

Ghost Peak: Practical Distance Reduction Attacks Against HRP UWB Ranging

Patrick Leu^{1,*}, Giovanni Camurati^{1,*}, Alexander Heinrich², Marc Roeschlin¹, Claudio Anliker¹, Matthias Hollick², Srdjan Capkun¹, and Jiska Classen²

¹ETH Zurich

²TU Darmstadt

*Authors contributed equally to this research

Abstract

We present the first over-the-air attack on IEEE 802.15.4z High-Rate Pulse Repetition Frequency (HRP) Ultra-Wide Band (UWB) distance measurement systems. Specifically, we demonstrate a practical distance reduction attack against pairs of Apple U1 chips (embedded in iPhones and AirTags), as well as against U1 chips inter-operating with NXP and Qorvo UWB chips. These chips have been deployed in a wide range of phones and cars to secure car entry and start and are projected for secure contactless payments, home locks, and contact tracing systems. Our attack operates without any knowledge of cryptographic material, results in distance reductions from 12 m (actual distance) to 0 m (spoofed distance) with attack success probabilities of up to 4%, and requires only an inexpensive (USD 65) off-the-shelf device. Access control can only tolerate sub-second latencies to not inconvenience the user, leaving little margin to perform time-consuming verifications. These distance reductions bring into question the use of UWB HRP in security-critical applications.

1 Introduction

Ultra-Wide Band chips that measure distance are being massively deployed in smartphones, cars, and other products [5, 32, 52]. Applications range from entry and start systems in cars to mobile payments, contact tracing, spatial awareness, and indoor localization. In addition to enhanced precision compared to more traditional signal strength based ranging, UWB aims to provide security against relay and distance reduction attacks [23], which have been used in practice for car thefts and attacks on contactless payments [16, 33, 59].

The recently adopted IEEE 802.15.4z standard [4] aims to address known distance reduction attacks. It introduces two ranging modes: Low-Rate Pulse Repetition Frequency (LRP) and High-Rate Pulse Repetition Frequency (HRP). Although both modes are used in automotive applications, primarily for Passive Keyless Entry and Start (PKES) systems [5, 11, 17, 56], HRP has seen adoption in Apple iPhones and AirTags, as well

as Samsung phones and SmartTags [10, 51, 55]. Despite its standardization and deployment, no public example implementations or standardized algorithms for security-relevant functionality exist. IEEE 802.15.4z focuses on message formats without mandating in detail how ranging is done and protected at the endpoints.

This paper demonstrates the first practical over-the-air distance reduction attack against the UWB IEEE 802.15.4z HRP mode. Even though HRP security has been recently studied, these studies were done in simulations [58]. We refine existing attacks, introduce a new one, and demonstrate their feasibility in practical settings with Apple U1 (iPhone/AirTag/HomePod), NXP Trimension SR040/SR150, and Qorvo DWM3000 chips. Our attack enabled a successful distance reduction of up to 12 m with an overall success rate of 4%, which is higher than what is generally accepted for relevant applications. Typically, false acceptance rates are $1/2^{20}$ for gate access control and $1/2^{48}$ for mobile payments, such that it would take days to years until a fake measurement gets accepted.

Manufacturers advertise some of the evaluated chips as secure ranging capable [38]. We performed our tests using the configurations that are openly accessible on these chips. Since security algorithms and parameters are not public in the chips that we tested (Apple, NXP, Qorvo), it is hard to determine if these systems can be configured differently and if these alternative configurations would be vulnerable to our or other attacks. Additionally, the past has shown that undisclosed wireless protocols can signify security-by-obscurity solutions [28, 60]. Prior work [58] further suggests that making HRP ranging both secure and reliable is likely hard.

The deployment and use of UWB will presumably increase in the future. The FiRa consortium [18] has been founded to contribute to the development and widespread adoption of UWB technologies in the context of *secured fine ranging and positioning*. The Car Connectivity Consortium recently published Digital Key Release 3.0, enabling PKES via UWB in combination with Bluetooth Low Energy [13]. At least one car manufacturer has already announced that it will support the iPhone as an access token for PKES, citing UWB as a

ranging mechanism [11]. Since UWB as an access system is a new protocol, it might take time until malicious actors can fully understand and bypass security checks [61]. However, systems in cars and other areas related to access control have to be secure for decades after initial deployment. Therefore, we see this work as another step towards a better understanding of the security of UWB HRP.

In summary, we make the following contributions:

- We introduce the first practical distance reduction attack on IEEE 802.15.4z HRP. This amendment defines cryptographically generated high-rate pulse sequences for Time of Arrival (ToA) measurement, whose unpredictability is supposed to prevent distance reduction by preventing the attacker from transmitting valid signals earlier than the victim. Our attack operates in a black-box manner and assumes neither knowledge of cryptographic material shared between the attacked devices nor access to (randomized) ranging message content before messages are transmitted. This attack not only validates observations from simulation-based studies of HRP but also introduces a novel attack dimension—it selectively varies the power of the injected packet per packet field. The power level is independently adjusted for different fields so that the injected signal is neither perceived as an additional packet nor as jamming the legitimate one. Our attack can therefore also be seen as a type of selective overshadowing.
- We implement our attack on inexpensive (USD 65), commercial off-the-shelf components and demonstrate it on Apple iPhones and AirTags (U1 chip) and on iPhones interoperating with NXP SR040/SR150 and Qorvo DWM3000 UWB chips. We evaluate our attack through a series of experiments and show that the attacker can reduce the measured distances from 12 m to 0 m (measured distance). During normal execution, the measurement error is between 10 cm and 20 cm. With a success rate as high as 4 %, our attack suffices to deceive ranging systems that rely on single HRP measurements.
- We discuss the implications of our results to different applications and use cases and the applicability of different mitigation techniques in practical settings.
- We responsibly disclosed our findings to Apple, and NXP, and are in the process of disclosing to Qorvo.

The rest of the paper is organized as follows. In [Section 2](#), we provide background on UWB secure distance measurements. In [Section 3](#), we present our attack. We discuss our experimental results in [Section 4](#). Finally, we reflect on the security of HRP UWB in [Section 5](#) and compare it to related work in [Section 6](#) before concluding in [Section 7](#).

2 Background

In this chapter, we provide the necessary background on Time-of-Flight (ToF) HRP UWB. We first introduce the concept of ToF ranging and show how HRP uses cross-correlation to determine the ToF before explaining security considerations behind HRP. Finally, we provide a brief overview of available HRP chips and products.

2.1 UWB Secure Ranging

The simplicity and practicality of relay attacks on PKES systems [16, 23, 59] urged a paradigm shift in secure ranging. Utilizing a signal’s ToF is promising since a relay can only increase the ToF and, thus, the measured distance. However, research in this field has shown that such systems can still be vulnerable to more sophisticated attacks, such as Cicada [43] or Early Detect/Late Commit (ED/LC) [41].

UWB aims to implement secure ranging, including physical-layer security [19]. IEEE 802.15.4 proposes two modes for UWB ranging named LRP and HRP. They are both subject to stringent power limitations, as their channels overlap with frequency bands used by existing technologies, such as Wi-Fi or cellular networks. While LRP approaches the power limit by using fewer but stronger pulses (each individually ‘visible’ to the receiver), HRP relies on a larger number of weaker pulses (which cannot be individually decoded in most environments by the receiver). This difference in design has consequences; while the security of LRP is easy to demonstrate, the resilience of HRP against reduction attacks is an open research question. Recent in-simulation analysis has shown that HRP might be hard to configure to be both performant and secure [58].

2.1.1 Two-Way Ranging

The IEEE 802.15.4 standard defines three different ranging and localization methods, namely Single-Sided Two-Way Ranging (SS-TWR), Double-Sided Two-Way Ranging (DS-TWR), and Time Difference of Arrival (TDOA); our work focuses on SS-TWR and DS-TWR.

SS-TWR is depicted in [Figure 1a](#), which shows how ToF for the distance calculation can be determined by subtracting T_{reply} , the processing time of the responder, from T_{round} , the total round trip time measured by the initiator. Dividing the result by two yields an estimation of the propagation delay \hat{T}_{prop} , or the ToF required by the signal to cover one way. However, this result may be affected by a possible clock frequency offset between initiator and responder. If the initiator can measure this offset, it can compensate for it and improve the measurement.

DS-TWR, as shown in [Figure 1b](#), mitigates the clock offset by transmitting more messages. DS-TWR comprises two SS-TWR exchanges in opposite directions. T_{reply} and T_{round}

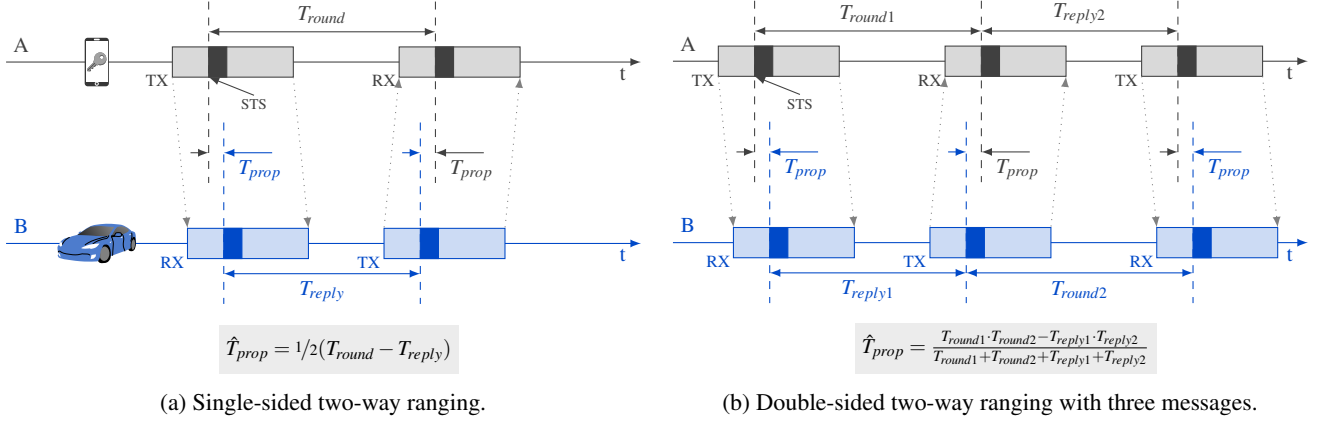


Figure 1: The principle of two-way ranging [4].

are measured with both devices/clocks, significantly reducing errors induced by clock offset and drift. DS-TWR is optimized by simultaneously using the response message of the first exchange as the request message of the second, thus reducing the procedure to three ranging messages. The derivation of the propagation time formula can be found in [36].

2.1.2 Receiver Design and Cross-Correlation

Most RF communication technologies rely on cross-correlation to detect the presence of an incoming message. In UWB, the receiver constantly scans the acquired signal for a static (pre-negotiated) preamble using a local template. The received signal is digitized and recorded as I/Q samples fed into a correlator. If the output exceeds the level for noise by a certain amount, the receiver concludes that a packet must be present and analyzes the signal further.

In practice, this process has to be optimized to cope with channel distortions, most notably multi-path fading. During transit, objects in the vicinity reflect the signal, which creates copies of the signal that are slightly delayed in time, as shown in Figure 2. Those copies are superimposed onto the original signal, causing constructive or destructive interference. There-




Figure 2: In a Non-Line-of-Sight (NLoS) scenario, the receiver needs to detect the arrival time of the early Line-of-Sight (LoS) copy (leading edge).

fore, a momentary output of the correlation is non-conclusive, and instead, the Channel Impulse Response (CIR), i.e., the correlation output over time, must be inspected. The CIR greatly supports the search for a known template in the received signal and can determine the precise arrival time of a packet. The CIR can be estimated as follows:

$$\text{CIR}[t] = (g_{\text{loc}} * s)[t] = \sum_{m=0}^{|g_{\text{loc}}|-1} g_{\text{loc}}[m] \cdot s[m+t]$$

where $s[\cdot]$ is the complex and time-discrete received signal, $g_{\text{loc}}[\cdot]$ is the template of the expected signal, and $*$ denotes cross-correlation.

As shown in Figure 3, HRP UWB ranging relies heavily on cross-correlation, to detect and determine the arrival times of preamble and a Scrambled Timestamp Sequence (STS). We explain STS in the next section. A UWB receiver cross-correlates the incoming signal with a template (e.g., a known sequence) for the preamble and, if present, also with a known template for the STS. High correlation values imply similarities between the template and the received signal. However, the CIR only shows a single distinct peak in perfect conditions. Due to multi-path, the CIR often shows a profile containing several peaks, and it is not straightforward to identify the *first* peak/path that reflects the actual physical distance. Construc-



Figure 3: The CIR is calculated based on the received signal and a local template of the expected signal. A correlation peak indicates high similarity. However, with multi-path effects, there are multiple peaks. The receiver identifies the first LoS path, e.g., by searching back in time from the strongest peak.

tive and destructive interference can lead to a CIR where the first path emerges as a peak with an amplitude significantly below the maximal value. To be precise, even in absence of multi-path, an additional source of noise in the CIR is the non-ideal (auto)correlation of the STS. Ideally, we would expect a CIR peak only when incoming STS and local template are perfectly aligned in time. In practice, if the two copies are shifted by a multiple of the pulse rate, they might still exhibit some similarity (some of the bits will randomly be the same), causing additional noise in the form of (significantly) smaller side lobe peaks. Channel and receiver noise make the search for the first path and thus the correct distance even more challenging.

2.1.3 High-Rate Pulse Repetition (HRP)

HRP mode of IEEE 802.15.4 uses a high pulse repetition frequency of 64 MHz. The spacing between pulses is narrow and, to meet stringent restrictions on power spectral density (-41.3 dBm/MHz) [1], the power per pulse is low, in the order of -80 , instantaneous dBm (at the antenna port). The information elements of a packet are either encoded with Burst-Position Modulation (BPM) using Binary Phase Shift Keying (BPSK) or just BPSK symbols. In BPM-BPSK, a symbol can encode two bits by varying the position of the burst and the polarity of the pulses, while in BPSK, a positive polarity pulse encodes a bit of value zero, and a negative polarity pulse (180° phase shift) encodes a bit of value one. Most UWB channels are 499.2 MHz wide, which is the bandwidth used by all our tested devices. At 499.2 MHz, the duration of a pulse is in the order of 2 ns.

HRP PHY Packet It is essential for the attacks described in this paper to understand the HRP packet construction and pulse sequence. Figure 4 shows the different segments that constitute an HRP ranging message using an STS. The packet preamble is used to detect the presence of a ranging message. The STS contains a cryptographically-secure pseudo-random bit sequence for security purposes, and the data segment may be used to transmit additional information. The Start-of-Frame Delimiter (SFD) should be taken as a reference to calculate the propagation delay, and the PHR carries the physical header of the packet. We refer interested readers to the official release of the IEEE 802.15.4 standard for a more detailed description of the PHY [2].

Scrambled Timestamp Sequence (STS) The preamble is a pre-defined and static sequence of pulses representing -1 , 0 , and 1 , modulated using a ternary code, i.e., positive, negative and no pulse. In contrast, the STS consists of BPSK-



Figure 4: Example format of an HRP packet [4]. The field lengths and their order depend on the configuration.

modulated pulses representing -1 and 1 . The bit sequence in the STS is the output of a pseudo-random generator and derived as outlined in Figure 5. The ranging devices need to agree on a 128-bit key, e.g., by using an out-of-band channel, before the UWB ranging operation can commence. A fresh STS is generated for every ranging message, as the *STS V Counter* increases with every packet. The ranging devices know the expected STS bit sequence in advance, and they can create a local template to detect the incoming STS using cross-correlation as described in Section 2.1.2. Since the STS contents cannot be predicted, it is theoretically impossible for an external source to emit a signal that arrives at the targeted device earlier in time and still contains the legitimate STS; only the legitimate device knows which data/signal to send. As a result, ranging devices can base the ToA of the packet on the arrival time of the STS and thereby guarantee that no external adversary reduces the measured distance by advancing the received signal in time. Moreover, it is also impossible to react to isolated pulses and send those earlier in time since the BPSK pulses as part of the STS are only about 2 ns long. An adversary cannot acquire the polarity of a single pulse in sufficient time, which makes any replay or ED/LC attack physically impossible. In any case, advancing pulses would only yield a 2 ns reduction at the maximum, translating to less than 60 cm in distance.

Channel Distortion and Multi-Path Fading UWB ranging packets are subject to channel noise and multi-path fading, rendering the (direct) demodulation of single pulses of the STS intricate and in some channel conditions impossible. At a 64 MHz pulse repetition frequency, the pulse spacing is in the order of 16 ns, which is shorter than the typical channel delay spread. As a consequence, inter-pulse interference and multi-path fading effects make separate pulses unrecognizable. To work around this, HRP detects STS by cross-correlating the received signal with the expected STS, similar to preamble detection. Although cross-correlation is a powerful tool to determine the presence of the STS, the computed Channel Impulse Response (CIR) often shows a profile that contains multiple correlation peaks, and pinpointing the exact arrival time remains challenging. The CIR is a superposition of cross-correlation side peaks and weak early path correlation peaks.




Figure 5: Cryptographically secure STS generation with AES in counter mode. Each iteration results in a random 128-bit block. The *STS V Counter* is incremented for every iteration. The entire STS comprises 32 blocks, or 4096 bits [4].

Figure 3 shows two pulses after reception (in red) and the template used by the receiver (in grey). The resulting CIR (in blue) exhibits multiple peaks. The highest peak does not necessarily correspond to the LoS path of the signal. Even before the strongest correlation value, any HRP receiver must check for additional peaks within a specific time window. Such a peak might suggest an earlier but weaker copy of the signal, which belongs to a shorter path. By using this path as a reference, the receiver can compute a more accurate ranging result. Details on how the time of arrival of the STS is determined are not specified in the standard for HRP. At the time of writing, the exact procedure remains protected intellectual property for all commercially available HRP transceiver chips we have evaluated.

2.1.4 Ideal versus Real Security Guarantees

If every pulse contained in the STS would be demodulated, absent of noise and channel effects, the receiver could verify every single bit in the sequence. However, in HRP UWB, the 4096 bit long STS does not result in 4096 verifiable bits. First, the entropy of the key used for AES in counter mode is only 128 bits, see Figure 5. Second, since the STS is verified by correlation instead of single pulse demodulation, any security guarantee is given by the significance level of the early peak compared to the overall cross-correlation profile. Non-ideal cross-correlation properties of random sequences, such as the STS, can cause side-lobes in the correlation and play a minor role.

A bit-wise STS comparison instead of cross-correlation would have to allow for transmission errors, which naturally happen in NLoS scenarios. IEEE 802.15.4z does not specify whether the STS should be compared bit-wise after the correlation operation. Even if a vendor implements additional checks, they need to account for bit flips and choose a threshold that significantly impacts on the security provided by the STS.

2.2 Commercial HRP UWB Chips

As of now, only a few vendors offer HRP transceiver chips, despite the fact that HRP-based location and tracking tags have entered the consumer market at scale [32] and automotive manufacturers are planning to release cars featuring PKES systems built on top of HRP chips, such as the BMW iX and the Genesis GV60 models [11, 52]. The FiRa consortium considers HRP viable for both consumer-grade and security-critical applications alike [19].

Apple has a diverse UWB software and hardware stack. Different versions of the Apple U1 chip have been released in recent products, such as the iPhone (since iPhone 11), the HomePod mini, the Apple Watch (since Series 6), and even the USD 30 AirTag. On the iPhone, Apple integrated UWB into AirDrop with iOS 13 [40], using Angle of Arrival (AoA)

measurements to simplify the location of devices and enhance user experience. With iOS 14, they introduced the Nearby Interaction framework, exposing a selected set of UWB-based ranging functionality to application developers [30]. A compatibility mode for third-party accessory support has been available since the release of iOS 15 [31]. However, details about the compatibility mode configuration parameters are only available to Made for Apple (MFi) program members.

NXP advertises their Trimension chip series for secure ranging and precise positioning [38]. Development kits exist for the SR150 and SR040 [56]. Our analysis showed that several Samsung products, for example, the SmartTag+ and phones starting from Samsung Note20 Ultra [55], contain NXP chips to enable ranging and improve Point to Share [50] data transfers. Examples for cars that comprise NXP chips are upcoming BMW and VW models [53, 54], whereas VW seems to incorporate LRP chips for PKES use cases [5].

Qorvo, also known as Decawave before their acquisition [47], manufactures the DW3000 chip series. These chips are interoperable with the Apple U1 chip [44]. Nevertheless, to the best of our knowledge, there are no commercially available products that use the DW3000 series and are compatible with Samsung or Apple consumer devices. Qorvo also offers two development kits: DWM3000EVB, an Arduino-based development board [45], and DWM3001CDK, an integrated board that contains an nRF52833 with Bluetooth 5.2 [46].

3 A Practical Distance-Reduction Attack

In the following, we explain our attacker model, the theoretical working principle of our attack, including boundaries of distance reduction, and the attack algorithm and setup.

3.1 Attacker Model and Attack Overview

We consider an attacker that is trying to reduce the distance measured between two HRP UWB devices.¹ E.g., an attacker trying to unlock and start a car by tricking it into believing that the legitimate owner’s car keyfob is near. Even a distance reduction in the order of a few meters can have a severe impact, e.g., if the car is parked in front of the legitimate owner’s house.

We consider a *black-box* attacker with the following limitations. The attacker has no access to any secrets shared between victim devices and cannot predict message field contents that are assumed to be unpredictable in HRP UWB; i.e., the attacker cannot predict the Scrambled Timestamp Sequence (STS). Unable to guess the STS, our attacker cannot simply send a valid packet to advance the message time of arrival and therefore reduce the distance. The attacker can

¹We do not focus on bearing, which is not covered by IEEE 802.15.4z, is not protected in current implementations, and would likely be vulnerable to other physical layer attacks.

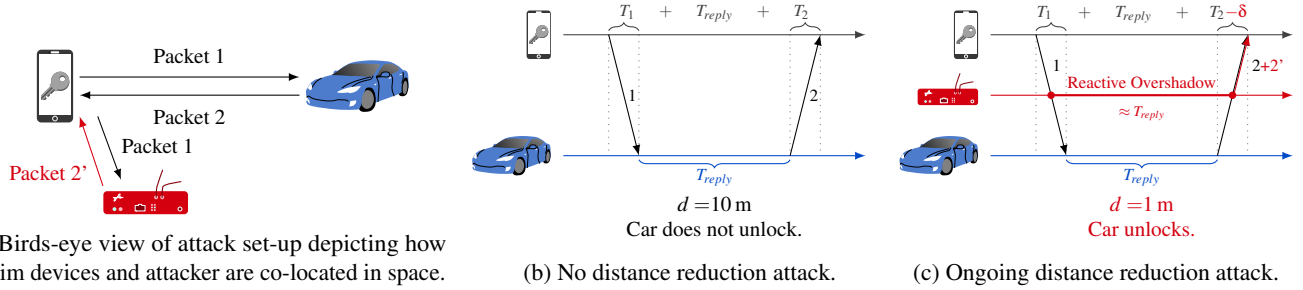


Figure 6: Overview of a distance reduction attack (a). Two devices determine their relative distance with UWB HRP by measuring the ToF of packets 1 and 2 (b). The attacker cannot transmit 1 or 2 in advance because their content is unpredictable (STS). Instead, the attacker transmits a different packet 2' with random STS' coarsely aligned over packet 2 (c). At the receiver, the additional noise caused by 2' is mistaken for a low-power copy of 2 arrived earlier ($T_2 - \delta$), decreasing distance.

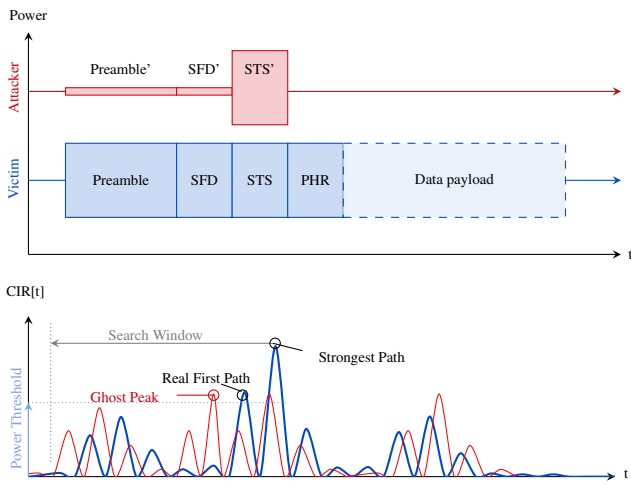


Figure 7: The attacker transmits a carefully crafted packet (red), coarsely synchronized with the legitimate signal (blue), and with power sufficiently low to avoid jamming. The STS is secret and STS' is randomly chosen by the attacker. Therefore, at the receiver the correlation peaks caused by the attacker (red) are lower than those of the strongest path (blue). However, one of them ('ghost peak') is higher than the threshold for accepting peaks that correspond to legitimate paths, and it falls inside the back-search window. Therefore, it is mistakenly classified as an early path, shorter than the real one.

place its devices in physical proximity to one of the victim devices but has no physical access and cannot tamper with these devices. The attacker can receive and inject signals on the wireless communication channel. Specifically, they can craft and transmit UWB messages based on the HRP standard.

We illustrate our attack in Figure 6 and Figure 7. During an initial observation phase, the attacker device behaves like an HRP UWB packet analyzer and resolves the sequences of packets exchanged by the victim devices. Once the ranging sequence and its timings have been identified, the attacker

device reactively injects a signal over selected packet components (*overshadowing*), as shown in Figure 6c. The receiver mistakes the noise induced by overshadowing for an early copy of the legitimate signal, reducing distance (Figure 7).

Injected overshadowing signals follow the structure of an HRP PHY packet but are crafted such that different packet fields are transmitted at different power levels. The attacker needs to synchronize with packet transmissions. Packets are sent every few ms, depending on the ranging implementation. Based on the first packet sent during a ranging sequence, which triggers the attack, the attacker can adjust its timing, and then overshadow the following packets belonging to the same ranging. During overshadowing, the attacker's synchronization accuracy only needs to be in the order of μs , despite attacking a protocol that measures timings with ps resolution [15]. Since the effect of distance on timings (approximately 3.3 ns per meter) is significantly smaller than the synchronization accuracy, fine-grained tuning is not necessary, even when distance changes. To establish the necessary transmission power, the attacking device can initially send a strong signal, then reduce power until distance reductions start to occur, while not causing any jamming. Tuning happens over a small set of values, and it requires some coarse adjustments only if the distance among devices changes considerably (to account for the quadratic decrease of power caused by propagation). Due to the simplicity of the described tuning process and the small search space, this is not a problem in practice (see Section 3.2 and Section 4).

The term *overshadowing* distinguishes our attack (i.e., overlapping an attack signal with the legitimate packet to decrease the perceived ToA at the receiver) from spoofing (i.e., sending a forged valid copy of the packet in advance, which is impossible because of the unpredictable STS).

Our attack can be implemented and executed using a simple and inexpensive off-the-shelf HRP UWB device. Therefore, no complex laboratory equipment is needed, making the attack practical and easy to implement. Figure 7 shows an injected packet aligned with a legitimate packet, and the

Table 1: Comparison with previous work exploiting leading edge detection.

	UWB Standard	Security	Victim	Attacked Field	Attacker	Attack signal
<i>Ghost Peak</i>	IEEE 802.15.4z	STS	Apple U1	STS	Off-the-Shelf	Weak Preamble + Strong Random STS
Adaptive [58]	IEEE 802.15.4z	STS	Simulation	STS	Simulation	Pulses at lower rate for shorter time
Cicada++ [58]	IEEE 802.15.4z	STS	Simulation	STS	Simulation	Pulses at lower rate
Cicada [43]	IEEE 802.15.4a	None	Simulation	Preamble	Simulation	Pulses

corresponding CIR at the receiver. The injected packet is composed of Preamble, SFD, and STS, where the STS is randomly generated without any knowledge of the legitimate STS. Consequently, the correlation peaks caused by the attacker are smaller than the peak corresponding to the legitimate strongest path. However, one of the peaks (ghost peak) is high enough to be (mis)classified as a legitimate early peak, corresponding to a shorter path. The power of each field is independently adjusted to obtain optimal results, as explained in more detail in Section 3.2.

In many practical cases, HRP UWB devices use DS-TWR and possibly exchange additional synchronization or data packets. This information can also be exchanged out-of-band (e.g., using Bluetooth [13, 14], NFC, UHF). However, this pre-negotiation does not impact the attack, which only targets ToA of packets in the ranging sequence [34]. As shown in Figure 8, the attack can be easily generalized. By configuring the delay of reaction after the reception of the first packet, the attacker can attack any desired packet in the sequence. In the case of DS-TWR, this can be leveraged to select the device to attack. Alternatively, an attacker could also use two devices to attack both ends simultaneously, increasing the chances of success.

In summary, an attacker needs to configure which packet in the ranging sequence to attack (by selecting the delay from the reception of the first packet) and the power of the preamble, SFD, and STS to inject. In Section 3.2, we will explain why and how these parameters affect the distance measurement.

3.2 Working Principle

In this subsection, we provide details about why and how the attack works. Furthermore, we compare it to existing distance shortening attacks.

3.2.1 Secure Leading Edge Detection

Accurate timestamps require detecting the earliest copy of the received signal, also called *leading edge detection*. In the following, we explain the challenge of leading edge detection and describe how our attack selectively attacks specific fields of targeted packets in a ranging sequence by overshadowing the contents.

In a realistic environment with obstacles and reflections, the receiver will likely be presented with multiple copies of the

transmitted signal, arriving with different power from different paths. In HRP, the problem is exacerbated because these delays might cause self-interference among pulses that are spaced only by 16 ns (high repetition frequency of 64 MHz), see Section 2.1.3. For Time-of-Flight measurements used in Two-Way Ranging, the receiver must find the earliest copy, corresponding to the shortest path (Line-of-Sight). When receiver and transmitter are not in LoS, the copy corresponding to the direct path is likely to arrive at lower power than other NLoS reflections, as previously shown in Figure 2. When looking for the leading edge copy, any algorithm or implementation must decide whether it faces noise or a very low power early copy of the signal, which is challenging.

Suppose an attacker is able to inject noise that looks reasonably similar to a legitimate low-power copy to the reception algorithm. In that case, it might trick the receiver into accepting it as the leading edge, causing a distance reduction. This attack has been first proposed for IEEE 802.15.4a in [41]. In IEEE 802.15.4a there is no STS and the attacker can inject a UWB pulse to attack the preamble. A recent study [58] has made the hypothesis, confirmed by simulation, that variations of the Cicada attack can be used to attack the STS in IEEE 802.15.4z by injecting HRP pulses. Since HRP UWB reception algorithms are not publicly known, simulations are based on three main assumptions: (i) arrival time and quality of the STS are computed via time-domain cross-correlation, (ii) the leading edge is found by looking for a smaller correlation peak in a limited backsearch window before the strongest peak, and (iii) thresholds are set to evaluate the significance of correlation compared to noise. Simulations in [58] highlight that, given a reception algorithm, there is a fundamental trade-off between security and performance: lax thresholds are necessary to accept legitimate early copies in challenging multi-path environments, but this increases the chance of accepting attacker-induced noise.

In this paper, we take the opposite approach. Instead of hypothesizing a certain algorithm and design choice and studying it in simulation, we empirically analyze the behavior of the unknown algorithms deployed in real products (Apple U1) when subject to signal injection. Because of their closed-source nature, we do not know most of the design choices. For example, we are not aware whether they implement time-domain cross-correlation or take a frequency-domain approach, how they estimate the noise floor, how they

define, and configure thresholds and whether such thresholds are dynamically adjusted to the environment.

The only assumption we make when developing our attack is that the receiver is able to work in NLoS conditions, which we were able to confirm empirically. We then chose to transmit signals crafted from standard packets, to maximize the probability of generating noise that is misclassified for a legitimate copy and to make the attack practical to implement. Instead of injecting fine-grained aligned pulses at different power and repetition frequencies, we observe how the fields of standard packets affect reception. We adapt the structure of the packet and the power level of the fields to maximize the chances of reduction (by injecting STS pulses) while avoiding jamming and other errors. In general, differently from previous work, our attack handles many of those challenges due to the fact that it operates on real sequences of packets used in real exchanges.

It is worth noticing that the attacker does not have direct control over the amount of distance reduction. A method to gain partial control has been proposed in simulation in [58]. However, it requires to delay the legitimate copy, emulate the leading edge detection algorithm at reception, analyze its output in real-time, and interrupt the injection when the desired result is obtained. For these reasons, it is hard to implement in practice, in particular with off-the-shelf devices. As an alternative, in Section 3.2.3 we show how the choice of the victim packet(s) in a sequence can affect the distribution of reduction, and in Section 4 we empirically analyze it. In Table 1 we compare previous work on leading edge detection with our approach. Nevertheless, in our threat model the attacker is not interested in controlling the reduction but in causing practical distance reductions that will trick the victim into believing the legitimate user is close enough to grant access.

3.2.2 Selective Overshadowing to Avoid Jamming

An attack against leading edge detection can be successful in practice only if the injection of the attack signal does not accidentally produce other errors that invalidate a ranging sequence. To achieve this goal, our attack carefully crafts the timing, format, and power level of the attack signal. The attacker's transmission is not continuous but reactive. As opposed to the continuous transmission of Gaussian noise or UWB pulses, a reactive transmission allows targeting a specific packet in the ranging sequence, without affecting packets carrying data. Similarly, the attack packet does not contain any data field that could corrupt the content of the legitimate packet. The preamble is transmitted at low power so that it does not trigger a new receive event. Such an event would indeed lead to an error when the receiver determines the STS quality and the presence of expected data fields. The STS pulses are instead sent at higher power so that they overshadow the legitimate signal and produce noise that is misclassified as an early copy. Finally, both the power of preamble

and STS are adjusted based on the relative distance between devices. In particular, power is lowered to avoid jamming when the device that transmits the packet to overshadow is far away.

3.2.3 Selecting Victim Packet(s)

Typically, DS-TWR is used because it compensates for clock errors and asymmetric reply times. We have confirmed this in our analysis of many HRP UWB configurations. As mentioned in the standard, distance is computed with the method proposed in [36]:

$$\hat{d} = c \cdot \hat{T}_{prop} = c \cdot \frac{T_{round1} \cdot T_{round2} - T_{reply1} \cdot T_{reply2}}{T_{round1} + T_{round2} + T_{reply1} + T_{reply2}} \quad (1)$$

For simplicity we can neglect non-idealities and consider that distance is measured as the average of the two rounds [2]:

$$\hat{d} = c \cdot \hat{T}_{prop} = \frac{c}{4} \cdot (T_{round1} + T_{round2} - T_{reply1} - T_{reply2}) \quad (2)$$

Sometimes, a fourth message is used, likely for the transmission of additional data. If the ranging packets contain only preamble and STS but no data, additional data packets are sent earlier and/or later. In any case, the attacker can configure the delay from reception of the first packet to attack either the second or the third packet of the DS-TWR sequence. As shown in Figure 8, attacking the second packet corresponds to overshadowing a packet transmitted by the responder and received by the initiator, while attacking the third packet corresponds to the opposite. It is convenient for the attacker to be closer to the receiver to use less power for overshadowing, but it is not strictly necessary. It is worth noting that the choice between the second and third packet is not entirely symmetric. Attacking the third packet has the only effect of reducing the round time measured by the responder (T_{round2}) leading to:

$$\hat{d}' = \frac{c}{4} (T_{round1} + T_{round2} - \delta - T_{reply1} - T_{reply2}) = \hat{d} - c \cdot \frac{\delta}{4} \quad (3)$$

Instead, attacking the second packet reduces both T_{round1} (because the initiator receives the packet earlier) and T_{round2} (because the initiator consequently replies earlier), leading to:

$$\hat{d}'' = \frac{c}{4} (T_{round1} - \delta + T_{round2} - \delta - T_{reply1} - T_{reply2}) = \hat{d} - c \cdot \frac{\delta}{2} \quad (4)$$

Clearly, by attacking the second packet, the attacker can obtain reductions that are twice as big as those obtained by attacking the third packet. The reduction δ is a random variable not in control of the attacker and it is bounded by the maximum difference between LoS and NLoS path accepted by the receiver (width of the backsearch window). However, precise control over distance is not required. For example, any reduction below 2 m would break PKES and unlock a car.

As an alternative, the attacker can use devices to target both the second packet (near the initiator) and the third packet




Figure 8: Generalization to more complex sequences (e.g., DS-TWR). The attacker can choose which packet/side to target.

(near the responder). We can consider the two attacks as independent events. Therefore, the attacker will obtain reductions of $c \cdot \delta/4$, $2c \cdot \delta/4$, and $3c \cdot \delta/4$, each with decreasing probability.

We confirm these calculations in Section 4. We tested our targets in their operating range (≈ 15 m) achieving reductions of up to 12.45 m. A system designed for larger distances (e.g., 100 m) would likely have a larger backsearch window allowing longer reductions, but common use cases (e.g., car keys, item finder) only envision short ranges.

3.3 Implementation

We have presented a general approach to conduct distance-reduction attacks. In principle, it can be implemented with any off-the-shelf HRP UWB IEEE 802.15.4z compatible device that can be programmed to receive and transmit packets and that allows configuring individual power levels for each field. In practice, we have implemented the attack using a Qorvo DWM3000EVB [45], controlled by a Nordic Semiconductor nRF52 DK [37], for a total cost of around USD 65 only. These devices can be easily programmed with open-source firmware [25], they have limited size, and they can be powered by a portable USB battery.

The delay can be configured to be a multiple of the reply time used by the victims so that the attack signal is transmitted on top of one of the following packets (Figure 8). The attacker can find this and other reception parameters in an attack preparation phase. The preparation phase is only required once per protocol, e.g., parameters stay the same for every iPhone–AirTag distance measurement.

For this, we have developed a sniffer and packet analyzer based on a Qorvo DWM3000EVB attached to an STM32 Nucleo-F429ZI. Using Qorvo’s SDK, we implement a fast UWB receiver, which forwards frames over a USB connection to a host computer. Here, packets are analyzed with a custom Wireshark dissector [22] that also supports Apple’s proprietary UWB frame format. Multiple packet analyzers with different configurations can be connected, which is required to observe complex ranging procedures. The DWM3000EVB chip in our packet analyzer can receive timestamps with an accuracy of 15.65 ps [15]. These timestamps are recorded and forwarded to the Wireshark dissector. As we will show later in Section 4, accuracy in the order of μ s is sufficient to run

the attack. Since most protocols that are using UWB today are closed source, there is no option to analyze the protocols for potential privacy and security issues thoroughly. Besides ranging frames, the UWB packet analyzer also receives data frames. This allows us to inspect if any private data, static identifiers, key material, or similar is shared over UWB. Apple does not use the IEEE 802.15.4 MAC frame format, e.g., for iPhone–iPhone and iPhone–HomePod ranging. We implemented a Wireshark dissector that allows inspecting the parts of it that are not encrypted. The analysis of a new UWB sequence is simple. As long as a few initial parameters such as the channel number are set, the device starts reporting reception events and diagnostics such as preamble quality (any preamble number triggers a reception, but with different quality) and other error codes. The attacker can then proceed to adjust other parameters (e.g., packet structure, STS length) until correct reception of a full packet occurs).

3.4 Application to Real HRP UWB Chips

We successfully applied our distance-reduction attack against Apple U1 chips deployed in different products (iPhone, AirTag, HomePod). When the U1 is interoperated with chips from other vendors (NXP SR040, NXP SR150, and Qorvo DWM3000), attacking the U1 still results in distance reduction for both sides.²

Figure 9a shows a concrete example. One iPhone 11 Pro (Apple U1) is placed at 8 m distance from an NXP SR150 in line of sight. The two devices exchange a total of 6 messages, where 3 are used for DS-TWR. The iPhone is the initiator (and victim) and the NXP SR150 is the responder. A Qorvo DWM3000EVB acts as an attacker placed at around 30 cm from the victim iPhone. By hitting the second message of the DS-TWR sequence, the attacker causes distance reductions of up to 10 m. The application running on the iPhone shows 8 m when the attack is off and 0 m during a successful reduction.

Figure 9b shows another example of an attack targeting ranging between two identical iPhones. In this case, the total number of messages is 4, but the attack is similar. By targeting the second packet of the DS-TWR sequence, the

²Since all implementations are closed source, we do not know if different parameters of the attack would work also for other combinations of chips (it would require an exhaustive search).

attacker causes reductions from 10 m to less than 2 m in the raw measurements plotted on the laptop.

4 Experimental Evaluation

In this section, we demonstrate the feasibility of our attack, show the number of distance reductions possible, and determine the success rate.

4.1 Setup

We ran the attacks in an indoor LoS environment with two victim devices placed at various distances between 5 m and 15 m with antennas facing each other. This setup results in a relatively good baseline signal quality with a small ranging error (normally 10 cm to 20 cm) when the attack is turned off. We chose this setup to avoid measurement noise due to channel (e.g., excess paths) that would otherwise distort the outcome. Different, potentially worse, channel conditions do not pose an inherent challenge to an attacker, since an attacker can relay signals (e.g., by cable) and establish relatively good channel conditions this way. We evaluated the following device combinations: iPhone–iPhone (Nearby Interaction) [30], iPhone–AirTag (FindMy ranging) [7], iPhone–HomePod (Handoff music) [8], iPhone–NXP and iPhone–Qorvo (compatibility mode) [31].


The attacker places either one or two Qorvo DWM3000EVB in ca. 30 cm proximity to one or both ranging devices. The adversarial transceivers perform a reactive attack as introduced in Section 3, i.e., they are programmed to detect the initial frame of the ranging

exchange and then overshadow preamble and STS of one or two subsequent frames. It is important to note that, while the overall success rate of the attack and the maximum distance reduction increases when both sides are targeted independently, the result of the ranging procedure is synchronized among the devices, i.e., both legitimate devices eventually report the same measurement time series.

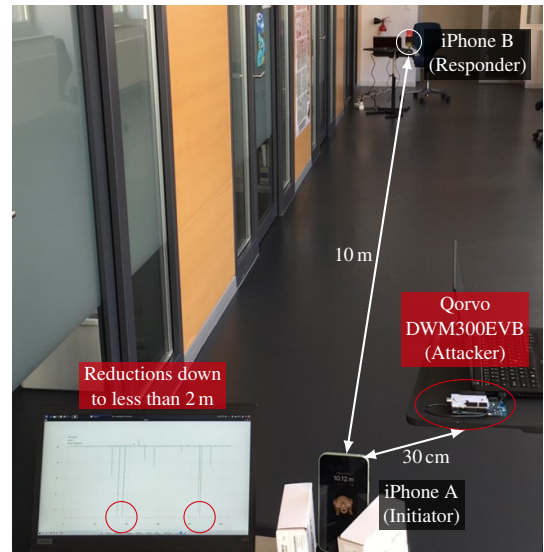
As shown in Table 2, different applications have different methods to start a sequence of ranging measurements. One of the main goals for UWB-based access control is to provide seamless access when the user is close, while turning on ranging only when necessary to save power. One method to achieve this is to use Bluetooth Low Energy (BLE) to detect the presence of the device and, if the BLE Received Signal Strength Indicator (RSSI) is high enough, start UWB ranging, as done by Handoff music in the iPhone–HomePod scenario. Note that using BLE RSSI is not a secure measurement and is vulnerable to simple physical attacks (e.g., relay). The iPhone and Apple Watch are projected to use a similar mechanism when used as PKES: using BLE to detect a car, generate and exchange keys to be used for the STS, and start ranging [6]. The choice for ranging initiation for the iPhone–iPhone and iPhone–NXP/Qorvo scenarios is left to the developer implementation. In our case, the user starts/stops ranging. AirTags also require user interaction.

4.2 Retrieving Raw Distance Measurements

UWB-based key solutions only need to determine if a distance is below or above a threshold. Thus, many applications do not display detailed distance information in the user interface.



(a) iPhone (initiator, victim) + NXP SR150 (responder): reduction from 8 m to 0 m visible on the screen of the iPhone.



(b) iPhone + iPhone: reduction from 10 m to less than 2 m visible in the raw measurements logs.

Figure 9: Two concrete examples of distance reduction attacks.




Figure 10: This figure shows a 25 min ranging experiment. Two iPhones are placed at a real distance of 15 m between each other, under attack with two devices. The right part shows the distances reported for each measurement in blue, with obvious reductions (i.e., reported distance less than 15 m). The success rate, which is calculated as a rolling average over 300 measurements, is plotted in orange. Over the entire experiment, the rate of reductions was 4.08 %. The histogram on the left side reflects the distribution of the reduced distances reported in the experiment.

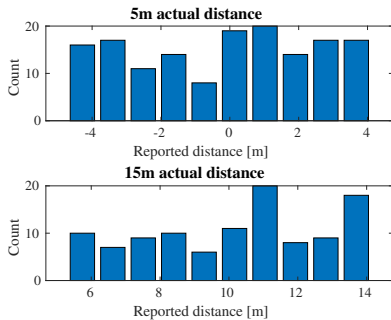


Figure 11: Distribution of reduced distance reports for the iPhone–iPhone (5 m and 15 m) setup, attacked with single device over a 15 min observation period. The overall rate of successful distance reductions (i.e., less than 5 meter and 15 meter, respectively) is ca. 2.2 % in both cases.

In contrast, we need precise distance measurement results without aggregation to evaluate the success rates of attacks. In the case of the Apple UWB implementation, the U1 chip reports the raw measurements to iOS drivers, which log them. Viewing these measurement logs requires the *Location Services* and *AirTag* debug profile, which can be installed on any iPhone without jailbreak [29]. Then, detailed measurement information appears in the logs, including the distance.

```
nearbyd #me,MeasEngMetricsCalculator::checkCirMetrics:
AOA CycleIdx 1497 RangeMsmt 3.27164 machAbsTime ...
nearbyd #sp,[Solution Provider] r1 range: 3.272 m
```

iOS will forward the raw measurement to the corresponding daemon or Nearby Interaction framework in most measurement modes. For example, attacking a single distance measurement between an iPhone and a HomePod mini is sufficient to show the HomePod’s music playback menu on the

iPhone immediately. We observe the same behavior in the compatibility mode, which is used for third-party integration like car keys. For example, Figure 9a shows a reduction to 0 m visible on the screen.

Apple noticed that UWB distance measurements are not always reliable. When using the Nearby Interaction framework with the example *Peekaboo* application [30], the current measurement is only published if it does not deviate more than 1 m from the median of the last 11 distance measurements. This filter is not applied when the phone interacts with third-party devices. We manually identify this boundary by replacing distance values reported by the U1 chip on an iPhone with FЯIDA [24], similar to previously published hooks [14]. iOS only discards measurements on the application layer, thereby hiding them from curious developers. In the following, we use the raw U1 chip measurements provided by the logs of an unmodified iPhone to get comparable results, irrespective of opaque application-layer filters.

4.3 Results

We quantify the success of the attack as the relative rate of ranging measurements (as read from the iOS logs) indicating a distance shorter than the baseline, averaged over an observation interval of at least 15 min. To separate benign measurement non-idealities from actual reductions, we only count measurements lower than two times the maximum benign (negative) deviation during a 100 s interval before running the attack.

For different device combinations, our attack causes distance reductions between ca. 2 m to 12 m with success rates in the range of 2 % to 4 %. An overview is provided in Table 2. Some of the differences in success rates, i.e., those between 2 % and 4 %, can be explained by the fact that either only one

or two packets of the ranging procedure are attacked. This means, for a given success rate per individual ToA measurement (i.e., by packet), we increase the chances that at least one ToA measurement of the ranging exchange is successfully reduced by targeting both the second and third packet. To exemplify the cumulative effect of multiple attacking devices on the overall reduction, Figure 10 shows the entire time series of measurements (iPhone–iPhone) over a 25 min observation interval with one attack device placed at each end. Due to the attack success rate changing over time, we also display the instantaneous success rate using a sliding window over 300 consecutive measurements. The attack results in an overall rate of reduced measurements of 4.1 %, whereas the rolling average over 300 consecutive packets can get as high as 7.7 %. The main uncontrollable source of such variations is likely the randomness of the STS and correlation noise caused by the adversarial transmission. The distribution of distance reductions is biased towards reductions ≤ 5 m because either of the devices, i.e., the one targeting the second packet and the one targeting the third packet, can cause those. In contrast, the device replying to the initiator (i.e., transmitting over the second packet) can solely have an effect up to 10 m. This observation is in line with the analysis provided in Section 3.2. The longest reduction observed over this interval is over 12.35 m, caused by successful reductions on both packets attacked during the same ranging procedure. Assuming independence of the effects on either side, these additive reductions (exceeding 10 m), while orders of magnitude less likely, are still frequent enough to occur within a realistic time window (25 min). This shows that in any scenario where a key is placed less than ca. 14 m away, an attack can be successful with high likelihood. A potential scenario is a car that is parked outside the main door of a house, whereas the key is placed somewhere close to the entrance³. In a configuration where only the responder is vulnerable, distance reductions are limited by ca. 5 m, because only the ToA of the third packet can be targeted. An example for this is the combination of iPhone and NXP SR040, since NXP SR040 can only be configured as initiator.

The range of possible relative distance reductions does not depend on the actual distance of the ranging devices, and the U1 chip even reports negative distances in case the distance reduction exceeds the nominal distance. Figure 11 highlights this, showing the distribution of reduced distance reports in the iPhone–iPhone setup with one attack device over two different distances, 5 m and 15 m, over a 15 min observation period. It becomes evident that the relative reduction is, irrespective of the nominal distance, bounded by 10 m.

Jamming Even though organizations claim that UWB is *immune to jamming* [20], we saw that it is perfectly possible to disturb ranging measurements through jamming. When

³Precisely this attack scenario has become an increasing concern for PKES that do not rely on signal ToF [16, 23, 59, 62].

setting the transmission power to a high level and the transmission time to match with an expected ranging frame the receiver will not be able to receive the frame. This is likely caused by disturbing the STS.

4.4 Parameter Tuning

The timings and power level of the adversarial signal have to be matched to the legitimate signal. To run a successful attack, a time delay that matches the actual frame has to be accurate in the order of 5 μ s to have success rate of above 1 percent. Whereas the timing is relatively static, the power level depends on the baseline signal quality. This means the attack power needs to be adjusted depending on the path loss, i.e., the distance between the legitimate devices. E.g., for victim devices at a nominal distance of 15 m, we adjusted the gain parameters to a fraction of the maximum possible value on the Qorvo, ca. $\frac{1}{4}$ for the preamble and $\frac{1}{3}$ for the STS. This adjustment is required to avoid sending a signal too strong (i.e., competing with the legitimate main peak) but strong enough to register a fake early peak. We found that changes of the communication distance between 15 m and 5 m did not change the requirement on output power in order to achieve success rates above 1 %. This means, while there is a dependency on the channel, there is also a significant window allowing for success. To address potentially stronger channel variations, e.g., due to movement, we highlight directions for calibration in Section 5.4. Other configurations, like the channel frequency, packet sequence, preamble, SFD, and delay between packets were found to be static for a given target.

4.5 Device Pairings

In Table 2 we show the results of performing the attack against different pairs of devices. In most cases, the iPhone has been the main victim, since its implementation seems to be most affected by this vulnerability. Our results have shown that one vulnerable device results in a distance reduction for both devices. This issue cannot be mitigated on one end only, since every UWB ranging algorithm requires both devices to report round-trip time T_{round} and reply delay T_{reply} to the other devices. This means that a user has to trust both devices, which can only be achieved through independent certification, including a review of the algorithms.

Additionally, we see that it is irrelevant with which device the iPhone performs ranging. Every device combination is vulnerable to distance reduction with a good success rate.

4.6 Confirming Results by Binary Analysis

Our attacks work without knowing UWB chip implementation details. Nonetheless, binary analysis of the UWB implementations helps understanding why attacks are feasible.

Table 2: Overview of attack scenarios against Apple U1 (primary victim) and results.

Scenario	Primary Victim	Secondary Victim	Roles	Initiation	Max. Reduction	Success Rate
Handoff Music	HomePod mini (Apple U1)	iPhone (Apple U1)	Init./Resp.	Proximity*	9.01 m	2.10 %
Nearby Interaction	iPhone (Apple U1)	iPhone (Apple U1)	Init./Resp.	Developer choice**	12.45 m	4.08 %
AirTag	AirTag (Apple U1)	iPhone (Apple U1)	Init./Resp.	User interaction	9.09 m	4.25 %
NXP Initiator	iPhone (Apple U1)	Tag (NXP SR040)	Resp./Init.	Developer choice**	4.80 m	1.87 %
NXP Responder	iPhone (Apple U1)	Tag (NXP SR150)	Init./Resp.	Developer choice**	9.68 m	2.15 %
Qorvo	iPhone (Apple U1)	Tag (Qorvo DWM3000)	Init./Resp.	Developer choice**	8.13 m	3.09 %

*Measured with BLE Received Signal Strength Indicator, which is not secure. **Defined by developer implementation, turned on/off by the user in our case.

All UWB chips analyzed in this paper are split into a main application and a low-level Digital Signal Processor (DSP). The DSP can be instrumented over a serial interface. The NXP and Qorvo chip have a documented Application Programming Interface (API) for this interface. However, the DSP itself is inaccessible on these platforms. Qorvo does not allow reprogramming the DSP to the best of our knowledge. NXP ships firmware files for the DSP, but they are encrypted and signed, even in the development kit, which prevents analysis.

In contrast, Apple’s U1 DSP firmware is part of the software updates for all UWB-enabled devices. It ships in a proprietary `ftab` format [14], and the firmware is not encrypted. However, it is a bare metal firmware without any symbols. It contains a few strings, including assertions about the distance measurement. One of these strings refers to the backsearch window and is part of a function used for STS correlation:

```
(inp->sum_window_right - inp->sum_window_left + 1) <= 16
```

Since the backsearch window sampling rate is unknown, we cannot calculate the maximum possible distance reduction.

5 Discussion

5.1 Strengths and Limitations

We proposed a practical attack that achieves distance reduction of several meters using only a simple and inexpensive off-the-shelf device.

Its practicality makes this attack particularly relevant for security-critical applications that are gaining more market traction, for example, PKES systems in cars.

The main limitation of the attack is the little control on the amount of distance reduction caused by selective overshadowing. The maximum reduction is set by the maximum difference between NLoS and LoS accepted by the victim, and how many packets in the DS-TWR sequence are targeted. In the case of Apple U1, this results in attacks that reduce distance by a maximum of 5 m when attacking the third packet of DS-TWR, 10 m when attacking the second, and 15 m when attacking both. A system designed to be operated at larger ranges would also require a larger backsearch window, resulting in larger reductions, too. It is worth noting that the attacker

is generally not interested in steering distance precisely. For example, to attack an access control system, it is enough for the attacker to cause a reduction below the threshold that grants access within a reasonable time frame.

A related limitation is that the attacker cannot control precisely which ranging sequence is affected by a reduction. Thus, the victim could try to identify outliers as a distance reduction attack. However, in a practical setting, the users move and have their car key in a pocket, leading to similar outliers. In addition, detecting outliers would generally require many measurements, while a smooth user experience requires short response times.

5.2 Reflections on HRP UWB Security

Our results are a clear call for research to improve the security of HRP. We have shown that current security properties hinge entirely on the quality of proprietary algorithms and black-box implementations and that mere compliance with IEEE 802.15.4z does not protect systems against distance reduction attacks. We argue that security should not be a distinguishing feature of individual implementations but publicly available and verifiable. We are convinced that a secure standard that withstands the scrutiny of the research community also benefits device vendors, as they can rely on its correctness guarantees and focus entirely on implementation challenges. Consequently, it is important to direct future work towards the development of a secure and open algorithm for first path detection, which can be integrated into an upcoming standard. The fact that AES is used to generate the STS might intuitively suggest a high level of security, since the probability of guessing the key used to derive the STS is only 2^{-128} . Unfortunately, this intuition is wrong because the receiver never verifies the correctness of the STS explicitly. While cross-correlation peaks may indicate a certain similarity between the expected STS and the received signal, looking at single values does not always provide sufficient information about the correctness of the received STS. Consequently, the probability of having a randomly injected STS accepted is around $1 - 4\%$ (see Table 2), which is far above the adversarial success rates accepted for PKES or payment systems, which are around 2^{-20} and 2^{-48} , respectively.

5.3 Countermeasures

The necessity to distinguish legitimate early copies of the signal from the attacker’s induced noise creates a tension between HRP UWB security and performance. Our results indicate that Apple’s current U1 implementation does not perform any noticeably advanced checks on the STS. However, the chip does perform some basic checks, since we cannot simply achieve 100 % success rate by transmitting packets in advance. The receiver must apply advanced statistics on the incoming signal to detect attacks while still performing well in challenging NLoS conditions. For example, the receiver could compare the consistency between the channel response of the early low-power copy and the main copy. The STS quality, which represents the similarity of the incoming signal with the expected one, should be checked independently for both early and late copies. In other words, the receiver should not assume that an early copy is acceptable just because another valid late copy appears just afterwards. Such countermeasures require increasing the complexity of the receiver, which might not be feasible for battery-powered devices like the AirTags, and for devices with strict reaction time constraints like PKES.

Reducing the maximum accepted difference between LoS and NLoS copies would reduce the maximum reduction that an attacker can achieve, but it would also limit the use of the product in realistic scenarios, e.g., a car key in a pocket.

More countermeasures can be applied at the upper layers, for example, detecting reductions as outliers. However, real-time applications do not have a margin for accumulating more than a few measurements before reacting based on the measurement value.

5.4 Future Work

On the theoretical side, our results are a call for researching whether it would be possible to assure the HRP UWB security level based on cross-correlation or other techniques. This is challenging because its security is intertwined with proprietary algorithms and design choices. On the practical side, the analysis and attack phases of our attack could be combined in a feedback loop, achieving automated calibration of the attack setup for varying scenarios. E.g., the power of the attack packet could be automatically chosen based on some observations at reception. A chip such as the Qorvo DWM3000EVB offers many reception diagnostics. E.g., when receiving the packets sent by the victims, the attacker could estimate the quality of their preamble, obtain a rough estimate of its relative distance from the victims, and adjust the power used for overshadowing. Similarly, the attacker could measure the rate of packets exchanged by the victims and lower its power if it detects a lower rate due to jamming. Other attack parameters could be adaptively configured in a similar automated fashion.

6 Related Work

The UWB IEEE 802.15.4 standard is described in [3,4]. Chips following the HRP mode of the standard have been implemented by several vendors, such as Apple (U1) [9], NXP (SR040, SR150, SR100T) [38], and Qorvo (DWM3000) [45]. Chips implementing the standard LRP mode have been implemented by Microchip (ATA8352, ATA8350) [35] and Renesas [48]. To the best of our knowledge, these LRP mode chips are not available in consumer electronic devices.

The first implementation-independent security evaluation of HRP UWB at the physical layer has been conducted in [58]. That work proposed two attacks on HRP, derived from the Cicada attack [42, 43], and shows in simulations that even conservative receiver implementations could be susceptible to distance reductions. In contrast to our paper, the authors neither conducted experiments with real UWB chips, nor prove that attacks are practical with off-the-shelf hardware. Furthermore, they did not consider other aspects of the UWB ranging protocols, i.e., the sequence of messages, significance of different message fields, or their power.

Further research on UWB ranging has been done in [57]. This work proposes improvements to LRP that aim at securely extending the range of the LRP through pulse interleaving.

Previously documented attacks against UWB [21], which applied to earlier standards (IEEE 802.15.4a), cannot be used against HRP because of the high frequency of the pulses. For example, an attacker cannot acquire the polarity of a 2 ns pulse in time to advance it to conduct an ED/LC [41] attack.

The Apple U1 chip and secure ranging in iOS have been studied recently in [14], with a focus on the overall software architecture rather than physical-layer aspects. After the release of AirTags, the hardware hacking community has discovered that its main firmware can be easily extracted and modified by glitching [39, 49]. Although the mentioned work provides an interesting overview on Apple’s usage of HRP UWB in its products, it has not analyzed physical-layer attacks or features of the firmware that have an impact on the security of the physical layer (e.g., the DSP code or the logic that decides whether to accept early peaks in the backsearch window).

Other studies [27] focused on the security of the Apple FindMy network that allows locating devices including AirTags. It is possible to add custom BLE devices to the FindMy network [26] or to leverage FindMy to upload data from devices without Internet connection [12]. These works focus on BLE and the architecture of FindMy and are mostly unrelated to the UWB technology that some of the devices deploy.

The security of time of arrival measurements has been formalized in the form of Message Time of Arrival Codes (MTAC) in [34].

7 Conclusion

We demonstrated for the first time a practical distance reduction attack against HRP UWB (IEEE 802.15.4z) secure ranging, implemented in Apple U1 chips and widely deployed in Apple products. We demonstrate that the impact reaches beyond the Apple ecosystem, showing attacks when ranging is performed between an Apple U1 chip in an iPhone and development kits with chips by NXP and Qorvo. Distance reduction is a considerable concern in many applications, from access control (e.g., opening cars, doors) to mobile payments and indoor positioning for industrial plants. Our attack is practical, and it can be implemented with a cheap off-the-shelf device. Our results raise the awareness on the pitfalls of HRP UWB technology. On the one hand, HRP UWB promises a nominally high security level based on a cryptographically secure STS sequence that cannot be guessed by an attacker. On the other hand, the actual security level depends on obscure design choices at the receiver. No independent experimental evaluation and certification framework exists either. Our results show that distance-reduction attacks are practical. To improve the state of HRP UWB security, we have proposed and discussed several countermeasures.

Acknowledgments

This research has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under grant agreement No 726227. This research has received funding from the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40_180545. This project has been partially funded by Fondation Botnar. This work has been co-funded by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

Availability

Source code for the UWB sniffer, packet analyzer, and Wireshark dissector will be made available on

<https://securepositioning.com/ghost-peak/>

References

- [1] UWB Regulations. https://www.decawave.com/sites/default/files/apr001_uwb_worldwide_regulations_summary_ev1.2.pdf. Accessed 2021-10-08.
- [2] IEEE Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs). *IEEE Std 802.15.4-2011 (Revision of IEEE Std 802.15.4-2006)*, pages 1–314, 2011.
- [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] 3db Access. VW Adopts UWB for Secure Car Access. <https://www.3db-access.com/article/18>. Accessed 2021-10-09.
- [6] Apple. Explore UWB-based car keys. <https://developer.apple.com/videos/play/wwdc2021/10084/>. Accessed 2022-02-02.
- [7] Apple Inc. Find your keys, wallet, and more with AirTag. <https://support.apple.com/en-us/HT210967>. Accessed 2021-10-12.
- [8] Apple Inc. Hand off audio to HomePod. <https://support.apple.com/en-nz/guide/homepod/apdfb81a72e4/homepod>. Accessed 2021-10-12.
- [9] Apple Inc. Ultra Wideband security in iOS. <https://support.apple.com/guide/security/ultra-wideband-security-sec1e6108efd/web>. Accessed 2021-10-08.
- [10] Brian Barrett. The biggest iPhone news is a tiny new chip inside it. <https://www.wired.com/story/apple-u1-chip/>. Accessed 2021-10-05.
- [11] 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 2021-10-08.
- [12] Fabian Bräunlein. Send My: Arbitrary data transmission via Apple’s Find My network. <https://positive.security/blog/send-my>. Accessed 2021-10-05.
- [13] Car Connectivity Consortium. Digital Key Release. <https://carconnectivity.org/press-release/car-connectivity-consortium-publishes-digital-key-release-3-0/>. Accessed 2021-10-11.
- [14] Jiska Classen and Alexander Heinrich. Wibly Wobbly, Timey Wimey – What’s Really Inside Apple’s U1 Chip. Presentation at Black Hat USA 2021, August 2021.
- [15] Decawave Ltd. DW3000 Family User Manual. <https://www.decawave.com/wp-content/uploads/2021/05/DW3000-User-Manual-1.pdf>.
- [16] Electrek. Tesla warns of theft risk through relay attacks, shares ‘tips’ to help prevent. <https://electrek.co/2018/07/31/tesla-theft-tips-help-prevent-relay-attacks/>. Accessed 2021-10-09.
- [17] Embedded. Ultra-wideband (UWB) adoption picks up pace. <https://www.embedded.com/ultra-wideband-uwband-adoption-picks-up-pace/>. Accessed 2021-10-09.
- [18] FiRa Consortium Inc. FiRa Consortium Website. <https://www.firaconsortium.org/>. Accessed 2021-10-11.
- [19] FiRa Consortium Inc. Technical FAQ - How secure is UWB positioning? <https://www.firaconsortium.org/discover/technical-faq>. Accessed 2021-10-11.
- [20] FiRa Consortium Inc. UWB Technology Comparison. <https://www.firaconsortium.org/discover/comparison>. Accessed 2021-10-12.
- [21] Manuel Flury, Marcin Poturalski, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. Effectiveness of distance-decreasing attacks against impulse radio ranging. In *Proceedings of the Third ACM Conference on Wireless Network Security, WiSec 2010, Hoboken, New Jersey, USA, March 22-24, 2010*, pages 117–128. ACM, 2010.
- [22] Wireshark Foundation. Wireshark. <https://www.wireshark.org>. Accessed 2021-10-11.

- [23] 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.
- [24] FRIDA. Dynamic instrumentation toolkit for developers, reverse-engineers, and security researchers. <https://frida.re/>. Accessed 2021-10-10.
- [25] GitHub user foldedtoad. Nordic nRF52-series + Decawave DWM3000 on Zephyr v2.5. <https://github.com/foldedtoad/dwm3000>. Accessed 2021-10-12.
- [26] Alexander Heinrich, Milan Stute, and Matthias Hollick. OpenHaystack: A framework for tracking personal Bluetooth devices via Apple's massive Find My network. In *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 374–376. ACM, 2021.
- [27] Alexander Heinrich, Milan Stute, Tim Kornhuber, and Matthias Hollick. Who Can Find My Devices? Security and Privacy of Apple's Crowd-Sourced Bluetooth Location Tracking System. *Proc. Priv. Enhancing Technol.*, 2021(3):227–245, 2021.
- [28] Dennis Heinze, Jiska Classen, and Felix Rohrbach. MagicPairing: Apple's Take on Securing Bluetooth Peripherals. In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec '20*, page 111–121, New York, NY, USA, 2020. Association for Computing Machinery.
- [29] Apple Inc. Apple Profiles and Logs. <https://developer.apple.com/bug-reporting/profiles-and-logs/>. Accessed 2021-10-08.
- [30] Apple Inc. Nearby Interaction. <https://developer.apple.com/documentation/nearbyinteraction>. Accessed 2021-10-08.
- [31] Apple Inc. Nearby Interaction Accessory Protocol Specification, Developer Preview, Release R1.
- [32] Allison Johnson. The search is over: smart trackers from Apple, Samsung, and Tile compared. <https://www.theverge.com/22570161/apple-airtag-samsung-smarttag-tile-pro-bluetooth-tracker-review-test-comparison>. Accessed 2021-10-01.
- [33] Steffen Klee, Alexandros Roussos, Max Maass, and Matthias Hollick. Nfcgate: Opening the door for NFC security research with a smartphone-based toolkit. In *14th USENIX Workshop on Offensive Technologies (WOOT 20)*. USENIX Association, August 2020.
- [34] 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.
- [35] Microchip Technology. ATA8352 Impulse-Radio Ultra-Wideband (IR-UWB) Transceiver Data Sheet. https://www.microchip.com/content/dam/mchp/documents/RFA/ProductDocuments/DataSheets/ATA8352_Datasheet_RevA_FEB2021_70005450A.pdf. Accessed 2021-10-08.
- [36] Dries Neirynek, 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.
- [37] Nordic Semiconductor. nRF52 DK - Nordic Semiconductor. <https://www.nordicsemi.com/Products/Development-hardware/nRF52-DK>. Accessed 2021-10-12.
- [38] NXP Semiconductors. Secure Ultra-Wideband (UWB) Positioning and Ranging Optimized for IoT Use Cases. https://www.nxp.com/products/wireless/secure-ultra-wideband-uwband/trimension-sr040-secure-uwband-solution-for-iot-tags:SR040?tab=Documentation_Tab. Accessed 2021-10-08.
- [39] Colin O'Flynn. Apple AirTag Teardown & Test Point Mapping. <https://colinoflynn.com/tag/airtag/>. Accessed 2021-10-05.
- [40] Malcom Owen. AirDrop and device tracking only the beginning of Ultra Wideband in the iPhone. <https://appleinsider.com/articles/19/10/17/airdrop-device-tracking-only-the-beginning-of-ultra-wideband-in-the-iphone>. Accessed 2021-10-08.
- [41] 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.
- [42] 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.
- [43] 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.
- [44] Qorvo Inc. DW3110 - Qorvo. <https://www.qorvo.com/products/p/DW3110>. Accessed 2021-10-12.
- [45] Qorvo Inc. DWM3000EVB - Qorvo. <https://www.qorvo.com/products/p/DWM3000EVB>. Accessed 2021-10-12.
- [46] Qorvo Inc. DWM3001CDK - Qorvo. <https://www.qorvo.com/products/p/DWM3001CDK>. Accessed 2021-10-12.
- [47] Qorvo Inc. Qorvo Completes Acquisition of Decawave. <https://www.qorvo.com/newsroom/news/2020/qorvo-completes-acquisition-of-decawave>. Accessed 2021-10-12.
- [48] Renesas Electronics. Renesas Electronics and 3db Access to Collaborate and Bring Secure Ultra-Wideband Solutions to Market. <https://www.renesas.com/us/en/about/press-room/renesas-electronics-and-3db-access-collaborate-and-bring-secure-ultra-wideband-solutions-market>. Accessed 2021-10-08.
- [49] Thomas Roth. Hacking the Apple AirTags. Presentation at DEF CON 2021, August 2021.
- [50] Samsung Electronics Co., Ltd. . How do I use Point to Share? <https://www.samsung.com/global/galaxy/what-is/uwband/>. Accessed 2021-10-11.
- [51] 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 2021-10-07.
- [52] 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 2021-10-09.
- [53] NXP Semiconductors. NXP and VW share the wide possibilities of Ultra-Wideband's (UWB) fine ranging capabilities. <https://www.nxp.com/company/about-nxp/nxp-and-vw-share-the-wide-possibilities-of-ultra-widebands-uwband-fine-ranging-capabilities:NW-VOLKSWAGEN-SHOWCASES-UWB>. Accessed 2021-10-12.
- [54] NXP Semiconductors. NXP Announces New Automotive Ultra-Wideband Chip Capable of Turning Smartphones into Car Keys. <https://www.nxp.com/company/about-nxp/nxp-announces-new-automotive-ultra-wideband-chip-capable-of-turning-smartphones-into-car-keys:NW-AUTOMOTIVE-ULTRA-WIDEBAND>. Accessed 2021-10-12.

- [55] NXP Semiconductors. NXP Secure UWB deployed in Samsung Galaxy Note20 Ultra Bringing the First UWB-Enabled Android Device to Market. <https://www.nxp.com/company/about-nxp/nxp-secure-uwband-deployed-in-samsung-galaxy-note20-ultra-bringing-the-first-uwband-enabled-android-device-to-market:NW-SECURE-UWB-SAMSUNG-GALAXY>. Accessed 2021-10-05.
- [56] NXP Semiconductors. Secure Ultra-Wideband Developer Resources. <https://www.nxp.com/products/wireless/secure-ultra-wideband-uwband:UWB-TRIMENSION>. Accessed 2021-10-08.
- [57] Mridula Singh, Patrick Leu, and Srdjan Capkun. UWB with pulse reordering: Securing ranging against relay and physical-layer attacks. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019.
- [58] 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 *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.
- [59] Sky News. Police warn of rise in keyless car thefts as CCTV shows thieves stealing Mercedes in 60 seconds. <https://news.sky.com/story/police-warn-of-rise-in-keyless-car-thefts-as-cctv-shows-thieves-stealing-mercedes-in-60-seconds-12361152>. Accessed 2021-10-09.
- [60] Milan Stute, Sashank Narain, Alex Mariotto, Alexander Heinrich, David Kreitschmann, Guevara Noubir, and Matthias Hollick. A Billion Open Interfaces for Eve and Mallory: MitM, DoS, and Tracking Attacks on iOS and macOS Through Apple Wireless Direct Link. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 37–54, Santa Clara, CA, August 2019. USENIX Association.
- [61] Roel Verdult, Flavio D. Garcia, and Josep Balasch. Gone in 360 seconds: Hijacking with hitag2. In *21st USENIX Security Symposium (USENIX Security 12)*, pages 237–252, Bellevue, WA, August 2012. USENIX Association.
- [62] West Mercia Police. Car theft prevention advice. <https://www.westmercia.police.uk/news/west-mercia/news/2022/january/car-theft-prevention-advice/>. Accessed 2022-02-03.