

Effectiveness of Distance-Decreasing Attacks Against Impulse Radio Ranging

Manuel Flury, Marcin Poturalski,
Panos Papadimitratos, Jean-Pierre Hubaux, Jean-Yves Le Boudec
Laboratory for Computer Communications and Applications, EPFL, Switzerland
firstname.lastname@epfl.ch

ABSTRACT

We expose the vulnerability of an emerging wireless ranging technology, impulse radio ultra-wide band (IR-UWB), to distance-decreasing attacks on the physical communication layer (PHY). These attacks violate the security of secure ranging protocols that allow two wireless devices to securely estimate the distance between them, with the guarantee that the estimate is an upper-bound on the actual distance. Such protocols serve as crucial building blocks in security-sensitive applications such as location tracking, physical access control, or localization.

Prior works show the theoretical possibility of PHY attacks bypassing cryptographic mechanisms used by secure ranging protocols. They also demonstrate that for physical layers used in ISO 14443 RFID and wireless sensor networks, some PHY attacks are indeed feasible. IR-UWB was proposed as a possible solution, but we show that the *de facto* standard for IR-UWB, IEEE 802.15.4a, does not automatically provide security against such attacks. We find that with the mandatory modes of the standard an external attacker can decrease the measured distance by as much as 140 meters with a high probability (above 99%).

Categories and Subject Descriptors

C.2.0 [Computer - Communication Networks]: General—Security and Protection

General Terms

Performance, Security

1. INTRODUCTION

Secure ranging [1–13] allows a (wireless) device to estimate, in a *secure* manner, the distance from itself to another device. More specifically, the measured distance provides an *upper-bound* on the actual distance (hence the name *distance bounding* used by many authors), even if the protocol is executed in the presence of an adversary that is trying to in-

terfere with the ranging process. This makes secure ranging a crucial building block for many security-sensitive services and applications. One example is the tracking of goods and people [14]. Consider a valuable item, such as a Swiss watch, equipped with a wireless-enabled (RFID) tag. The watch is displayed in a store equipped with a monitoring system that executes secure ranging with the tag every few seconds. If anyone tries to move the watch beyond some distance from the monitoring system, an alarm is triggered. Other examples include RFID access control [15], secure neighbor discovery [16], secure time synchronization [17], and secure localization [6].

In previous work, Clulow *et al.* [18] point out physical layer (*PHY*) attacks as a possible attack vector against secure ranging. These attacks make it possible even for an *external* adversary to decrease the estimated distance *without* breaking any cryptographic primitives or protocols. In [19] Hancke and Kuhn demonstrate with a proof-of-concept implementation that some of these attacks are indeed feasible for the ISO 14443 PHY and a compliant RFID receiver, as well as for 433MHz ASK/FSK modulation and a super-heterodyne receiver, common in wireless sensor networks.

Nevertheless, PHY attacks are, by nature, PHY-specific. Therefore, the results of [18, 19] cannot be mapped easily to other PHYs – separate investigations to quantify the effect of PHY attacks are required. One PHY particularly worth such an investigation is the emerging, yet promising ranging technology: *Impulse Radio Ultra-Wideband (IR-UWB)*. The defining feature of IR-UWB is the use of nanosecond pulses, which gives it unmatched capabilities of high (sub-meter) precision indoor ranging, even in dense multi-path environments [20]. Indeed, the recently proposed standard for IR-UWB, IEEE 802.15.4a [21, 22], is the only wireless standard that has been specifically designed for ranging applications. Further, this standard even includes an optional *private ranging mode* meant to enable ranging in the presence of an adversary. Using IR-UWB is also proposed in [18] as a way to mitigate PHY attacks, because of its potential for use of very short symbols. From the application perspective, IR-UWB is envisioned to be used for tracking and access control with a new generation of RFID [14], as well as high-precision localization [20], including secure localization [23]. Undoubtedly, other security sensitive applications will emerge as the technology matures. All of this makes IR-UWB a natural candidate for the PHY of a secure ranging protocol, as is pointed out in [4]. Understanding the impact of PHY attacks on IR-UWB clearly is of utter importance.

This is precisely the problem we address in this paper.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WiSec'10, March 22–24, 2010, Hoboken, New Jersey, USA.
Copyright 2010 ACM 978-1-60558-923-7/10/03 ...\$10.00.

Our main contributions and findings are:

► **The first investigation and quantification of the security of IR-UWB ranging against PHY attacks.**

We investigate the IEEE 802.15.4a [21, 22] standard, focusing on its mandatory modes, and a non-coherent energy-detector [24, 25], which is a realistic receiver for the RFID applications we are interested in (cost- and complexity-wise). We devise a *distance-decreasing relay attack*, which an *external* adversary can use to decrease the measured distance by up to 140m.¹ By increasing the signal-to-noise-ratio (SNR), the attack success rate can be made arbitrarily large. In particular, a success rate of around 99% requires the adversary to operate only at a slightly higher SNR (a few dB) than needed for normal system operation. This is easily achievable by using a high-gain antenna, by transmitting at a power exceeding the regulatory limit, or by simply moving the adversarial devices closer to the victim devices.

► **The first study of PHY attacks that encompassed the complete process of packet reception.** PHY packets typically consist of two distinct parts: the *payload* that carries the actual information bits and a *preamble* that is used to acquire the packet and determine the packet time-of-arrival. We show that in order to mount a distance decreasing attack, it is not enough to attack the payload part of a PHY packet but that the attack also has to be extended to the preamble. To the best of our knowledge, this has not been considered in any previous works.

The remainder of the paper is structured as follows. In Section 2 we introduce the IR-UWB PHY of IEEE 802.15.4a as well as the energy-detection receiver that we use in our evaluation. In Section 3, we design physical layer attacks tailored to this PHY and receiver. In Section 4 we evaluate the effectiveness of the proposed attacks with detailed physical layer simulations (due to very limited availability of IR-UWB hardware). In Section 5 we discuss possible countermeasures, elaborate on some of the assumptions, and explain why the private ranging mode of the standard is not resilient against the proposed attack. Finally, we conclude in Section 6.

2. SYSTEM MODEL

We assume that devices engaging in a secure ranging protocol share the necessary cryptographic material, and that they are equipped with an IEEE 802.15.4a compliant receiver and transmitter. The choice of transmitter is of little consequence to our investigation, any standard-compliant transmitter is acceptable [26, 27]. The architecture of the receiver used by honest devices is described in Section 2.3.

The IEEE 802.15.4a standard contains a myriad of details. For the reader's convenience, whenever applicable, we first provide information that is **essential** for understanding the paper. Then, for completeness, we provide **details** (denoted as such throughout the paper) that can be skipped

¹Revisiting the tracking example, we can clearly see the threat this attack poses: An adversary can use it to convince the store monitoring system that the watch is within the store premises, whereas in reality it is in the adversary's pocket more than one block away. This attack is also quite easy to mount, as the adversary does not need to detach the tag from the item, or to corrupt the tag, or to break any cryptographic primitives or protocols. Further, jamming or destroying the tag are not an option, as the monitoring system will set off an alarm if it cannot hear from the tag.

mode	T_{psym}	T_{sync}	T_{sfd}	T_{pre}	T_{sym}
LPRF	3968ns	254 μ s	31.8 μ s	285.8 μ s	1024ns
HPRF	992ns	63.5 μ s	7.9 μ s	71.4 μ s	1024ns

Table 1: IEEE 802.15.4a mandatory modes.




Figure 1: IEEE 802.15.4a preamble structure.

by a casual reader; they are necessary only to understand some subtle points in the attack design and performance. Note that we omit some less relevant information altogether (in particular, features of the standard designed with only coherent receivers in mind).

2.1 IEEE 802.15.4a

IEEE 802.15.4 [21] is a standard for low-rate wireless personal area networks (WPAN). The 802.15.4a amendment [22] defines an IR-UWB PHY allowing for low-rate communication and high precision ranging, using a number of 500MHz or 1.5GHz bandwidth channels from approximately 3GHz to 10GHz. Because of the ultrawide-band nature of this PHY, the transmitting power is significantly limited by regulation. This results in a relatively low communication range (20-30m). The standard does not specify a particular transmitter or receiver implementation, nevertheless it provides recommendations on the suitability of various modes for different receivers and channels.

Mandatory modes. The IEEE 802.15.4a standard is very flexible, allowing for many combinations of parameter values. However, only two of these combinations need to be implemented by a standard compliant device: one LPRF and one HPRF mode (*high/low pulse repetition frequency*). We argue that the majority of devices are likely to implement only these mandatory modes and we therefore focus on these. The parameters of the mandatory modes most important for our investigation are summarized in Table 1.

2.1.1 Preamble

Essentials. An IEEE 802.15.4a packet is composed of a preamble followed by a payload part. The preamble (Fig. 1), of duration T_{pre} , is further composed of a SYNC part of duration T_{sync} and the start frame delimiter (SFD) of duration T_{sfd} . The SYNC part is used for packet detection, timing acquisition and channel estimation; the SFD marks the start of the payload. The SYNC part consists of $N_{\text{psym}} = 64$ identical preamble symbols S_i . The SFD is a sequence of $N_{\text{sfd}} = 8$ preamble symbols $[0, S_i, 0, -S_i, S_i, 0, 0, -S_i]$. The duration of a preamble symbol is $T_{\text{psym}} = 3968\text{ns}$ (LPRF) or 992ns (HPRF).

Details. A preamble symbol S_i is built from a ternary preamble code $C_i = \{-1, 0, 1\}^{N_{\text{pcode}}}$, of length $N_{\text{pcode}} = 31$. More precisely, S_i is composed of N_{pcode} code symbols, each consisting of $L = 64$ chips in the case of LPRF




Figure 2: BPPM modulation from the receiver's perspective: The integrator output samples are shown and the channel delay spread T_{spread} is highlighted. The channel is NLOS, as indicated by the strongest path component being delayed with respect to the symbol start.

or $L = 16$ chips in the case of HPRF: The first chip being C_i , and the remaining $L - 1$ being 0. Index i indicates which of the predefined codes to use. The duration of S_i is $T_{\text{psym}} = N_{\text{pcode}}LT_c$, where T_c is the duration of a chip, fixed to $T_c = 2.0032\text{ns}$; in the rest of the paper, for convenience, we assume that $T_c = 2\text{ns}$. This, the duration of a code symbol equals $LT_c = 128\text{ns}$ (LPRF) and $LT_c = 32\text{ns}$ (HPRF). Assuming $p(t)$ is the pulse shape of a single pulse, and $d_j = [0, 1, 0, -1, 1, 0, 0, -1]$ is the sequence defining the SFD, the signal transmitted in the preamble is:

$$s(t) = \sum_{j=1}^{N_{\text{psym}}} \sum_{k=1}^{N_{\text{pcode}}} C_i^{(k)} p(t - kLT_c - jT_{\text{psym}}) + \sum_{j=1}^8 \sum_{k=1}^{N_{\text{pcode}}} d_j C_i^{(k)} p(t - kLT_c - jT_{\text{psym}} - T_{\text{sync}}) \quad (1)$$

2.1.2 Payload

Essentials. The payload is composed of symbols, each carrying one bit of information. The modulation format is *Binary Pulse Position Modulation (BPPM)*: for a 0 bit, a pulse is sent in the first part of the symbol (the *0-block*); for a 1 bit, a pulse is sent in the second part (the *1-block*), see Fig. 2. The duration of a payload symbol is $T_{\text{sym}} = 1024\text{ns}$.

Details. More specifically, a symbol (Fig. 3) consists of $N_c = 512$ chips and has a duration of $T_{\text{sym}} = N_c T_c = 1024\text{ns}$. It is divided into 4 blocks of equal duration: the 0-block, the 1-block and two guard blocks. The guard blocks serve as a means to prevent inter-symbol interference (ISI). For a 0 bit, a burst of $N_{\text{cpb}} = 4$ (LPRF) or 16 (HPRF) pulses is sent in the 0-block; for a 1 bit, the burst is sent in the 1-block. In addition, the burst is shifted by a time-hopping offset $\delta_n < N_c/4$, according to a publicly known time-hopping sequence. Finally, the signal is scrambled according to a publicly known scrambling sequence $\beta_{n,j} \in \{-1, 1\}$. Thus, the signal transmitted for the n th bit $b_n \in \{0, 1\}$ is:

$$s(t) = \sum_{j=1}^{N_{\text{cpb}}} \beta_{n,j} p(t - jT_c - \delta_n T_c - b_n T_{\text{sym}}/2) \quad (2)$$

The maximum allowable packet size is 1016 data bits. Before modulation, the data is encoded using a (55,63) Reed-Solomon (RS) code, yielding a maximum number of 1208 coded bits per packet. The raw data bitrate is 0.85Mb/s .

2.1.3 Private Ranging Mode

The standard includes a private ranging mode whose main purpose is to prevent an adversary from learning sensitive ranging information. To this end the private ranging mode allows for the encryption of timestamp information that is




Figure 3: IEEE 802.15.4a payload symbol structure.




Figure 4: Energy-detection receiver architecture.

exchanged during the ranging process. More importantly, the preamble codes used in the ranging packets are secretly agreed on by the legitimate participants. However, this provides a very minor increase in security, because the nodes are only allowed to choose from a set of 8 predefined preamble codes. We explain this in detail in Section 5.

2.2 Multipath Channel

Essentials. We adopt the multiple propagation path channel model, which is commonly used for UWB. Under this model, a number of time-delayed, phase-shifted and attenuated copies of the transmitted signal arrive at the receiver. More specifically, we make use of two channel models appended to the IEEE 802.15.4a standard [28]: the residential non-line-of-sight (NLOS) model, and the office line-of-sight (LOS) model. For our purpose, the main characteristic of interest is the channel delay spread T_{spread} – the duration over which the channel distributes most of the energy (Fig. 2). The channel delay spreads are $T_{\text{spread}}^{\text{NLOS}} \approx 60\text{ns}$ and $T_{\text{spread}}^{\text{LOS}} \approx 30\text{ns}$, respectively.

Details. The transmitted signal $s(t)$, defined by (1) (for the preamble) or (2) (for the payload) is transformed by the channel into the received signal:

$$r(t) = s(t) * h(t - \nu) + n(t) \quad (3)$$

$n(t)$ accounts for thermal noise assumed to be additive white Gaussian (AWGN), ν is the unknown propagation delay (corresponding to line-of-sight propagation), $*$ denotes convolution, and $h(t)$ captures the response of the multipath channel and is defined according to the tapped-delay-line model as:

$$h(t) = \sum_{l=1}^L \alpha_l \delta(t - \tau_l) \quad (4)$$

where L is the number of propagation paths, α_l is the attenuation coefficient and τ_l is the delay induced by the l th propagation path, and δ is the Dirac delta function.

2.3 Baseline Receiver

Essentials. We consider a non-coherent energy-detection receiver composed of an antenna, a bandpass filter of bandwidth B , followed by a squaring device and an integrator (Fig. 4). The integrator outputs a discrete time sample every $T = T_c = 2\text{ns}$.

Details. The discrete samples at the integrator output of the receiver are denoted by $y_{m,n}$. During the reception of the preamble, $y_{m,n}$ denotes the m th sample of the n th code symbol of the preamble. During reception of the pay-

load, $y_{m,n}$ denotes the m th sample of the n th symbol of the payload (Fig. 2). These samples are given by:

$$y_{m,n} = \int_{mT+\alpha(n)}^{(m+1)T+\alpha(n)} [r(t)]^2 dt, \quad (5)$$

where $\alpha(n) = nLT_c + \nu$ during the preamble, $\alpha(n) = nT_{\text{sym}} + \delta_n T_c + \nu$ during the payload, and ν is the propagation delay.

For noise-only parts of $r(t)$, the samples $y_{m,n}$ are distributed independently and according to a chi-square distribution with $2BT$ degrees of freedom. If $r(t)$ contains contributions from the transmitted signal, then $y_{m,n}$ is distributed according to a non-central chi-square distribution with $2BT$ degrees of freedom and a non-centrality parameter $p_{m,n}$ that depends on the channel response [24, 25].

Motivation. We deem this receiver to be a realistic solution for RFID in terms of complexity, cost and performance. We chose a non-coherent energy-detection receiver [24, 25, 29] over a complex coherent rake receiver because of its reduced complexity. [The large number of resolvable multipath components of the UWB channel impose fine channel estimation and stringent timing requirements [30] on coherent rake receivers, which makes them hard to implement on the small and cheap active tags that we consider.] Nevertheless, the considered receiver has a sufficiently high sampling rate to allow for precise ranging. It has also been shown to have optimal performance in this class of receivers [24] and to be quite robust to multi-user interference [25].

Alternative Receivers. We also investigate resilience of sub-optimal energy-detector receivers: performing SFD detection based on a correlation with a template derived from the SFD sequence (similar to the timing acquisition method explained in Section 2.3.1), and demodulating without weighting by the channel energy-delay profile (see Section 2.3.2). They are also vulnerable to the attack, as we briefly report in Section 4.

2.3.1 Synchronization

Essentials. Synchronization allows the receiver to detect the presence of a packet on the wireless channel and to determine the beginning of the packet. This is necessary both for decoding the payload and for establishing the packet time-of-arrival, based on which the measured distance is estimated. Infrastructure-less packet based wireless networks typically lack global synchronization. Thus synchronization has to be performed on a packet-per-packet basis. In IEEE 802.15.4a, this is done with the help of the known preamble sequence (Section 2.1.1).

To detect the presence of a packet on the wireless channel, the receiver employs a process called *timing acquisition*. This allows the receiver to discover where the boundaries of the S_i symbols fall and thus learn the time-of-arrival *modulo* T_{psym} . Due to random factors, such as the channel impulse response and the noise, some S_i symbols might be corrupted. Hence, based on timing acquisition alone, the receiver cannot be certain about the number of S_i symbols needed to acquire timing and thus the exact time-of-arrival.

After the timing acquisition phase, the receiver can perform *channel estimation*: Based on a small number of S_i symbols, it estimates the energy-delay profile of the channel. This profile allows for the decoding performance, as well as ranging accuracy, to significantly improve.

After the channel estimation phase, the receiver goes into

the *SFD detection* mode, and begins to look for the special signal sequence given by the SFD. The SFD marks the end of the preamble and the beginning of the payload. It thus allows for the time-of-arrival to be determined exactly, eliminating the uncertainty remaining after timing acquisition. Once the SFD is found, the receiver starts demodulating and decoding the payload. The exact method the receiver uses to detect the SFD is described below: It is essential that the receiver utilizes the full length of the SFD sequence, for maximum performance.

Details. During *timing acquisition*, the receiver correlates the received signal with a template derived from the known preamble code sequence. The presence of a preamble on the channel is detected if the correlation exceeds a certain threshold. In practice, this process usually takes about one third of the N_{psym} preamble symbols. The algorithm used in the receiver is the baseline algorithm from [31]; for further details the interested reader is referred there. *Channel estimation* involves estimating the energy-delay profile of the channel (from which the non-centrality parameters $p_{m,n}$ can be derived [25]). The energy-delay profile of the channel is estimated by averaging the received signal over a small number of S_i symbols.

During *SFD detection*, the receiver looks at a sliding window containing the last $N_{\text{sfd}} = 8$ received preamble symbols. The length of this window thus corresponds to the length of the SFD sequence. Given the preamble structure of IEEE 802.15.4a, only 9 different sequences of 8 preamble symbols can possibly be observed². The first sequence ($S_i, S_i, S_i, S_i, S_i, S_i, S_i, S_i$) occurs if the window is fully inside the SYNC part of the preamble. The second sequence ($S_i, S_i, S_i, S_i, S_i, S_i, S_i, 0$) is observed if the window contains 7 preamble symbols from the SYNC part and the first symbol of the SFD. The third sequence ($S_i, S_i, S_i, S_i, S_i, 0, S_i$) is observed if the window contains 6 preamble symbols from the SYNC part and the first two symbols from the SFD. In total we thus have 9 possible observations, the last one corresponding to the SFD sequence ($0, S_i, 0, S_i, S_i, 0, 0, S_i$). The receiver employs the maximum likelihood criterion (based on the channel estimation phase) in order to determine which of the 9 possible sequences is the most likely for the observation currently within the sliding window. If the most likely sequence corresponds to the SFD sequence, detection of the SFD is declared and reception of the payload starts.

2.3.2 Payload Demodulation

Essentials. Payload symbol demodulation is performed by weighting the received signal with the estimated channel energy-delay profile and then by comparing the energies in the 0-block and in the 1-block of the symbol.

Details. More precisely, demodulation of the n -th payload bit, b_n , is done according to the decision rule:

$$\sum_{m=0}^{M-1} y_{m,n} \cdot p_{m,n} \underset{b_n=1}{\overset{b_n=0}{\gtrless}} \sum_{m=0}^{M-1} y_{m+\frac{N_c}{2},n} \cdot p_{m,n} \quad (6)$$

where the parameter M defines the number of chips that the receiver takes into consideration to account for the channel delay spread. Such a decision rule corresponds to a linear approximation of the maximum likelihood ratio test [24, 25].

²Note that with a non-coherent receiver, the negative symbols are flipped to positive due to the squaring operation.




Figure 5: Example secure ranging protocol. Device A estimates the distance to device B with the formula $d_{AB} = c(t_{RTT} - t_{TA})/2$, where c is the channel propagation speed. MAC_{AB} stands for Message Authentication Code with a symmetric key shared between A and B, N_A and N_B are freshly generated nonces, t_{TA} is a constant turn-around time that A and B know, and t_{RTT} is the round-trip-time measured by A.

3. DISTANCE-DECREASING ATTACK

Secure ranging (distance bounding) protocols [1–13] are, in general, developed without a particular physical layer in mind. Essentially, they are cryptographic protocols built on top of “traditional” ranging that measures the distance between two devices based on message time-of-flight (Fig. 5). Secure ranging protocols augment ranging messages with bit-strings cryptographically bound to a secret shared by the devices. (The binding is typically sealed with additional messages.) This prevents an external adversary from injecting bogus messages (spoofing) and thus interfering with the ranging process.

Many distance bounding protocols go one step further, and attempt to thwart internal attacks: a misbehaving ranged device (traditionally called a *prover*) that convinces the ranging device (*verifier*) that it is closer than it actually is. Such protocols include, for the most part³, a *rapid-bit-exchange* phase (RBE): The verifier sends a number of single bit challenges, to which the prover must respond instantly. Such unusual requirements make these protocols difficult to implement.⁴ In particular, an IEEE 802.15.4a implementation of RBE would not only be extremely inefficient, as every bit would have to be prefixed with a (relatively long) preamble – it would also be open to packet-level attacks considered in [18].

In contrast, protocols that attempt only to prevent external attacks [6, 7] rely on a small number (typically 2) of ranging messages that are several bits long. An example of such a protocol is the protocol in Fig. 5. Such protocols can be easily implemented on IEEE 802.15.4a compliant devices. Furthermore, security against external attacks is sufficient in many applications (e.g., the tracking scenario from the Introduction). Finally, an external PHY attack is more challenging to mount than an internal one: The *early detection* (ED) and *late commit* (LC) components of the *distance-decreasing relay attack* [18], which we devise in this section, can be used individually by a malicious prover to

³A recent proposal [10] shows that the RBE can be replaced by a full-duplex transmission in protocols that provide security against internal attacks

⁴To this date, no implementation of RBE for wireless networks exists.




Figure 6: Distance-decreasing relay attack setup.

mount an internal attack. For all these reasons, we focus on external PHY attacks and on secure ranging protocols that employ several-bits-long ranging messages unpredictable by the adversary (like the protocol in Fig. 5).

3.1 Threat Model

We consider an external adversary mounting a *distance-decreasing relay attack* between two honest devices, A and B, that execute a secure ranging protocol (Fig. 6). The adversary uses two devices M_A and M_B , where A can communicate directly only with M_A and B can communicate directly only with M_B .⁵ Furthermore, M_A and M_B exchange information using an out-of-band adversarial channel. The IEEE 802.15.4a channel propagation speed is c , the speed of light. The same speed is assumed for the adversarial channel.

The adversarial devices are equipped with transmitters similar to the honest devices, but able to send non-standard-compliant pulse sequences, and to ignore regulatory transmission limits. Their receivers extend the baseline architecture (Section 2.3) and can be equipped with high gain antennas. Hence, the adversary is able to increase the SNR observed by both adversarial and honest devices. Note that such an increase in received SNR can also be achieved by the adversary moving its devices closer to the honest devices.

We focus on the exchange of a single ranging message, during which one of the honest devices acts as a transmitter (HTX) and the other one as a receiver (HRX), as depicted in Fig. 6. Accordingly, the adversarial devices act as a receiver (ARX), and as a transmitter (ATX). It is easy to extend this attack to an entire secure ranging protocol. The adversary simply mounts the distance-decreasing relay attack on all ranging messages. Any non-ranging messages of the protocol, which are not time critical, can be relayed in an arbitrary fashion.

3.2 Attack Principle

For the relay attack to be distance-decreasing, the adversary needs to “shift” the relayed message back in time by some offset $t_{\text{relay}} > 0$, which we call the *relay time-gain* (Fig. 7). The amount by which the measured distance is decreased is $c \cdot t_{\text{relay}}$, assuming an optimal (for the adversary) configuration where ARX and ATX lie on a line between HTX and HRX.⁶ This distance decrease is reduced by any additional processing delays the adversary introduces (Section 3.5).

The difficulty in mounting this attack is twofold: (1) ATX needs to begin the transmission of the preamble before it learns from ARX **when** HTX started the transmission. (2) ATX needs to transmit the payload before it learns **what** bits the payload carries. Existing work has focused exclusively on the second problem [18, 19], but the first one is

⁵We elaborate on this assumption in Section 5.

⁶In other configurations the distance decrease will be smaller. Note that the choice of the configuration rests with the adversary.




Figure 7: Overview of the distance-decreasing relay attack. ARX and ATX are assumed to lie on a line between HRX and HTX. The thick dotted arrow indicates time-of-arrival corresponding to the actual distance between HTX and HRX.

equally important: Without shifting the time-of-arrival at HRX, attacks on the payload are in vain.

The first problem might seem simpler, because the adversary knows the content of the preamble. However, in general, the adversary cannot know when HTX will begin the transmission. Although ATX can always choose to begin the transmission at a random time, there is a good chance it either does this too late, resulting in a distance increase, or too early, such that the distance-decreasing attack on the payload fails (see paragraph below). Even if neither is the case, the achieved time-gain is random, which might lead to undesired results, e.g., negative distance estimates.

Naturally, the preamble and the payload must be relayed with the same time-gain. This implies that the upper-bound on the achievable time-gain is the minimum of (1) the upper-bound on the time-gain for the preamble and (2) the upper-bound of the time-gain for the payload. As we will see shortly, the payload upper-bound is more strict and determines the achievable relay-gain.

3.3 Attack on the Preamble

The attack on the preamble is depicted in Fig. 8. For clarity of presentation, we assume the distance between ARX and ATX to be 0. When HTX sends a packet, ARX performs the synchronization procedure, exactly as the baseline receiver. Upon timing acquisition, ARX learns (modulo T_{psym}) the time at which it began receiving the packet, which we denote by t_0 . Deviating from the baseline receiver, ARX performs *early SFD detection*: ARX chooses an early SFD detection delay $t_{\text{ED}}^{\text{SFD}}$ and acts on the first $t_{\text{ED}}^{\text{SFD}}/T_c$ samples of every received preamble symbol. [Details: ARX performs a maximum likelihood criterion-based test, with two hypotheses: an S_i symbol plus noise, indicating that the SFD has not started yet, versus only noise, indicating that the SFD is starting. ARX completes early SFD detection when the latter is more likely.] Note that early SFD detection is necessary, as ARX cannot know how many received S_i symbols were necessary for timing acquisition (see Section 2.3.1).

In the mean time, ATX chooses a late SFD commit delay $t_{\text{LC}}^{\text{SFD}}$ and remains silent until ARX signals that timing acquisition was successful, thus providing t_0 modulo T_{psym} . Then, after an appropriately chosen (we explain how shortly) delay $\tau < T_{\text{psym}}$, ATX begins transmitting a sequence of preamble




Figure 8: Distance-decreasing relay attack on the preamble.

symbols S_i . This is repeated until ARX signals that the SFD was detected, providing the exact value of t_0 . Immediately afterwards, ATX switches to the transmitting of a standard compliant SFD, beginning from $t_{\text{LC}}^{\text{SFD}}$ into the SFD.

In contrast to a standard-compliant preamble (Fig. 1), the preamble generated by ATX (Fig. 8) begins with a number of 0 preamble symbols, and the SFD begins with a $t_{\text{LC}}^{\text{SFD}}$ -long part of an S_i symbol, instead of the standard-compliant 0. As we show in Section 4, this difference is small enough for HRX to synchronize correctly.

The relay time-gain achieved by this attack is $t_{\text{relay}} = t_{\text{LC}}^{\text{SFD}} - t_{\text{ED}}^{\text{SFD}}$. This determines the choice of τ , as $T_{\text{psym}} - \tau = (t_{\text{LC}}^{\text{SFD}} - t_{\text{ED}}^{\text{SFD}}) \bmod T_{\text{psym}}$. Note that nothing prevents the adversary from choosing $t_{\text{LC}}^{\text{SFD}} > T_{\text{psym}}$. Rather, the hard upper-bound on the time-gain achievable by the preamble attack is T_{sfd} .

3.4 Attack on the Payload

The attack on the payload is performed on a symbol by symbol basis. ARX performs an *early detection* attack, first choosing an early detection delay $t_{\text{ED}} \ll T_{\text{sym}}$. It acts only on the first t_{ED} nanoseconds of the symbol (Fig. 9), effectively performing on-off keying demodulation (OOK) in place of BPPM demodulation, ignoring deliberately most of the BPPM symbol. [Details: ARX employs a maximum likelihood test with the hypotheses: signal plus noise, in which case the symbol is demodulated to 0, versus only noise, in which case the symbol is demodulated to 1.] The time t_{ED} can be made arbitrarily small, it determines the attack's performance. The value optimal for performance is dictated by the channel delay spread T_{spread} , as we show in Section 4. After demodulation, ARX signals the result to ATX.

In the mean time, ATX performs a *late commit* attack. ATX begins the transmission of a symbol $T_{\text{sym}}/2$ before the symbol's bit value is received from ARX. At first, ATX does not know what bit the symbol should carry, thus it always begins the symbol transmission with a pulse of energy E_0 . Once the bit value is received, ATX acts accordingly: If it is a 0, it transmits nothing in the 1-block of the symbol; if it is a 1, it transmits a pulse with energy $E_1 > E_0$. (The optimal ratio between E_0 and E_1 , $\gamma = E_0/E_1$ is determined in Section 4.) This attack exploits the fact that HRX performs a simple energy comparison to demodulate. The *late commit delay* is $t_{\text{LC}} = T_{\text{sym}}/2$. The relay time-gain due to this attack is $t_{\text{relay}} = t_{\text{LC}} - t_{\text{ED}} \leq T_{\text{sym}}/2$, considerably less than the upper-bound due to the preamble part of the attack.

Details. For clarity, above we assumed a symbol with time-hopping offset $\delta_n = 0$. However, the time-hopping sequence does not affect the time-gain of the attack. Indeed, assume that the early detection and late commit delays with




Figure 10: SER versus SNR for LPRF comparing benign performance to ED with varying ED delays t_{ED}^{SFD} .




Figure 11: SER versus SNR for LPRF comparing benign performance to LC with varying LC delays t_{LC}^{SFD} .




Figure 12: Performance loss ΔSNR with respect to benign case at fixed $\text{SER} = 10^{-2}$ versus t_{ED}^{SFD} and t_{LC}^{SFD} for LPRF and HPRF.

4.1 Attack on the Preamble

An honest receiver performing SFD detection takes the entire length T_{sfd} of the SFD into account. For LPRF this equals $T_{\text{sfd}} = 31.8 \mu\text{s}$, for HPRF $T_{\text{sfd}} = 7.95 \mu\text{s}$.

Fig. 10 shows the SER for an honest receiver, as well as for an adversary that performs early SFD detection with different early SFD detection delays t_{ED}^{SFD} . The curves shown here are for LPRF. Not surprisingly, the earlier an adversary performs SFD detection, the more additional received power with respect to an honest receiver it is going to cost him to reach a given level of SER. If we fix $\text{SER} = 10^{-2}$, detecting the SFD at $t_{ED}^{SFD} = 3.712 \mu\text{s}$ costs the adversary $\Delta\text{SNR} = 2.8\text{dB}$ in additional received power, detecting at $t_{ED}^{SFD} = 0.128 \mu\text{s}$ entails a cost of $\Delta\text{SNR} = 11.2\text{dB}$.

For t_{ED}^{SFD} , we only consider values shorter than the length of the first SFD symbol. Larger values for t_{ED}^{SFD} do not make much sense for the adversary because they also force him to commit after the first SFD symbol, which is only possible at a considerable additional cost. This can be seen in Fig. 11, which shows the SER of an adversary that commits late, at time t_{LC}^{SFD} into the SFD. Again, the results shown are for LPRF. Committing at $t_{LC}^{SFD} = 8 \cdot 128\text{ns} = 1.02 \mu\text{s}$, or earlier is within 0.6dB of the benign case and thus comes at practically no additional cost at a target SER of 10^{-2} . Committing later comes at an ever increasing cost: Committing at $t_{LC}^{SFD} = 29 \cdot 128\text{ns} = 3.712 \mu\text{s}$, already costs $\Delta\text{SNR} = 7.5\text{dB}$. [Details: According to the preamble and SFD codes, no pulse is sent between the 29th code symbol of the first SFD symbol and the first code symbol of the third SFD symbol. So committing anywhere between $t_{LC}^{SFD} = 3.712 \mu\text{s}$ and $t_{LC}^{SFD} = 63 \cdot 128\text{ns} = 8.064 \mu\text{s}$ is equivalent to committing at $t_{LC}^{SFD} = 8.064 \mu\text{s}$, which costs more than $\Delta\text{SNR} = 9\text{dB}$.]

Results for HPRF are generally close to those of LPRF shown so far. Performing ED at $t_{ED}^{SFD} = 928\text{ns}$, for example, costs the adversary about $\Delta\text{SNR} = 2.9\text{dB}$, compared to 2.8dB for LPRF. This can be seen in Fig. 12 where we show the additional cost ΔSNR with respect to an honest receiver versus t_{ED}^{SFD} and t_{LC}^{SFD} for both LPRF and HPRF and a fixed SER of 10^{-2} . The corresponding SNR values were found via interpolation of curves such as those shown in Figures 10 and 11. Results for ED are close and within 0.5dB. Late commit generally costs about 1dB more in the case of HPRF. Note the different time scales that are due

to the fact that a preamble symbol in HPRF is four times shorter.

An important observation is that none of the curves showing the performance under attack exhibits an error floor. This indicates that by increasing the SNR, the attack success rate can be made arbitrarily large. The same holds for the payload, as we will see shortly.

Alternative Receivers. We also evaluated a receiver that performs SFD detection using a sub-optimal correlation method (2dB performance loss compared to the baseline receiver). This receiver is also vulnerable to the attack, and the attack's cost in terms of ΔSNR is close to the cost in the case of the baseline receiver: more precisely, up to 1dB greater (for t_{LC}^{SFD} in the order of T_{psym}).

4.2 Attack on the Payload

We now look at the effect of ED and LC on the payload. The following results do not contain effects of synchronization: We assume here that the receiving party, ARX in the case of ED and HRX in the case of LC, is able to perfectly synchronize to each packet. Perfect synchronization here means that an oracle returns the exact packet time-of-arrival. (Hence, there are no false alarms or missed detections.) The channel energy-delay profile is still estimated; but the estimation is performed under the assumption that the packet boundaries are perfectly aligned. In the case of LC, we further assume that the packet sent by ATX does not contain any errors due to a preceding ED. For all of the results presented here, we further assume a payload of 128 bits. We consider 128 bits a conservative upper-bound on the length of a ranging message in a secure ranging protocol.

Fig. 13(a) shows the PER at different SNRs for the LPRF mode. We show the performance curves for a benign receiver and an adversary performing ED at different ED delays t_{ED} . The optimal ED delay for the adversary is in the order of the channel delay spread and found to be $t_{ED}^{\text{OPT}} = 68\text{ns}$ in the present example that uses the NLOS channel model. Deciding on the symbol at t_{ED}^{OPT} introduces a loss of about 1.7dB with respect to the benign curve at a packet error rate of $\text{PER} = 10^{-2}$. This can also be seen in Fig. 13(b). Here we show the loss in SNR, ΔSNR , with respect to the benign case versus the ED delay t_{ED} for a target packet error rate of $\text{PER} = 10^{-2}$. The curve has been obtained from curves such as those shown in Fig. 13(a) via interpolation. Detecting





Figure 13: (a): PER versus SNR for the payload comparing benign performance to ED with varying ED delays t_{ED} . The optimal t_{ED} is in the order of the channel delay spread and gives a loss of about 1.7dB. (b): More compact representation of the data in (a), showing the loss ΔSNR with respect to the benign case versus t_{ED} for a fixed PER of 10^{-2} .

after t_{ED}^{OPT} gives a slightly worse performance because the adversary then merely integrates more noise instead of useful signal. Performing ED much earlier than t_{ED}^{OPT} results in substantially larger loss because a large part of the useful signal energy is lost: Deciding at $t_{ED} = 32\text{ns}$, for example, introduces a loss of 5.6dB.

Fig. 14 shows the performance of LC on the payload in the case of LPRF. As explained in Section 3.4, the LC delay t_{LC} is fixed to $t_{LC} = T_{\text{sym}}/2 = 512\text{ns}$. We show the PER for different ratios γ of the energies E_0 and E_1 corresponding to the signal energies transmitted by the adversary during the 0-block and 1-block, respectively. E_1 here corresponds to the energy a benign receiver would transmit and E_0 is typically smaller (see also Section 3.4). A ratio of $\gamma^{\text{OPT}} = 0.35$ gives optimal performance throughout the whole operating range, thus this is the energy ratio we will use in all subsequent simulations. The optimal ratio gives a loss of about 4dB with respect to the benign case.

For HPRF, we do not show any curves because the results are very similar. With HPRF and the LOS channel, the optimal ED delay is $t_{ED}^{\text{OPT}} = 48\text{ns}$. [Details: Note that this is significantly larger than the channel delay spread. The reason is that in the HPRF mode, a burst of 16 pulses is sent during the payload, spreading the received signal wider in time.] The difference in SNR, with respect to the benign case, is 2dB versus 1.7dB with LPRF. For LC, we find the optimal energy ratio to be $\gamma^{\text{OPT}} = 0.35$ as well, and the corresponding loss of 3.9dB is close to the 4dB found for LPRF.

Uncoded Transmission. We also look at transmissions that do not use the RS code. We did so to make sure that it is not the RS code that allows the attack to succeed (by masking the errors introduced by the attack). It turns out that this is not the case: Without the RS code, the ΔSNR between the benign and attack case is practically the same as for the coded case.

Alternative Receivers. We also evaluated a simplistic receiver that demodulates without weighting with the estimated energy-delay profile (2dB performance loss compared to the weighted decision). Such a receiver is vulnerable to the attack as well, and the attack's cost in terms of ΔSNR is within 0.5dB of the cost for the baseline receiver.




Figure 14: PER for LC on the payload with varying energy ratios γ . The optimal ratio at $\gamma^{\text{OPT}} = 0.35$ gives a loss of about 4dB with respect to the benign setting.

4.3 Overall Performance of the Attack

We now establish the overall performance of the distance-decreasing relay attack. As the relay attack involves two transmissions, ARX and HRX potentially have different received SNRs, which we will denote by SNR_{ED} and SNR_{LC} . This difference can be a result of the topology, but it can also be introduced by the adversary. Depending on his abilities, an adversary can, for example, send with a higher power in order to increase SNR_{LC} , or move closer to HTX, or use a directive antenna to increase SNR_{ED} . Combined with the observation that the same relay time-gain, t_{relay} , can be obtained with different combinations of ED and LC delays, this gives the adversary room for a trade-off: Depending on the SNR values achievable for SNR_{ED} (SNR_{LC} , respectively) the adversary can choose to perform ED earlier or later (commit earlier or later, respectively). If SNR_{LC} is high with respect to SNR_{ED} , the adversary will prefer to commit late in order to be able to detect late as well. If SNR_{LC} is low with respect to SNR_{ED} , the adversary will prefer to detect early in order to be able to commit early.

In our analysis, Fig. 15 will serve as a benchmark. It shows the PER in the benign case for both LPRF and HPRF. Packet sizes of 128 and 1016 data bits are shown. A packet size of 1016 bits is the maximum packet size allowed by the standard; as stated earlier, 128 bits correspond to a conservative length of a ranging message. In LPRF the factor limiting performance is the payload. This can be seen by observing that the LPRF curve for the shorter packet size is almost identical to the benign curve in Fig. 13 (which assumes perfect synchronization). For HPRF the opposite is true, the limiting factor is the synchronization. This can be seen in Fig. 15 where the size of the packet hardly influences the PER. [Details: The reason is that in HPRF 16 times more energy is sent in a payload symbol compared to a preamble pulse, whereas in LPRF it is only 4 times more.]

Fig. 16(a) shows the probability of success for LPRF and an attack that tries to gain 480ns when relaying a 128bit packet between HTX and HRX. This relay time-gain is equivalent to a 144m distance decrease between HTX and HRX.⁸

⁸Here we assume for simplicity that the processing delays at the adversarial transceivers are zero. Processing delays are discussed in Section 3.5.




Figure 15: Reference curves showing PER in the benign case for LPRF and HPRF with different packet sizes.

The results shown are for different combinations of SNR_{ED} and SNR_{LC} . For every SNR combination, the probability of success that is reported corresponds to the tuple of $(t_{\text{ED}}^{\text{SFD}}, t_{\text{LC}}^{\text{SFD}}, t_{\text{ED}})$ ⁹ yielding best performance among all the tuples that achieve the given relay time-gain of 480ns. In the benign case we achieve a PER of approximately 10^{-2} at an SNR of around 8dB, see Fig. 15. In Fig. 16(a), a probability of success of $P_s = 0.9869$ is achieved for the pair $(\text{SNR}_{\text{ED}} = 12\text{dB}, \text{SNR}_{\text{LC}} = 14\text{dB})$. For all pairs above (12dB, 14dB) the probability of success is above 99%. With respect to an honest transmitter-receiver pair, an adversary thus needs an additional 4dB in SNR for ED and an additional 6dB for LC, in order to reduce the distance by 144m with a probability of success in the order of 99%. Attaining SNR values in this range would not pose much of a challenge to the adversary.

The corresponding HPRF results, decreasing the distance by 144m, are shown in Fig. 16(b). A probability of success of $P_s = 0.9875$ is reached at (11dB, 11dB). Compared with Fig. 15, the additional cost is 7dB for both ED and LC. Compared with LPRF, we thus see that decreasing the distance by the same amount costs a bit more in HPRF. This was to be expected for several reasons. First of all we have seen that, contrary to LPRF, the performance is not limited by the payload but by the synchronization. We can thus not hope to achieve the ED/LC performance of the payload-only attacks shown in Figures 13 and 14. [Details: Second, to obtain a given relay time-gain on the preamble is more costly for HPRF because of the closer spacing of the pulses. We have seen in Fig. 12 that detecting early at the i th code symbol or committing late at the j th code symbol costs roughly the same for both LPRF and HPRF. The distance decrease achieved corresponds to $(j - i) \cdot L \cdot T_c$ which depends on the length of a code symbol $L \cdot T_c$. At the same cost, the distance decrease achieved by HPRF is thus four times shorter than for LPRF.]

Increasing the packet length to its allowable maximum of 1016 bits decreases slightly the probability of success: To reach $P_s > 99\%$ we see virtually no cost increase for HPRF. For LPRF the cost in SNR_{ED} and SNR_{LC} increases by about 1.5dB each. A smaller distance decrease obviously comes at a lower cost: e.g., for an attack decreasing the distance by

⁹Recall that $t_{\text{LC}} = 512\text{ns}$ is fixed, thus limiting to some extent the degrees of freedom on the payload part.

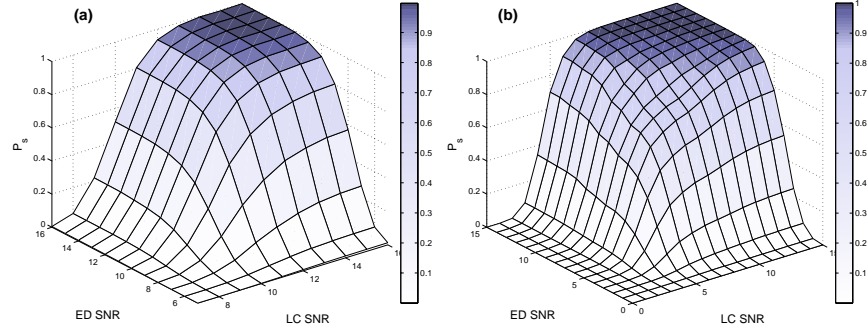


Figure 16: Probability of success, P_s , for an attack trying to achieve a distance decrease of 144m. Packet length is 128bits. (a): With LPRF, $P_s > 99\%$ is reached at a cost of $(\Delta\text{SNR}_{\text{ED}} = 4\text{dB}, \Delta\text{SNR}_{\text{LC}} = 6\text{dB})$. (b): For HPRF (7dB, 7dB) gives $P_s > 99\%$.

100m for HPRF with 128bit packets, we found the additional cost for $P_s \approx 99\%$ to be 5dB for ED and 4dB for LC. Compared to the corresponding attack achieving a decrease of 144m, this signifies a 2dB smaller cost in SNR_{ED} and a 3dB smaller cost in SNR_{LC} .

5. DISCUSSION

Private Ranging Mode. As explained in Section 2.1.3, the standard includes a private ranging mode that allows the legitimate participants to secretly agree on the preamble codes used in the ranging packets. Hence, the adversary does not know the exact structure of the preamble, which makes the attack on the preamble harder. Nevertheless, the honest devices can only choose among 8 allowable preamble codes, which offers little security. First, the adversary could simply guess the codes with a decent success probability $\frac{1}{64}$. Second, the adversary could detect a packet using, in parallel, all 8 allowable codes. This can be done entirely in the digital domain¹⁰ by correlating the received signal with each of the 8 codes and choosing the one with the highest correlation output. What additionally helps the adversary is the fact that these codes were designed to have minimum cross-correlation. In summary, the private ranging mode only moderately increases the complexity of the distance-decreasing relay attack, and cannot be considered a valid countermeasure. Furthermore, it seems the private ranging mode was designed for more complex coherent receivers only: The preamble parameters that the private ranging mode employs imply strong inter-symbol interference (ISI), which a non-coherent receiver, such as the one used in our investigation, cannot cope with well.

Possible Countermeasures. Two factors determine the quality of a countermeasure. The first is the effectiveness: the maximum distance by which the adversary can decrease the distance with the countermeasure in place. The second factor is the cost: how much the countermeasure degrades benign case performance (if no attack is taking place), compared to a system without countermeasures deployed.

The simplest countermeasure is to decrease payload symbol duration [18], as the distance-decreasing attack cannot decrease the distance by more than one symbol duration.

¹⁰This is much easier and cost-efficient than, e.g., adding circuitry to the analog part of the receiver.

This applies to the BPPM modulation: if the symbol duration is T_{sym} , the time-gain of the the attack we proposed in this paper is at most $T_{\text{sym}}/2$.¹¹ This solution can be implemented even within the IEEE 802.15.4a standard, some non-mandatory modes have symbols as short as 32ns. However, reducing T_{sym} is not without effect on the benign performance. The first problem is ISI, which manifests itself if the symbol duration is close or below the channel delay spread. Low-complexity non-coherent receivers cannot cope well with ISI and even if some solutions exist, they entail a loss of 5–10dB in the benign case [34]. Furthermore, shorter symbols have less resilience to multi-user interference.

Alternatively, the symbol duration can be preserved, but the honest receiver can choose to only take into account the beginning of the symbol [19], essentially performing early detection. This is particularly attractive in our case, as switching from BPPM demodulation to OOK demodulation significantly reduces the achievable time-gain: Indeed, $t_{\text{LC}} \geq t_{\text{relay}}$ can be reduced from 512ns to a value in the order of the channel delay spread $T_{\text{spread}} \approx 60\text{ns}$ (for optimal OOK performance). This corresponds to *at most* 20m distance-decrease (assuming an unrealistic instant ED and no processing delays). This solution does not induce any additional ISI and is compliant with the mandatory modes of the standard. However, our simulations evaluating payload early detection show that OOK decreases the benign case performance by roughly 1.5dB because half of the available information is discarded. Coding could potentially compensated for this degradation.

We also investigate another countermeasure, whose main idea is to detect the non standard signal sent by the adversary during the payload LC attack. With this countermeasure, the receiver records, for every bit, the energy in the 0-block of the symbol, and compares the distribution of these energies for the 0 bits (bits that were decoded as a 0) with the 1 bits (decoded as a 1). In the benign case, the 0-blocks of the 0 bits carry more energy (Fig. 2), whereas under attack these energies are the same (Fig. 9). To distinguish these cases, one can use a robust statistical test, such as the Mann-Whitney-Wilcoxon test. This countermeasure prevents the attack presented in Section 3.4 with virtually no degradation of benign case performance. However, an adversary can modify the attack (vary the energy levels between symbols) to severely degrade the performance of this countermeasure.

Finally, another direction might be to use secret preamble codes, known only to the communicating honest nodes. This could make (early) preamble detection infeasible, at least within the constrained time budget available to the adversary to mount the relay attack. It is uncertain, but worth investigating, how such random codes without nice auto-correlation properties would affect the benign case performance. Alternatively, secret time-hopping sequences could be used to make early detection of payload symbols more difficult. This would also require further investigation.

Isolating HRX from HTX. In our threat model, we assume that the honest receiver, HRX, cannot receive sig-

nals sent by the honest transmitter, HTX. This is inherent in some scenarios, e.g., picking virtual pockets [15]. In other scenarios, such as the watch-tracking example given in the Introduction, the adversary has to make an extra effort to prevent the tag from communicating with the monitoring system. A simple way of achieving this is to place the stolen item in a Faraday cage with one adversarial transmitter connected via a wired link to the second transmitter placed outside the Faraday cage.

Although constructing such a Faraday cage is simple (even today, shoplifters use “booster bags” coated with aluminium foil [35]), an interesting question is whether shielding HRX from HTX is really necessary? The answer is yes, as long as ARX, the adversarial receiver, takes the same or more time to acquire the packet as HRX. In this case, by the time ATX starts transmitting the preamble (Fig. 8), HRX is already synchronized to HTX’s signal. Even a much stronger signal from ATX cannot undo this – the attack fails. However, if ARX acquires the packet first, ATX can start sending the preamble while HRX is still in the process of acquiring HTX’s signal. The attack would succeed, provided that the signal level of ATX exceeds the one of HTX by a large enough margin. The performance of this attack depends on the details of the synchronization algorithm. We plan to investigate this as part of future work.

6. CONCLUSION AND FUTURE WORK

We have investigated the vulnerability of the IR-UWB standard, IEEE 802.15.4a, to physical layer distance-decreasing relay attacks. We demonstrated, for the mandatory modes of the standard, that an attack decreasing the measured distance by 140m achieves an impressive success rate of 99% at a cost of just a few dB in SNR with respect to normal system operation; a further increase in SNR allows the adversary to make the success rate arbitrarily large. In terms of future work, we plan to further investigate possible countermeasures and alternative PHY attacks.

7. ACKNOWLEDGMENTS

We would like to thank Yannick Do and Florence Le Goff for their extensive help with the simulations underlying Section 4.

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.

8. REFERENCES

- [1] S. Brands and D. Chaum. Distance-bounding protocols. In *EUROCRYPT*, 1993.
- [2] S. Čapkun, L. Buttyán, and J.-P. Hubaux. SECTOR: secure tracking of node encounters in multi-hop wireless networks. In *SASN*, 2003.
- [3] L. Bussard. *Trust establishment protocols for communicating devices*. PhD thesis, 2004.
- [4] G. Hancke and M. Kuhn. An RFID distance bounding protocol. In *SecureComm*, 2005.
- [5] J. Reid, J. M. Gonzalez Nieto, T. Tang, and B. Senadji. Detecting relay attacks with timing-based protocols. In *ASIACCS*, 2007.

¹¹We note that the attack proposed in this paper can be further improved (in terms of the achieved distance-decrease), by employing late commit techniques in the fashion of [19]. However, the additional challenge is the weighting by the channel mask performed by the baseline receiver. We leave the investigation of such attacks for future work.

- [6] S. Čapkun and J.P. Hubaux. Secure positioning in wireless networks. *IEEE J. Sel. Areas Commun.*, 24(2), 2006.
- [7] C. Meadows, R. Poovendran, D. Pavlovic, L.-W. Chang, and P. Syverson. 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.
- [8] D. Singelée and B. Preneel. Distance bounding in noisy environments. In *ESAS*, 2007.
- [9] Y.-J. Tu and S. Piramuthu. Rfid distance bounding protocols. In *First International EURASIP Workshop on RFID Technology*, 2007.
- [10] K. B. Rasmussen and S. Čapkun. Location privacy of distance bounding protocols. In *CCS*, 2008.
- [11] Jorge Munilla and Alberto Peinado. Distance bounding protocols for RFID enhanced by using void-challenges and analysis in noisy channels. *Wireless Communications and Mobile Computing*, 8(9), 2008.
- [12] C. H. Kim, G. Avoine, F. Koeune, F.-X. Standaert, and O. Pereira. The swiss-knife RFID distance bounding protocol. In P.J. Lee and J.H. Cheon, editors, *ICISC*, 2008.
- [13] C. H. Kim and G. Avoine. RFID distance bounding protocol with mixed challenges to prevent relay attacks. In *CANS*, 2009.
- [14] R. Mulloy. Ultrawide-band RFID technology. FCC Radio Frequency Identification Workshop, 2004.
- [15] Z. Kfir and A. Wool. Picking virtual pockets using relay attacks on contactless smartcard. In *SecureComm*, 2005.
- [16] P. Papadimitratos, M. Poturalski, P. Schaller, P. Lafourcade, D. Basin, S. Čapkun, and J.-P. Hubaux. Secure neighborhood discovery: A fundamental element for mobile ad hoc networking. *IEEE Communications Magazine*, Vol.46, No.2, 2008.
- [17] K. B. Rasmussen, S. Čapkun, and M. Gaalaj. SecNav: secure broadcast localization and time synchronization in wireless networks. In *MobiCom*. ACM, 2007.
- [18] J. Chulow, G. P. Hancke, M. G. Kuhn, and T. Moore. So near and yet so far: Distance-bounding attacks in wireless networks. In *ESAS*, 2006.
- [19] G. P. Hancke and M. G. Kuhn. Attacks on time-of-flight distance bounding channels. In *WiSec*, 2008.
- [20] S. Gezici, Z. Tian, G.B. Giannakis, H. Kobayashi, A.F. Molisch, H.V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *Signal Processing Magazine, IEEE*, 22(4), 2005.
- [21] *IEEE Std 802.15.4-2006 (Rev. of IEEE Std 802.15.4-2003)*, 2006.
- [22] *IEEE Std 802.15.4a-2007 (Amd. to IEEE Std 802.15.4-2006)*.
- [23] Y. Zhang, W. Liu, Y. Fang, and D. Wu. Secure localization and authentication in ultra-wideband sensor networks. *IEEE J. Sel. Areas Commun.*, 24(4), 2006.
- [24] A. A. D'Amico, U. Mengali, and E. Arias-De-Reyna. Energy-detection UWB receivers with multiple energy measurements. *IEEE Trans. Wireless Commun.*, 6(7), 2007.
- [25] M. Flury, R. Merz, and J.-Y. Le Boudec. An energy detection receiver robust to multi-user interference for IEEE 802.15.4a networks. In *ICUWB*, 2008.
- [26] C. Duan, P. Orlik, Z. Sahinoglu, and A. F. Molisch. A non-coherent 802.15.4a uwb impulse radio. In *ICUWB*, 2007.
- [27] J. Ryckaert, G. Van der Plas, V. De Heyn, C. Desset, G. Vanwijnsberghe, B. Van Poucke, and J. Craninckx. A 0.65-to-1.4nJ/burst 3-to-10GHz UWB digital TX in 90nm CMOS for IEEE 802.15.4a. *ISSCC*, 2007.
- [28] IEEE 802.15.4a channel model - final report, 2004.
- [29] M. Weisenhorn and W. Hirt. Robust noncoherent receiver exploiting UWB channel properties. In *Joint UWBST & IWUWBS*, 2004.
- [30] W.M. Lovelace and J.K. Townsend. The effects of timing jitter on the performance of impulse radio. In *UWBST*, 2002.
- [31] M. Flury, R. Merz, and J.-Y. Le Boudec. Robust non-coherent timing acquisition in IEEE 802.15.4a IR-UWB networks. In *PIMRC*, 2009.
- [32] S. Paquelet, L.-M. Aubert, and B. Uguen. An impulse radio asynchronous transceiver for high data rates. *IWUWB*, 2004.
- [33] P. F. Driessen and L. J. Greenstein. Modulation techniques for high-speed wireless indoor systems using narrowbeam antennas. *IEEE Trans. Commun.*, 1995.
- [34] F. Trösch and A. Wittneben. MLSE post-detection for ISI mitigation and synchronization in UWB low complexity receivers. In *VTG*, 2007.
- [35] <http://www.realtechnews.com/posts/1366>.