

On Secure and Precise IR-UWB Ranging

Marcin Poturalski, Manuel Flury,
Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec

Abstract—To provide high ranging precision in multipath environments, a ranging protocol should find the first arriving path, rather than the strongest path. We demonstrate a new attack vector that disrupts such precise Time-of-Arrival (ToA) estimation, and allows an adversary to decrease the measured distance by a value in the order of the channel spread (10-20 meters). This attack vector can be used in previously reported physical-communication-layer (PHY) attacks against secure ranging (or distance bounding). Furthermore, it creates a new type of attack based on malicious interference: This attack is much easier to mount than the previously known external PHY attack (distance-decreasing relay) and it can work even if secret preamble codes are used.

We evaluate the effectiveness of this attack for a PHY that is particularly well suited for precise ranging in multipath environments: Impulse Radio Ultra-Wideband (IR-UWB). We show, with PHY simulations and experiments, that the attack is effective against a variety of receivers and modulation schemes. Furthermore, we identify and evaluate three types of countermeasures that allow for precise and secure ranging.

Index Terms—Security, ranging, distance bounding, impulse radio, ultra-wideband.

I. INTRODUCTION

RANGING, i.e., the estimation of distance between two wireless devices based on message time-of-flight, serves as a building block in many applications and services. A number of applications requires ranging to be both *precise* and *secure*. Secure and precise localization [1], used e.g., for robot fleet navigation, is one example. Physical access control is another: As an example, consider a system for unlocking a car door with a wireless-enabled key that combines two modes of operation in a single physical layer: The door is unlocked if: (1) the key is in communication range of the car (tens of meters) and the user presses the “unlock” button, or (2) the car determines that the key is in proximity of 1-2 meters (conveniently located in the drivers pocket), not requiring the “unlock” button to be pressed (like in Passive Keyless Entry and Start (PKES) systems used in some modern cars). Clearly, the proximity verification should be secure [2], have a meter precision, and provide a range of tens of meters.

For ranging to be *precise* (to achieve meter precision) multipath propagation has to be taken into account. Essentially,

Manuscript received March 17, 2011; revised September 13, 2011 and December 7, 2011; accepted December 18, 2011. The associate editor coordinating the review of this paper and approving it for publication was D. I. Kim.

M. Poturalski, M. Flury, J.-P. Hubaux, and J.-Y. Le Boudec are with the Laboratory for Computer Communications and Applications, EPFL, Switzerland (e-mail: {Marcin.Poturalski, Manuel.Flury, Jean-Pierre.Hubaux, Jean-Yves.LeBoudec}@epfl.ch).

P. Papadimitratos is with the School of Electrical Engineering, KTH, Stockholm, Sweden (e-mail: papadim@kth.se).

Digital Object Identifier 10.1109/TWC.2012.011812.110483

under multipath propagation, a receiver receives a sum of multiple distorted and attenuated copies of the transmitted signal, each shifted in time by the propagation delay of the path it traversed. In typical indoor environments, the *channel spread* of the time-delays is in the order of tens of nanoseconds [3], which translates to at least a few meters of ranging error. To be precise, a receiver needs to accurately identify the *time-of-arrival* (ToA) of the first path. This is especially difficult under *weak non-line-of-sight* (weak NLOS) conditions, when the first path is not the strongest path. A wireless technology that is particularly well suited for this task is *Impulse Radio Ultra-wideband* (IR-UWB or simply IR), because the very narrow pulses it transmits allow for high ranging precision and make the task of resolving the channel easier. We focus on this technology in our investigation.

To make ranging *secure*, i.e., protect the *integrity* of ranging, a cryptographic protocol should be employed. Such protocols are known as *distance bounding* (DB) [4]. More precisely, DB protocols allow a *verifier* to securely measure an *upper-bound* on the distance to a *prover* (honest or malicious). However, because of their cryptographic nature, these protocols abstract away lower-layer details. This makes their implementations potentially vulnerable to physical layer (PHY) attacks [5]. These attacks completely bypass cryptographic mechanisms and decrease the measured distance. The achievable distance-decrease depends on the physical layer, and it can be in the order of kilometers (ISO 14443 RFID, [6]) or hundreds of meters (IR PHY of IEEE 802.15.4a, [7], [8]). In general, IR PHY has the potential to mitigate such attacks: With appropriate countermeasures in place, the achievable distance-decrease can be limited to values in the order of the channel spread, i.e., 10-20 meters [8] (or even less, as we show in Sec. V). This is another reason why we focus on IR.

In this paper, we explore a new PHY attack vector that manifests itself when ranging is designed to be precise, i.e., when the ranging protocol attempts to find the first arriving path. This allows an adversary to *spoof* the ToA estimation, and thus decrease the measured distance and violate ranging integrity. The distance-decrease achievable by this attack is in the order of the channel spread. Although not as effective as the attacks mentioned above, a distance-decrease in the order of 10-20 meters remains a threat in some scenarios. For example, the attack could allow an adversary to gain access to a car parked in front of a restaurant, while the car’s owner is enjoying his dinner, with his key still in the car’s communication range.

We show, with detailed PHY simulations (Sec. IV) and simple experiments (App. A), that the ToA attack vector can be effectively exploited by an external adversary that

introduces malicious interference or mounts a relay attack, or by an internal adversary (malicious prover). The malicious interference attack, although it offers more modest distance-decrease, has two advantages compared to the previously reported external PHY attack (distance-decreasing relay attack, [5], [8]): It is significantly easier to mount, and it works even if a secret preamble code is used. In the case of the relay attack and the malicious prover attack, the ToA attack vector gives an adversary the means to decrease the measured distance by the maximum amount allowed by the countermeasures advocated in [8]. Furthermore, we identify three types of effective countermeasures and provide a detailed performance and cost comparison (Sec. V).

II. RELATED WORK

Distance bounding was first proposed by Brands and Chaum in [4]. A number of other DB protocols are proposed, addressing aspects such as distance bounding over noisy channels [9], mutual ranging [1], [10], resilience to the terrorist fraud [11]–[13], efficiency [14], privacy [15], and formal verification [16], [17]. Most of these proposals ignore PHY issues. Physical layer attacks against distance bounding were first introduced in [5], which included the malicious prover attacks and distance-decreasing relay attacks. The effectiveness of PHY attacks against concrete PHYs is studied in [6] (ISO 14443 RFID and wireless sensor networks) and [7], [8] (IEEE 802.15.4a). It should be noted that the attacks in [5]–[8] focus on the packet payload (some preamble attacks are developed in [7], [8], but only to make payload attacks possible). In contrast, the attacks considered in our paper and in [18] (where we first introduced the malicious prover attack) focus on ToA estimation, hence on the preamble.

An IR-UWB architecture for implementing DB protocols is proposed in [19]. The maximum distance-decrease an adversary can gain against this PHY is 3–6m. This is achieved with a short symbol duration of 20ns, which limits the applicability of this PHY in dense multi-path environments. An ID-based distance bounding protocol is implemented on proprietary IR radios in [20]. In [21], the authors design a time-hopping-based IR PHY that is secure against external, but not internal attacks. Beyond IR-UWB, DB PHYs tailored to narrow-band RFID systems are proposed in [12], [22], [23], and a DB PHY for smartcards (wire-line) is introduced in [24]. These RFID systems are designed for very short communication (centimeters, decimeters or at most a few meters). Hence the multipath problem is much less pronounced than for IR-UWB, and a ToA estimation based on the strongest path is sufficient. The authors of [25] introduce a fast analog method of implementing the rapid response used in DB protocols (< 1ns delay). Their proposal can be applied with a variety of PHYs, including IR-UWB, and it can work with any ToA estimation method, which means its implementations can be potentially vulnerable to the ToA attack vector.

III. SYSTEM MODEL

The distance-decreasing physical-layer attacks that we consider work on the packet level. Hence, we focus on the exchange of a single ranging packet. Extending these attacks

to the protocol level is straightforward: The adversary simply mounts the attack and some or all ranging packets exchanged by the ranging protocol, be it two-way, one-way or pseudo-ranging [8].

A. Modulation Scheme

A ranging packet is composed of: 1) the preamble, 2) the data. The transmitted signal is:

$$s(t) = \sum_{i=1}^{N_P} a_i^P p(t - iT_f^P - \tau_i^P) + \sum_{i=1}^{N_D} a_i^D p(t - iT_f^D - \tau_i^D - T_{\text{tot}}^D) \quad (1)$$

where a P, D index stands for preamble, or data, respectively; N is the number of frames, T_f is the duration of a frame, $a_i \in \{-1, 0, 1\}$ is the amplitude of the i th frame, $\tau_i < T_f$ is the time-hopping offset of the i th frame, $T_{\text{tot}}^D \geq N_P T_f^P$ is the time-offset between the preamble and data, and $p(t)$ is the pulse shape. We assume that the frame duration and time-hopping sequences are such that there is no inter-frame interference.

The data is modulated by a sequence a_i^D known only to the *transmitting party*, whereas for the preamble a_i^P is known to both parties, or is public. The sequences τ_i are known to both parties for the preamble and for the data; they can also be public.

B. Receivers

We consider two basic classes of receivers: a low-complexity non-coherent *energy-detection receiver* (we also use the term *energy detector*), and a more sophisticated *rake receiver*. The former is composed of an antenna, a 500MHz bandpass filter, followed by a squaring device and an integrator that outputs a discrete time sample every $T_{\text{int}} = 2\text{ns}$. The rake receiver is composed of an antenna, a 500MHz bandpass filter and a filter matched to the pulse shape $p(t)$.

We further distinguish between different algorithms implemented by these receivers: vanilla (basic algorithms), PID (Power Independent Detection [26], [27]), and MINF (min-filter [28], [29]) for the energy-detection receiver, and vanilla and PID for the rake receiver. The PID and MINF algorithms are robust to interference. We consider them because they have the potential to prevent malicious interference attacks. In all cases, the receivers operate in the following stages:

- *Coarse synchronization* – the receiver detects a packet, and achieves a rough synchronization (in the order of the frame duration).
- *Fine synchronization / ToA estimation* – following coarse synchronization that provides a rough ToA estimate, the receiver finds a more precise ToA estimate.
- *Channel estimation* (optional) – the receiver estimates the channel delay profile to improve the performance of data demodulation.
- *SFD detection* – the receiver detects a special *Start Frame Delimiter* sequence at the end of the preamble, indicating that data demodulation should starting.
- *Data demodulation*.

The last 3 stages use classic maximum likelihood algorithms for the vanilla energy detector and the rake (maximal ratio combining), or variants of these algorithms that are robust to multi-user interference [30] for the PID and the MINF energy detector.

1) *Coarse Synchronization*: This stage is performed using a traditional synchronization algorithm, based on correlating the received signal with the known preamble *template*. More precisely, the baseline method from [27] is implemented by the vanilla and MINF energy detectors, and the conventional method from [26] is implemented by the vanilla rake. Coarse synchronization locks on the strongest path component of the received signal.

The PID receivers rely on the PID method [26], [27]. In this method the received signal is first compared to a noise-based threshold, and converted into a binary (energy detector) or ternary (rake) sequence (we call this the *PID filter*). In other words, the PID filter performs binary (ternary) quantization of the input signal. The output of the PID filter is then correlated with the preamble template. This has a particular effect relevant to our investigation: When a path is strong enough, the pulses in all frames are converted to 1 (or -1) by the PID filter, and the output of the correlator is maximized. Thus, the PID receiver, in contrast to the vanilla receiver, *cannot distinguish a strong path from the strongest path*.

2) *Fine Synchronization*: ToA estimation performs a *back-search* [31], [32] in a window of duration T_{BS} preceding the rough synchronization point found by coarse synchronization. The back-search window duration T_{BS} should be in the order of the channel spread. The back-search identifies the first time-offset for which a noise-based threshold is reached – this is considered to be the first arriving path, i.e., the ToA. All receivers except MINF perform the back-search on the output of the correlator that correlates the received signal with the template. MINF [28], [29] uses an average of N preamble frames, filtered with a moving min filter before averaging (min-window length $W_{min} = 8$). The min filter removes interference based on the assumption that the interference is present in at most $W_{min} - 1$ consecutive frames.

3) *Data Detection Window*: This parameter, denoted by t_{det} , determines the window that the receiver uses to demodulate one data frame (Fig. 1). For maximum reliability, it should be at least as long as the channel delay spread. It plays an important role in countering the attack.

C. Adversary Model

We distinguish three classes of attacks that aim to decrease the measured distance: An internal *malicious transmitter (prover) attack*, an external *distance-decreasing relay attack* and an external *malicious interference attack*. We focus on the latter attack. In this attack, an honest receiver (HRX) receives a ranging packet (1) from an honest transmitter (HTX). This signal interferes at HRX with the adversarial signal generated by an adversarial transmitter (ATX):

$$s^A(t) = \sum_{i=-\infty}^{\infty} a_i^A p(t - iT_f^A - \tau_i^A) \quad (2)$$

where T_f^A is the duration of a frame, a_i^A is the amplitude, $\tau_i^A < T_f^A$ is the time-hopping offset of the i th frame, and $p(t)$ is the same pulse shape as used by the honest transmitter. Normally, the adversarial signal is spread by the channel. However, we also consider a powerful adversary that transmits over a single-tap channel (i.e., no multipath) by using a highly directive antenna and/or moving close to HRX.

We distinguish between different types of malicious interferers. The simplest *blind* adversary is not equipped with a receiver. Hence, it transmits its signal blindly, without synchronizing with HTX. A *reactive* adversary is equipped with a receiver, and synchronizes its transmission to the transmission of HTX. The level of synchronization can vary. If it is in the order of nanoseconds, we speak of a *precisely synchronized* adversary. Note that even precise synchronization is quite feasible to achieve: If the adversary places its adversarial receiver close to HTX, it can reliably and precisely detect the first pulse transmitted by HTX.

Furthermore, an adversary can be oblivious to the codes used by the honest devices. Or, the adversary can know the preamble codes (and the data time-hopping codes). We assume the adversary does not know the amplitude data code (the payload of the ranging packet) and we assume the adversary always knows the frame lengths in the honest signal. Finally, the adversary can target all devices in range. Or, it can target a specific device, or a specific location. We term these attacks *broadcast* or *targeted*, respectively.

The most sophisticated malicious interference attack we consider is targeted, reactive and precisely synchronized with the HTX; it is mounted by an adversary that knows the preamble codes and the data time-hopping code. Such an attack is – from the perspective of HRX – essentially equivalent to the two other attack classes: In a distance-decreasing relay attack, ATX and an adversarial receiver (ARX) relay (with small modifications on the PHY) the ranging packet transmitted by HTX to HRX. In a malicious transmitter attack, ATX sends a malformed ranging packet to HRX. For more information, see [5], [8].

D. Simulation Setup

We simulate malicious interference attacks only, because the the malicious prover attack and the relay attack are essentially equivalent to a sophisticated malicious interference attack. HRX is exposed to the adversarial signal transmitted by ATX at *signal-to-noise ratio* SNR_A . HRX receives, at random times, ranging packets transmitted by HTX with SNR_H . (In both cases, the SNR is defined as $\frac{E_p}{N_0}$, where E_p is the energy of a single pulse, and the power spectral density is $N_0/2$.) We simulate the entire packet reception process (from synchronization to demodulation) in Matlab, with a self-developed simulator. The physical layer is simulated with an accuracy of 100ps. We use the (weak) NLOS channel model number 2 from [3] with a channel spread of roughly 60ns. Unless otherwise stated, we assume the parameter values summarized in Table I.

Metrics: We consider that *distance-decrease* occurs if a packet is received and the data is recovered without errors, but the estimate ToA is at least $T_{dd} = 4$ ns below the actual

TABLE I

DEFAULT PARAMETER VALUES USED IN SIMULATIONS. THE DEFAULT IEEE 802.15.4A PREAMBLE LENGTH IS USED. THE BACK-SEARCH WINDOW IS CHOSEN TO MATCH THE CHANNEL SPREAD.

frame length	T_f^P, T_f^D	256ns
preamble length	N_P	64 · 31
data length	N_D	32
preamble code	a_i^P	IEEE 802.15.4a code 5
back-search window	T_{BS}	64ns
data detection window	t_{det}	128ns
energy detector sampling time	T_{int}	2ns




Fig. 1. The cicada attack (blind constant 1-attack) mounted against a vanilla energy-detection receiver. HTX sends a ranging preamble modulated with a preamble code $[-1, 0, 1, -1, 0, 0, 1, 1, \dots]$. ATX simultaneously sends the adversarial signal. Both signals propagate through the multipath environment before they are received by HRX. HRX aggregates the received signal over a number of pulses, and finds the strongest path (1). It then searches back for the first path (2), but instead selects the bogus path introduced by the adversary (3). However, even shifted back, the data detection window t_{det} is large enough to capture a significant fraction of the honest signal, allowing for correct data demodulation.

ToA. (T_{dd} is chosen such that the probability of obtaining such a ToA in benign conditions is negligible.) We consider that *denial* occurs, if the packet is not correctly received – either due to failure to detect the packet in coarse synchronization or due to failure of subsequent reception stages. We measure the percentage, or the *rate* of packets subject to distance-decrease or denial. As an additional metric, we measure the amount of distance-decrease (*ToA error*) for the packets for which distance-decrease occurs.

IV. ATTACKS

In this section, we provide an overview of the attack space. We look at various variants of the attack and evaluate their effectiveness against different modulation schemes and receivers. We investigate the attacks in increasing complexity. We first explore the simplest *cicada attack*, and demonstrate that it is effective against most receivers, with and without preamble time-hopping. Next, we generalize the attack to the *coded cicada attack* and show how this attack can defeat the vanilla rake receiver that is robust to the basic cicada attack. Then, we discuss how a *reactive* adversary can improve the attack effectiveness. Finally, we consider the most sophisticated version of the attack, where ATX is tightly synchronized and knows the preamble code. We show that this attack is highly effective, even against strongest-path ToA estimation algorithms.




Fig. 2. Performance of constant 1-attack mounted against a vanilla energy-detection receiver: distance-decrease and of denial (due to coarse synchronization failure or failure of other reception stages) rates at $\text{SNR}_H = 20\text{dB}$.

A. Blind Attacks: The “Cicada” Attack

We start with the simplest attack, the *cicada attack*: A blind adversary transmitting an infinite sequence of identical equally spread pulses: $a_i^A = 1$, $\tau_i^A = 0$. The adversarial frame duration is $T_f^A = \frac{1}{\rho} T_f^P$, where the *attack rate* ρ is an integer. We term this the *constant ρ -attack*. Note that it does not require the knowledge of any codes. Without synchronization with HTX, adversarial pulses are shifted randomly with respect to the honest signal. This is a broadcast attack.

This attack is tailored to modulation schemes without preamble time-hopping. The principle is illustrated in Fig. 1. The signals of HTX and ATX interfere at HRX. If the adversarial signal is weaker than HTX’s signal, HRX should correctly detect HTX’s signal. However, there is a good chance that the fine synchronization algorithm will incorrectly find the “first arriving path” in the adversarial signal. The estimated ToA is then significantly lower than the actual ToA, resulting in a distance-decrease. Furthermore, the distance-decrease is typically low, such that the data detection window still contains a large fraction of the honest signal. The adversarial signal is weaker, hence HRX can demodulate the payload intended by HTX.

If either the honest or the adversarial signal or both use (random) time-hopping, the honest and adversarial pulses will not be aligned. Rather, from the perspective of HRX, the adversarial pulses are randomly spread over time. Still, if the adversary transmits with appropriate power, such random interference turns out to be sufficient to spoof fine synchronization, and not disrupt the other reception stages. We show this in Sec. IV-A2.

1) *Preamble without Time-hopping*: The main factor determining the attack outcome is SNR_A . This can be seen in Fig. 2(a), which shows the performance of the constant 1-attack against the vanilla energy-detection receiver at $\text{SNR}_H = 20\text{dB}$. For low SNR_A , the adversarial signal is too weak to influence the receiver operation. From $\text{SNR}_A \approx 0\text{dB}$ distance-decrease begins, and it reaches its maximum of around 36% for $\text{SNR}_A \approx 15\text{dB}$. Beyond the maximum point, denial begins to take over, and for $\text{SNR}_A \approx 25\text{dB}$, it reaches 100% – partially due to coarse synchronization failure ($\text{Denial}_{\text{SYNC}}$), and partially due to failure of subsequent reception stages ($\text{Denial}_{\text{postSYNC}}$). More generally, the distance-decrease begins at $\text{SNR}_A \approx 0\text{dB}$ and ends at $\text{SNR}_A \approx \text{SNR}_H$, as shown in Fig. 2(b). We observe this general performance pattern for all receivers vulnerable to the attack.

With $\rho = 1$, the adversarial signal, even though spread by




Fig. 3. Distance-decrease rate achieved by constant ρ -attack mounted against a vanilla energy-detection receiver at $\text{SNR}_H = 20\text{dB}$

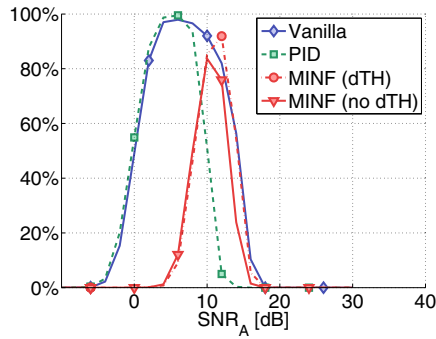


Fig. 4. Distance-decrease rate achieved by constant 8-attack mounted against a vanilla, PID, and MINF energy-detection receivers $\text{SNR}_H = 20\text{dB}$; for the latter we shown performance against modulation with and without data time-hopping (dTH).




Fig. 5. Distance-decrease rate achieved by constant 8-attack mounted against the vanilla rake receiver at $\text{SNR}_H = 10\text{dB}$ with: no time-hopping (“no TH”), data-only time-hopping (“no pTH”), preamble time-hopping with a non-zero-sum template (“not-0-sum”) and a zero-sum template (“0-sum”).

the channel (Fig. 1), is not always present in the back-search window. To increase the probability of distance-decrease, the adversary can increase ρ (Fig. 3). Note that increasing ρ also results in distance-decrease ending at lower SNR_A – this is because there is more interference that disrupts other (than fine synchronization) stages of the receiver operation, causing denial. In subsequent experiments, unless otherwise stated, we fix $\rho = 8$, which strikes a balance between achieving a high maximum distance-decrease rate and not interfering too much with other stages of the receiver operation.

We observe a similar attack performance for the MINF and PID energy-detection receivers (Fig. 4). Both methods were designed with benign interference in mind but, as expected, neither can prevent the attack. In the case of the MINF receiver, the min filter cannot remove the adversarial signal present in *every* frame. Distance-decrease begins at SNR_A about 7dB larger than for vanilla, but only because of a more conservative back-search threshold (inherent to MINF); it ends at the same SNR_A as vanilla. In the case of the PID receiver, distance-decrease starts at the same SNR_A as vanilla, but ends at SNR_A approximately 5dB lower than for vanilla and is due to coarse synchronization failure. The difference in performance occurs because the PID method cannot distinguish the strong enough adversarial pulses from the strongest honest pulses. (Note: In [18] we show how the adversary can use this effect to improve attack performance.) For the vanilla and MINF receivers, the attack performance improves slightly if the modulation scheme uses time-hopping for data. This is because, for these receivers, the attack fails at high SNR due to interference with the data part of the packet. Data time-hopping mitigates some of this interference.

For the vanilla rake receiver, the effectiveness of the attack depends on the sum of the amplitudes of the fine synchronization template. On one hand, if the sum is non-zero, distance-decrease will occur. This is confirmed in Fig. 5, which shows the attack performance at $\text{SNR}_H = 10\text{dB}$ when the template amplitudes follow the IEEE 802.15.4a preamble code 5. In this code, for every 10 frames with $a_i = 1$, there are only 6 frames with $a_i = -1$. The attack performance follows the familiar pattern, and the distance-decrease reaches 80%

– 100% depending on the data modulation scheme (not shown in figure) – data is the main factor limiting attack performance. On the other hand, a fine synchronization template that sums to zero cancels out the constant cicada code, and no distance-decrease occurs.

2) *Preamble with Time-hopping*: The attack works well against the vanilla and PID energy-detection receivers (Fig. 6). Compared to the case without preamble time-hopping, distance-decrease begins at higher SNR_A , because the adversarial pulses are not aligned, rather spread from the perspective of the receiver – hence it takes more power to raise them above the fine synchronization threshold. For vanilla, distance-decrease ends at the same SNR_A with and without preamble time-hopping, because the limiting factor is data demodulation. However, for PID, distance-decrease ends significantly later for the case *with* preamble time-hopping. This is because for PID *without* time-hopping, the limiting factor is coarse synchronization (see Sec. IV-A1). Preamble time-hopping circumvents this limitation.

For the vanilla rake receiver, the attack performance is shown in Fig. 5, and it follows the same pattern as for the vanilla energy detector, with data demodulation being the limiting factor. However, with preamble time-hopping the attack now works with a zero-sum template.

For the MINF receiver, the min filter is somewhat effective in stopping the attack, once the adversarial pulses are randomly spread due to the time-hopping. In Fig. 6 we show the performance of a 16-attack, which manages to reach almost 20% probability of distance-decrease. This is significantly lower than for the other receivers, but still far from negligible.

3) *Ranging Error*: In Fig. 7 we show the median absolute error of the ToA under the 8-attack (for packets for which distance-decrease occurs). It increases with SNR_A , because more adversarial peaks rise above the fine-synchronization threshold. For a high SNR_A , an adversarial pulse in the start of the back-search window is almost always detected as the ToA, hence the distance-decrease tends to T_{BS} . However, the starting point of the back-search is the strongest path in HTX’s signal, which can be greater than the actual distance. In such cases, the distance-decrease is necessarily smaller than T_{BS} .

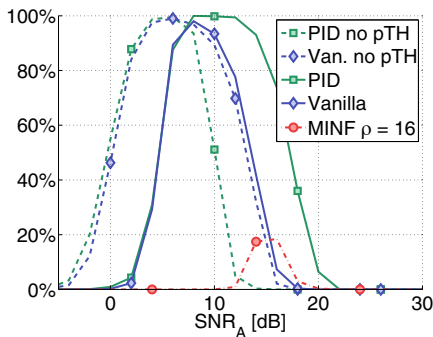


Fig. 6. Modulation with preamble time-hopping: Distance-decrease rate achieved by constant 8-attack (16-attack for MINF) mounted against energy detection receivers at $\text{SNR}_H = 20\text{dB}$. Performance without preamble time-hopping (“no pTH”) is shown for comparison.




Fig. 7. ToA median absolute error under constant 8-attack against the vanilla energy detection receiver at $\text{SNR}_H = 20\text{dB}$ (over packet for which distance-decrease was successful), modulation without time-hopping. 5% and 95% percentiles are also shown.




Fig. 8. Distance-decrease rate achieved by a coded 16-attack against a vanilla rake receiver at $\text{SNR}_H = 20\text{dB}$. The number of frames during which adversary transmits (out of 120) is denoted by n .

Nevertheless, assuming a fixed channel between HTX and HRX (e.g., ranging repeated in a short time interval), a high-rate attack achieves a low distance-decrease variance. This ToA error pattern is universal to all receivers.

4) *Clock Drift*: An unsophisticated or blind adversary will not have the means to adjust its clock speed to exactly match the clock of HTX. We verify that clock drift of 40ppm (the largest allowed by IEEE 802.15.4a) has a negligible effect on the attack effectiveness for energy detectors, and a minor effect (5dB shift of attack start) for rake receivers.

B. Blind Attacks: Coded “Cicada” Attack

We now consider an “coded” version of the cicada attack. It is identical to the basic cicada attack, except that the adversary modulates the amplitudes of the pulses using a non-constant code. This allows the adversary to mount a successful attack against a vanilla rake receiver with zero-sum fine synchronization template (without time-hopping).

Assume that the number of non-zero elements in the fine synchronization template is N and that the time-offset t contains only adversarial signal. Then, fine synchronization considers the offset t as a valid ToA candidate if:

$$\left| \sum_{i=0}^{N-1} a_i^P a_i^A x(t) \right| > N\theta \quad (3)$$

where $a_i^P, a_i^A \in \{-1, 0, 1\}$ are the honest and adversarial codes, respectively, $x(t)$ represents the adversarial signal power at offset t , and θ is a noise-based threshold.

The preamble code a_i^P can be chosen to be pseudo-random, which is the worst case for the adversary. This implies $\mathbb{P}(a_i^P a_i^A = -1 \mid a_i^P \neq 0) = \mathbb{P}(a_i^P a_i^A = 1 \mid a_i^P \neq 0) = 0.5$. Then, the probability of spoofing the ToA for a single time-offset t is:

$$\begin{aligned} \mathbb{P}(A(t)) &= \mathbb{P}(|2\mathcal{B}(n, 0.5) - n| > N\theta x^{-1}(t)) \\ &= 2 \cdot \mathbb{P}(2\mathcal{B}(n, 0.5) - n > N\theta x^{-1}(t)) \\ &= 2 \cdot F_{\text{BIN}}(0.5n - 0.5N\theta x^{-1}(t) \mid n, 0.5) \\ &\leq 2 \exp(-N^2\theta n^{-1} x^{-1}(t)) \leq 2 \exp(-N\theta x^{-1}(t)) \end{aligned} \quad (4)$$

where $n \leq N$ is the number of non-zero elements in the adversarial code a_i^A , $\mathcal{B}(n, 0.5)$ follows the binomial distribution with parameters n and 0.5 and $F_{\text{BIN}}(\cdot \mid n, 0.5)$ is the binomial cdf; the first bound follows from the Hoeffding’s inequality.

Although the probability of spoofing decreases exponentially fast with N , the transmission power $x(t)$ is under the control of the adversary. By increasing $x(t)$, the adversary can achieve a reasonable spoofing probability for practical values of N and θ . Indeed, Fig. 8 shows practical instances of this attack. To increase the attack success probability, the attack rate is $\rho = 16$, meaning that 4 to 5 adversarial frames fall into the back-search window. Although (4) suggests to set $n = N$, this does not take into account the interference created by the adversarial signal during other reception stages. Indeed, Fig. 8 shows that the adversary achieves better results with $n = 12$, rather than $n = 120$, where $N = 120$ in both cases.

C. Reactive Attacks

A reactive adversary can leverage the synchronization with HRX to increase the upper limit on SNR_A . Knowing the algorithms run by the honest receiver, the adversary will know (roughly) when the receiver is performing which stage. Then, it can transmit only when fine synchronization is performed, with an arbitrary high power and high rate ρ , which would guarantee that the ToA is spoofed. If no other stage is performed on this signal, then the adversary will not interfere with these stages. This guarantees correct packet reception.

However, a receiver can perform other reception stages on the same signal as fine synchronization. If this stage is coarse synchronization, SNR_A cannot exceed SNR_H , or the honest signal will be overshadowed by the adversarial signal. If this stage is channel estimation, SNR_A can be increased, but not indefinitely. Intuitively, the ratio between SNR_A and SNR_H in channel estimation determines how much weight is put on the bogus part versus how much is put on the honest information-bearing part of the data symbols.

D. Known Code Attacks

We now look at a targeted version of the malicious interference attack, in which the adversary knows the preamble code




Fig. 9. Distance-decrease rate achieved by known-code attack mounted against a vanilla energy detection receiver at $\text{SNR}_H = 20\text{dB}$. We show the performance under two adversarial channel models: the default one and a single-tap channel (“1-tap”). We distinguish between the usual success in spoofing the 1st path (“1st”), and spoofing of the strongest path (“Str”). In the data, we fix $\text{SNR}_A^D = -\infty\text{dB}$.

and is tightly synchronized with HTX. Under these assumptions, the adversary can spoof the ToA easily, by transmitting a copy of the preamble that arrives at HRX k nanoseconds before the honest preamble. In the data part, the adversary knows the time-hopping sequence, but not the amplitudes – they are determined by the unpredictable challenge/response bits. Hence, ATX transmits a pulse k nanoseconds before each honest pulse, with SNR_A^D . Without loss of generality, we assume that the adversarial pulse has a constant amplitude, whereas the honest pulse has an amplitude of 0/1 for OOK modulation, or of ± 1 for BPSK modulation. Note that in the case of a distance-decreasing relay attack or a malicious prover attack, the adversary would transmit a signal that is the combination of the honest signal and the adversarial signal described above. Hence, we focus on the malicious interference attack.

We can see in Fig. 9 that this attack achieves virtually 100% distance-decrease across a relatively broad range of SNR_A . (We set $\text{SNR}_A^D = -\infty$.) Due to tight synchronization, the distance-decrease is much less variable than with the blind attacks. For $k = 20\text{ns}$, the 5- and 95-percentiles of absolute ToA error are 10ns and 22ns at $\text{SNR}_A = 10\text{dB}$ and 17ns and 22ns at $\text{SNR}_A = 15\text{dB}$. With the single-tap channel, they are between 19ns and 22ns. Furthermore, the adversary can succeed in spoofing not only the first arriving path, but even the strongest path, albeit with a lower probability of success. This circumvents a “countermeasure” that estimates ToA based on the strongest path (Sec. V).

V. COUNTERMEASURES

In this section, we use the findings of Sec. IV to propose and evaluate countermeasures that can thwart the ToA attack vector. After briefly covering naive countermeasures, we look at three different types of stronger countermeasures: (1) *secure ToA estimation*, (2) *early data detection* (EDD) that prevents the attack by reducing the data detection window t_{det} and (3) a *variance countermeasure*, that detects the attack based on the high variance of the ToA estimates observed under multiple measurements. The first two countermeasures are universally applicable, whereas the variance countermeasure should be

coupled with secret preamble time-hopping and a large back-search window. For the variance countermeasure and secure ToA estimation preamble-only packets are sufficient if both ranging devices are honest. However, if the prover can be malicious, additional data unpredictable to the prover is necessary.

A. Naive Countermeasures

An obvious way of countering the attack is to disable fine synchronization, and estimate ToA based on the strongest path. This has, however, the significant disadvantage of decreasing the ranging precision. Furthermore, an adversary that knows the preamble codes and that is precisely synchronized with the honest transmitter can circumvent this method (Sec. IV-D).

In Sec. IV-A2 we have seen that the MINF energy-detection receiver, if coupled with secret preamble time-hopping, can offer some degree of protection against the attack. However, the adversary still has a non-negligible probability of success, thus making this a relatively weak countermeasure.

B. Secure ToA Estimation

The key to securing the ToA estimation is the PID method. Indeed, we have seen in Sec. IV-B that a rake receiver that does not use PID is vulnerable to an attack by an adversary that transmits with high power.

1) *PID Rake*: Assuming that the ± 1 amplitudes of non-zero preamble frames are random, and the rake receiver uses the PID method, we can apply an almost identical reasoning as in (4) to estimate the probability of spoofing the ToA for a single time-offset t :

$$\begin{aligned} \mathbb{P}(A(t)) &= \mathbb{P}(|2\mathcal{B}(n, 0.5) - n| > N\theta) \\ &= 2 \cdot F_{\text{BIN}}(0.5n - 0.5N\theta | n, 0.5) \\ &\leq 2 \exp(-N\theta) \end{aligned} \quad (5)$$

where N is the number of non-zero frames in the fine synchronization template, $n \leq N$ is the number of non-zero frames in the adversarial code a_i^A , and θ is a noise-based threshold. The difference with (4) is that the power factor x disappears. This is crucial, because now the adversary can no longer increase the transmission power to compensate for an exponentially fast decline of $\mathbb{P}(A(t))$. Hence, the PID rake receiver is secure against the attack, if the values of N and θ are chosen appropriately. This can be done in the same fashion as for the PIDH countermeasure, and the effect on ranging precision will be similar, hence we omit the details.

2) *PIDH Countermeasure*: The PIDH ToA estimation relies on the PID method, but it uses the Hamming distance in place of correlation. Hence, we call it *Power Independent Detection with the Hamming distance* (PIDH).

More specifically, the PIDH fine synchronization algorithm takes as input: (1) the known preamble template $a_{i=1,\dots,N} \in \{0, 1\}$, and (2) a sequence of samples $y_{i=1,\dots,N} \in \{0, 1\}$ (after applying the PID filter) corresponding to some time-offset t . Then, the time-offset t is considered a valid ToA candidate if:

$$d(y_i, a_i) \leq T_{\text{PIDH}} \quad (6)$$

where d is the Hamming distance and T_{PIDH} is a threshold.

For optimal security, the template a_i should be a random binary sequence. Under this assumption, for a single time-offset t , the adversary can spoof the ToA with probability:

$$\mathbb{P}(A(t)) = F_{\text{BIN}}(T_{\text{PIDH}}|N, 0.5) \quad (7)$$

Given that there are N_{BS} time-offsets in the back-search window that a receiver evaluates, the total probability that the adversary achieves a distance-decrease can be upper-bounded by:

$$\mathbb{P}(A) = \mathbb{P}\left(\bigcup A(t_j)\right) \leq \sum \mathbb{P}(A(t_j) = N_{\text{BS}} \cdot F_{\text{BIN}}(T_{\text{PIDH}}|N, 0.5)) \quad (8)$$

For a desired security level P_{attack} , we can invert (8) and obtain a threshold T_{PIDH} that achieves this security level:

$$T_{\text{PIDH}} = F_{\text{BIN}}^{-1}(P_{\text{attack}}|N, 0.5) \cdot N_{\text{BS}}^{-1} \quad (9)$$

C. Early Data Detection

The idea of the *early data detection* (EDD) countermeasure is simple: shrink the data detection window drastically, such that any distance-decrease removes the honest signal from the window, making data demodulation fail. More specifically, the countermeasure works as follows:

- 1) **ToA estimation.** Perform non-secure fine synchronization to find the first arriving path. Denote the preamble template length by N_{toa} .
- 2) **Verification.** Perform data demodulation with detection window $t_{\text{det}} = T_{\text{int}}$ (on the time-offset determined by fine synchronization); reject the packet if the number of errors is above N_{err} . The data length is N_{nonce} .

The EDD countermeasure offers poor benign-case performance if the first path is weak: The ToA estimation can detect such a path, but it is likely that this path is too weak to perform reliable data demodulation, resulting in a rejection of a valid ranging packet. To address this, we propose the following extension:

- 1) **ToA estimation.** Perform non-secure fine synchronization that selects N_{off} time-offsets t_i :
 - a) t_1 is selected as the first offset in the back-search window above the noise-based threshold (regular ToA estimation).
 - b) t_{i+1} is selected as the first offset in the back-search window above the i th time-offset. The last offset $t_{N_{\text{off}}}$ corresponds to the strongest path.
- 2) **Verification.** Perform data demodulation with detection window $t_{\text{det}} = T_{\text{int}}$ for every time-offset t_i ; choose the final ToA estimate to be the first t_i for which the number of errors is below N_{err} . The data length is N_{nonce}

This countermeasure finds the first path that is strong enough for secure verification. Hence, in some cases it returns a ToA that is slightly larger than the true ToA, rather than rejecting the packet as EDD does. We hence term it *EDD with Graceful degradation* (EDDG).

An obvious drawback of this type of countermeasures is a significant reduction of data demodulation reliability. However, this can be compensated by increasing the data

length. For a desired security level P_{attack} , the values of N_{nonce} and N_{err} can be determined by a formula equivalent to (9):

$$N_{\text{err}} = F_{\text{BIN}}^{-1}(P_{\text{attack}}|N_{\text{nonce}}, 0.5) \cdot N_{\text{off}}^{-1} \quad (10)$$

Note that a countermeasure based on shrinking the data detection window was first proposed in [6]. The name of EDD is, ironically, based on the *early detection* attack, in which the adversary decreases the data detection window to obtain a distance-decrease [5]. A version of this countermeasure is also advocated in [8]. See [8] for the derivation of (10).

D. Variance Countermeasure

The variance countermeasure prevents the attack by performing multiple ToA measurements and by detecting a high variance of the ToA estimates caused by the attack. In its basic form, this countermeasure can mitigate some naive adversaries. To be truly effective, it should be coupled with a secret preamble time-hopping and (perhaps counter-intuitively) a large back-search window.

ToA estimation with the basic *variance countermeasure* proceeds in the following Steps:

1. Perform M independent ToA measurements, thus obtaining ToA estimates t_1, \dots, t_M .
2. If the ratio of failed measurements is above R_{fail} , reject. Consider a ToA measurement that is “too small” (see below) as failed.
3. If $\text{Var}(t_i) > T_{\text{var}}$, reject.
4. Compute the final ToA estimate $t = \text{median}(t_i)$.

R_{fail} and T_{var} are parameters that can be set based on the desired benign- and adversarial-case performance of the system.

The variance countermeasure should be coupled with secret preamble time-hopping. Otherwise, the adversary (notably a tightly synchronized one) can achieve a constant distance-decrease on all ToA estimates and mitigate the variance test in Step 3. Furthermore, even with secret preamble time-hopping, a high-rate attack achieves a low variance (in the order of the square of the adversarial inter-pulse spacing). However, under such an attack the final ToA estimate falls into the beginning of the back-search window. Hence, we can chose a *large* T_{BS} , forcing a high-rate attack to result in a “too small” ToA estimate that will be rejected in Step 2. A “too small” ToA estimate can be detected if it results in a negative distance estimate (in case of ranging). Or, it can be detected implicitly if the back-search window is set to be noticeably larger than the data detection window t_{det} , thus making data demodulation fail under a high distance-decrease. In addition, Step 2 prevents an adversary from deliberately failing a significant number of ToA estimates t_i , thus achieving a higher probability of circumventing the variance test in Step 3.

Determining the security level of this countermeasure is much less straightforward than for the other countermeasures. We provide an estimate of this probability under a specific set of parameters: We set $T_{\text{BS}} = 2t_{\text{det}}$ and set the maximum time-hopping offset to t_{det} , where $t_{\text{det}} = 64\text{ns}$. We set $R_{\text{fail}} = 0$ and $T_{\text{var}} = (5\text{ns})^2$, which balances good benign-case performance with security. Larger values of R_{fail} and T_{var} improve

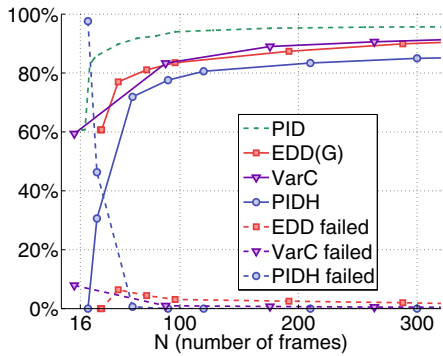


Fig. 10. Countermeasures benign case performance: We show the precise-ToA rate (“PIDH”, “EDD(G)”, “VarC”), as a function of the number of frames N at $\text{SNR}_H = 20\text{dB}$. For the vulnerable PID algorithm, $N = N_{\text{toa}}$ (the template length), likewise for PIDH. For EDD(G), $N = N_{\text{toa}} + N_{\text{nonce}}$, the latter being the length of the data (roughly 5 to 7 times larger than N_{toa}). For the variance countermeasure, N is the length of the fine synchronization template times M . In addition, we show the failure rate (“failed”). Note that we ignore the data frames necessary for all countermeasures except EDD(G); such data frames are necessary to defend against a malicious prover.

benign case performance slightly, but they degrade security significantly.

Furthermore, we assume (for simplicity) that there is only one frame per ToA estimation, and we consider the following, powerful adversary. The adversary is tightly synchronized, and knows exactly when a honest preamble pulses with time-hopping offsets 0 are sent; we denote this time by 0. (Obviously, the adversary does not know the secret time-hopping offsets). Finally, the adversary wants to achieve a distance-decrease of at least 3m.

We determine numerically that the best strategy for this adversary is to transmit a pulse at time 0. Indeed, a pulse transmitted at time $t < 0$ causes measurement rejection if the preamble time-hopping offset is above $t_{\text{det}} - t$. This leads to crossing the R_{fail} threshold. A pulse transmitted at time $t > 0$ results in $\frac{t}{t_{\text{det}}}$ of the ToA estimates having no distance-decrease (all preamble frames with time-hopping offset below t). If this fraction is large, only a small distance-decrease can occur (lower than 3m). If the fraction is small, the variance is actually increased, because all the ToA estimates t_i without distance-decrease are far from the mean ToA. Given the $t = 0$ strategy, the distance-decrease for the M measurements can be modeled as M iid uniform random variables ranging from 0 to $t_{\text{det}} \cdot c$. We obtain numerically that to limit $P_{\text{attack}} \leq 2^{-16}$, we need to set $M \geq 11$.

E. Performance Evaluation

We compare the benign case performance of the countermeasures tuned to provide the same security level, which we fix at $P_{\text{attack}} = 2^{-16}$. We derive the thresholds for PIDH, EDD and EDDG according to (9) and (10); for the variance countermeasure, M was estimated in Sec. V-D. Our primary metric pertains to ranging precision: We count the percentage of ToA estimations for which the receiver finds the first path, i.e., the ranging error is below 2ns. We term this metric *precise-ToA rate*. We find it to be more informative than the rather coarse-grained mean absolute ranging error. Furthermore, we count the percentage of ToA estimations that




Fig. 11. Countermeasures performance under constant 8-attack (at $\text{SNR}_A = 8\text{dB}$). We show the precise-ToA rate (“PIDH”, “EDD”, “EDDG”), as a function of the number of frames N at $\text{SNR}_H = 20\text{dB}$. For PIDH, N is the template length. For EDD/EDDG, $N = N_{\text{toa}} + N_{\text{nonce}}$, where N_{toa} is the template length and N_{data} is length of the data. We also show the failure rate of all methods (“failed”).

fail, i.e., the *failure rate*. The exact values that we obtain are specific to the channel model (the NLOS model number 2 from [3]), but we are interested in the relative performance of different countermeasures.

The precise-ToA and failure rates of the countermeasures is show in Fig. 10. In terms of ranging precision, EDD and EDDG display a virtually identical precise-ToA rate. Both outperform PIDH for all N . EDD/EDDG and PIDH work in a similar fashion, and the performance difference is mostly due to N_{BS} being much higher than N_{off} in (9) and (10) – resulting in a more conservative threshold for PIDH. EDD also outperforms the variance countermeasure for low N ; for higher N , the variance countermeasure perform slightly better than EDD. Compared to the basic PID method, which is not secure against the ToA attack vector, all countermeasures experience a significant performance drop: the cost of precision and security is quite high.

In terms of failure rate, PIDH and EDDG (latter not shown in figure) display negligible failure for $N > 100$. This is in contrast to EDD that experiences a failure rate of 2 – 3% (as expected). The variance countermeasure shows a failure rate of below 1%. Under attack (Fig. 11), EDDG and PIDH achieve a precise-ToA rate in the order of 80%, and a failure rate in the order of 5%. In contrast, the failure rate for EDD is quite high, and for the variance countermeasure is it virtually 100% (not in figure). In the latter case, this is caused by reliable attack detection on subsequent packet rejection. Although the attack robustness shown by EDDG and PIDH can be desirable in some cases, is not a crucial feature – an adversary aiming at denial can obtain a 100% failure rate by simply increasing the transmission power.

Taking both metrics into account, we can recommend EDDG as the most robust and precise countermeasure. Although the variance countermeasure can offer a slightly higher precision, it comes at a cost of non-negligible failure rate, notably under attack, and longer frames to accommodate time-hopping and a long back-search window. (For our choice of parameters, the frame duration is doubled).

VI. CONCLUSION

We have identified a novel attack vector against IR-UWB ranging. The attack allows an adversary to decrease the distance measured by ranging protocols that are designed to be precise (by searching for the first arriving path). The attack disrupts the time-of-arrival (ToA) estimation procedure, exploiting a fundamental difficulty in distinguishing the signal of interest from interference. We have demonstrated that even a simple-to-mount variant of this attack (the “cicada” attack) is effective against a number of modulation schemes and receivers.

Furthermore, we have shown that the attack can be thwarted by a three types of countermeasure. First, we identify a *secure ToA estimation* algorithm (PID) for rake receivers, and propose a new algorithm (PIDH) for energy-detection receivers. Second, we revisit the *early data detection* countermeasure proposed in [6], and propose a more robust extension (EDDG). Third, we propose a *variance countermeasure* that prevents the attack if coupled with a secret preamble time-hopping and a long back-search window. All three countermeasures allow for ranging that is both *precise* and *secure*, meaning that they prevent the adversary from decreasing the measured distance by more than the ranging precision of the receiver. This is an improvement on the state of the art for IR-UWB, in which the achievable distance-decrease was limited to approximately 10m [8]. We have demonstrated that security comes at a cost: Achieving a good ranging precision requires packets of length an order of magnitude larger than ranging algorithms vulnerable to the ToA attack vector.

VII. ACKNOWLEDGMENTS

We would like to thank Alexander Feldman, Alexander Bahr, James Colli-Vignarelli, Stephan Robert, Catherine Dehollain, and Alcherio Martinoli for giving us access to the IR test-bed [33] and for their invaluable help in performing the experiments reported in the Appendix.

The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

APPENDIX EXPERIMENTS ON AN IR TEST-BED

We evaluate the cicada attack performance on an Impulse-Radio test-bed [33]. The test-bed is designed for TDOA localization. It comprises a set of static receivers and a set of mobile transmitters. The transmitter is composed of a front-end (described in [34]) and of a simple modulator module. The transmitter can send a train of pulses with pulse repetition frequency (PRF) of 10MHz. The train of pulses can on/off modulated. Pulses have a bandwidth of 500MHz at 4.25GHz central frequency. The receiver front-end is a energy-detection receiver with a twist inconsequential for our investigation. The receiver back-end is composed of analog-to-digital converter (ADC) which sampling of 2.842GS/s and 8 bit resolution and a Field-Programmable Gate Array (FPGA), connected to a PC. The FPGA has relatively limited memory of 140Kb, which

influences the experiment design. The receiver is described in detail in [33].

At the time we performed the experiments, the test-bed was still under development and it did not provide synchronization between any pairs of devices (i.e., no common clock base). This causes the following difficulty: When the receiver estimates the ToA, there is no *ground truth* ToA that the estimate could be compared with. Hence, we cannot directly determine if a distance-decrease has occurred.

A. Experiment Design

We resolve the limitations of the test-bed with a new experiment design. We involve the usual set of devices: HTX, ATX and HRX. HTX transmits a series of ranging packets back-to-back. Each ranging packet is composed of one preamble symbol, composed of 31 frames modulated according to the IEEE 802.15.4a preamble code no 5. (We refrain from using complete ranging packets due to the modest memory on the FPGA.) ATX transmits periodically a sequence of 4·31 frames modulated with a sequence of 2·31 ones followed by 2·31 zeros. Note that this is a simple variant of the cicada attack with rate $\rho = 1$. The frame duration is 100ns. Examples of the HTX and ATX signals are shown in Fig. 13.

This transmission pattern has the following result: Some of the honest ranging packets are affected by the cicada signal, and some are not. We use the latter ranging packets to estimate the *ground truth ToA*.

In a single *experiment*, we capture $12\mu\text{s}$ of samples (maximum that the FPGA allows). The duration of a preamble symbol is $3.1\mu\text{s}$, hence in one experiment we obtain a ToA estimate for 3 ranging packets (dashed lines in Fig. 13(a)).

We consider two *configurations*: In configuration (A), there are no obstacles between either transmitter and the receiver. In configuration (B) we attenuate the ATX signal by putting a tinfoil obstacle between ATX and HRX. In both configurations, the distances between HTX and HRX and ATX and HRX are in the order of one meter. The test-bed is located in the corner of a 10m by 10m electrical engineering lab room. For each configuration, we perform 1000 experiments.

We define SNR as the ratio between the signal power and the noise power. (In contrast to the other section of the paper, where we define SNR as the ratio between the pulse energy and the noise spectral density, E_p/N_0 . Converting E_p/N_0 to $P_{\text{signal}}/P_{\text{noise}}$ entails subtracting roughly 24dB.) In both configurations, $\text{SNR}_H \approx 22\text{dB}$. In configuration (A), $\text{SNR}_A \approx 21\text{dB}$. In configuration (B), SNR_A varies between 6dB and 17dB with mean 11dB. Note that these signals are significantly stronger than the signals used in the simulations. This is due to the short distances between devices in the experiment setup.

B. Metrics

We define the *coverage* of a ranging packet received by HRX as the percentage of the packet that is covered by the ATX signal (Fig. 13). In practice, we compute the coverage based on the position of the cicada signal, which we detect with the PIDH method. (We chose this method for its robustness.) We note that the detection is not extremely accurate (it




Fig. 12. Distance-decrease rate and denial rate as a function of coverage. Configuration (A) results are shown in (a) for vanilla, (b) for PID and (c) for MINF. Configuration (B) results are shown in (d) for vanilla, (e) for PID and (f) for MINF. We show 95% confidence intervals for distance-decrease rate.

can be off by a few frames), but this is good enough for our purposes.

Based on the coverage, we define the *ground truth* ToA estimate as the HRX ToA estimate for the preamble symbol with coverage 0. If no preamble symbol has such coverage, we discard the experiment. We measure the ToA estimation error by comparing the other two ToA estimates (coarse or fine) to the ground truth ToA estimate.

For a preamble symbol we consider that *denial* occurs (synchronization fails) if the coarse synchronization error is greater than 100ns. We consider that *distance-decrease* occurs if denial does not occur and if the fine synchronization ToA estimate is lower by more than 4ns than the ground truth estimate. We then measure the denial rate and distance-decrease rate as the percentage of preamble symbols for which denial or distance-decrease occurs, respectively.

Note that contrary to simulations performed in Sec. IV, denial can occur due to a benign failure. This is because we are working with prototype-grade hardware, which sometime results in imperfections in the received signals. (Note the pulse amplitude variability in Fig. 13, notably (c)).

C. Honest Receiver Operation

HRX operation is adjusted to the signal transmitted by HTX. Most notably, the receiver only performs synchronization, but no channel estimation, SFD detection or data demodulation. *Coarse synchronization* can be significantly simplified compared to Sec. III, as HRX can a priori assume that the HTX signal is always present. HRX simply correlates the received signal with a template that corresponds to one preamble symbol; it then chooses the maximum in the correlator output of duration $3.1\mu\text{s}$, which is the “ranging packet” duration (i.e., one preamble symbol).

Fine synchronization is performed on the same signal as coarse synchronization. It follows the description in Sec. III, searching back from the coarse ToA estimate for a time-offset above a threshold. The only differences are the duration of the back-search window (32ns in place of 64ns, corresponding to the observed channel spread) and the MINF window size (4 in place of 8, to account for honest signal imperfections).

In our implementation, we use the ADC and FPGA to capture the samples, and send them to the PC, where all




Fig. 13. Samples from the HRX’s ADC for a single experiment: HTX signal only is shown in (a), the dashed lines show the start of preamble symbols (ranging packets). In (b) and (c) ARX signal only is shown. If (a) and (b) interfere at HRX, the coverage of the 3 consecutive ranging packets is roughly equal to 1, 0.9 and 0.

subsequent receiver operations are implemented in Matlab. This minimizes the implementation overhead. It also allows us to process the samples off-line and to run different receiver algorithms on the same input.

D. Experimental Results

We show the distance-decrease rate and the denial rate as a function of coverage in Fig. 12. The results met our expectations. Consider the first configuration (A), where the ATX signal is roughly as strong as the HTX signal. For low coverage, the vanilla receiver shows some distance-decrease. The distance-decrease increases up to roughly 20% as the coverage increases, and for coverage beyond 0.3 denial starts to take over. To explain this behavior, recall that the vanilla receiver *sums* the samples from the frames indicated by the template. Hence, the coverage effectively works as a

multiplying factor for the total adversarial signal power; e.g., at coverage 0.5, SNR_A can be considered approximately 3dB lower than at coverage 1. Thus, the distance-decrease rate and denial rate follow the same pattern that in Sec. IV. A similar coverage interpretation can be applied to MINF receiver, although it should be noted that the min filter causes some non-linear distortion. In particular, at low coverage, the min filter removes the cicada signal completely, and thus no distance-decrease is observed. In general, the performance pattern of the attack is similar as for vanilla, but the distance-decrease rate reaches only roughly 10% due to a conservative MINF threshold value.

For the PID receiver, there is no simple parallel between the coverage and SNR_A , but we can find an equally simple interpretation. In noiseless conditions, the output of the PID correlator at a time-offset corresponding to the true ToA (the start of the “ranging packet”) is equal to N , the number of non-zero frames in the template. With coverage x , the output of the correlator at a time-offset with adversarial signal contribution is roughly $x \cdot N$. This explains why in Fig. 12(b) denial becomes dominant only at coverage close to 1. This allows the distance-decrease rate to reach roughly 35%.

In configuration (B) the ATX signal is weaker than the HTX signal. Hence, for the vanilla and MINF filters, Fig. 12 can be considered to show the performance for a low SNR_A only. This explains why denial occurs marginally, and the distance-decrease rate increases with coverage. In contrast, for the PID receiver the power of the adversarial signal plays a smaller role, as long as it is above the PID filter threshold. Hence, the attack performance in configuration (B) is close to the performance in configuration (A).

Finally, we verify that no distance-decrease occurs with the PIDH fine synchronization algorithm. The other two countermeasures are not applicable, as there is no payload and not enough packets to get meaningful variance estimates.

REFERENCES

- [1] S. Čapkun and J. Hubaux, “Secure positioning in wireless networks,” *IEEE J. Sel. Areas Commun.*, vol. 24, no. 2, 2006.
- [2] A. Francillon, B. Danev, and S. Čapkun, “Relay attacks on passive keyless entry and start systems in modern cars,” in *2011 Netw. Distributed Syst. Security Symp.*
- [3] A.-F. Molisch, K. Balakrishnan, C.-C. Chong, S. Emami, A. Fort, J. Karedal, J. Kunisch, H. Schantz, U. Schuster, and K. Siwiak, “IEEE 802.15.4a channel model - final report, document 04/662r1,” 2004.
- [4] S. Brands and D. Chaum, “Distance-bounding protocols,” in *1993 EUROCRYPT*.
- [5] J. Clulow, G. P. Hancke, M. G. Kuhn, and T. Moore, “So near and yet so far: distance-bounding attacks in wireless networks,” in *2006 ESAS*. Available: <http://www.crysys.hu/ESAS2006/cfp.html>
- [6] G. P. Hancke and M. G. Kuhn, “Attacks on time-of-flight distance bounding channels,” in *2008 WiSec*.
- [7] M. Flury, M. Poturalski, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec, “Effectiveness of distance-decreasing attacks against impulse radio ranging,” in *2010 ACM Conf. Wireless Netw. Security*.
- [8] M. Poturalski, M. Flury, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec, “Distance bounding with IEEE 802.15.4a: attacks and countermeasures,” *IEEE Trans. Wireless Commun.*, to appear.
- [9] G. Hancke and M. Kuhn, “An RFID distance bounding protocol,” in *2005 SecureComm*.
- [10] D. Singelée and B. Preneel, “Distance bounding in noisy environments,” in *2007 ESAS*.
- [11] L. Bussard, “Trust establishment protocols for communicating devices,” Ph.D. dissertation, 2004.
- [12] J. Reid, J. M. G. Nieto, T. Tang, and B. Senadji, “Detecting relay attacks with timing-based protocols,” in *2007 ASIACCS*.
- [13] C. H. Kim, G. Avoine, F. Koeune, F.-X. Standaert, and O. Pereira, “The swiss-knife RFID distance bounding protocol,” in *ICISC*, P. Lee and J. Cheon, editors, 2008.
- [14] C. H. Kim and G. Avoine, “RFID distance bounding protocol with mixed challenges to prevent relay attacks,” in *2009 CANS*. Available: <http://www.uclouvain.be/sites/security/download/papers/KimA-2009-cans.pdf>.
- [15] K. B. Rasmussen and S. Čapkun, “Location privacy of distance bounding protocols,” in *2008 CCS*.
- [16] C. Meadows, R. Poovendran, D. Pavlovic, L.-W. Chang, and P. Syver-son, “Distance bounding protocols: authentication logic analysis and collusion attacks,” in *Secure Localization and Time Synchronization for Wireless Sensor and Ad Hoc Networks*. Springer-Verlag, Series: Advances in Information Security, Vol. 30, 2007.
- [17] P. Schaller, B. Schmidt, D. Basin, and S. Čapkun, “Modeling and verifying physical properties of security protocols for wireless networks,” in *Proc. 2009 IEEE Comput. Security Foundations Symp.*
- [18] M. Poturalski, M. Flury, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec, “The cicada attack: degradation and denial of service in IR ranging,” in *2010 IEEE International Conf. Ultra-Wideband*.
- [19] M. Kuhn, H. Luecken, and N. Tippenhauer, “UWB impulse radio based distance bounding,” in *Proc. 2010 Workshop Positioning, Navigation Commun.*
- [20] N. O. Tippenhauer and S. Čapkun, “ID-based secure distance bounding and localization,” in *Proc. 2009 European Symp. Research Comput. Security*.
- [21] A. Benfarah, B. Miscopein, J. Gorce, C. Lauradoux, and B. Roux, “Distance bounding protocols on TH-UWB radios,” in *2010 IEEE Global Telecommun. Conf.*
- [22] J. Munilla and A. Peinado, “Distance bounding protocols for RFID enhanced by using void-challenges and analysis in noisy channels,” *Wireless Commun. Mobile Comput.*, vol. 8, no. 9, 2008.
- [23] G. Hancke, “Design of a secure distance-bounding channel for RFID,” *Elsevier J. Netw. Comput. Appl.*, 2010.
- [24] S. Drimer and S. Murdoch, “Keep your enemies close: distance bounding against smartcard relay attacks,” in *Proc. 2007 USENIX Security Symp.*
- [25] K. B. Rasmussen and S. Čapkun, “Realization of RF distance bounding,” in *Proc. 2010 USENIX Security Symp.*
- [26] A. El Fawal and J.-Y. Le Boudec, “A robust signal detection method for ultra wide band (UWB) networks with uncontrolled interference,” *IEEE Trans. Microwave Theory Techniques*, vol. 54, no. 4, part 2, pp. 1769–1781, 2006.
- [27] M. Flury, R. Merz, and J.-Y. Le Boudec, “Robust non-coherent timing acquisition in IEEE 802.15.4a IR-UWB networks,” in *2009 PIMRC*. Available: <http://www.pimrc2009.org/>
- [28] Z. Sahinoglu and I. Guvenc, “Multiuser interference mitigation in noncoherent UWB ranging via nonlinear filtering,” *EURASIP J. Wireless Commun. Netw.*, 2006.
- [29] D. Dardari, A. Giorgetti, and M. Win, “Time-of-arrival estimation of UWB signals in the presence of narrowband and wideband interference,” in *2007 ICUBW*.
- [30] M. Flury, “Interference robustness and security of impulse-radio ultra-wide band networks,” Ph.D. dissertation, Lausanne, 2010. Available: <http://library.epfl.ch/theses/?nr=4698>
- [31] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Win, “Ranging with ultrawide bandwidth signals in multipath environments,” *Proc. IEEE*, vol. 97, no. 2, 2009.
- [32] I. Guvenc, Z. Sahinoglu, P. Orlik, and H. Arslan, “Searchback algorithms for TOA estimation in non-coherent low-rate IR-UWB systems,” *Wireless Pers. Commun.*, vol. 48, no. 4, 2009.
- [33] A. Feldman, A. Bahr, J. Colli-Vignarelli, S. Robert, C. Dehollain, and A. Martinoli, “Toward the deployment of an ultra-wideband localization test bed,” in *Proc. 2011 Veh. Technol. Conf.*
- [34] J. Colli-Vignarelli and C. Dehollain, “A discrete-components impulse-radio ultrawide-band (IR-UWB) transmitter,” *IEEE Trans. Microwave Theory Techniques*, vol. 59, no. 4, Apr. 2011.



Marcin Poturalski earned his Master of Science degree in Computer Science and Bachelor Degree in Mathematics from Warsaw University in 2005. During his studies he completed internships Microsoft Research Asia in Beijing and Microsoft Corp. in Redmond, USA. In 2006 he joined the Laboratory for Computer Communication and Applications (LCA) at EPFL, School of Computer and Communication Sciences. He earned his PhD in 2011. His research interests are in security of wireless communication, notably neighbor discovery, ranging

and localization.



Manuel Flury earned his Master of Science degree in Communication Systems Engineering from Ecole Polytechnique Federale de Lausanne (EPFL) in 2005. During his studies he completed internships at Nokia Research Center in Helsinki, Finland and at Qualcomm Inc. in San Diego, USA. In 2005 he joined the Laboratory for Computer Communication and Applications (LCA) at EPFL, School of Computer and Communication Sciences, and began working on his PhD thesis. There, he participated in the National Center of Competence in Research

project on Mobile Information and Communication Systems (NCCR-MICS) and he earned his PhD in 2010. His research interests are in wireless communication and computer networks.



Panos Papadimitratos earned his PhD from Cornell University, Ithaca, NY, USA, in 2005. After being a post-doctoral fellow at Virginia Tech, Blacksburg, VA, USA, and a scientist at EPFL, Lausanne, Switzerland, he joined the faculty of KTH, Stockholm, Sweden, where he is now an Associate Professor in the School of Electrical Engineering. His research is concerned with security and networked systems, with more than 70 related technical publications. He has delivered a number of invited talks and tutorials, including tutorials at the ACM

CCS and the ACM MobiCom conferences. He is an Area Editor for the ACM MC2R journal and he has served in several technical program committees, including IEEE INFOCOM and ACM WiSec, ASIACCS, and MobiHoc. His webpage is: <http://www.ee.kth.se/~papadim/>



Jean-Pierre Hubaux joined the faculty of EPFL in 1990. He became full professor in 1996. He is a Fellow of both ACM and IEEE. His current research activity is focused on privacy preservation mechanisms in pervasive communications.

In 1991, he designed the first curriculum in Communication Systems at EPFL. He was promoted to full professor in 1996. In 1999, he defined some of the main ideas of the National Competence Center in Research named "Mobile Information and Communication Systems" (NCCR/MICS); this center (still very active) was initially nicknamed "the Terminodes Project". In this framework, he has notably defined, in close collaboration with his students, novel schemes for the security and cooperation in wireless networks; in particular, he has devised new techniques for key management, secure positioning, and incentives for cooperation in such networks. In 2003, he identified the security of vehicular networks as one of the main research challenges for real-world mobile ad hoc networks. In 2008, he completed a graduate textbook entitled "Security and Cooperation in Wireless Networks", with Levente Buttyan. Some of his current research activities are funded by Nokia.

He is co-founder and chairman of the steering committee of WiSec (the ACM Conference for Wireless Network Security). He has served on the program committees of numerous conferences and workshops, including SIGCOMM, INFOCOM, MobiCom, MobiHoc, SenSys, WiSe, WiSec and VANET. He is one of the seven commissioners of the Federal Communications Commission (ComCom), the "Swiss FCC".

He was born in Belgium, but spent most of his childhood and youth in Northern Italy. After completing his studies in Electrical Engineering at Politecnico di Milano, he worked 10 years in France with Alcatel, primarily in the area of switching systems architecture and software. He held visiting positions at the IBM T.J. Watson Research Center and at UC Berkeley.



Jean-Yves Le Boudec is full professor at EPFL and fellow of the IEEE. He graduated from Ecole Normale Supérieure de Saint-Cloud, Paris, where he obtained the Agregation in Mathematics in 1980 (rank 4) and earned his doctorate in 1984 from the University of Rennes, France. From 1984 to 1987 he was with INSA/IRISA, Rennes. In 1987 he joined Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department. In 1988, he joined the IBM Zurich Research Laboratory where

he was manager of the Customer Premises Network Department. In 1994 he joined EPFL as associate professor.

His interests are in the performance and architecture of communication systems. In 1984, he developed analytical models of multiprocessor, multiple bus computers. In 1990 he invented the concept called "MAC emulation" which later became the ATM forum LAN emulation project, and developed the first ATM control point based on OSPF. He also launched public domain software for the interworking of ATM and TCP/IP under Linux. He proposed in 1998 the first solution to the failure propagation that arises from common infrastructures in the Internet. He contributed to network calculus, a recent set of developments that forms a foundation to many traffic control concepts in the internet, and co-authored a book on this topic. He is also the author of the book *Performance Evaluation* (2010). He received the IEEE Millennium Medal, the Infocom 2005 Best Paper award, the CommSoc 2008 William R. Bennett Prize and the 2009 ACM Sigmetrics Best Paper award.

He is or has been on the program committee or editorial board of many conferences and journals, including Sigcomm, Sigmetrics, Infocom, Performance Evaluation and ACM/IEEE TRANSACTIONS ON NETWORKING.