



Time for Change: How Clocks Break UWB Secure Ranging

Claudio Anliker, Giovanni Camurati, and Srdjan Čapkun, *ETH Zurich*

<https://www.usenix.org/conference/usenixsecurity23/presentation/anliker>

**This paper is included in the Proceedings of the
32nd USENIX Security Symposium.**

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

**Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.**

Time for Change: How Clocks Break UWB Secure Ranging

Claudio Anliker
ETH Zurich

Giovanni Camurati
ETH Zurich

Srdjan Čapkun
ETH Zurich

Abstract

Due to its suitability for wireless ranging, Ultra-Wide Band (UWB) has gained traction over the past years. UWB chips have been integrated into consumer electronics and considered for security-relevant use cases, such as access control or contactless payments. However, several publications in the recent past have shown that it is difficult to protect the integrity of distance measurements on the physical layer. In this paper, we identify transceiver clock imperfections as a new, important parameter that has been widely ignored so far. We present *Mix-Down* and *Stretch-and-Advance*, two novel attacks against the current (IEEE 802.15.4z) and the upcoming (IEEE 802.15.4ab) UWB standard, respectively. We demonstrate *Mix-Down* on commercial chips and achieve distance reductions from 10 m to 0 m. For the *Stretch-and-Advance* attack, we show analytically that the current proposal of IEEE 802.15.4ab allows reductions of over 90 m. To prevent the attack, we propose and analyze an effective countermeasure.

1 Introduction

UWB has experienced a revival over the last years as one of the leading technologies for short-distance wireless ranging. Because of its wide bands, it can achieve sub-decimeter accuracy. Furthermore, UWB ranging computes the distance between two devices based on the Time of Flight (ToF) of exchanged messages. This approach is inherently secure against *simple* relay attacks, which were conducted successfully against Received Signal Strength Indication (RSSI)-based ranging schemes, like those deployed in Passive Keyless Entry and Start (PKES) systems [15].

Due to these attacks, the standardization of UWB ranging progressed at a high pace, and several car vendors have already integrated UWB chips into their vehicles and keys [8, 10, 22, 24]. Companies like Apple, Samsung, and Google have released smartphones, tags, and other gadgets with embedded UWB transceivers [2, 5, 6, 40, 41]. In addition, BMW and Apple announced a collaboration enabling

customers to unlock their cars using an UWB-equipped iPhone, indicating that UWB will be used in this domain in the future [8]. However, protecting the integrity of distance measurements on the physical layer has proven to be difficult: attacks such as Cicada [34, 35], Early-Detect/Late-Commit [33], and, most recently, Ghost Peak [27] show that distance reductions are still possible.

This work introduces *Mix-Down* (MD) and *Stretch-and-Advance* (S&A), which represent a new class of attacks on UWB ranging. In these attacks, we exploit the receivers' tolerance towards natural clock errors and show how subsequent synchronization and compensation render them vulnerable to distance reductions. Our contributions can be summarized as follows:

Discussion of the security implications of clocks: We explain the functional consequences of non-ideal transceiver clocks, namely the carrier frequency offset (CFO) and the sampling frequency offset (SFO), and show how they are mitigated in practice. We discuss how these mitigations can be abused to manipulate distance measurements over the air.

Two novel attacks against UWB ranging: We present *Mix-Down* (MD), an attack against Single-Sided Two-Way Ranging (SS-TWR), and *Stretch-and-Advance* (S&A), which affects both SS-TWR and Double-Sided Two-Way Ranging (DS-TWR). Both distance reduction attacks depend heavily on design parameters, such as clock accuracy and the length of ranging messages. We show that unfavorable parameter choices can lead to distance reductions of up to dozens of meters and demonstrate MD practically. While the impact of S&A is currently negligible, the attack may lead to considerable reductions in the future standard 802.15.4ab. [Table 1](#) provides an overview of the standards and modes affected by either attack.

Practical evaluation of MD: We conduct experiments to demonstrate the MD attack against the SS-TWR mode of 802.15.4z-compliant ranging chips, showing that even conservative parameters allow distance reductions of several meters. Furthermore, we conclude that SS-TWR with real-

Attack	4z SS	4z DS	4ab SS	4ab DS
MD	✓	✗	~✓	✗
S&A	✗*	✗*	~✓	~✓

Table 1: Overview on standards / ranging modes affected by each attack. Notation: "✓" means affected and experimentally verified, "~✓" potentially affected future standard, "✗*" theoretically affected but negligible attack margin, "✗" not affected.

time compensation of clock errors is fundamentally insecure and should not be used in security-sensitive applications.

Countermeasure against Stretch-and-Advance: We introduce a countermeasure to secure DS-TWR against the S&A attack. We recommend that this proposal be considered in the upcoming IEEE 802.15.4ab standard.

Responsible disclosure: We contacted leading chip vendors and informed them about both attacks prior to publication.

2 Background: UWB Ranging

This section provides an overview of UWB ranging and discusses effects of clocks errors that are relevant to UWB security.

2.1 UWB Ranging Standards

The UWB channels considered in this paper are roughly 500 MHz wide, which enables transmitters to generate radio frequency signals consisting of short pulses (*i.e.*, 2 ns). A signal's Time of Arrival (ToA) is easier to measure with such pulses than, for example, with WiFi or Bluetooth waveforms. This advantage makes UWB an attractive candidate for wireless ranging. The current UWB standard is IEEE 802.15.4z, published in 2020 [3, 4]. A new version, 802.15.4ab, is currently under development [23].

2.1.1 Single-Sided and Double-Sided Two-Way Ranging

On the protocol layer, the current UWB standard supports three modes for ranging: SS-TWR, DS-TWR, and ranging based on Time Difference of Arrival (TDoA). This paper focuses on SS-TWR and DS-TWR, which are shown in Figure 1. As the names suggest, both modes require a back-and-forth or two-way exchange of ranging messages to estimate the distance between the *initiator*, who sends the first ranging message, and the *responder*. We will stick to these names, whenever their roles are relevant, to be consistent with the nomenclature of the 802.15.4z standard. When we use the terminology *transmitter* and *receiver*, we only indicate in which direction a message is being sent, regardless of the device's role in SS-TWR/DS-TWR.




Figure 1: SS-TWR (dashed box) and DS-TWR between the initiator and responder in 802.15.4z.

The ToA of a UWB packet is measured at a predefined position of the packet called the RMARKER. Since this is just a reference point, it can be put in any position if done so consistently. For example, in 802.15.4z, the RMARKER is at the end of the preamble, roughly in the middle of the packet (see Figure 1).

SS-TWR: In the simpler case of SS-TWR, only two UWB messages are sent. The distance between the devices is then estimated based on the ToF of these messages, which can be computed by subtracting the known processing time T_{reply} from the entire duration of the exchange:

$$T_{ToF} = \frac{1}{2} \cdot (T_{round1} - T_{reply1}(1-c)) \quad (1)$$

Note that the correction term c is related to the clock drift between the two devices. This difference in clock speed means that, from the initiator's perspective, the responder's clock is either too fast or too slow. Therefore, it sent the reply either too early or too late. The reply time deducted by the initiator would be incorrect in this case, and the ToF calculation would be inaccurate. The initiator compensates for this clock drift and rectifies the ranging measurement by scaling T_{reply} . This correction is the enabler of the Mix-Down attack we present in Section 3.

DS-TWR: In this mode, the transceivers perform two complete exchanges, one in each direction. The ToF can be computed as:

$$T_{ToF} = \frac{T_{round1} \cdot T_{round2} - T_{reply1} \cdot T_{reply2}}{T_{reply1} + T_{reply2} + T_{round1} + T_{round2}} \quad (2)$$

This formula implicitly corrects errors introduced by the clock drift mentioned above, but it forces the responder to send T_{round2} to the initiator. This formula is used (instead of taking the average of the two rounds) because it works even when the reply times of the two devices are not symmetric [31]. The formula also averages out differences in the ToF of the messages. Generally speaking, the DS-TWR mode requires longer media access, and the additional messages result in higher latency. Conversely, it is less prone to inaccuracies introduced by the clock drift between devices.

We have looked at an assortment of gadgets in the Apple and Samsung ecosystems, such as the iPhone 11 Pro, the

AirTag, the HomePod, the Samsung Galaxy S10, and the Samsung Smart Tag. By default, all of them use DS-TWR. We could not find official statements as to why this is the case. We assume that the motivation is the ToF computation, whose accuracy does not rely on a clock drift estimation. Furthermore, the additional delay and media access incurred by this mode do not seem to hamper performance in existing use cases. Other products, such as the Decawave Positioning and Networking Stack R2 in MDEK1001 [37], use SS-TWR, as its lower power consumption and latency are advantageous in Real Time Location Service (RTLS) systems. Furthermore, several open-source projects [17, 19, 20] employ SS-TWR, and SS-TWR is considered in the context of one-to-many ranging in the upcoming standard [25]. The Software API Guide of the Qorvo DWM3000 [36] (v2.2, section 6.3.26) mentions the newly available clock drift measurement in the driver as a feature that makes SS-TWR more accurate, and explicitly states that it no longer recommends DS-TWR over SS-TWR for accuracy reasons. To the best of our knowledge, there is no indication in the publicly accessible documentation of any vendor that one mode should be preferred over the other due to security reasons.

2.1.2 High and Low Pulse Repetition Frequency modes

UWB generally does not have exclusive access to its frequency bands and is subject to stringent power regulations. This is a challenge for UWB transceivers, as it limits the link budget of ranging messages. 802.15.4z comprises two physical layer designs that approach the power limits differently. First, the High Rate Pulse Repetition Frequency (HRP) mode uses a high number of narrowly spaced pulses, which have to be sent at low power. On the other hand, the Low Rate Pulse Repetition Frequency (LRP) mode opts for higher peak power at the cost of sending fewer, more widely spaced pulses. The two approaches have led to distinct transceiver designs. However, in the context of this paper, the differences between LRP and HRP are largely irrelevant.

2.1.3 802.15.4ab: Future of UWB ranging

The link budget limitations are a major drawback of the current UWB standard. UWB chips have found adoption in small gadgets without high-gain antennas, such as smartphones, smartwatches, and tags [1, 6, 40, 41]. Consequently, increasing the link budget is a central goal of the upcoming 802.15.4ab standard. To that end, Apple proposed Narrowband-Assisted Multi-Millisecond Ultra-Wide Band (NBA-MMS-UWB) ranging [14]. The main idea behind the proposal is to use long ranging messages that comprise several fragments at one-millisecond intervals. This increases the signal's energy and, therefore, its link budget. Since the fragments are separated by transmission breaks, the mean Power Spectral Density of NBA-MMS-UWB messages does not exceed

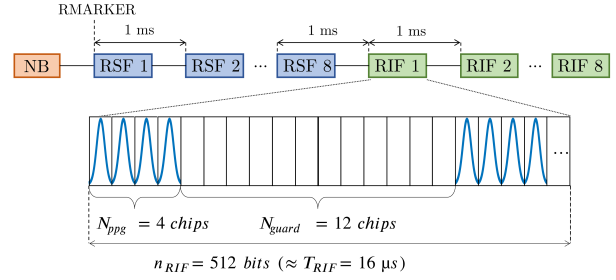


Figure 2: The proposed NBA-MMS-UWB frame consists of a narrowband (NB) fragment followed by Ranging Sequence Fragments (RSFs) for measuring the ToA (RMARKER), and Ranging Integrity Fragments (RIFs) for ToA integrity verification. Following [12], each RIF contains n_{RIF} decodable bits. Using multiple fragments (e.g., here $N_{RSF} = N_{RIF} = 8$) increases the link budget at the cost of a longer transmission time.

the limits imposed on UWB communication. However, the multi-millisecond ranging messages also have a drawback: they are more susceptible to clock-related errors, which we discuss in detail in Section 2.2.

NBA-MMS-UWB consists of narrow-band (NB) and UWB fragments. Before the UWB ranging starts, NB fragments are exchanged between the devices to transmit necessary data. The UWB fragments comprise Ranging Sequence Fragments (RSFs) and Ranging Integrity Fragments (RIFs). The purpose of the RSFs is to establish the ToA of the ranging message, whereas the RIFs enable ToA verification to prevent distance reduction attacks. Nevertheless, the RIFs could be used to establish and verify the ToA simultaneously. The standard envisions ranging messages with between 1 and 8 RSFs and RIFs. Figure 2 shows an example configuration. The standard puts the RMARKER, which represents the ToA internally, at the beginning of the first RSF [47].

At the time of writing, the design of the RIF was still a matter of ongoing discussion, and two alternatives had been submitted to the task group. The first proposal retains the Scrambled Timestamp Sequence (STS) waveform from 802.15.4z HRP [47]. This approach would indicate a ToA verification based on an approximated Channel Impulse Response (CIR), which can be obtained by cross-correlating the incoming signal with a local STS template.

The second proposal from Barras et al. [12] suggests modulating the pseudo-random bits as bursts of $N_{ppg} \in \{1, 2, 4, 8\}$ pulses with a peak Pulse Repetition Frequency (PRF) of 499.2MHz, which corresponds to one chip (≈ 2 ns) per pulse. The bursts are separated by guard intervals that reduce the mean PRF to 124.8MHz. Consequently, their duration depends on N_{ppg} and is three times the one of a burst. For $N_{ppg} = 1$, the scheme coincides with the legacy STS waveform mentioned before. However, according to the

authors, choosing $N_{pg} \in \{4, 8\}$ leads to guard intervals long enough to sufficiently reduce inter-burst interference. This should enable receivers to decode and verify bits individually.

From a security perspective, this is a major advantage over the first, CIR-based approach, whose security guarantees depend heavily on the receiver implementation and are neither well understood nor easy to prove [27]. Therefore, we assume throughout this work that the second proposal is used. Where we provide concrete examples, we choose the configuration shown in Figure 2. With $N_{pg} = 4$ pulses per burst/bit, it balances inter-burst interference and symbol length. This compromise simplifies decoding of individual bits while limiting the attack window for Early-Detect/Late-Commit [33]. As a result, we obtain $T_{bit} = 32\text{ns}$ (16 chips) per bit/burst, of which 8ns (4 chips) are active. We further assume that a RIF is $N_s = 32$ units of 512 chips long. One RIF contains $n_{RIF} = 512$ bits, and the entire ranging message $N_{RIF} = 8$ RIFs, resulting in $n_{tot} = 4096$ bits in total.

2.2 Clocks

A digital transceiver uses a clock derived from a crystal oscillator to measure time intervals, sample incoming signals, and generate radio frequencies. Because of hardware imperfections and environmental factors, the clock frequency of a device and all derived quantities differ from their nominal values. For example, an initiator with a nominal clock frequency f_{nom} would actually run at a different clock frequency f_{init} . We call this error, which can be conveniently expressed in parts per million (ppm), the *clock drift*¹ c_{init} :

$$c_{init} = \frac{f_{init} - f_{nom}}{f_{nom}} \quad (3)$$

The UWB standard [4] defines the accepted range of clock drifts as $\pm 20\text{ppm}$. While clock drifts may vary over time, they can be considered constant during the short duration of a ranging exchange.

A device with a negative clock drift generates frequencies below their nominal values. If asked to wait for a nominal time interval, it waits longer, and when measuring an absolute time interval between two events, it reports a smaller value. Opposite effects occur for devices with positive clock drifts. The larger the times and frequencies involved, the larger the absolute errors in s and Hz caused by the same clock drift in ppm.

In general, transmitter and receiver exhibit different clock drifts, causing the following problems. First, the carrier frequencies they generate differ by a Carrier Frequency Offset (CFO) expressed in Hz. The CFO causes a rotation of the symbols on the constellation plane, which can lead to decoding errors. Second, the receiver is affected by a Sampling Frequency Offset (SFO) causing it to sample the incoming symbols at the wrong intervals. As a result, the

¹We refrain from using the term *clock frequency offset* to avoid confusion with the Carrier Frequency Offset (CFO), which is caused by the clock drift.

receiver accumulates less energy per symbol, which degrades performance. Finally, transceivers have different estimates of time intervals, which impacts the accuracy of distance measurements based on the ToF.

Upon reception, the receiver can derive the clock drift between the two from the measured CFO [42], and apply necessary corrections.

3 The Mix-Down Attack (MD)

In this section, we describe the Mix-Down (MD) attack, the first of two novel attacks related to device clocks in UWB ranging. The underlying vulnerability affects the SS-TWR modes of the current standard (802.15.4z) and the current proposal of the upcoming one (802.15.4ab).

3.1 Threat Model

The attacker aims to reduce the distance measured by two UWB transceivers. We assume a Dolev-Yao attacker that controls the wireless channel and can manipulate, block, inject, and eavesdrop on all wireless signals. Furthermore, the attacker is able to configure the power and timing of their transmission precisely if required. Following Kerckhoff's principle, we also assume that the attacker knows everything about the victim's devices except their cryptographic secrets.

This threat model decouples the security of the system from factors related to specific scenarios or use cases, such as the spatial arrangement of devices, and from obscure implementation details, which cannot be openly analyzed or verified. A secure ranging scheme should provide a high and verifiable level of protection against distance reduction attacks under this strong threat model.

In practice, attacks on wireless ranging succeed under far less ideal conditions: many instances of car thefts have shown that parked vehicles are a viable target, even if the exact position of the key fob is not known.^{2,3,4}

3.2 General Idea

Figure 1 depicts a Single-Sided Two-Way Ranging (SS-TWR) exchange. The initiator computes the ToF by measuring the round-trip time of two messages with the responder, and subtracting the reply time:

$$T_{ToF} = \frac{1}{2}(T_{round} - T_{reply}(1 - c)) \quad (4)$$

The corrective term $c = c_{resp} - c_{init}$ measured by the initiator denotes the *relative* drift between the clocks of the two devices. It is the difference between the *absolute* clock drifts (*i.e.*, in reference to an ideal clock) of the responder and

²<https://www.youtube.com/watch?v=PEMOWPj2i-0>

³<https://www.youtube.com/watch?v=bR8RrmEizVg>

⁴<https://www.youtube.com/watch?v=bR8RrmEizVg>

the initiator, respectively. Due to these drifts, the transceivers disagree about how long, for example, 1 ms or T_{reply} last in absolute time.

The formula above implies that the SS-TWR initiator assumes its own clock to be correct. From the initiator's perspective, a relative clock drift c means that the responder clock is imprecise, and that the reply time it waited is incorrect. Since a deviation of a single nanosecond results in a ranging error of $\pm 15\text{cm}$ in Eq. 4, the initiator must estimate the relative clock drift and compensate accordingly. As described in Section 2.2, clock drift estimation can be done by measuring the incoming signal's Carrier Frequency Offset (CFO).

In the MD attack, we exploit that this estimation can be controlled by the attacker. We achieve this by replaying the response message on a different carrier frequency corresponding to a strong negative clock drift c_{att} . The initiator receives the replayed message and interprets c_{att} as c_{resp} . It concludes that the responder is too slow and, consequently, must have replied to late. To compensate for this apparent delay, the initiator deducts it from what would be the correct ToF. The final, shorter ToF results in a distance reduction.

This attack only changes the carrier frequency of the ranging message but not its content. Therefore, it cannot be prevented by existing security mechanisms relying on the pseudo-randomness of bits/pulses, such as the STS in 802.15.4z HRP. Furthermore, the MD attack is fully deterministic and theoretically works on every ranging round. Since the attacker can choose c_{att} , they can also steer the distance reduction within certain limits. This is in stark contrast to attacks like Ghost Peak [27], which are probabilistic and whose measurement outcomes are beyond the attacker's control. Finally, the MD attack could also be used to increase the ToF by faking a faster responder clock. We do not discuss this in detail since distance enlargement is usually trivial to achieve through delayed retransmission.

3.3 Analysis

We can express the result of the attack as

$$T_{ToF} + \delta_{ToF} = \frac{1}{2} (T_{total} - T_{reply}(1 - c')) \quad (5)$$

where δ_{ToF} denotes the ToF difference caused by the attack. Note that the attacker's interference changed $c = c_{resp} - c_{init}$ to $c' = c_{att} - c_{init}$, so the initiator measures the clock drift between its clock and the attacker's. This equation shows that a ToF reduction is a negative quantity, which is the convention we will follow throughout this paper. For simplicity, and because they are proportional, we will treat the terms "ToF reduction" and "distance reduction" synonymously. Subtracting Eq. 4 from Eq. 5 results in

$$\delta_{ToF} = \frac{T_{reply}}{2} \cdot \delta_c \quad (6)$$

where

$$\delta_c = c' - c = c_{att} - c_{resp} \quad (7)$$

The value of δ_{ToF} , therefore, depends on the length of the reply time and the difference $\delta_c = c_{att} - c_{resp}$ between the clock drift chosen by the attacker and the real clock drift of the responder. To cause a distance reduction ($\delta_c < 0$) the attacker must use a clock drift c_{att} less than the clock drift c_{resp} of the responder. The maximal distance reduction is then bounded by the range of clock drifts the initiator accepts.

An acceptable range for c_{init} and c_{resp} is $[-c_{std}, c_{std}]$, where c_{std} is defined in the UWB standard (e.g., $c_{std} = 20\text{ppm}$ in 802.15.4z). For simplicity, we first consider an initiator who has an ideal clock ($c_{init} = 0\text{ppm}$) and is aware of its precision. The latter enables the initiator to limit $c = c_{resp} \in [-20\text{ppm}, +20\text{ppm}]$, which covers all standard-compliant responder clocks. Under these conditions, the attacker is also limited to $c_{att} \in [-20\text{ppm}, +20\text{ppm}]$. They are generally interested in the highest negative c_{att} , which causes the largest reduction. In the best case for the attacker, with $c_{resp} = +20\text{ppm}$ and $c_{att} = -20\text{ppm}$, we have $\delta_c = -40\text{ppm}$. For an assumed reply time of $T_{reply} = 2\text{ms}$, this results in a distance reduction of -12m .

Transceiver clocks are not ideal in general, since hardware imperfections and environmental factors cause them to diverge. In the worst case, initiator and responder could *unknowingly* be affected by any $c_{init}, c_{resp} \in [-20\text{ppm}, +20\text{ppm}]$. Such deviations would force the initiator to accept incoming signals with $c \in [-2c_{std}, 2c_{std}]$. Under this condition, the attacker may choose $c_{att} \in [c_{init} - 2c_{std}, c_{init} + 2c_{std}]$. The best-case scenario for the attacker is $c_{init} = -20\text{ppm}$ and $c_{resp} = +20\text{ppm}$. Choosing $c_{att} = -60\text{ppm}$ leads to $\delta_c = -80\text{ppm}$, and, for $T_{reply} = 2\text{ms}$, to a distance reduction of -24m . In the average case of $c_{init} = c_{resp} = 0\text{ppm}$, the reduction limit is -12m .

3.4 Experimental Evaluation

In this section, we demonstrate distance reductions against commercial HRP UWB devices in an indoor environment. The MD attacker is built of simple off-the-shelf components and two software-defined radios. It can measure the clock drift between any two devices over-the-air, which we leverage to tune its own clock drift before an attack.

3.4.1 Setup and Practical Aspects

Test devices: We use two Qorvo DWM3000EVB boards [36] as initiator and responder. The firmware is based on the example provided in the chip's Software Development Kit (SDK) [18]. The devices perform two SS-TWR rangings per second and use the dynamic STS mode of 802.15.4z HRP for security: with a cryptographic key that was shared out-of-band, they generate an unpredictable, pseudo-random Scrambled Timestamp Sequence (STS) for every new packet.

The receiver’s code verifies the correctness of the STS by calling the `dwt_readstsquality` function available in the API. Since the attacker does not know the STS in advance, they cannot advance the packet to reduce the measured distance. Before MD, the only known physical-layer attack possible against 802.15.4z was GhostPeak, which can cause random distance reductions with a limited success rate [27]. We minimally adapted the firmware such that: (i) all parameters can be monitored, (ii) the clock drift of the devices can be trimmed, and (iii) the reply time T_{reply} can be configured.

Attacker: We implement our attack with a simple analog circuit consisting of off-the-shelf RF components. The clock drift c_{att} can be precisely and flexibly controlled via software. Figure 3 shows our attacker device: it consists of an antenna to receive the responder’s reply at $f_c = 6.4896\text{GHz}$ (HRP channel 5), a mixer to convert it down to an intermediate frequency of 1.6 GHz, and another mixer to convert it up to the shifted carrier $f'_c = f_c + f_{att} = f_c + f_c \cdot c_{att}$.⁵ For example, if an initiator with $c_{init} = 0\text{ppm}$ should measure $c = -10\text{ppm}$, the attacker applies a frequency shift $f_{att} = -10 \cdot 6.4896\text{ kHz}$ to bring the carrier down to $f'_c \approx 6.4895\text{GHz}$. The resulting signal is then transmitted over a second antenna in the initiator’s direction. We used several amplifiers and filters to improve the quality of the signal. The combined delay caused by the cables and the analog circuitry is below 5 ns. The frequencies for down- and upconversion are generated with two USRP B210 Software-Defined Radios (SDRs) and an additional mixer. Apart from the SDRs (\approx USD 1000), which could be replaced by cheaper alternatives, the entire hardware’s worth is approximately USD 400.

Initial clock drift estimation: In Appendix A.1 we present a sniffer that can estimate $c = c_{resp} - c_{init}$ (and/or $c' = c_{att} - c_{init}$) over-the-air, using the messages received from the other devices. This enables tuning c_{att} , for example, for a smooth takeover in which c' is initially close to c .

Environment and placement: We conduct our experiments in an indoor environment. The test devices are stationary during the experiments and up to 10 m apart. They are either in line-of-sight (experiments 1-3) or in different rooms separated by a wall or glass (experiment 4).

In experiments 1-3, we place the input antenna close to the responder and the output antenna close to the initiator. We connect the antennas with a coaxial cable to the attacker device. We also relay the first message from the initiator to the responder to facilitate their communication. This setup matches our threat model, which assumes complete attacker control over the wireless channel. In experiment 4, we demonstrate that the attack works even if the attacker cannot get close to the responder. The sniffer that monitors c and c' is placed in the attacker’s vicinity.

The input and output antennas of the attacker are directional and should generally point toward the victims.

⁵Using a single mixer for down-conversion is not possible because the 500 MHz bandwidth of the signal is much larger than the offset.




Figure 3: Main components of the attacker. The analog mixing applies the attacker’s offset c_{att} to the message from the responder in real-time. The offset can be configured in software since the carriers are generated with USRP B210 SDRs. Antennas can be connected via cables and amplifiers.




Figure 4: Example of distance reduction from 10 m to 0 m.

Furthermore, placing them in opposite directions helps to avoid undesired feedback loops between output and input. We can control the power of the attacker’s signal by adjusting the gain of the USRPs in software. Naturally, increasing the gain causes the initiator to lock onto the attacker’s signal.

Monitoring and control: In order to perform systematic and reproducible experiments, we have implemented a control and monitoring system based on the Avatar² [29] framework. Its most important functions are: (i) programming and configuring the victims, (ii) controlling c_{att} programmatically, and (iii) monitoring and storing reception and ranging diagnostics (e.g., measured distance, measured clock drift, STS verification result, and reception errors).

3.4.2 Experiments and Results

Experiment 1. Reduction from 10m to 0m. In this experiment, we position the devices in line-of-sight and ten meters apart ($d_{real} = 10\text{m}$). Furthermore, we set the reply time to




Figure 5: Distance reductions with respect to the attacker’s clock drift c_{att} for different reply times. Each star represents a successful measurement. The initiator’s clock drift c_{init} was constant during the experiment.

$T_{reply} = 5$ ms. Figure 4 shows that once the attack starts at ≈ 40 s, the measured distance decreases to ≈ 0 m and remains stable at the new value. As the MD attacker has complete control over the measured distance, they can also reduce the distance smoothly (e.g., to simulate a person approaching a car).

Experiment 2. Systematic evaluation of T_{reply} . Neither the DWM3000EVB manual nor 802.15.4z mention any security implications of a long T_{reply} . While it generally makes sense to keep T_{reply} as short as possible for measurement accuracy, a well-calibrated initiator (i.e., $c_{init} \approx 0$ ppm) can compensate for arbitrary errors caused by the responder’s clock drift, regardless of T_{reply} . Furthermore, use cases such as one-to-many ranging [25] or Real Time Location Service (RTLS) systems with dozens of tags may require longer reply times since tags reply one after the other. In this experiment, we positioned the devices at a distance of ≈ 2.5 m. For every value of $T_{reply} \in \{1 \text{ ms}, \dots, 5 \text{ ms}\}$ and c_{att} we conduct 10 rangings, and increase c_{att} as long as at least one reception succeeds. The results are shown in Figure 5.

Note on proprietary security mechanisms: The Qorvo DWM3000EVB provides a set of security flags (e.g., "peak growth rate warning" and "ADC count warning"), which are not documented in detail. The MD attack still works when the receiver is configured to heed them and discard affected measurements. This result matches our expectation since it is unlikely that the warnings were designed to detect our manipulation. However, we observed that the lower bounds of $c' \in [-42\text{ppm}, -32\text{ppm}]$ (see Figure 5) were reduced to $[-25\text{ppm}, -20\text{ppm}]$. We presume that fine-tuning the attack (e.g., through power adjustments) would mitigate this effect.

Experiment 3. Systematic evaluation of δ_c . We postulate in Section 3.3 that the attacker’s clock drift advantage $\delta_c = c_{att} - c_{resp}$ depends on the clock drifts c_{init} and c_{resp} of the two devices. This is because the initiator should accept




Figure 6: Experiment 4. The responder is placed at ≈ 5 m from the initiator and ≈ 4 m from the attacker, separated by a glass window. The distance estimated by the devices can be reduced to 0 m.

an incoming signal if its clock drift estimate c is in the range $[c_{init} - 2c_{std}, c_{init} + 2c_{std}]$. Consequently, both attacker and responder have to be within this range.

To test this assumption, we trim the two transceivers to different clock frequencies using a dedicated device register and run the MD attack for varying c_{att} . We observe that the chip reports reception events for $c_{att} \in [c_{init} - 60\text{ppm}, c_{init} + 60\text{ppm}]$ for any c_{init} . This result shows that the initiator’s clock drift tolerance is even higher than expected, and that the interval is centered around its own clock drift.

However, the initiator only produces a distance measurement if $\delta_c = c_{att} - c_{resp} \in [-25\text{ppm}, +25\text{ppm}]$ (with minor deviations), regardless of c_{init} and c_{resp} . Any other δ_c leads to a reception error. Consequently, for the device model we used, the maximum distance reduction does not depend on the clock drifts of the ranging devices.

This result can be explained with the sampling frequency offset (SFO, see Section 2.2): the initiator adjusts the sampling rate based on its clock drift estimate c . This measure ensures that, under benign conditions, the signal can be received correctly, even for large clock drifts. Under attack, however, the initiator adjusts the sampling based on c_{att} , whereas the signal was generated using the responder’s clock. Therefore, the "correction" by the initiator causes itself an SFO that is proportional to $c_{resp} - c_{att}$. If this difference exceeds a certain threshold, the reception fails.

Experiment 4. No physical access to responder. To demonstrate the feasibility of the MD attack in a more challenging environment, we put the responder into a different room that the attacker cannot access. In two separate runs, the responder is either behind a glass window (see Figure 6) or behind a wall. The distance between the attacker’s antenna and the responder is approximately 5m. In this setting, the cable relay used in the previous experiments is not possible, and the received signals have to be amplified to compensate for the additional pass loss. The results show that the distance can still be reduced from 5m to 0m in this scenario.

3.5 Countermeasures

The main problem of SS-TWR is that the initiator has to estimate the responder’s clock drift and compensate for it.

Without this compensation, it could not provide the accuracy required from UWB ranging. To the best of our knowledge, there is no secure way of verifying the integrity of a CFO under our threat model. Therefore, we argue that MD is hard to prevent, and should instead be mitigated by reducing the reply time and clock drift tolerances. In particular, the following approaches do not promise to solve the problem:

Monitor clock drifts over time: The devices could track clock drifts to detect attacks. However, a changing clock frequency may have benign causes (*e.g.* device warm-up or temperature changes), leading to a tradeoff between false negatives and false positives.

Exchange and compare clock drifts: During normal operation, the clock drift estimates of the two devices should be almost identical, but with opposite signs. Thus, deviations caused by MD could be detected. However, the attacker could manipulate the CFO of both messages accordingly: since the clock drift of the first message is not used in the distance computation, altering it has no consequences, apart from bypassing the countermeasure.

Check for presence of a second signal: If the receiver can detect a similar signal at a slightly higher frequency, it could detect the attack. However, an attacker could bypass this countermeasure by, for example, significantly increasing the power of their own transmission or, if they have sufficient control over the wireless channel, by blocking the genuine signal.

4 Attack 2: Stretch-and-Advance (SaA)

In this section, we present our second attack called Stretch-and-Advance (S&A). This attack affects the draft of the upcoming 802.15.4ab standard, for which no hardware implementation exists yet. Therefore, our evaluation is analytical; for implementation aspects, we refer the reader to [Appendix A.2](#). We assume the same threat model as for the MD attack, described in [Section 3.1](#).

4.1 Introduction

A transmitter whose clock frequency is lower than the nominal frequency (*i.e.*, with a clock drift $c < 0$ ppm) generates a signal that is "stretched" in time, as illustrated in [Figure 7](#). A receiver whose clock operates at the nominal frequency and processes this signal will experience a Sampling Frequency Offset (SFO), since it is in disagreement with the transmitter about the duration of and the delay between UWB pulses. This error grows over time and, without compensation, makes it impossible to acquire and decode long signals. In fact, this phenomenon is responsible for the result of experiment 3 in the [Section 3.4](#) of the MD attack. To avoid this issue, the receiver estimates the clock drift c between the devices by measuring the CFO during the Narrow Band (NB) part of the NBA-MMS-UWB ranging message. Based on the result, it may compensate for the drift in software or in hardware.




Figure 7: Comparison of two ranging messages generated with an ideal clock ($c = 0$ ppm) and a slower clock ($c < 0$ ppm).

For simplicity, we describe the attack with a basic NBA-MMS-UWB configuration containing a single RSF and a single RIF, as shown in [Figure 7](#). The figure illustrates that the delay between the RMARKER and the start of the RIF depends on the transmitter's clock frequency. We will show how this dependency can be exploited to manipulate the ToA of ranging messages and, thus, to reduce the distance between the devices.


4.2 Attack Description

The attacker injects a specially crafted ranging message, which is advanced (*i.e.*, it reaches the receiver before the genuine one) and stretched in time. Because of the stretching, the attacker gains enough time to learn and replay the unpredictable RIF bits from the legitimate transmitter. The resulting message will look to the receiver as if it had been sent by a legitimate transmitter with a slower clock, and its earlier ToA shortens T_{round} (see [Eq. 4](#) and [Eq. 2](#)), which leads to a distance reduction.


[Figure 8a](#) illustrates S&A from the attacker's perspective. The upper part shows a ranging message m received from the transmitter, and the lower one the message m' transmitted by the attacker. The attacker constructs m' by first generating the NB and RSF. This is possible because we assume these parts to be deterministic (we discuss non-deterministic NB fragments in [Appendix A.2](#)). The attacker uses a clock with an artificial drift c_{att} in the process, which mimics a slower transmitter and results in the stretching mentioned above. [Figure 8a](#) also shows the advancement or left shift by δ_{RSF} . The attacker chooses this shift such that the starting times of the RIFs of the two messages (t_{replay}) coincide. From t_{replay} on, the attacker can learn and replay the signal and data from the real message m . By stretching and replaying the entire RIF, the attacker finalizes m' .

In the attack description above, we made two implicit assumptions:

Receiver decodes the attacker's message, not the genuine one: Given that the attacker sends before the genuine transmitter and can choose the transmission power, the receiver will synchronize to the attacker's message. In [Section 3.4](#) we



(a) The attacker advances the message by δ_{RSF} , which is the maximum that still allows them to learn and reply all RIF bits in time (the messages are aligned at t_{replay}). The advancement depends on the clock drift (the stretch) and the interval between the RMARKER and t_{replay} .



(b) Advancing m' by an additional δ_{RIF} is possible, but the leading n_{att} RIF bits (left from t_{replay}) become unpredictable.

Figure 8: Illustration of the Stretch-and-Advance (S&A) on a ranging message with a single RIF and a single RSF.

have demonstrated that this is possible in practice for the MD attack, even for attacker signals with ≈ 5 ns delay.

Negligible processing delay: For simplicity, we assume an attacker capable of receiving and retransmitting RIF pulses instantly, without processing delay. Consequently, the ToF or distance reductions we compute in this section are theoretical upper bounds. In practice, with the analog implementation outlined in [Appendix A.2](#), we expect negligible delays in the order of nanoseconds, which minimally impact the achievable distance reduction.

We note that the attacker could advance m' by more than δ_{RSF} , as depicted in [Figure 8b](#). However, this additional shift δ_{RIF} might come at a cost: it moves t_{replay} into the RIF, which renders the leading RIF bits unpredictable for the attacker. Depending on the value of δ_{RIF} and the Bit-Error Rate (BER) accepted in the RIF decoding, the success probability of the attack might drop. We leverage this observation in the following security analysis.

4.3 Analysis

The S&A vulnerability affects SS-TWR and DS-TWR equally. Since both ranging modes average the ToF of all messages (see [Eq. 4](#) and [Eq. 2](#)), we can write the distance reduction introduced by the S&A attack as

$$\delta_{ToF} = \frac{a}{N}(\delta_{RSF} + \delta_{RIF}) \quad (8)$$

where N is the number of messages of the ranging scheme and a the number of messages attacked. Consequently, we have $N = 2$ for SS-TWR and $N = 4$ for DS-TWR⁶. In the case of $a \neq 1$, this equation assumes that the attacker uses the same δ_{RSF} and δ_{RIF} across all attacked messages.

In the remainder of the section, we will first quantify δ_{RSF} and δ_{RIF} . We will then show how p_{RIF} , the attacker's probability to pass the RIF integrity check, is coupled to δ_{RIF} .

4.3.1 δ_{RSF}

We consider the case shown in [Figure 8a](#). The advancement of the RMARKER can be written as

$$\delta_{RSF} = \delta_c \cdot T_{RSFs} \quad (9)$$

Here, T_{RSFs} is the time duration of all RSFs combined (N_{RSF} fragments at 1ms intervals), as shown in [Figure 2](#). Like in [Section 3](#), δ_c denotes the difference in clock drifts between the attacker and the legitimate transmitter. We show in [Section 4.4](#) that an advancement of δ_{RSF} alone can, when applied to all messages, lead to a distance reduction of over 90m.

4.3.2 δ_{RIF}

If the attacker aims at a larger reduction, they can advance the ranging message by an additional δ_{RIF} , as shown in [Figure 8b](#). However, this means the leading RIF bits ($\approx 50\%$ in the figure) become unpredictable for the attacker. If the receiver expects all bits of the RIF to be correct, a successful attack would require the attacker to guess these unknown bits. Thus, by increasing δ_{RIF} , the attacker would attempt to shorten the distance further but simultaneously reduce the probability p_{RIF} of passing the RIF verification. We first focus on the advancement δ_{RIF} and analyze the impact on p_{RIF} subsequently.

We assume that the attacker chooses a desired distance reduction, which requires them to apply a given advancement δ_{RIF} to the ranging message. We denote by n_{att} the number of leading RIF bits made unpredictable by this δ_{RIF} . The greater δ_{RIF} , the higher the number of unpredictable bits the attacker has to accept. We can write this as

$$\delta_{RIF} = \delta_c \cdot \mathcal{T}(n_{att}) \quad (10)$$

$\mathcal{T}(n_{att})$ is the time interval corresponding to the n_{att} bits, from the beginning of the first RIF to t_{replay} . For our example with a single RIF, we have

$$\mathcal{T}(n_{att}) = n_{att} T_{bit} \quad (11)$$

where T_{bit} is the length of a single bit. Generally speaking, $\mathcal{T}(n_{att})$ is a monotonically non-decreasing function of n_{att} that depends on the RIF design and configuration. For ranging messages with several RIFs, \mathcal{T} can be found in [Appendix A.3](#).

⁶For the three-message version of DS-TWR, $N = 4$ but an attack on the second message counts double, as it affects both ranging rounds.

4.3.3 Bit errors in RIFs

For optimal security, a receiver could verify that every single RIF bit is correct. However, multi-path effects, path loss, and transmission power constraints inevitably lead to flipped bits, even in benign environments. Receivers will therefore have to accept RIFs with a non-zero bit error rate. We denote by BER_{max} the highest accepted bit error rate, calculated over all RIFs, below which the receiver accepts RIFs as correct. This assumption is based on the very motivation behind NBA-MMS-UWB: the link budget or the SNR of a single RIF is insufficient to support reliable ranging. Furthermore, the channel might not be coherent over the duration of a long message, especially if the transceivers or objects in their vicinity are in motion.

The number n_{BER} of bit errors accepted by the receiver depends on BER_{max} and n_{tot} , the total number of bits in all RIFs:

$$n_{BER} = \lfloor BER_{max} \cdot n_{tot} \rfloor \quad (12)$$

In contrast to standard-compliant transceivers, the attacker can mitigate detrimental channel and noise effects by exploiting their power advantage, favorable positioning, and custom hardware. This enables them to detect and retransmit the *predictable* RIF bits, *i.e.*, those after t_{replay} , so that the receiver can decode all of them correctly. Therefore, in the best case, the attacker can make up to n_{BER} guessing errors among the unpredictable bits and still pass the RIF integrity check. n_{BER} is thus a key parameter that links the reduction δ_{RIF} to p_{RIF} , the probability of passing the check.

4.3.4 Probability p_{RIF}

Since the attacker can retransmit all predictable bits without errors, they can achieve $p_{RIF} = 1.0$ for all $n_{att} \leq n_{BER}$. However, for $n_{att} > n_{BER}$ the attacker is forced to guess at least $n_{att} - n_{BER}$ bits correctly, otherwise the attack fails. The number of bit errors can be modeled with a binomial random variable $X \sim \mathcal{B}(n_{att}, 0.5)$, and the probability that the attacker makes at most n_{BER} bit errors in n_{att} guesses corresponds to the CDF of X at the position n_{BER} , *i.e.*,

$$p_{RIF} = P(X \leq n_{BER}) = F_X(n_{BER}) \quad (13)$$

p_{RIF} only describes the probability of passing the RIF integrity check of a *single* message for a chosen n_{att} . Since an attacker has to pass the integrity check of all a attacked messages (see Eq. 8) and because the outcomes of these checks are statistically independent events, we can define the overall success probability of the S&A attack, p_{succ} , as

$$p_{succ} = \prod_{i=1}^a p_{RIF_i} \quad (14)$$

4.4 Impact on UWB standards

The maximum distance reduction through S&A depends on several parameters. To estimate the effect of the attack on 802.15.4ab, we use the configuration introduced in Section 2.1.3. For 802.15.4z we use the default configuration of the HRP mode.

4.4.1 802.15.4ab

We have conducted the evaluation for ranging messages with 8 RSFs and 8 RIFs, each at 1ms intervals. Every RIF comprises $n_{RIF} = 512$ bits, which adds up to $n_{tot} = 4096$ bits. Within a RIF, bits are sent at $T_{bit} = 32ns$ intervals. Like in Section 3.3, we assume $\delta_c = -40ppm$ to compute reduction limits. We assume that the devices use SS-TWR, and that the attacker advances both messages equally to optimize p_{succ} .

Results: Our evaluation shows that the distance reductions possible with the S&A attack are large enough to defeat any proximity-based access control system, irrespective of the tolerated bit error rate.

Figure 9 shows the success probability of the S&A attack in dependence of the distance reduction. The attack is deterministic ($p_{succ} = 1.0$) for reductions of at least $\delta_d = -93.95m$, which can be achieved through advancing both messages by δ_{RSF} only (see Eq. 9). Longer reductions require $\delta_{RIF} < 0$ and result in unpredictable RIF bits. Nevertheless, the distance between two transceivers tolerating $BER_{max} = 0.1$ can be reduced by $\delta_d = -107m$ while still keeping $p \approx 1.0$. For $BER_{max} = 0.25$, reductions of $\delta_d \approx -132m$ ($p_{succ} = 1.0$) or $\delta_d = -143m$ ($p_{succ} \approx 0.5$) are possible. For simplicity, we rounded all distance reductions to meters.

4.4.2 802.15.4z

The current 802.15.4 standard is practically unaffected by the S&A attack: generally speaking, 4z messages are too short for exploitable attack margins.

Out of the four message formats defined for HRP (see [4], Figure 15-2a), two put the RMARKER only microseconds before the start of the STS, which is the HRP equivalent of the RIF. Hence, the stretchable timeframe in Eq. 9 is too short and only results in a negligible $\delta_{RSF} < 1ns$. Another message format puts an encrypted payload field between the RMARKER and the STS, which is unpredictable and cannot be advanced, effectively leading to the same result. Finally, the last format contains no STS field, has no security claims, and can be attacked trivially.

For δ_{RIF} , the situation is similar: the STS is comparably short (*e.g.*, $\approx 131\mu s$ for $n_{tot} = 8192$ pulses). Even in the extreme case, where the attacker chooses t_{replay} such that $n_{att} \approx n_{tot}$, δ_{RIF} would be negligible: $\delta_{RIF} = 131\mu s \cdot -40ppm = -5.24ns$.

More exotic, optional configurations exist (*i.e.*, multiple and longer STS segments with additional RMARKERS), but




Figure 9: The success probability p_{succ} depends on the distance reduction δ_d (i.e., δ_{ToF} times the speed of light). Every curve corresponds to the results for a specific BER_{max} . The attacker can always achieve $\delta_d = -95.93\text{m}$, even if the receiver tolerates no bit errors. The discontinuity for $BER_{max} = 0.25$ is related to the pauses between RIFs, and enables the attacker to increase the distance reduction, to some extent, without sacrificing success probability (see the definition of \mathcal{T} in [Appendix A.3](#)). A distance reduction of about -192m requires an advancement δ_{RIF} that renders all RIF bits unpredictable.

the standard does not detail how they should be processed. To the best of our knowledge, we therefore conclude that distance reductions with S&A are negligible in 4z HRP.

4.5 Countermeasure

In this section, we propose a countermeasure that allows to bound the success probability of S&A distance reductions to an arbitrarily low value.

4.5.1 Working Principle

The S&A attack can be mitigated if the two devices verify that their clock drift estimates are consistent. Under benign conditions, a relative clock drift c measured by the responder causes the initiator to report approximately $-c$. In contrast, the S&A attack described before would cause a conspicuous deviation. The countermeasure we propose prevents S&A and cannot be bypassed if receivers set $BER_{max} \approx 0.2$.

We assume that both devices estimate their clock drift based on the CFO of the NB fragment before the UWB exchange begins, and use this estimate for the duration of the entire reception. We do not detail how the clock drift estimates are to be exchanged, but we require their integrity to be cryptographically protected. Both conditions should be easy to fulfill: the devices need the means to exchange timing information for DS-TWR in any case, and the generation of RIF sequences already requires a shared secret.

4.5.2 Security Analysis

For the sake of simplicity, we analyze the countermeasure for SS-TWR and the basic NBA-MMS-UWB configuration with one RIF. We transfer the result to DS-TWR at the end of the section and to the case of multiple RIFs in [Appendix A.3](#).

The following analysis is based on the publicly available draft documents and proposals of 802.15.4ab at the time of

writing. Since they may differ from the final specification, the proposal should go through the full standardization process.

In SS-TWR, initiator and responder send one message each (i.e., m_1 and m_2) and produce a clock drift estimate on the received message. In order to achieve a distance reduction, the combined ToF change of both messages must be negative:

$$\delta_{ToF} = \frac{1}{2}(\delta_{ToF1} + \delta_{ToF2})^7 \quad (15)$$

Without loss of generality, we assume that the attacker successfully conducts the S&A attack on m_1 . In the process, they change the clock drift estimate of the responder by $\delta_c < 0$ and achieve a ToF reduction of

$$\delta_{ToF1} = \delta_c(T_{RSFs} + \mathcal{T}(n_{att1})) \quad (16)$$

In order to bypass the countermeasure, the attacker also has to manipulate the initiator's clock drift estimate by $-\delta_c > 0$. In this case, the estimates are consistent and the transceivers accept the ranging result. However, this causes the initiator to expect a compressed message because it concludes that the responder's oscillator is too fast. Because of the resulting sampling frequency offset (SFO, see [Section 2.2](#)), the reception of m_2 inevitably fails. This issue already occurs in 802.15.4z for much shorter messages (see experiment 3 in [Section 3.4](#)). To prevent this, the attacker has to forge a compressed m'_2 . Consequently, they have to send the RIF pulses at a higher rate than they can receive them from the genuine transmitter. To replay all RIF bits correctly, the attacker has to delay the transmission of m'_2 long enough to learn the last RIF pulse just in time for retransmission. This is depicted in [Figure 10](#), which illustrates how compressing and delaying m'_2 increases its ToF by $\delta_{RSF2} = -\delta_c(T_{RSFs} + T_{RIFs})$.

If the attacker is willing to compromise on the RIFs correctness, they can advance m'_2 and mitigate the enlargement. This advancement can be written as $\delta_{RIF2} = -\delta_c \mathcal{T}(n_{att2})$,

⁷The division by 2 is due to the averaging in [Eq. 4](#).




Figure 10: To replay all RIF bits, the messages have to be aligned at the end of the RIF, which shifts the RMARKER by δ_{RSF2} to the right and increases the ToF compared to m_2 .

where n_{att2} is the number of *trailing* RIF bits that become unpredictable. This results in

$$\delta_{ToF2} = -\delta_c(T_{RSFS} + T_{RIFS} - \mathcal{T}(n_{att2})) \quad (17)$$

Since we consider $N_{RIF} = 1$, we can use $T_{RIFS} = n_{tot}T_{bit}$ and \mathcal{T} as defined in Eq. 11, which finally results in

$$\delta_{ToF} = \frac{\delta_c T_{bit}}{2} (n_{att1} + n_{att2} - n_{tot}) \quad (18)$$

Therefore, the attacker can reduce the distance if they choose to have $n_{att1} + n_{att2} > n_{tot}$ unpredictable RIF bits. Figure 11 illustrates how $n_{att1} = n_{att2} = n_{tot}/2$ pushes t_{replay} into the middle of the RIFs of both messages, which compensates for the delay δ_{RSF2} .

Since any further shifts would result in a distance reduction, we want to ensure that setting $n_{att1} + n_{att2} \geq n_{tot}$ makes it hard for the attacker to pass the RIF integrity checks⁸. Specifically, we want to ensure that choosing $n_{att1} = n_{att2} = n_{tot}/2$ is infeasible, since this case maximizes the attacker's success probability p_{succ} for any sum $n_{att1} + n_{att2} = n_{tot}$ (see Eq. 13 and Eq. 14).

Results: Figure 12 shows how a receiver can reduce the success probability by reducing the accepted bit-error rate. We see that the success probability is sensitive to small changes in BER_{max} . For $n_{tot} = 4096$, $BER_{max} = 0.22 \Rightarrow p_{succ} \leq 2^{-40}$, and $BER_{max} = 0.25 \Rightarrow p_{succ} = 2^{-2}$. The latter is because an attacker with $n_{att1} = n_{att2} = n_{tot}/2$ has a $p_{RIF} = 0.5$ of passing the RIF check for each message, resulting in $p_{succ} = 0.25$ for two messages.

We have shown that the clock drift consistency check protects SS-TWR against S&A. Since DS-TWR comprises two SS-TWR rounds, it is equally secure if the clock drifts of all involved messages are compared. In fact, it is even more challenging for an attacker to reduce the distance in DS-TWR because they have to pass the RIF integrity check of four messages instead of two.

⁸for $N_{RIF} > 1$ the equality already leads to a reduction, but any $n_{att1} + n_{att2} < n_{tot}$ does not. We refer readers to Appendix A.3 for more information.

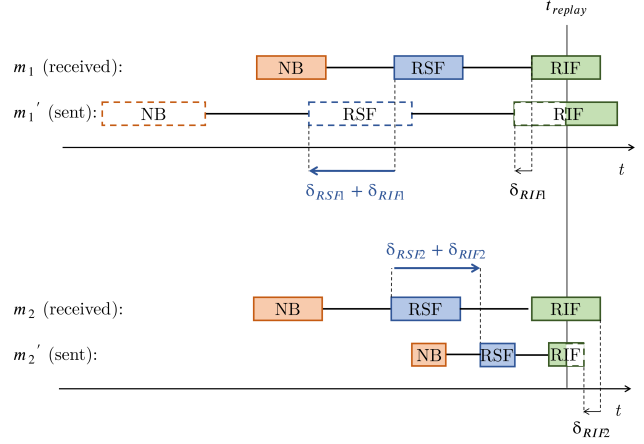


Figure 11: Despite the countermeasure, the attacker attempts to reduce the distance with the S&A attack. They stretch m'_1 and compress m'_2 , and advance both of them as shown. As a result, half of the RIF bits become unpredictable. However, the advancement of m'_1 and the delay of m'_2 (blue arrows) have the same value, *i.e.*, the attack does not change the computed ToF/distance. Further advancements would result in a distance reduction, and must be prevented by the RIF integrity checks.




Figure 12: Probability p_{succ} of passing the RIF integrity checks in both SS-TWR messages for $n_{att1} = n_{att2} = n_{tot}/2$. The probability depends on the overall number of bits all RIFs combined. $n_{tot} = 512$ corresponds to one RIF, $n_{tot} = 4096$ to 8 RIFs.

4.5.3 Further countermeasures and future work

Executing S&A with $\delta_{RIF} < 0$ inevitably results in an uneven distribution of bit errors between the predictable and unpredictable parts of the RIF (*i.e.* up to $BER = 0.0$ vs $BER \approx 0.5$ for an ideal attacker, respectively). Receivers could analyze this distribution as an additional hardening measure. However, this could result in a classification problem of distinguishing between malicious behavior and bad channels, and would require a thorough analysis. Another simple countermeasure worth evaluating is to move the RMARKER's position into the middle of the RIFs. This change would categorically prevent $\delta_{RSF} < 0$ and, thus, limit the maximum distance reduction

significantly. However, distance reductions due to $\delta_{RIF} < 0$ would still be possible and likely be subject to similar BER_{max} requirements like the countermeasure we have proposed.

5 Related Work

Distance bounding: Distance bounding protocols [9, 26, 45] establish an upper bound of the distance between a prover and a verifier under various threat models (*e.g.*, external attacker, honest/dishonest prover). Practical realizations of distance bounding protocols based on wireless ranging have been proposed (*e.g.*, based on UWB [39] or NFC [21]). In practice, secure ranging in the UWB standard [3, 4] assumes honest initiators and responders that protect ranging integrity using pseudo-random sequences in the ranging frames. A survey on distance bounding protocols and attacks can be found here [7].

Our work is orthogonal to the design of distance bounding protocols because it focuses on the manipulation of time references at the physical layer.

Physical-layer attacks: Guaranteeing the integrity of distance measurements against physical-layer attacks is not trivial. Systems based on signal strength or phase are insecure [15, 32]. Chirp-based systems are vulnerable to Early-Detect/Late-Commit attacks [38], which also affect Ultra Wide Band Impulse Radio (802.15.4a) [33]. In Cicada attacks [34, 35], the receiver mistakes the attacker's random pulses for a weak legitimate signal. The HRP mode of 802.15.4z is vulnerable to similar and improved attacks [27, 44].

Time-of-Arrival Integrity: Leu et al. introduced Message Time-of-Arrival Codes (MTACs) [28], a framework to formalize the security requirements of ToA measurements, taking into account physical layer attacks such as Early-Detect/Late-Commit (ED/LC). However, they do not consider the possibility of an attacker manipulating the time references of involved devices.

Clock imperfections: From a functional point of view, clock non-idealities in wireless ranging systems are well understood, and compensation techniques exist [13, 31]. The formulas described in 802.15.4z [3, 4] minimize the effect of clock drift for both SS-TWR and DS-TWR. The 802.15.4ab standard [23] will require highly accurate CFO estimates to support multi-millisecond frames.

Clulow et al. [11] propose an attack on ToF distance bounding in which a malicious prover leverages packet-level latencies to reply earlier to the challenge sent by the verifier. The authors also mention overclocking the prover to further shorten its reply time, relying on the tolerance of the receiver for clock drifts. In this attack, the malicious prover actually runs faster to shorten its reply time. In contrast, in MD, an external attacker convinces the initiator that the responder was slower, so that the initiator compensates for a larger reply time and computes a shorter distance. With clock drift compensation in place, overclocking the responder would have no effect.

In the context of OFDM-based ranging, Singh et al. [43] propose a "carrier frequency offset attack" for distance enlargement. The attacker overshadows the reference signal with a wrong carrier frequency offset, preventing decoding of legitimate data, and replays the legitimate signal with a delay, causing a distance enlargement.

Non-ideal clocks are a problem also in sensor networks where each device has a defined time slot to transmit. In this context, algorithms for secure time synchronization have been proposed in literature [16].

6 Conclusion

Because of their non-ideal clocks, different electronic devices have different notions of time and frequencies. Distance measurement based on the time of flight works under the implicit assumption that initiator and responder can deal with such differences through various forms of compensation. While these aspects are well known from a functional point of view, they have been largely overlooked in the security analysis of ranging. This is a topical issue in the definition of the upcoming 802.15.4ab standard, whose long frame duration exacerbates the problems of non-ideal clocks, but it is already relevant to the current 802.15.4z standard.

In this paper we provided an extensive analysis of the security impact of clock non-idealities and their compensation. We have presented two over-the-air attacks, Mix-Down and Stretch-and-Advance, which compromise the integrity of UWB measurements and lead to substantial distance reductions. This is a grave concern in security-critical applications, such as proximity-based access control. Our proposed countermeasures can defend against S&A, while MD does not have a straightforward solution.

Our insights prove that the implications of physical layer aspects on ranging integrity are still not completely understood, and should be analyzed in the design of new standards and protocols.

7 Acknowledgements

This research has received funding from the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40_180545.

References

- [1] Ultra Wideband availability. <https://support.apple.com/en-us/HT212274>.
- [2] Xiaomi will invest more in UWB. <https://xiaomiplanets.com/xiaomi-will-invest-more-in-uwbb/>.

- [3] IEEE Standard for Low-Rate Wireless Networks. *IEEE Std 802.15.4-2020 (Revision of IEEE Std 802.15.4-2015)*, pages 1–800, 2020.
- [4] IEEE Standard for Low-Rate Wireless Networks—Amendment 1: Enhanced Ultra Wideband (UWB) Physical Layers (PHYs) and Associated Ranging Techniques. *IEEE Std 802.15.4z-2020 (Amendment to IEEE Std 802.15.4-2020)*, pages 1–174, 2020.
- [5] 9To5Google. Pixel 6 Pro gets settings toggle to disable Ultra-Wideband (UWB). <https://9to5google.com/2022/01/02/pixel-6-pro-uw-b-setting/#jp-carousel-475456>. Accessed 2022-10-08.
- [6] Apple Inc. Find your keys, wallet, and more with AirTag. <https://support.apple.com/en-us/HT210967>. Accessed 2022-10-09.
- [7] Gildas Avoine, Muhammed Ali Bingöl, Ioana Boureanu, Srdjan Capkun, Gerhard P. Hancke, Süleyman Kardas, Chong Hee Kim, Cédric Lauradoux, Benjamin Martin, Jorge Munilla, Alberto Peinado, Kasper Bonne Rasmussen, Dave Singelee, Aslan Tchamkerten, Rolando Trujillo-Rasua, and Serge Vaudenay. Security of distance-bounding: A survey. *ACM Comput. Surv.*, 51(5):94:1–94:33, 2019.
- [8] BMW. BMW announces BMW Digital Key Plus with Ultra-Wideband technology coming to the BMW iX. <https://www.press.bmwgroup.com/global/article/detail/T0324128EN/bmw-announces-bmw-digital-key-plus-with-ultra-wideband-technology-coming-to-the-bmw-ix>. Accessed 2022-10-12.
- [9] Stefan Brands and David Chaum. Distance-bounding protocols (extended abstract). In Tor Hellesest, editor, *Advances in Cryptology - EUROCRYPT '93, Workshop on the Theory and Application of Cryptographic Techniques, Lofthus, Norway, May 23-27, 1993, Proceedings*, volume 765 of *Lecture Notes in Computer Science*, pages 344–359. Springer, 1993.
- [10] Car Connectivity Consortium. Digital Key Release. <https://carconnectivity.org/press-release/car-connectivity-consortium-publishes-digital-key-release-3-0/>. Accessed 2022-10-09.
- [11] Jolyon Clulow, Gerhard P. Hancke, Markus G. Kuhn, and Tyler Moore. So near and yet so far: Distance-bounding attacks in wireless networks. In Levente Buttyán, Virgil D. Gligor, and Dirk Westhoff, editors, *Security and Privacy in Ad-Hoc and Sensor Networks, Third European Workshop, ESAS 2006, Hamburg, Germany, September 20-21, 2006, Revised Selected Papers*, volume 4357 of *Lecture Notes in Computer Science*, pages 83–97. Springer, 2006.
- [12] David Barras, Boris Danev. Performance analysis of ranging integrity fragment (RIF) with distance commitment. <https://mentor.ieee.org/802.15/dcn/22/15-22-0413-00-04ab-performance-analysis-of-ranging-integrity-fragment-rif-with-distance-commitment.pptx>. Accessed 2022-10-11.
- [13] Igor Dotlic, Andrew Connell, and Michael McLaughlin. Ranging methods utilizing carrier frequency offset estimation. In *15th Workshop on Positioning, Navigation and Communications, WPNC 2018, Bremen, Germany, October 25-26, 2018*, pages 1–6. IEEE, 2018.
- [14] Ersen Ekrem, Apple inc. More on NBA-MMS. <https://mentor.ieee.org/802.15/dcn/21/15-21-0593-00-04ab-more-on-nba-mms.pptx>. Accessed 2022-10-11.
- [15] Aurélien Francillon, Boris Danev, and Srdjan Capkun. Relay attacks on passive keyless entry and start systems in modern cars. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2011, San Diego, California, USA, 6th February - 9th February 2011*. The Internet Society, 2011.
- [16] Saurabh Ganeriwal, Christina Pöpper, Srdjan Capkun, and Mani B. Srivastava. Secure time synchronization in sensor networks. *ACM Trans. Inf. Syst. Secur.*, 11(4):23:1–23:35, 2008.
- [17] GitHub user Amila Abeygunasekara. stm32-uw-range-finder. <https://github.com/amila93/stm32-uw-range-finder>. Accessed 2022-10-12.
- [18] GitHub user foldedtoad. Nordic nRF52-series + Decawave DWM3000 on Zephyr v2.5. <https://github.com/foldedtoad/dwm3000>. Accessed 2022-10-10.
- [19] GitHub user Woosang Yoo. Development of UWB ANCHOR using nRF52840 & DWM3000 for Hyundai E&C project. https://github.com/yws94/HDEC_ANC. Accessed 2022-10-12.
- [20] GitHub user Woosang Yoo. SS-TWR Multiranging. https://github.com/yws94/UWB_Multi-Ranging-DW3000. Accessed 2022-10-12.
- [21] Gerhard P. Hancke and Markus G. Kuhn. An RFID distance bounding protocol. In *First International Conference on Security and Privacy for Emerging Areas in Communications Networks, SecureComm 2005, Athens, Greece, 5-9 September, 2005*, pages 67–73. IEEE, 2005.

- [22] Hyundai. Introducing Hyundai Digital Key. <https://owners.hyundaiusa.com/us/en/resources/technology-and-navigation/introducing-all-new-digital-key.html>. Accessed 2022-10-08.
- [23] IEEE. IEEE 802.15 Documents Group TG4AB. https://mentor.ieee.org/802.15/documents?is_dcn=NBA&is_group=04ab. Accessed 2022-10-11.
- [24] Jaguar UK. Keyless Entry. <https://www.jaguar.co.uk/keylessentry/index.html>. Accessed 2022-10-11.
- [25] Jinjing Jiang et. al. One-to-Many Ranging using MMR Technical Framework Proposal. <https://mentor.ieee.org/802.15/dcn/23/15-23-0015-04-04ab-one-to-many-ranging-using-mmr-technical-framework-proposal.docx>. Accessed 2023-03-15.
- [26] Marc Kuhn, Heinrich Luecken, and Nils Ole Tippenhauer. UWB impulse radio based distance bounding. In *7th Workshop on Positioning Navigation and Communication, WPNC 2010, Dresden Germany, 11-12 March 2010, Proceedings*, pages 28–37. IEEE, 2010.
- [27] Patrick Leu, Giovanni Camurati, Alexander Heinrich, Marc Roeschlin, Claudio Anliker, Matthias Hollick, Srdjan Capkun, and Jiska Classen. Ghost peak: Practical distance reduction attacks against HRP UWB ranging. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1343–1359, 2022.
- [28] Patrick Leu, Mridula Singh, Marc Roeschlin, Kenneth G. Paterson, and Srdjan Capkun. Message time of arrival codes: A fundamental primitive for secure distance measurement. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*, pages 500–516. IEEE, 2020.
- [29] Marius Muench, Dario Nisi, Aurelien Francillon, and Davide Balzarotti. Avatar²: A Multi-target Orchestration Platform. In *Workshop on Binary Analysis Research (colocated with NDSS Symposium)*, BAR 18, February 2018.
- [30] Aravind Nagulu, Ahmed Mekawy, Mykhailo Tymchenko, Dimitrios Sounas, Andrea Alù, and Harish Krishnaswamy. Ultra-wideband switched-capacitor delays and circulators—theory and implementation. *IEEE Journal of Solid-State Circuits*, 56(5):1412–1424, 2021.
- [31] Dries Neiryneck, Eric Luk, and Michael McLaughlin. An alternative double-sided two-way ranging method. In *13th Workshop on Positioning, Navigation and Communications, WPNC 2016, Bremen, Germany, October 19-20, 2016*, pages 1–4. IEEE, 2016.
- [32] Hildur Ólafsdóttir, Aanjhan Ranganathan, and Srdjan Capkun. On the security of carrier phase-based ranging. In Wieland Fischer and Naofumi Homma, editors, *Cryptographic Hardware and Embedded Systems - CHES 2017 - 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings*, volume 10529 of *Lecture Notes in Computer Science*, pages 490–509. Springer, 2017.
- [33] Marcin Poturalski, Manuel Flury, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. Distance bounding with IEEE 802.15.4a: Attacks and countermeasures. *IEEE Trans. Wirel. Commun.*, 10(4):1334–1344, 2011.
- [34] Marcin Poturalski, Manuel Flury, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. On secure and precise IR-UWB ranging. *IEEE Trans. Wirel. Commun.*, 11(3):1087–1099, 2012.
- [35] Marcin Poturalski, Manuel Flury, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. The cicada attack: Degradation and denial of service in ir ranging. In *2010 IEEE International Conference on Ultra-Wideband*, volume 2, pages 1–4, 2010.
- [36] Qorvo Inc. DWM3000EVB - Qorvo. <https://www.qorvo.com/products/p/DWM3000EVB>. Accessed 2022-10-10.
- [37] Qorvo Inc. MDEK1001 - Qorvo. <https://www.qorvo.com/products/p/MDEK1001>. Accessed 2022-10-11.
- [38] Aanjhan Ranganathan, Boris Danev, Aurélien Francillon, and Srdjan Capkun. Physical-layer attacks on chirp-based ranging systems. In Marwan Krunz, Loukas Lazos, Roberto Di Pietro, and Wade Trappe, editors, *Proceedings of the Fifth ACM Conference on Security and Privacy in Wireless and Mobile Networks, WISEC 2012, Tucson, AZ, USA, April 16-18, 2012*, pages 15–26. ACM, 2012.
- [39] Kasper Bonne Rasmussen and Srdjan Capkun. Realization of RF distance bounding. In *19th USENIX Security Symposium, Washington, DC, USA, August 11-13, 2010, Proceedings*, pages 389–402. USENIX Association, 2010.
- [40] Samsung Electronics Co., Ltd. . Introducing the New Galaxy SmartTag+: The Smart Way to Find Lost Items. <https://news.samsung.com/us/introducing-the-new-galaxy-smarttag-plus/>. Accessed 2021-10-12.
- [41] Samsung Electronics Co., Ltd. . Unlock a New Experience: Galaxy Users Can Now Use Secure Digital Key

With the Genesis GV60. <https://news.samsung.com/global/unlock-a-new-experience-galaxy-users-can-now-use-secure-digital-key-with-the-genesis-gv60>. Accessed 2022-10-09.

- [42] Juri Sidorenko, Volker Schatz, Norbert Scherer-Negenborn, Michael Arens, and Urs Hugentobler. Decawave UWB clock drift correction and power self-calibration. *Sensors*, 19(13):2942, 2019.
- [43] Mridula Singh, Marc Roeschlin, Aanjhan Ranganathan, and Srdjan Capkun. V-range: Enabling secure ranging in 5g wireless networks. In *29th Annual Network and Distributed System Security Symposium, NDSS 2022, San Diego, California, USA, April 24-28, 2022*. The Internet Society, 2022.
- [44] Mridula Singh, Marc Roeschlin, Ezzat Zalzal, Patrick Leu, and Srdjan Capkun. Security analysis of IEEE 802.15.4z/hrp UWB time-of-flight distance measurement. In Christina Pöpper, Mathy Vanhoef, Lejla Batina, and René Mayrhofer, editors, *WiSec '21: 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks, Abu Dhabi, United Arab Emirates, 28 June - 2 July, 2021*, pages 227–237. ACM, 2021.
- [45] Nils Ole Tippenhauer, Heinrich Luecken, Marc Kuhn, and Srdjan Capkun. UWB rapid-bit-exchange system for distance bounding. In *Proceedings of the 8th ACM Conference on Security & Privacy in Wireless and Mobile Networks, New York, NY, USA, June 22-26, 2015*, pages 2:1–2:12. ACM, 2015.
- [46] Mykhailo Tymchenko, Dimitrios Sounas, Aravind Nagulu, Harish Krishnaswamy, and Andrea Alù. Quasielectrostatic wave propagation beyond the delay-bandwidth limit in switched networks. *Phys. Rev. X*, 9:031015, Jul 2019.
- [47] Xiliang Luo et. al. NBA-UWB Technical Framework. <https://mentor.ieee.org/802.15/dcn/23/15-23-0100-02-04ab-nba-uwbt-technical-framework-for-draft0.docx>. Accessed 2023-04-17.

A Appendix

A.1 Over-the-Air Estimation of Clock Drifts

We show that an attacker can estimate clock drifts over the air, using only the messages it receives from other devices, and leverage this information to fine-tune its own clock drift. In particular, this would enable smooth takeover attacks in which $\delta_c = c' - c = c_{att} - c_{resp}$ is initially set to ≈ 0 ppm and then gradually decreased.

The attacker estimates the natural drift $c = c_{resp} - c_{init}$ between initiator and responder by simply sniffing their messages with a receiver, and computing

$$c = c_{resp} - c_{init} = (c_{sniffer} - c_{init}) - (c_{sniffer} - c_{resp}) \quad (19)$$

where $c_{sniffer} - c_{init}$ ($c_{sniffer} - c_{resp}$) is the drift estimated by the sniffer when receiving a message from the initiator (responder). The location of the sniffer with respect to the other two devices does not affect the computation as the arrival time is not used. An additional rolling median filter over a sliding window can be applied to filter noise and outliers.

Assuming that the initiator is well calibrated and has a negligible drift, the attacker can choose $c_{att} = c \approx c_{resp}$. If this is not the case, the attacker could use the same method to estimate $c' = c_{att} - c_{resp}$, by placing a second sniffer directly at the output of its mixing stages. The attacker could then start with a low power, adjust the mixing frequency until $c' = c$ (i.e., $c_{att} = c_{resp}$) and then start increase the power to takeover. The advantage of this method is that it is a self calibration of the attacker to the responder over-the-air.

To demonstrate the ability of the attacker to estimate clock drifts over-the-air with good accuracy, we have implemented a sniffer with a QorvoDWM3000 transceiver, and we have evaluated it in a scenario without physical (the same as experiment 4 described in Section 3.4).

Figure 13 shows how the value estimated by the attacker over the air closely follows the value actually estimated by the initiator. The larger variations at around 200 s - 600 s and 800 s - 400 s were caused by a sudden increase in the temperature of the initiator that we induced with an off-the-shelf hair-dryer. The sniffer was able to follow the variations.

A.2 Implementation of S&A

In contrast to MD, S&A requires signals to be stretched in time. This has to be done differently for predictable and unpredictable fragments, but the two approaches can be combined into one attacker device.

RSFs and NB: The RSF and NB fragments can be generated and stretched using an off-the-shelf UWB transceiver. The stretching can be achieved by trimming the transceiver's clock in hardware or in the firmware.⁹

⁹For the current standard, the Qorvo DMW3000EVB offers clock trimming from software, and we expect future devices to have similar features.




Figure 13: The MD attacker uses a sniffer to estimate the natural clock drift between initiator and responder. The estimate of the sniffer (orange) follows closely the one of the initiator (blue) shown as ground truth.

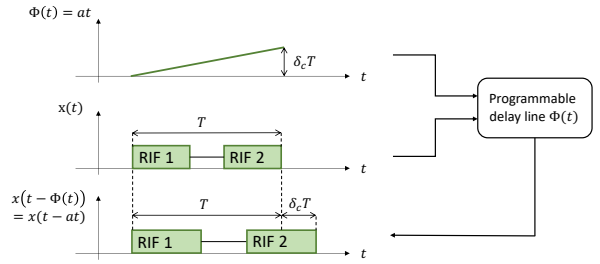


Figure 14: Proposed architecture to replay the unpredictable RIF bits with low latency in the S&A attack: the attacker stretches a signal in real-time by applying a linearly increasing delay with a programmable delay line (e.g., based on [30,46]).

RIFs: Since they are unpredictable, the attacker has to receive the RIF pulses and retransmit them with increasing delay. This could be implemented with programmable delay lines, based on UWB switched-capacitor delay elements [30,46]. Figure 14 shows how a delay line works: an analog circuit applies the delay $\Phi(t)$ to an input signal. Although $\Phi(t)$ is increased in small digital steps, the resolution is high enough to assume it to be linear. The authors have built a working prototype that delays 500MHz UWB signals several dozens of nanoseconds, which is the order of magnitude required for the RIFs.

Unpredictable NB: If the NB fragment contains unpredictable data (e.g., cryptographic nonces), the previous attack strategy has to be revised. Here, the PHY layer properties of the NB fragment can be helpful: according to the standard draft, the NB channels only have a bandwidth of 2.5MHz, which results in a pulse length of $\approx 1\mu$ s. The attacker could exploit this twofold:

First, note that we exaggerated the effects of stretching and advancing in our figures. They would be imperceptible if shown to scale: stretching a 17ms-long ranging message with $\delta_c = -40$ ppm only increases its length by $0.000'68$ ms. This is the maximum by which the NB fragment has to be advanced throughout all configurations. We expect that, after receiving

the NB fragment, the receiver accepts the first RSF to arrive a few hundreds of nanoseconds earlier or later, since its long waveforms make it difficult to derive an exact ToA (which is the reason UWB is used in the first place). The attacker might exploit this tolerance in order to handle NB like the RIFs: instead of sending an advanced NB, they can stretch and replay the one sent by the genuine transmitter. Then, the attacker transmits all the following fragments as if the NB had been advanced, which results in a shorter pause between the NB and the first RSF. The resulting message should still be processed correctly by the attacked receiver and result in a successful attack. In the worst case, the attacker would have to reduce the drift δ_c to be within the accepted margins.

Second, long symbols can be advanced using the ED/LC attack [11]. Without knowing the symbol, an attacker could send an arbitrary waveform to trigger an earlier reception event in the receiver. As soon as learn the symbol from the actual NB fragment, they can use their power advantage to set the advanced NB symbol to the correct value.

A.3 Analysis of S&A

A.3.1 $\mathcal{T}(n_{att})$ for multiple RIFs

For message formats with arbitrary RIFs, we can write:

$$\mathcal{T}(n_{att}) = \underbrace{\lfloor n_{att}/n_{RIF} \rfloor \cdot 1\text{ms}}_{\text{entire RIFs}} + \underbrace{(n_{att} \bmod n_{RIF}) \cdot T_{bit}}_{\text{remaining bits}} \quad (20)$$

Here, n_{att} is the number of unpredictable bits, n_{RIF} the number of bits in a RIF, and T_{bit} the duration of a single bit. The two contributions are a result of the pauses between the RIFs: as long as n_{att} is less than n_{RIF} , it grows linearly with n_{att} . However, for $n_{att} \geq n_{RIF}$ the attacker accepts that an entire RIF becomes unpredictable, which shifts t_{replay} (see Figure 8) to the end of the corresponding fragment. In this case, the attacker may just as well increase δ_{RIF} further and move t_{replay} to the start of the *next* RIF, since the pause between them does not contain any bits. These "free" advancements dominate δ_{RIF} , since the duration of a RIF ($\approx 16\mu\text{s}$) is small compared to the pause ($\approx 1\text{ms}$).

A.3.2 Countermeasure for multiple RIFs

We assume an attacker who aims at optimizing the probability p_{succ} of achieving any distance reduction $\delta_{ToF} < 0$. For $N_{RIF} \in \{2, 4, 8\}$, we have to use \mathcal{T} as defined in Eq. 20. If the messages fulfill the consistency requirement regarding the clock drifts, the smallest sum $n_{att1} + n_{att2}$ that results in a distance reduction $\delta_{ToF} < 0$ is n_{tot} , the total number of RIF bits in one ranging message:

$$\underset{n_{att1} + n_{att2} \in (0, 2n_{tot})}{\text{argmin}} (\delta_{ToF} < 0) = n_{tot} \quad (21)$$

The maximum ToF reduction for such n_{att1} and n_{att2} is

$$\min_{n_{att1} + n_{att2} = n_{tot}} \delta_{ToF} = \delta_c T_{pause} \quad (22)$$

Where $T_{pause} \approx 984\mu\text{s}$ corresponds to the pause between fragments. Because of \mathcal{T} , this result only holds when both n_{atti} are an integer multiple of the bits in a single RIF, *i.e.*, $n_{att1} = x \cdot n_{RIF}$, $n_{att2} = (N_{RIF} - x)n_{RIF}$. The reason are the inter-RIF pauses: for these values, the attacker can exploit the trailing pause in ranging messages. This results in an overall reduction despite the enlargement caused by δ_{RSF} in Eq. 17. In contrast, any other $n_{att1} + n_{att2} = n_{tot}$ only compensate for the enlargement, resulting in $\delta_{ToF} = 0$.

In the last step, we show that there is an x for which p_{suc} is maximal, which is the optimal solution we are interested in:

$$\begin{aligned} \underset{x \in \{1, \dots, N_{RIF}\}}{\text{argmin}} (p_{succ} = p_{RIF1} p_{RIF2}) &= \frac{N_{RIF}}{2} \\ \Rightarrow n_{att1} = n_{att2} &= \frac{n_{tot}}{2} \end{aligned} \quad (23)$$

Choosing any $n_{atti} > n_{tot}/2$ increases the ToF reduction, but at the cost of a lower probability. Therefore, $n_{atti} = n_{tot}/2$ maximizes the probability of any distance reduction. We conclude that setting BER_{max} such that p_{suc} is negligible prevents the S&A attack.