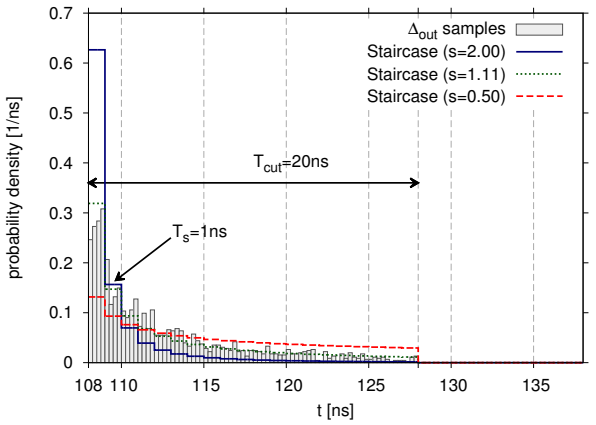

Fig. 12. Histogram of the outcome delay falling into Case 2 (JBSF), compared to Staircase *pdf*'s, with $T_{cut} = 20\mathrm{ns}$ and different values of s. The histogram comprises 1,800 outcome delays, with $\mathrm{SNR}_h = 30\mathrm{dB}$ and $\Delta_{ovrs} = 166\mathrm{ns}$. The Staircase with $s = 1.11$ is the one that best fits the histogram.

Such a distribution is fitted by a (properly shifted) "staircase"-shaped function, i.e., a function with horizontal steps of decreasing height (see Fig. 12). The width of the steps is equal to the sampling period, while their height follows a discrete power law, with exponent $-s$. The analytical form of the Staircase is:

$$\mathrm{Staircase}(t) \triangleq \frac{1}{H} \cdot \sum_{i=1}^{N_{cut}} (i^{-s}) \cdot \mathrm{rect}(t - (i-1)T_s), \quad (12)$$

where H is a normalization factor, and N_{cut} is the number of steps:

$$H = \sum_{i=1}^{N_{cut}} i^{-s} \quad (13)$$

$$N_{cut} = \left\lfloor \frac{T_{cut}}{T_s} \right\rfloor. \quad (14)$$

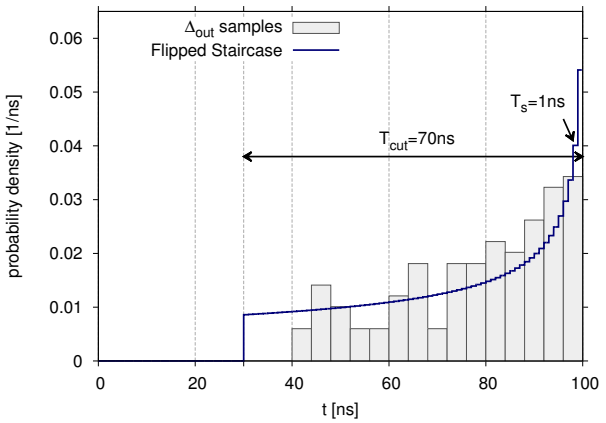

Fig. 13. Histogram of the outcome delay falling into Case 2 (SB), compared to a flipped Staircase *pdf*, with $T_{cut} = 70ns$ and $s = 0.43$. The histogram comprises 125 outcome delays, with $\mathrm{SNR}_h = 30\mathrm{dB}$ and $\Delta_{ovrs} = 100\mathrm{ns}$. The flipped Staircase with $s = 0.43$ is the one that best fits the histogram.

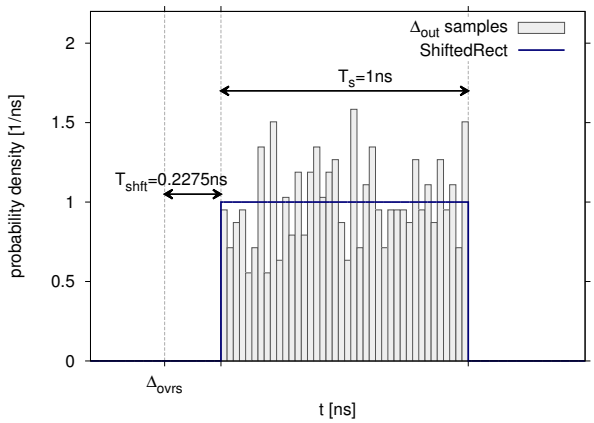

Fig. 14. Histogram of the TOA estimation error of a receiver that receives *only* the malicious signal, compared to a ShiftedRect *pdf* with $T_{shft} = 0.2275\mathrm{ns}$. The histogram comprises 10,000 outcome delays, with $\mathrm{SNR}_m = 30\mathrm{dB}$.

The *cutting time* T_{cut} represents the length of the Staircase. The decreasing trend is due to the rightward search of JBSF, because long searches are less probable than short ones. We heuristically found that a power law fits the histogram better[2] than other laws (e.g., exponential). The horizontal steps are due to the time sampling, which introduces a uniform error component.

On the other hand, the outcome delay of SB has a distribution exemplified by the histogram in Fig. 13. Such a distribution is fitted by a (properly shifted) "flipped" Staircase function (see Fig. 13). This shape is due to the fact that SB searches a sample above the noise threshold like JBSF, but with a leftward direction. Also here, long leftward searches are less probable than short ones. This is modeled by the increasing shape of the flipped Staircase.

C. *Basic* pdf *for Case 3*

In Case 3 (for both JBSF and SB), the outcome delay is equivalent to an ordinary TOA estimation error of a receiver that receives *only* the malicious signal. It does not follow a SampledSharpYu *pdf* as in Case 1, because the adversarial channel is Gaussian. In a Gaussian channel, the energy is concentrated in one pulse only, whose amplitude is not random. Thus, the randomicity of the TOA estimation error stems from the time sampling only.

The TOA estimation error of a receiver that receives only the malicious signal follows a distribution exemplified by the histogram in Fig. 14. Such a distribution is fitted by a T_s-wide shifted rectangular function, whose analytical form is:

$$\mathrm{ShiftedRect}(t) \triangleq \mathrm{rect}(t - T_{shft}), \tag{15}$$

where T_{shft} is a shift factor which depends on the adversarial signal-to-noise ratio. The shift is due to the non-zero rise time of the malicious pulse, which passes the noise threshold after T_{shft} seconds. The lower the adversarial signal-to-noise ratio is, the longer T_{shft}, because the malicious pulse passes the noise threshold later in time. The trend of the T_{shft} parameter with respect to the adversarial signal-to-noise ratio will be studied in the parametrization section (Section VII).

[2]Again, in terms of mean log-likelihood (see Section VII).

D. *Probabilistic Model of the Attacked JBSF Algorithm*

We model the outcome delay of an attacked JBSF algorithm with the following composite *pdf*:

$$\begin{aligned} f(\Delta_{out}) \;&=\; a_1 \cdot \mathrm{SampledSharpYu}_{\alpha}(\Delta_{out} - \left\lceil \tfrac{\Delta_{land}}{T_{pr}} \right\rceil \cdot T_{pr}) \\ &+\; a_2 \cdot \mathrm{Staircase}_{s,T_{cut}}(\Delta_{out} - \Delta_{land}) \\ &+\; a_3 \cdot \mathrm{ShiftedRect}_{T_{shft}}(\Delta_{out} - \Delta_{ovrs}). \end{aligned} \tag{16}$$

The multiplicative factors a_1, a_2, a_3 are non-negative, and their sum equals 1. They represent the probabilities that the attack falls into, respectively, Case 1, Case 2, Case 3. The Staircase of Case 2 starts from Δ_{land}, which corresponds to the instant that the jump-back lands on:

$$\Delta_{land} \triangleq \Delta_{ovrs} + T_r - T_{JB}. \tag{17}$$

Equation 17 takes into account that the jump does not take off from the beginning of the malicious pulse (i.e., from Δ_{ovrs}), but from its peak (i.e., from $\Delta_{ovrs} + T_r$). The cutting time of the Staircase is fixed to:

$$T_{cut} = T_{pr} - |\Delta_{ovrs} + T_r - T_{JB}|_{T_{pr}}. \tag{18}$$

This value is such that the Staircase always ends at the position of the malicious pulse, starting from which the attack falls into Case 3.

E. *Probabilistic Model of the Attacked SB Algorithm*

We model the outcome delay of an attacked SB algorithm with a *long-search error probability* (a_{err}) and a composite *pdf* of the outcome delay ($f(\Delta_{out})$). An overshadowing attack will incur in a long-search error with a_{err} probability. Otherwise, it will produce an outcome delay following the $f(\Delta_{out})$ *pdf*, which is given by:

$$\begin{aligned} (1-a_{err}) \cdot f(\Delta_{out}) \;&=\; a_2' \cdot \mathrm{Staircase}_{s',T_{cut}'}(-\Delta_{out} + \left\lfloor \tfrac{\Delta_{ovrs}}{T_{pr}} \right\rfloor \cdot T_{pr}) \\ &+\; a_1 \cdot \mathrm{SampledSharpYu}_{\alpha}(\Delta_{out} - \left\lfloor \tfrac{\Delta_{ovrs}}{T_{pr}} \right\rfloor \cdot T_{pr}) \\ &+\; a_2 \cdot \mathrm{Staircase}_{s,T_{cut}}(-\Delta_{out} + \Delta_{ovrs}) \\ &+\; a_3 \cdot \mathrm{ShiftedRect}_{T_{shft}}(\Delta_{out} - \Delta_{ovrs}). \end{aligned} \tag{19}$$

The multiplicative factors a'_2, a_1, a_2, a_3 are non-negative, and their sum equals $1-a_{err}$. They represent the probabilities that the attack falls into Case 1 (a_1), Case 2 (a'_2 and a_2), and Case 3 (a_3). Case 2 is split into components a'_2 and a_2, corresponding to separate flipped Staircases. This is because the SB algorithm can stop somewhere in the same period of the malicious pulse, or it can continue searching, pass a replica of the leading peak, and then stop somewhere in the precedent period. The former case is captured by the a_2 Staircase, the latter one by the a'_2 Staircase. The cutting times of the Staircases are fixed as follows:

$$T'_{cut} = \left\lfloor \frac{\Delta_{ovrs}}{T_{pr}} \right\rfloor \cdot T_{pr} - \Delta_{ovrs} + T_{pr}, \tag{20}$$

$$T_{cut} = \Delta_{ovrs} - \left\lfloor \frac{\Delta_{ovrs}}{T_{pr}} \right\rfloor \cdot T_{pr} - T_{SB}. \tag{21}$$

These values are such that the a_2 Staircase always ends at the position of (a replica of) the leading peak, while the a'_2 Staircase always ends where the long-search error occurs.

VII. Model's Parametrization and Evaluation

We parametrized our probabilistic models for a typical residential environment (standard channel model CM1 [14]) and for a typical office environment (standard channel model CM3 [14]). In order to determine the parameters, we proceeded as follows. We generated a number of outcome delay samples by signal-level simulations of the TOA estimation algorithm. Then, we tailored the parameters on these samples, following a maximum goodness-of-fit criterium. In particular, we determined those parameters that maximizes the *mean log-likelihood* ($\mathcal{L}$), defined as follows:

$$\mathcal{L} \triangleq \frac{1}{N} \sum_{i=1}^{N} \log f(\Delta_{out,i}), \tag{22}$$

where N represents the number of samples, and $\Delta_{out,i}$ represents the i-th sample. In the following, we show the complete parametrization for the channel model CM1. The parametrization for both channel models CM1 and CM3 is included in the Matlab tool we make publicly available (see below).

The parameter α can be studied separately, since the outcome delay in Case 1 is equivalent to an ordinary TOA estimation error of a non-attacked receiver. As a consequence, it depends only on the legitimate signal-to-noise ratio. We tailored the α parameter on 10,000 TOA estimation errors of a non-attacked JBSF algorithm, with randomly generated UWB channels following the standard statistical model of a typical residential environment (CM1) [14]. Fig. 15 shows the best-fitting α with respect to the legitimate signal-to-noise ratio. As anticipated in Section VI, lower signal-to-noise ratios correspond to higher α's, which in turn make the SampledSharpYu *pdf* spread towards the right. This catches the fact that, with a lower signal-to-noise ratio, the first pulse passes the noise threshold later in time.

Note that, as we said in Section IV, Case 1 captures also the case in which a *replica* of the first pulse is identified as the leading peak. This case happens only if $\Delta_{ovrs} \geq T_{pr}$ (both

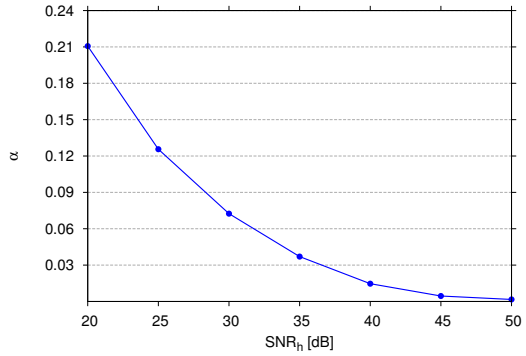

Fig. 15. Trend of α wrt the legitimate signal-to-noise ratio (CM1).

with JBSF and SB). However, pulse replicas after the first period are attenuated (see Fig. 4). Accordingly, Case 1 with $\Delta_{ovrs} \in [T_{pr}, 3T_{pr}]$ must be parametrized with a greater α, corresponding to an equivalent signal-to-noise ratio of $\mathrm{SNR}_h - 6.0\mathrm{dB}$.

Also the parameter T_{shft} can be studied separately, since the outcome delay in Case 3 is equivalent to a TOA estimation error of a receiver that receives *only* the malicious signal. As a consequence, it depends only on the adversarial signal-to-noise ratio. We tailored the T_{shft} parameter on 10,000 TOA estimation errors of a JBSF algorithm receiving only the malicious signal. Fig. 16 shows the best-fitting T_{shft} with respect to the adversarial signal-to-noise ratio.

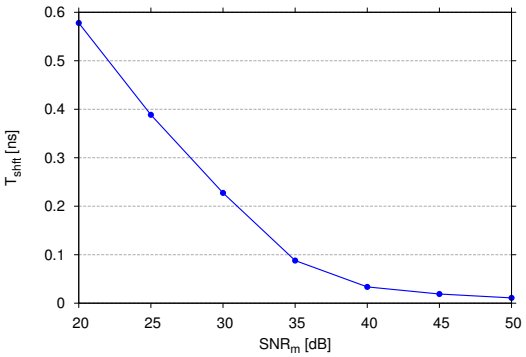

Fig. 16. Trend of T_{shft} wrt the adversarial signal-to-noise ratio (CM1).

As anticipated in Section VI, lower adversarial signal-to-noise ratios correspond to higher T_{shft}'s. This catches the fact that the malicious pulse passes the noise threshold later in time.

The other parameters (a_1, a_2, a_3, s for the JBSF model, and a_{err}, a'_2, a_1, a_2, a_3, s', s for the SB model) depend on the legitimate signal-to-noise ratio and the overshadow delay. Notably, they do not depend on the adversarial signal-to-noise ratio, as long as it is greater than the honest one. We tailored these parameters on 10,000 outcome delays of attacked TOA estimation algorithms, with randomly generated UWB channels following the CM1 statistical model. We tested different overshadow delays, namely $\Delta_{ovrs} = 1\mathrm{ns}, 2\mathrm{ns}, \ldots, 384\mathrm{ns}(= 3T_{pr})$. We fixed the adversarial signal-to-noise ratio to $\mathrm{SNR}_m = \mathrm{SNR}_h + 1\mathrm{dB}$. Fig. 17 shows the best-fitting parameters of the JBSF and SB models. They all have been estimated with the maximum goodness-of-fit

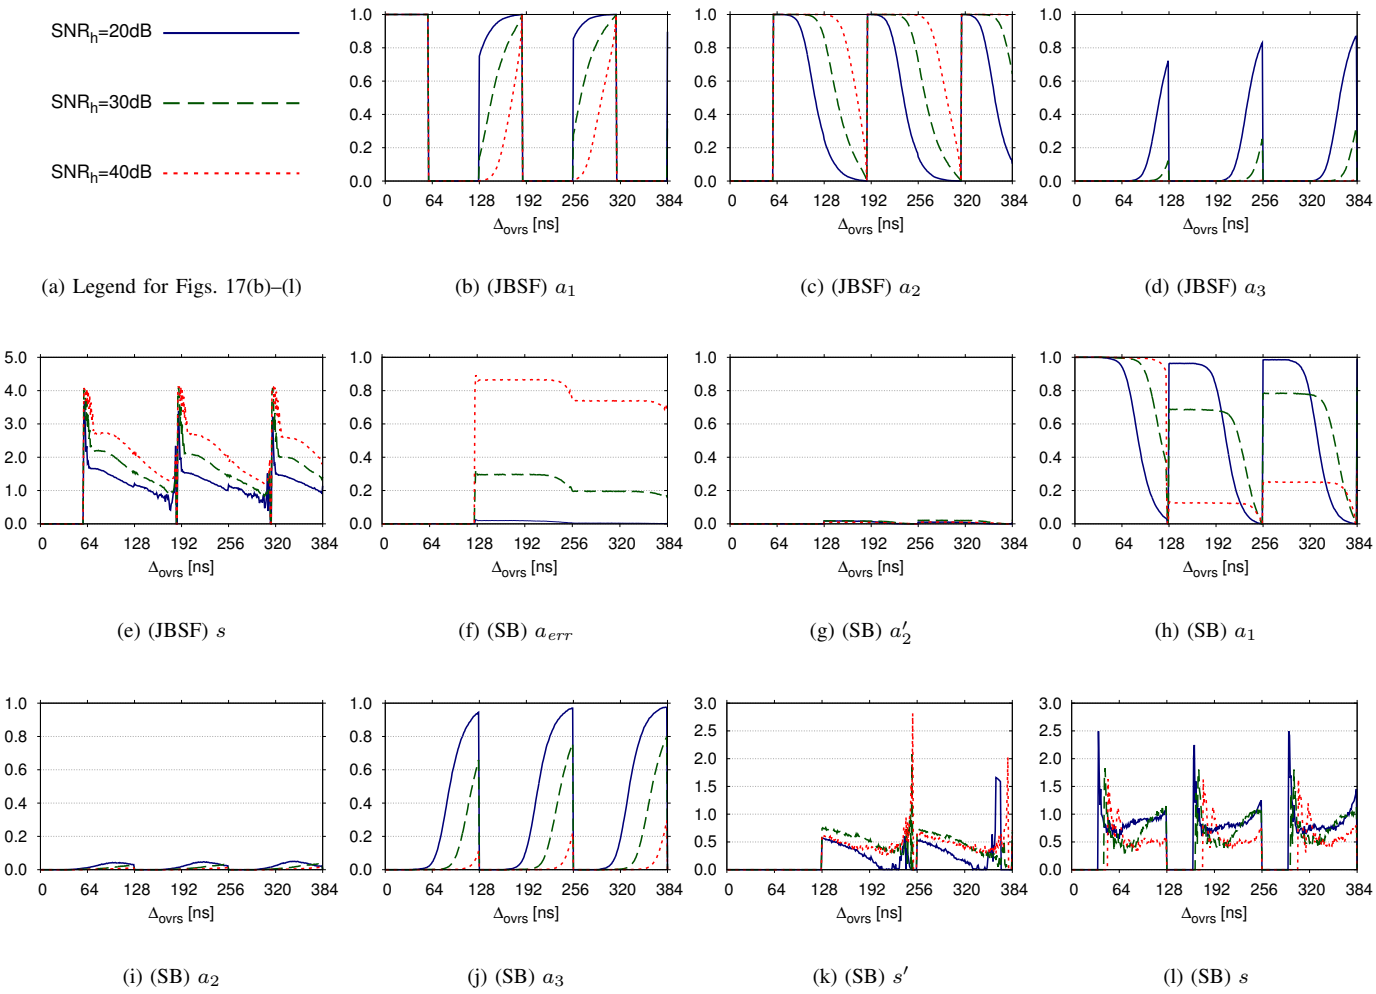

(a) Legend for Figs. 17(b)–(l) (b) (JBSF) a_1 (c) (JBSF) a_2 (d) (JBSF) a_3

(e) (JBSF) s (f) (SB) a_{err} (g) (SB) a'_2 (h) (SB) a_1

(i) (SB) a_2 (j) (SB) a_3 (k) (SB) s' (l) (SB) s

Fig. 17. Parameters for the JBSF and SB models (CM1).

criterium, except of a_{err}, which has been estimated as the percentage of long-search errors on the total TOA estimations.

A. Goodness-of-fit Tests

In order to check the soundness of our probabilistic models, we estimated their goodness of fit (in terms of mean log-likelihood) on 10,000 outcome delays of attacked TOA estimation algorithms. We compared it with the goodness of fit of some "naive" models, namely a Gaussian *pdf* (with the same μ and σ of the samples) and a uniform *pdf* (uniform between the min and the max samples). Fig. 18 shows such a comparison for the JBSF and the SB model and for CM1 and CM3. It can be seen that our probabilistic models better describes the outcome of an overshadowing attack than reference models.

The likelihood comparison is a general purpose goodness-of-fit test. In order to further evaluate the validity of our probabilistic model we study its equivalence to the samples in terms of security. More precisely, we define a metric, namely the *enlargement control probability* (P_{ctrl}), which expresses the probability that an adversary is able to introduce a controlled enlargement. Given the *objective delay* (Δ_{obj}) and the *objective precision* (δ_a) of the adversary, the enlargement control probability is defined as follows:

$$P_{ctrl} \triangleq \max_{\Delta_{ovrs}} \Pr\left[|\Delta_{out} - \Delta_{obj}| \leq \delta_a\right]. \tag{23}$$

The control probability is referred to the overshadowing attack with the *most convenient* overshadow delay (i.e., the one that maximizes the control probability). It can be computed from the samples:

$$P_{ctrl} = \max_{\Delta_{ovrs}} \frac{\text{samples in } [\Delta_{obj} - \delta_a,\ \Delta_{obj} + \delta_a)}{\text{total samples}}, \tag{24}$$

as well as from the probabilistic models:

$$P_{ctrl} = \max_{\Delta_{ovrs}} \int_{\Delta_{obj}-\delta_a}^{\Delta_{obj}+\delta_a} f(\Delta_{out}) d\Delta_{out} \quad \text{for JBSF,} \tag{25}$$

$$P_{ctrl} = \max_{\Delta_{ovrs}} (1-a_{err}) \int_{\Delta_{obj}-\delta_a}^{\Delta_{obj}+\delta_a} f(\Delta_{out}) d\Delta_{out} \quad \text{for SB.} \tag{26}$$

We compared the control probability computed from the samples, to the one computed from the probabilistic models. Fig. 19 shows such a comparison for CM1 and CM3. We see that the control probabilities computed by our models closely follow those estimated by the samples. This corroborates

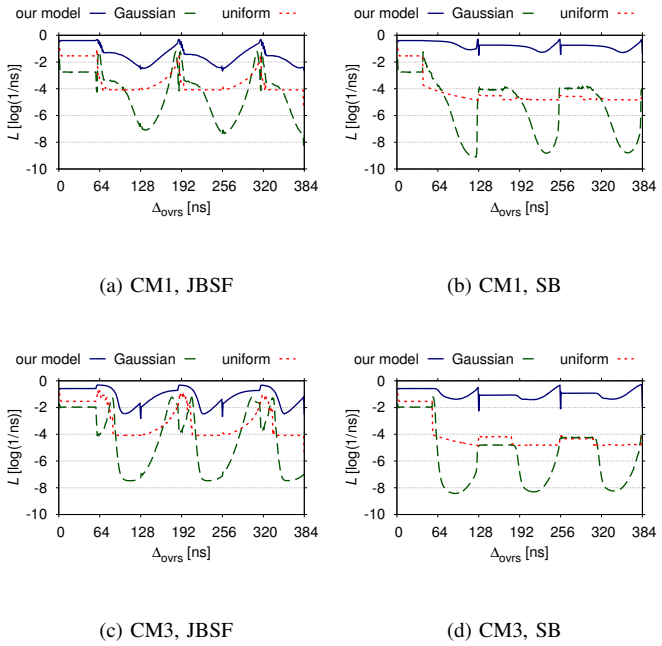

(a) CM1, JBSF (b) CM1, SB

(c) CM3, JBSF (d) CM3, SB

Fig. 18. Mean log-likelihood of our probabilistic models compared to reference models, with $\mathrm{SNR}_h = 30\mathrm{dB}$.

the validity of our probabilistic models for security-focused simulations.

Finally, we developed a Matlab tool[3] based on our probabilistic models. The tool is capable of generating efficiently the random outcome of an overshadowing attack, without performing burdensome signal-level simulations. It takes as input the TOA estimation algorithm, the legitimate and adversarial signal-to-noise ratios, the overshadow delay, and the channel model. The tool can simulate non-attacked TOA estimations as well, which is useful to evaluate the non-malicious scenario.

VIII. Conclusions

In this paper, we provided a probabilistic model of the outcome of an overshadowing attack against a distance bounding protocol realized with IEEE 802.15.4a UWB. Our model takes into consideration several variables, like the propagation environment, the signal-to-noise ratio, and the TOA estimation algorithm. We evaluated the soundness of our model by comparing it to attack outcomes generated by physical-layer simulations, and by performing goodness-of-fit tests. The results showed that our model is sufficiently realistic to replace physical-layer simulations. We finally developed a Matlab tool based on our model, capable of simulating attacked and non-attacked TOA estimations. The tool allows researchers to evaluate the security of the ranging/positioning solutions that can be subject to enlargement attacks. We made such a tool available to the research community.

Acknowledgment

We thank the anonymous reviewers for their insightful comments and suggestions, which significantly contributed to improving the quality of the paper. Mauro Conti is supported by a Marie Curie Fellowship funded by the European Commission (agreement PCIG11-GA-2012-321980). This work is partially supported by the EU TagItSmart! project (agreement H2020-ICT30-2015-688061), the EU-India REACH project (agreement ICI+/2014/342-896), the Italian MIUR-PRIN TENACE project (agreement 20103P34XC), the project "Analisi di dati sensoriali: dai sensori tradizionali ai sensori sociali" funded by "Progetti di Ricerca di Ateneo - PRA 2016" of the University of Pisa, and the project "Tackling Mobile Malware with Innovative Machine Learning Techniques" funded by the University of Padua.

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

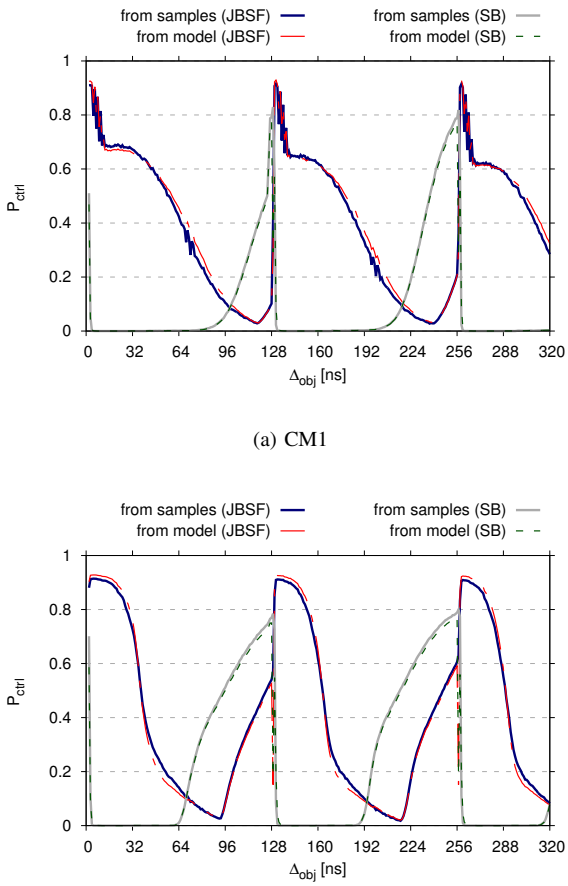

(a) CM1

(b) CM3

Fig. 19. Control probability of the adversary, computed by means of the model and by means of the samples. The adversarial signal-to-noise ratio is $\mathrm{SNR}_m = 30\mathrm{dB}$. The objective precision of the adversary is $\delta_a = 0.5\mathrm{ns}$, which is a reasonable precision for UWB ranging systems.

References

[1] A. Abu-Mahfouz and G. P. Hancke, "Distance bounding: a practical security solution for real-time location systems," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 1, pp. 16–27, February 2013.

[2] J. T. Chiang, J. J. Haas, J. Choi, and Y.-C. Hu, "Secure location verification using simultaneous multilateration," *IEEE Transactions on Wireless Communications*, vol. 11, no. 2, February 2012.

[3] G. Dini, F. Giurlanda, and P. Perazzo, "SecDEv: Secure distance evaluation in wireless networks," in *Proceedings of ICNS'13*, 2013, pp. 207–212.

[4] L. Taponecco, P. Perazzo, A. D'Amico, and G. Dini, "On the feasibility of overshadow enlargement attack on IEEE 802.15.4a distance bounding," *IEEE Communications Letters*, vol. 18, no. 2, pp. 257–260, February 2014.

[5] T. Wang, Y. Liu, and J. Ligatti, "Fingerprinting far proximity from radio emissions," in *ESORICS'14*. Springer, 2014, pp. 508–525.

[6] IEEE Computer Society, "IEEE Std 802.15.4a-2007 (Amendment 1: Add Alternate PHYs)," 2007.

[7] M. Poturalski, M. Flury, P. Papadimitrios, J.-P. Hubaux, and J.-Y. Le Boudec, "Distance bounding with IEEE 802.15.4a: Attacks and countermeasures," *IEEE Transactions on Wireless Communications*, vol. 10, no. 4, pp. 1334–1344, April 2011.

[8] S. Brands and D. Chaum, "Distance bounding protocols," in *Proceedings of EUROCRYPT'93*, 1993, pp. 344–359.

[9] G. P. Hancke and M. G. Kuhn, "An RFID distance bounding protocol," in *Proceedings of SecureComm'05*, 2005, pp. 67–73.

[10] M. Fischlin and C. Onete, "Terrorism in distance bounding: modeling terrorist-fraud resistance," in *Applied Cryptography and Network Security*. Springer, 2013, pp. 414–431.

[11] I. Boureanu, A. Mitrokotsa, and S. Vaudenay, "Practical & provably secure distance-bounding," in *Proceedings of ISC'13*, 2013.

[12] J. Clulow, G. P. Hancke, M. G. Kuhn, and T. Moore, "So near and yet so far: Distance-bounding attacks in wireless networks," in *Proceedings of ESAS'06*. Springer, 2006, pp. 83–97.

[13] M. Poturalski, M. Flury, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec, "On secure and precise IR-UWB ranging," *IEEE Transactions on Wireless Communications*, vol. 11, no. 3, pp. 1087–1099, March 2012.

[14] A. F. Molisch, D. Cassioli, C.-C. Chong, S. Emami, A. Fort, B. Kannan, J. Karedal, J. Kunisch, H. G. Schantz, K. Siwiak *et al.*, "A comprehensive standardized model for ultrawideband propagation channels," *IEEE Transactions on Antennas and Propagation*, vol. 54, no. 11, pp. 3151–3166, November 2006.

[15] I. Sharp and K. Yu, "Indoor TOA error measurement, modeling, and analysis," *IEEE Transactions on Instrumentation and Measurement*, vol. 63, no. 9, pp. 2129–2144, September 2014.

[16] S. Čapkun and J.-P. Hubaux, "Secure positioning in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 2, pp. 221–232, feb 2006.

[17] P. Perazzo and G. Dini, "Secure positioning with non-ideal distance bounding protocols," in *Proceedings of ISCC'15 (to appear)*, 2015.

[18] DecaWave Ltd., http://www.decawave.com.

[19] Nanotron Technologies GmbH, http://nanotron.com.

[20] A. A. D'Amico, U. Mengali, and L. Taponecco, "Energy-based TOA estimation," *IEEE Transactions on Wireless Communications*, vol. 7, no. 3, pp. 838–847, March 2008.

[21] ——, "TOA estimation with the IEEE 802.15.4a standard," *IEEE Transactions on Wireless Communications*, vol. 9, no. 7, pp. 2238–2247, July 2010.

[22] I. Guvenc, Z. Sahinoglu, A. F. Molisch, and P. Orlik, "Non-coherent TOA estimation in IR-UWB systems with different signal waveforms," in *Proceedings of UWBNETS'05*, 2005, pp. 245–251.

Alberto Compagno Alberto Compagno received the MSc degree in Computer Science from the University of Padua, Italy, in 2012. He is currently pursuing the Ph.D. degree in Computer Science at Sapienza University of Rome, Italy, under the supervision of Prof. L. V. Mancini and the co-supervision of Prof. M. Conti from University of Padua. His main research interests include network security and privacy.

Mauro Conti Mauro Conti is an Associate Professor at the University of Padua, Italy. He obtained his Ph.D. from Sapienza University of Rome, Italy, in 2009. After his Ph.D., he was a Post-Doc Researcher at Vrije Universiteit Amsterdam, The Netherlands. In 2011 he joined as Assistant Professor the University of Padua, where he became Associate Professor in 2015. He has been Visiting Researcher at GMU (2008), UCLA (2010), UCI (2012, 2013, and 2014), and TU Darmstadt (2013). He has been awarded with a Marie Curie Fellowship (2012) by the European Commission, and with a Fellowship by the German DAAD (2013). His main research interest is in the area of security and privacy. In this area, he published more than 120 papers in topmost international peer-reviewed journals and conferences. He is Associate Editor for several journals, including IEEE Communications Surveys & Tutorials and IEEE Transactions on Information Forensics and Security. He was Program Chair for TRUST 2015 and ICISS 2016, and General Chair for SecureComm 2012 and ACM SACMAT 2013. He is Senior Member of the IEEE.

Antonio Alberto D'Amico Antonio A. D'Amico received the Dr. Ing. Degree in Electronic Engineering in 1992 and the Ph.D. degree in 1997, both from the University of Pisa, Italy. He is currently an Assistant Professor at the Department of Information Engineering of the University of Pisa. His research interests are in digital communication theory, with emphasis on synchronization algorithms, channel estimation and detection techniques.

Gianluca Dini Gianluca Dini is an associate professor in the Department of Information Engineering, University of Pisa. His research interests include the field of distributed computing systems, with particular reference to security. He has published 100+ papers and has participated in many projects funded by the Commission of the European Community, the Italian Government and private companies.

Pericle Perazzo Pericle Perazzo received the Dr. Ing. degree (cum laude) in Computer Engineering in 2010 and the Ph.D. degree in Information Engineering in 2014, both from the University of Pisa, Italy. During his Ph.D. studies, he has been Visiting Researcher in the Institute for Parallel and Distributed Systems (IPVS) of Stuttgart, Germany. Since 2014, he has been Research Fellow at the Department of Information Engineering at the University of Pisa. His research interests include the area of security and privacy, with special emphasis on location privacy and secure localization.

Lorenzo Taponecco Lorenzo Taponecco received the Dr. Ing. degree (cum laude) in Telecommunications Engineering in 2002 and the Ph.D. degree in Information Engineering in 2007, both from the University of Pisa, Italy. Since 2007, he has been with the Department of Information Engineering at the University of Pisa where he is currently a Research Assistant of telecommunications. His research interests include the area of digital communication theory, with special emphasis on ultra-wideband (UWB) communications, parameter estimation, synchronization and localization techniques.