DISS. ETH NO. 29036

Secure Ranging: Physical-Layer Attacks and Countermeasures

Patrick Leu

Doctoral Thesis

2023

DISS. ETH NO. 29036

Secure Ranging: Physical-Layer Attacks and Countermeasures

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH Zürich

(Dr. sc. ETH Zürich)

presented by

Patrick Leu

MSc ETH in Electrical Engineering and Information Technology
ETH Zürich

born 08.04.1991
citizen of Hohenrain LU

accepted on the recommendation of

Prof. Dr. Srdjan Čapkun, examiner
Prof. Dr. Ivan Martinovic, co-examiner
Prof. Dr. Aanjhan Ranganathan, co-examiner
Prof. Dr. Patrick Traynor, co-examiner

2023

Abstract

Over the past years, applications that use relative distance estimates between devices have seen widespread adoption. How devices can establish a notion of distance or relative proximity is an ongoing concern. This question is critical if supposed physical proximity between a user and a verifying device grants access to a facility or vehicle or releases a payment.

With the Message Time of Arrival Code (MTAC), we propose formal definitions to secure ranging against physical-layer attacks. These definitions encompass a detailed attacker model and requirements for secure signal modulation and reception. We then propose and evaluate a physical-layer procedure that combines security against physical-layer distance-reduction attacks with robust bit transmission.

The practical relevance of Impulse-Radio Ultra-Wide Band (IR-UWB) signals for these types of applications increased significantly with the recent rollout of dedicated chips (e.g., Apple's U1) for high-precision distance measurement using Ultra-Wide Band (UWB) signals.

We describe the current directions in IR-UWB ranging techniques in terms of their physical-layer security. We report and evaluate a practical attack on widely deployed ranging technology. Specifically, we demonstrate a practical distance reduction attack against pairs of Apple U1 chips (embedded, e.g., in iPhones and AirTags), as well as against U1 chips inter-operating with chips from different manufacturers. 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.

As opposed to IR-UWB, Orthogonal Frequency-Division Multiplexing (OFDM) systems have seen widespread use over the past decades, e.g., in WiFi and cellular systems. OFDM transmits data over multiple subcarriers in parallel, thus providing high resilience against frequency-

dependent channel drops (fading) and achieving high throughput.

This makes OFDM systems not a natural fit for secure ranging, as long symbols allow an attacker longer observation and reaction times to mount an Early Detect/Late Commit (ED/LC) attack. Despite these concerns, a recent IEEE standardization effort envisions the use of OFDM-based signals for secure ranging. The work presented in the last part of this thesis provides an analysis OFDM Time of Flight (ToF) measurements and studies the security guarantees of OFDM-based ranging against a physical-layer attacker.

Zusammenfassung

Anwendungen, welche auf Distanz- oder Näheinformation zwischen mobilen Endgeräten basiern, haben in den letzten Jahren stark an Bedeutung gewonnen. Die Sicherheit der Distanzmessung bleibt dabei eine Herausforderung. Insbesondere deren Integrität ist von spezieller Bedeutung, wenn auf Grund der vermeintlichen Nähe zweier Geräte ein Zutritt zu einem Fahrzeug oder Gebäude gewährt, oder eine Zahlung ausgeführt werden kann.

Mit dem Begriff des Message Time of Arrival Code (MTAC) definieren wir die Anforderungen an die Signalmodulation, welche Integrität Distanzmessung mit hohen Sicherheitsanforderungen gewährleisten kann.

Die praktische Relevanz von Ultrabreitbandsignalen für obige Anwendung hat in jüngerer Zeit stark zugenommen, seitdem Hersteller dedizierte Chips für hochpräzise Distanzmessung basierend auf Ultrabreitbandsignalen in mobilen Endgeräten verbauen.

Wir beschreiben den aktuellen und vergangenen Stand der Technik der sicheren Distanzmessung mit Ultrabreitbandsignalen. In der Tat ist der aktuell auf dem Apple U1 Chip implementierte Ansatz angreifbar, d.h., die vermeintlich gemessene Distanz durch einen Angreifer reduzierbar. Wir beschreiben und evaluieren einen praktischen Angriff auf Endgeräte mit U1 Chips (iPhones und AirTags), der kein Wissen über kryptographische Schlüssel voraussetzt. Diese Chips wurden in einem breiten Segment von mobilen Endgeräten und Automobilen verbaut für die Sicherung von passiven Zutritts- und Startsystemen und sind Kandidaten für künftigen Einsatz für kontaktloses Zahlen, Gebäudesicherung und Contact Tracing.

Im Gegensatz zu der Ultrabreitbandtechnologie ist das orthogonale Frequenzmultiplexverfahren schon über Jahrzehnte in breitem Einsatz, sei es für WiFi oder im Mobilfunk. In jüngerer Zeit ist auch in diesen Bereichen eine Entwicklung zur Integration von Mechanismen für die genaue und sichere Distanzmessung festzustellen. Das orthogonale Frequenzmul-

tiplexverfahren überträgt Daten über mehrere orthogonale Träger, was eine vereinfachte Kompensation von Kanaleffekten und somit hohe Datenraten erlaubt. Da diese Technologie primär für Performanz ausgelegt ist, wirft deren Anwendung für sichere Distanzmessung Fragen auf. Die relativ langen Symbolintervalle und zeitliche Redundanzen machen diese Technologie empfindlich gegenüber Angriffen auf der Bitübertragunsschicht. Im letzten Teil dieser Arbeit analysieren wir die Sicherheit dieser Technologie gegen solche Angriffe.

Acknowledgements

First and foremost, I would like to thank my thesis supervisor, Prof. Dr. Srdjan Čapkun, for his guidance and support.

I would also like to thank the members of my doctoral committee, Prof. Dr. Aanjhan Ranganathan, Prof. Dr. Ivan Martinovic, and Prof. Dr. Patrick Traynor.

My gratitude goes out to everyone I had the pleasure of working with during this process. In particular, Dr. Mridula Singh, Dr. Marc Röschlin, Dr. Giovanni Camurati, Prof. Dr. Kenneth Paterson, Claudio Anliker, Martin Kotuliak, Dr. Jiska Classen, Alexander Heinrich, and Prof. Dr. Matthias Hollick. It was my privilege to work with you on projects that impacted this work.

I would also like to thank all the members of the System Security Group for making this a memorable experience.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

We are witnessing a rapid increase in the widespread use of applications that rely in estimates of relative distance between devices [16, 65, 125]. These estimates can be secured against attacks that replay, relay or spoof signals by measuring the relative distance based on the signal time of flight in a two-way procedure. However, even then, the security depends on the integrity of the measured signal time of arrival. This raises the questions: Can the estimate of the message arrival time be manipulated by an attacker that controls the communication channel? Can an attacker modify this estimate in a way that devices appear closer than they actually are?

Applications that are based on relative position estimates go beyond navigation: The physical proximity between devices is increasingly used as a means for securing access to facilities and vehicles or for contactless payment. This not only creates different and potentially stronger incentives for an attacker, but these more recent use cases are also latency-sensitive and typically not based on continuous position estimates. This increases the need for quantifiable security of individual, short-term ranging procedures and limits the applicability of attack detection mechanisms based on temporal inconsistencies. Moreover, the tolerance for adversarially created discrepancies is only in the order of meters.

Attacks on both navigation and proximity-based systems are not new. Both Global Navigation Satellite System (GNSS) [57, 81, 120] and Passive Keyless Entry and Start (PKES) systems [18, 39, 47, 52, 92, 114, 116] have in the past been found vulnerable to attacks that either spoof or relay signals.

This created increased awareness for the need to put insights from the area of distance bounding and secure ranging into practice. Recent standardization efforts aim to address these types of attacks by using two-way ranging protocols and protected ranging sequences [8, 14].

While this means that the messages used to establish the signal time of flight contain cryptographically generated symbol sequences, the integrity of the Time of Arrival (ToA) also depends on both signal modulation and receiver operation, i.e., physical-layer design.

Application aspects also put constraints on the design of the physical layer and protocols. Navigation applications predominantly use broadcast signals due to a need for global availability and large distances between infrastructure (satellites) and end devices. This approach faces the problem against an attacker that either spoofs or records and replays (relays) the signals. In the case of broadcast positioning, besides signal spoofing [57, 120], as long as an attacker can receive the signals from individual satellites, there is a potential for delaying these signals individually, and modification of position estimates. With more direct integration of localization and operation of vehicles (autonomous driving, ACC, platooning), security guarantees on the location information need to be strong, creating a need for other approaches, i.e., using direct car-car communication or based on existing terrestrial infrastructure. However, these communication systems are typically designed for high data rates, meaning the modulation is not designed for secure and precise distance measurement. On the other hand, PKES, access and payment systems require only short communication distances and lower signal power, yielding more flexibility in the physical-layer design since they can use already licensed segments of the spectrum. This makes this scenario not only a premier candidate for a two-way approach, but also for the use of UWB technology.

There has been extensive research on protocols that actively involve both parties in the distance measurement [17,22,27,38,54,71,72,102]. The overall Round-Trip Time (RTT) measurement of such a procedure captures the ToF in both directions. This is required in critical applications, i.e., those for PKES systems [21].

Current implementations for the latter rely on UWB signals with operating distances in the order of tens of meters [7, 85]. The space of UWB communications has itself seen an interesting development over the past one and a half decades. Initially touted as a technique to be used in sensor networks for its ability to resist narrowband fading [127], its applications for that purpose remained limited. However, its core

benefit, precise (i.e., centimeter-level) time-of-arrival estimation, resulted in widespread adoption recently with a push by phone manufacturers to add UWB capability for low-range, high-precision ranging in mobile phones, small tags, and everyday appliances [9,100,105,106]. This creates a need for signal modulations and receiver operations that can guarantee the integrity of the ToA measurement.

While navigation and PKES are without doubt high-stakes use cases, the practical relevance of secure distance estimation may not remain limited to these. Quite recently, exposure notification protocols [59,61, 121] found widespread use in proximity-tracing applications. Moreover, high-precision spacial awareness could in the future also become useful for augmented reality applications [65].

1.1 Contributions

This work presents the following main contributions:

Message Time of Arrival Codes With MTAC we formalize the problem of secure ToA acquisition. We provide security definitions for both the problem of distance advancement and distance enlargement. With the concept of the MTAC, we are able to formally define the security requirements of physical-layer measures that protect ToA measurement systems against attacks. We use our perspective to systematically explore the trade-offs between security and performance that apply to all signal modulation techniques enabling ToA measurements.

Variance-based MTAC We propose a physical-layer procedure that addresses the fundamental problem of reconciling robust transmission and ToA security, targeted to scenarios without meaningful inter-pulse interference. The proposal essentially measures the quality of the (randomized) de-spreading operation. We show that this achieves the security level of the underlying bit transmission under a meaningful trade-off between security and performance.

Practical attack on IEEE 802.15.4z High-Rate Pulse Repetition Frequency (HRP) ranging We report and evaluate a practical distance-reducing attack on an IR-UWB ranging system. These systems are supposed to be protected by a cryptographically generated correlation sequence. We show that, during the attack, the attacker advances the

ToA of ca. 4 % of distance reports. We discuss underlying causes and point towards directions for improving the security level.

PHY Security Analysis of OFDM multicarrier ranging OFDM is a widely used modulation scheme. It transmits data over multiple subcarriers in parallel, which provides high resilience against frequency-dependent channel drops (fading) and achieves high throughput. Due to the proliferation of OFDM-enabled devices and the increasing need for location information, the research community has suggested using OFDM symbols for secure (ToF) distance measurements. However, a consequence of relying on multiple subcarriers is long symbols (time-wise). This makes OFDM systems not a natural fit for secure ranging, as long symbols allow an attacker longer observation and reaction times to mount a so-called early-detect/late-commit attack. Despite these concerns, a recent standardization effort (IEEE 802.11az [14]) envisions the use of OFDM-based signals for secure ranging. This chapter studies the security guarantees of OFDM-based ranging against a physical-layer attacker. We use Binary Phase Shift Keying (BPSK) and 4-Quadrature Amplitude Modulation (QAM), the most robust configurations, as examples to present a strategy that increases the chances for early-detecting the transmitted symbols. Our analysis shows that such OFDM systems are vulnerable to ED/LC attacks, irrespective of frame length and number of subcarriers. We identify the underlying causes and explore a possible countermeasure, consisting of orthogonal noise and randomized phase.

1.2 Thesis Organization

This thesis first covers the necessary background on both protocols for secure distance estimation (secure ranging) and their guarantees. This also means identifying the challenges on the underlying physical layer design, i.e., how the main requirements on modulations such that the properties of these protocols hold. Chapter 3 introduces the notion of MTAC, i.e., algorithms for signal generation and verification that provide the necessary security properties. In Chapter 4, we cover a current IR-UWB implementation. In Chapter 5, we investigate the security of ranging using multicarrier OFDM signals.

1.3 Collaboration

Some of the presented results were obtained in collaboration with co-authors of the respective publications. In particular, the method for obtaining the raw distance measurements from the relevant iOS logs, as presented in the evaluation of Chapter 4, was provided by Jiska Classen and Alexander Heinrich. Some of the attack results in the same chapter (device combinations) refer to measurements performed by Alexander Heinrich, based on an attack implementation provided by the author.

1.4 Publications

This thesis is based on the following publications:

- 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 USENIX Security, 2022 (*equal contribution)

- Patrick Leu, Martin Kotuliak, Marc Roeschlin, Srdjan Čapkun
  Security of Multicarrier Time-of-Flight Ranging
  In Annual Computer Security Applications Conference (ACSAC), 2021

- Patrick Leu, Mridula Singh, Marc Roeschlin, Kenneth G. Paterson, Srdjan Capkun
  Message Time of Arrival Codes: A Fundamental Primitive for Secure Distance Measurement
  In IEEE Symposium on Security and Privacy (S&P), 2020

In addition, during my PhD, I contributed to the following publications:

- Simon Erni, Martin Kotuliak, Patrick Leu, Marc Röschlin, Srdjan Capkun
  AdaptOver: Adaptive Overshadowing Attacks in Cellular Networks
  In Annual International Conference on Mobile Computing and Networking (MobiCom), 2022

- Martin Kotuliak, Simon Erni, Patrick Leu, Marc Röschlin, Srdjan Capkun
  LTrack: Stealthy Tracking of Mobile Phones in LTE
  In USENIX Security, 2022

- Mridula Singh*, Marc Röschlin*, Ezzat Zalzala*, Patrick Leu, Srdjan Čapkun
  Security Analysis of IEEE 802.15.4z/HRP UWB Time-of-Flight
  In ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), 2021 (*equal contribution)

- **Decentralized Privacy-Preserving Proximity Tracing**
  Carmela Troncoso, Mathias Payer, Jean-Pierre Hubaux, Marcel Salathé, James Larus, Edouard Bugnion, Wouter Lueks, Theresa Stadler, Apostolos Pyrgelis, Daniele Antonioli, Ludovic Barman, Sylvain Chatel, Kenneth Paterson, Srdjan Capkun, David Basin, Jan Beutel, Dennis Jackson, Marc Roeschlin, Patrick Leu, Bart Preneel, Nigel Smart, Aysajan Abidin, Seda Gürses, Michael Veale, Cas Cremers, Michael Backes, Nils Ole Tippenhauer, Reuben Binns, Ciro Cattuto, Alain Barrat, Dario Fiore, Manuel Barbosa, Rui Oliveira, José Pereira.
  In IEEE Data Engineering Bulletin 43, 2020

- Mridula Singh, Patrick Leu, AbdelRahman Abdou, Srdjan Capkun
  UWB-ED: Distance Enlargement Attack Detection in Ultra-Wideband
  In USENIX Security, 2019

- Mridula Singh, Patrick Leu, Srdjan Capkun
  UWB with Pulse Reordering: Securing Ranging against Relay and Physical Layer Attacks
  In Proceedings of the Network and Distributed System Security Symposium (NDSS), 2019

- Patrick Leu, Ivan Puddu, Aanjhan Ranganathan, Srdjan Čapkun
  I Send, Therefore I Leak: Information Leakage in Low-Power Wide Area Networks
  In ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), 2018

Chapter 2

Secure Ranging

2.1 Introduction

Mobile devices play an important role in our daily lives, to the extent that we increasingly rely on them to interact with the physical world. Numerous mobile applications, such as navigation, Passive Keyless Entry and Start systems, are based on a notion of relative position or proximity between devices. This notion can be established by cryptographic protocols that contain a timed protocol exchange. Based on the RTT of a correct exchange, a verifier can determine an upper bound on the distance to another entity. Ideally, this allows, for instance, the secure operation of access systems without the need for any user interaction.

All of these approaches have in common that they derive a physical statement (proximity, distance) from a physical measurement of the signal Time of Arrival. Protocols that use this measurement for distance bounding are well studied. However, how exactly this measurement is established is typically out of scope in the design of such protocols, and implementation non-idealities and performance constraints can break with some of the assumptions made.

With protocols for distance bounding and secure ranging in which both parties actively take part in the exchange finding their way into real-world applications, the focus shifted from protocol-level and relay attacks to attacks that target the physical measurement of the signal's ToA in more intricate ways. Such attacks typically exploit temporal redundancies of a signal, a technique to achieve performant operation. The requirements

for performance (i.e., high precision and range), and security are not straightforward to reconcile. While UWB is well suited for high precision, its operating range is limited due to its use of high bandwidth in licensed spectrum, consequently under tight regulatory constraints on signal power. This causes tension between robust bit transmission and physical-layer signal integrity required for secure verification of a signal's ToA. On the other side, technologies classically used for high throughput, namely OFDM systems, are increasingly also used for localization, bringing the question to which extent these modulations can provide ToA security.

This work addresses the problem of translating the security guarantees on the protocol level to the physical layer, also unveiling vulnerabilities in current implementations. The work presented here is very much aligned with recent and ongoing standardization and specification efforts. In the following, we will cover the relevant developments in this space. Then, we will cover the fundamentals of distance bounding and secure ranging. A secure ranging system needs to combine cryptographic and protocol aspects of distance bounding and authenticated ranging with communication and signal processing techniques that allow for precise ToF measurement. We will highlight the relevance of physical-layer concerns since the security and correctness of the output of any such protocol depends on the validity of the ToA established by the receiver. We will also cover different adversarial behaviors and known physical-layer attacks.

2.1.1 Industry developments

In the past years, we have seen accelerated adoption of applications that rely on secure proximity verification [40] with the potential use case of PKES, access systems, and contactless payment. Major mobile manufacturers, like Apple [9, 60] and Samsung [100], rolled out the capability for precise distance measurement in their latest devices. Compatibility with other chip manufacturers [106] ensures that this technology can be used for accessing cars in the future [21, 105]. These developments come in the wake of the IEEE 802.15.4z standardization effort [8] that specifies ranging procedures and modulations for secure ranging.

Especially if channel aspects require performance optimizations, physical-layer security heavily depends on the receiver design. IEEE considers this mostly to be out of the scope of a standard but subject to competition between vendors. However, in any application that offers interoperability between different manufacturers (e.g., phone-car-interaction [104, 105]),

Figure 2.1: Mafia Fraud attack scenario against secure ranging. The attacker is located between prover and verifier and does not possess any knowledge of the cryptographic material.

an insecure acquisition procedure on one side will lead to an insecure outcome on both sides of the protocol. I.e., the lack of security of one participating receiver cannot be compensated by a secure acquisition procedure on the other end.

2.2 Attacker Models

The goal of any attacker is to make the prover appear to be closer to the verifier than it actually is. In distance bounding and secure ranging, different attacker models capture varying trust assumptions on the prover [13]. Although later chapters exclusively focus on Mafia Fraud, we introduce other attacker models for context.

In general, the attacker is assumed to be able to eavesdrop on, intercept, modify or inject messages. The extension of this attacker model to the physical layer will be introduced in Chapter 3.

Mafia Fraud First described in [36], this attacker consists of a fraudulent prover and verifier that aim to convince the legitimate prover of a wrong (i.e., reduced) ToA of a legitimate prover. As illustrated in Figure 2.1, one can imagine this attacker to be in between the legitimate entities, i.e., to act as a Machine-in-the-Middle (MitM). The attacker possesses no knowledge of the challenge or response bits transmitted by the legitimate verifier and prover.

Distance Fraud A dishonest prover aims to convince the verifier of being closer than it is. The fact that the prover cannot be trusted, in particular, means that the prover cannot be trusted with responding after a pre-agreed delay, which requires a fast reply time to the incoming challenge.

Terrorist Fraud A dishonest prover collaborates with an external attacker to appear closer than it actually is (without disclosing the secret to the attacker) [19].

Distance Hijacking A dishonest prover replaces the signature of an honest prover with its own, thereby linking the distance claim of a prover to its own identity [32]. This attack only depends on protocol aspects and is, therefore, orthogonal to physical-layer design.

2.3 Protocols

Protocols for secure ranging and proximity verification can address different trust assumptions. Typically, the term distance bounding is used for protocols that resist a compromised prover as well as an external adversary. In the literature, the term authenticated ranging is sometimes used if only a Mafia Fraud model is considered. Both protocol types have in common that they securely establish proximity by exchanging cryptographically generated bit sequences and either establishing or validating the RTT (and, hence, the relative distance) based on the ToA of these sequences.

2.3.1 Distance Bounding

Distance bounding protocols securely establish whether a prover (responder) is within a certain distance to a verifier (initiator) by upper bounding the relative distance based on the RTT of a protocol exchange. While an adversary can always delay signals and, this way, make the distance appear larger, distance bounding protocols prevent an attacker from reducing the measured distance. Different distance bounding protocols have been proposed in the literature [22,38,54,102], some of them addressing an external adversary (Mafia Fraud), most of them, in addition, a dishonest prover (Distance Fraud), or more complex trust assumptions [63]. Typically, the term distance bounding is used for protocols that resist both an external and a compromised (internal) adversary [27]. These protocols all consist of an initialization phase, a timed (fast) challenge-response phase, and a verification phase. In the initialization phase, both the prover and verifier generate fresh, random nonces. The prover is required to send a commitment to this nonce to prevent an opportunistic choice after the timed exchange. After the timed challenge-response phase,

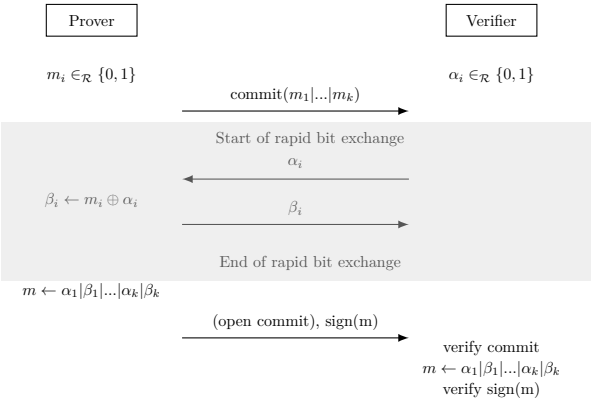

Figure 2.2: Distance bounding with rapid bit exchange, adapted from [22], with time-critical part highlighted

the prover opens the commitment. The verifier checks whether the fast reply sequence matches the committed value and verifies the authenticity of the prover's record of the fast exchange. For the RTT of the timed exchange to be a secure bound on relative proximity, operations of the receiver need to be executed either with a small delay (if the prover is not trusted) or the response has to occur some agreed reply time after the challenge (if the prover is trusted). On the physical layer, a protocol addressing an untrusted prover can be implemented based on a rapid bit exchange. This means the ranging procedure consists of a series of single bits that the prover responds to immediately after an XOR operation. This is illustrated in Figure 2.2. The verifier can then establish a secure upper bound on the relative distance based on the RTT, provided that verification was successful. A distance bounding protocol that is resistant to both Mafia Fraud and Distance Fraud and satisfies the requirement on fast reply time is challenging to implement based on standard demodulation techniques. Existing proposals that aim to minimize processing time are based on fully analog processing of the challenge pulses [94], or a fast reply logic based on energy detection [93]. Direct implementation of rapid bit exchange and integration into a ranging system was shown to be feasible but associated with more than 10 m potential distance reduction [118].

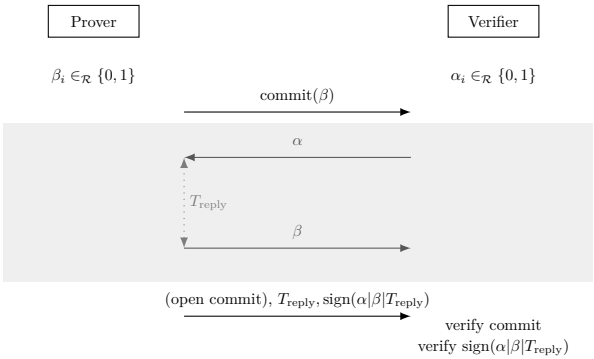

Figure 2.3: Authenticated Ranging protocol, with time-critical part highlighted

2.3.2 Authenticated Ranging

Protocols that address Mafia Fraud only, i.e., assume a trusted prover, are sometimes referred to as authenticated ranging protocols [17, 27, 71, 72]. This trust assumption essentially removes the requirement for a fast reply since the prover is entrusted with adhering to a given (either pre-agreed or reported) reply time. Since the verifier subtracts this interval from the RTT before calculating the distance, operations that occur within this interval are not time-critical in the sense that they have to be as fast as possible, but only faster than T_{reply}. Hence, this trust assumption allows for more flexibility in the choice of modulations, frame formats, and the design of receivers.

An example of an authenticated ranging protocol is shown in Figure 2.3. Both the verifier and prover generate fresh, unguessable nonces. Then, the prover commits to its nonce. For ranging, the verifier sends its nonce α to the prover. The prover registers the ToA and replies with its nonce β precisely T_{reply} later. For calculating the relative distance, the verifier subtracts the reply time from the measured RTT. In the final step, the verifier checks the commitment, as well as the authenticity of the reply time and the prover's observation of the nonces.

Real-world secure ranging protocols typically implement a variation of an authenticated ranging protocol, i.e., are designed against Mafia Fraud. This can be explained by the important use cases of PKES and contactless payment, where the concern is a theft or a fraudulent payment initiated

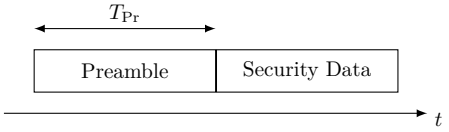

Figure 2.4: Distance Commitment: The publicly known preamble is followed by the security data at a pre-agreed time interval.

by an external adversary. In practice, the term distance bounding is often used to include protocols that technically (in the literature) fall under the term authenticated ranging.

2.4 Secure Ranging

Irrespective of the protocol choice, the core statement of any distance bounding or authenticated ranging protocol is based on a measurement of the signal ToA. Realizing the protocol guarantees in a practical ranging system is not straight-forward since other requirements, e.g., related to ranging performance and operation range play into it. Previous work has shown that careful physical-layer design is important for the resulting security [30, 92], otherwise a modulation can become vulnerable to physical-layer distance modification attacks [82]. In the following, we highlight the known design principles for secure and precise ToA establishment in the context of practical ranging systems.

2.4.1 Distance Commitment

Distance-bounding protocols typically assume a rapid bit exchange, i.e., a relatively high number of consecutive exchanges [94]. Combining such a protocol with precision ranging, requires a fine-grained acquisition process. A preamble, used for precise ToA acquisition, can precede the security bits at a precisely pre-agreed offset. This is a concept known as distance commitment [117, 118] and illustrated in Figure 2.4. Using a preamble before the cryptographic information improves the precision and efficiency of the ToA acquisition in a ranging system since it allows to separate signal and ToA acquisition from the validation of the cryptographic sequence. In the case of a Mafia Fraud attacker, the challenge and the response frame do not need to be interleaved. This also corresponds to

the typical design of the time-critical messages in authenticated-ranging protocols. Under this attacker model, the responding device is trusted, in particular, entrusted with adhering to a time schedule in both reception and transmission, be it for validation of the distance commitment in reception or to wait for a specific time interval before responding. This assumption allows more flexibility in the modulation of the signal since there is less emphasis on the receiver-side operations being time-critical.

2.4.2 Receiver Operation

Physical-layer security of ranging is concerned with the integrity of a ToA measurement. While this places constraints on how signals need to be modulated, irrespective of the signal modulation, the security also critically depends on the acquisition and detection logic. The receiver determines the ToA as a physical measurement based on a received signal. Therefore, receiver design is of crucial importance for the security of the resulting distance estimation. A receiver has to perform multiple tasks in order to provide a precise ToA measurement, robust data demodulation, and quantifiable (ToA) security. Typically, these tasks are associated with different parts of a frame that transmitted during a ranging procedure.

The receiver has to perform a series of operations for deriving a secure ToA estimate from an incoming ranging frame. Typical steps are the following:

Coarse-grained acquisition A receiver detects the existence of a frame on the medium. Typically, the receiver detects an initial, publicly known part of a frame by correlating it with a template. Since, in this initial step, there is no requirement on ToA security, the sequence (preamble) can be selected for good correlation properties which allows for efficient acquisition.

Fine-grained acquisition, ToA estimation Fine acquisition based on a known preamble yields the precise ToA, which translates to precise distance estimation. High signal bandwidth allows for high ToA precision. In addition, this step can be used for precise estimation of clock drift, which is of particular importance if the data is modulated coherently (i.e., if information is encoded in the phase of the signal).

(Robust) demodulation The receiver demodulates information bits that follow the preamble. If these bits are related to the validation of the ToA, it is important that the information is received at precise, pre-agreed time-intervals after the preamble (distance commitment). To this end, typical performance-enhancing techniques, such as channel equalization cannot be directly applied. These techniques typically seek to invert the dispersion profile of the channel, by combining different copies of the same pulse, which is at odds with the principle underpinning the distance commitment.

Validation of ToA The receiver checks whether a cryptographically generated symbol sequence occurs at the precise time consistent with the resulting ToA statement. Validation can be done by demodulation, which results in a directly quantifiable security level or based on a different similarity score that is computed over an entire symbol sequence.

2.5 Physical-layer attacks

We consider a Mafia Fraud attacker that aims to reduce the ToA of the signal. The attacker can be assumed to be "in between" the two communicating parties. The attacker can record legitimate signals and react to observations by transmitting its own signals. If the two devices are out of range, the attacker can establish a communication path with a relay. We will cover this attacker model in the physical-layer context in more detail in Chapter 3. While attacks that relay, replay and spoof entire frames have become a widespread concern, there are more intricate ways to interact with the physical representation of an information frame "in transit". Equipment to receive and generate arbitrary waveforms, e.g., Software Defined Radios (SDRs), have become more affordable and easier to use in recent years [97], increasing the importance of physical-layer concerns.

2.5.1 Relay

A relay allows the attacker to establish a channel between devices that otherwise might be out of range [55], as shown in Figure 2.5. An attacker can do this either with a wired connection [47], or using wireless relay devices [98], i.e., one device in proximity to each communicating party. If a system relies upon the ability to communicate as a notion of proximity,

Figure 2.5: Relay attack: The attacker establishes a communication path between prover and verifier. Using amplification and a wired or wireless connection, the attacker can bridge an arbitrary distance and circumvent any propagation obstacles of the actual channel.

this can be directly circumvented by a relay attack. In ToF systems, a relay allows the attacker to establish a communication path over the fastest possible medium (wormhole attack), which is of direct utility if the intended communication medium is slower than the speed of light (e.g., ultrasound [103]).

Moreover, a relay can be considered an auxiliary technique combined with any other attack. It allows the attacker to selectively pass through parts of the legitimate signal and impose a channel at will, e.g., perform various analog operations on the incoming signal—the simplest being amplification (to circumvent a system that relies on Received Signal Strength Indicator (RSSI)). Moreover, an attacker can change the signal phase [79] or apply an arbitrary (causal) filter, to impose a particular channel characteristic. This last point is critical regarding the security of potential performance-enhancing (channel equalizing techniques), resp., matched filtering based on a channel estimate. Relay attacks have become a practical concern in both PKES [47,114,116] and payment systems [48]. Proximity verification that relies on the signal ToF is not vulnerable to a simple relay. However, the ability of an attacker to selectively relay parts of a signal underpins the Dolev-Yao attacker model [37] on the physical layer and enables more complex attacks.

2.5.2 Early-detect/Late-commit

A distance-reducing attack against a ToF ranging system can succeed if the notion of ToA can be manipulated, i.e., reduced. An ED/LC attacker reduces the measured distance by preemptively injecting a non-committal waveform that triggers an early signal detection at the receiver [30,46]. The goal is to cause the receiver to register an earlier time of arrival, even if the attacker cannot generate the full symbol in advance due to a lack of knowledge about the signal content. We illustrate this attack in Figure 2.6. The attack works by exploiting non-idealities of the receiver,

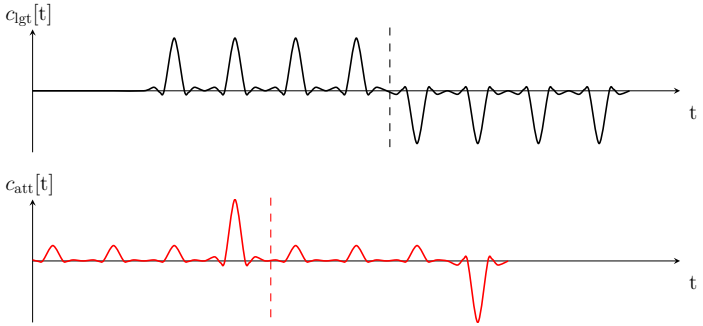

Figure 2.6: ED/LC attack: The attacker transmits its symbols earlier and compensates a guessing error with a correct ending once the value of the legitimate symbol is known.

requiring it to integrate signal power over time for each bit-wise decision, effectively limiting its temporal resolution. To compensate for early deviations from the legitimate symbol (i.e., guessing errors), the attacker significantly amplifies its signal towards the end of each symbol. For maximum effect (distance reduction), the attacker sends the committal, information-bearing part as late as possible after the start of the injected signal. Ideally, this is done to precisely coincide with the start time of the legitimate signal so that the attacker can "copy" its content (with amplification). An ED/LC attack, as shown in Figure 2.6 can be executed with 100% success probability, given a fast enough reaction time of the attacker, and can lead to a distance reduction up to the product of the symbol duration and the speed of light.

2.5.3 Guessing attacks

If the polarity of individual pulses constituting a modulated symbol does not only depend on the bit-value of the symbol, e.g., by being fully randomized, the attacker can resort to a probabilistic ED/LC attack, i.e., a guessing attack. Here, the attacker tries to guess signal components in advance in order to reduce the measured distance. As in an ED/LC attack, the attacker exploits signal redundancies that are required for robust signal demodulation. The basic idea is that the attacker can compensate for early guessing errors by using more power towards the end of the symbol and, in this way, increase its success probability. For

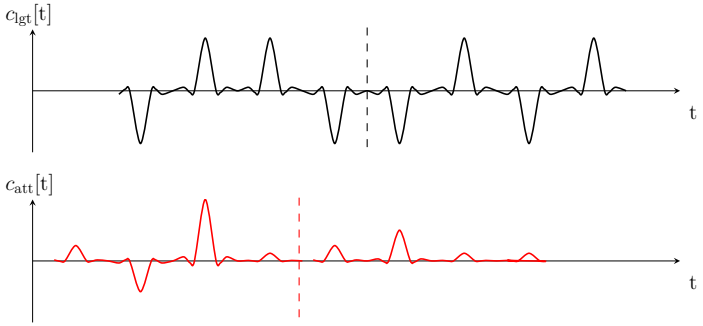

Figure 2.7: Guessing attack with power-increase strategy: The attacker transmits its symbol earlier and compensates wrong guesses by increasing the power until correct.

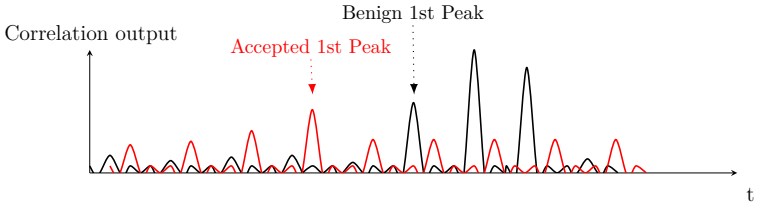

Figure 2.8: Cicada attack: An attacker transmits a random pulse train (coarsely time aligned with the legitimate signal) that creates side peaks after cross-correlation. One of these side peaks can be mistakenly classified as corresponding to an early path of the legitimate signal.

each symbol, the attacker can, for instance, double the power after wrong guesses and stop interfering as soon as a pulse is guessed correctly.

2.5.4 Cicada attacks

The wireless channel causes the signal to arrive at the receiver over different paths. This means a receiver encounters a superposition of multiple time-shifted copies of a signal. Under such multipath conditions, it can be difficult for a receiver to detect the leading edge, i.e., the ToA of the signal corresponding to the direct (i.e., first) path. In particular, if different parts of a frame occur close enough in time to collide under

channel effects. This is the case if symbols are separated in time by less than the power delay profile of the channel, causing Inter-Pulse Interference (IPI). An attacker can exploit this ambiguity by transmitting a pulse train over the legitimate signal to raise the noise floor and create fake (early) correlation peaks that are misclassified as the signal leading edge [83,84,113]. This attack principle was first introduced as the Cicada attack [84].

Proposed countermeasures against this type of attack encompass quantization of the incoming signal before correlation [42,45], or detection of the randomness of consecutive (disturbed) distance measurements (i.e., detection based on the variance of the inter-arrival time) [83].

A quantization approach is associated with information loss, which can be a problem if there is significant self-interference. Managing the tension between performance and ranging security under significant self-interference has not been fully addressed and is still an ongoing question.

In contrast to guessing attacks, Cicada attacks exploit the back-search algorithm, i.e., a receiver's decision based on the presumed path-delay profile, and are facilitated by the existence of a correct signal copy that occurs later on the medium.

2.6 Summary

Secure distance estimation and proximity verification are increasingly used in real-world applications. The security against physical-layer attacks needs to be considered independently of protocol aspects. Protocol design typically assumes that the ToA of a bit sequence can be unambiguously acquired and verified. Putting distance bounding and authenticated ranging into practice can lead to difficult trade-offs, mainly between the robustness of information transmission and performance (precision) of ToA acquisition, as well as security against various physical-layer attacks. In the context of IR-UWB, performance aspects are governed by strict regulatory requirements.

Chapter 3

Message Time of Arrival Codes

3.1 Introduction

When did the message arrive at the receiver? Can this estimate of the message arrival time be manipulated, and in particular by an attacker that controls the communication channel? In particular, can message advancement and delay attacks be prevented? This question is at the core of the problem that distance bounding protocols, secure positioning, and navigation systems are trying to solve: can we prevent the attacker from reducing or enlarging the distance that is measured between the devices? This problem is relevant in a number of application scenarios: contactless payments [48], PKES [18,47,52,92], GNSS (e.g., Galileo, GPS) security [57, 81, 120]. If we could prevent ToA and therefore distance manipulation attacks, we could enable many proximity-based applications, from location-based access control to secure navigation [24, 95].

As a result, many distance bounding protocols have been proposed and analyzed [22,23,63]. Implementations of distance bounding protocols have emerged that combine such protocols with distance measurement techniques [54, 91, 94, 118], in particular with IEEE 802.15.4 UWB radios [7, 85, 128].

The main idea behind these solutions is to prevent ToA manipulation by the randomization of message content. Namely, it was commonly believed that if the attacker cannot predict the bits of the messages,

then he will not be able to advance their time of arrival at the receiver. In [30] the authors argued this to be false – since bits are encoded into symbols, attackers can advance their arrival time. Different physical-layer attacks followed also validating this in practice [46, 82, 92]. This led to the conclusion that secure distance measurement systems can only be built with short symbols and using rapid bit exchange [30]. Given the limits on the output power, such a result would mean that only short-range systems could be made secure. This was shown to be incorrect in [111], which showed that longer symbols can be used if they are interleaved in transmission in a manner that is unpredictable to the attacker. This further demonstrated that secure, long-range distance measurement systems are possible. Recent works further show that, under certain conditions, distance enlargement can also be detected [110]. All these works showed that consideration of the details of how bits are encoded into symbols (i.e., modulation) is crucial in the design of secure distance measurement systems.

This discussion leads to the following questions:

Can we construct a generic message to symbol encoding that prevents any message advancement/reduction (and therefore distance delay/enlargement) for symbols of arbitrary lengths (and therefore arbitrary measurement ranges)?

Can we derive the main principles for the design of such encodings?

In this work, we show that answering these questions is indeed possible. To do so, we introduce Message Time Of Arrival Codes (MTACs), a new class of cryptographic primitive that allows receivers to verify if an adversary has manipulated the message arrival time. In a similar way that Message Authentication Codes protect message integrity, MTACs preserve the integrity of message arrival times. They are, therefore, fundamental to any protocol that relies on Time of Arrival information, such as clock synchronization [49], distance measurement [56] and positioning protocols [26, 27, 102, 109].

In the same sense that bits can be encrypted with a shared key, the shape of a signal can also be hidden by masking it with a random fast-changing sequence. However, to verify a signal shape, a receiver has to aggregate the signal over a considerable time interval in order to capture enough energy. This is especially so when sender and receiver are separated by longer distances. If the attacker knows the temporal alignment of those aggregations with the signal, the attacker can hide its guessing errors in the null space of the (linear) aggregation function. Simple signal masking is, therefore, not sufficient for the protection against

distance manipulation attacks. To address this problem, in addition to using cryptographically-secured modulation (i.e., signal generation), an MTAC also *performs cryptographic checks of the consistency of the modulation* at the receiver.

In the following, we give a formal definition of MTACs and their security and provide the main principles for the design of these codes. We then introduce a new *Variance-Based MTAC* that is inspired by our design considerations. We show that this allows protection against physical-layer distance-reducing attacks over a wide, realistic performance region and systematically explore the trade-off between performance and security.

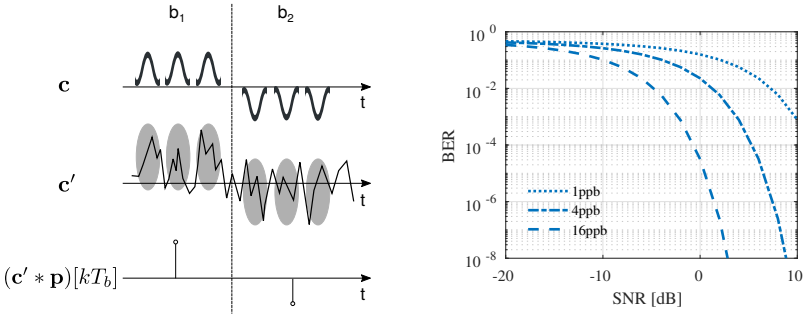

Figure 3.1: Repetition coding: Under noisy conditions, the receiver has to combine multiple short-term signal contributions (samples) to retrieve information. $(\mathbf{c}' * \mathbf{p})$ denotes linear aggregation, e.g., through a matched filter.

3.2 Definitions

3.2.1 System and Attacker Model

We can model any ranging signal as consisting of short-time signal contributions (i.e., pulses) that carry the information used for precise ranging. As shown in Figure 3.1, linear combinations of these pulses provide the statistics for the detection of information bits at the receiver. This model covers a broad range of modulation schemes.

3.2.2 Modulation

In the following, we state some assumptions on the modulation. Following Kerckhoffs' principle, we assume the attacker to be aware of all of these aspects of the modulation.

- The modulation consists of a series of elementary, short-time signal contributions called pulses. The effect of ED/LC attacks on such individual signal contributions is considered insignificant (say, less than 1 m) in a system of sufficient bandwidth. We refer to the amplitude level of such a pulse as a *sample*.
- For performance-related considerations, we assume the pulses to be sufficiently spaced such that there is no inter-pulse-interference in the given channel.

- Each information bit is encoded in a symbol consisting of n_{ppb} pulses (and samples). The value of n_{ppb} is chosen in compliance with a target performance level p within a performance region $\mathcal{P} = (0, BER_{max}] \times [0, d_{max}] \times [0, \Gamma_{max}]$, defined by intervals bounded by the maximally tolerated bit error rate BER_{max}, the maximum communication distance d_{max} as well as the maximum NLoS signal attenuation Γ_{max}.
- Bits are grouped together to form frames, and each frame consists of n_p pulses (and hence n_p/n_{ppb} bits).

3.2.3 Receiver Demodulation

We assume the receiver demodulates by aggregating n_{ppb} samples using correlation with a polarity[1] mask that fits the corresponding hypothesis for each possible value of the information bit. Then, a binary hypothesis test is applied to recover each bit. This is illustrated in Figure 3.1. We assume an Additive White Gaussian Noise (AWGN) channel model without inter-pulse interference. In general, the BER at the receiver is therefore given by the tail bound on the Gaussian distribution, i.e.

$$BER = Q\left(\sqrt{\frac{n_{ppb}P_{rx}}{\sigma_n^2}}\right), \tag{3.1}$$

under Gaussian thermal noise with variance $\sigma_n^2 = bW \cdot N_0$, where N_0 is the noise power spectral density at room temperature, bW is the system bandwidth and P_{rx} is the receiver-side signal power. Figure 3.1 highlights the effect of larger n_{ppb} (longer symbols) on BER. This is to highlight the beneficial effect of pulse repetition on performance. Although Equation 3.1 refers to a BPSK modulation, this effect extends to other modulation techniques. We note that, within this model, for any channel and target BER, there exists an adequate symbol length and assume that the receiver chooses the symbol length accordingly. In this work, we do not assume any (error-correcting) coding.

Attacker Model

We assume that the attacker fully controls the communication channel and has no limitations on how fast she can process messages and react

[1] Polarity refers to one of two possible phase values of the sample.

to them. She is, therefore, able to detect individual samples ideally. As a consequence, the attacker's information advantage increases as the channel for legitimate communication worsens, e.g., due to increased distance. We consider two distinct attack models capturing distance reduction (message advancement) and distance enlargement (message delay). In the case of the distance reduction attacker, we pose no restriction on the attacker's abilities regarding the speed of computation,[2] location, or control of the communication channel (e.g., we give the attacker the ability to record and reactively inject messages on the channel with negligible delay). The only restriction that we pose is that the attacker cannot transmit information faster than the speed of light. The attacker's sampling rate needs to be sufficient to recover the signal. For an attack to be effective, we don't need to assume that the attacker has a higher bandwidth since we assume the attacker can precisely synchronize to the start of the signal. For the distance enlargement attacker, we assume that the attacker is constrained in terms of location, computation and control of the environment such that she is only able to block the reception of samples if she can anticipate their polarity. However, this includes attackers that operate with multiple (smart) antennas or increase noise levels at the legitimate receiver.

3.2.4 Definitions

A Message Time of Arrival Code (MTAC) is intended to allow detection of any kind of physical-layer distance-modifying attack with high likelihood.

Definition 1 *A Message Time of Arrival Code (MTAC) is a tuple of probabilistic polynomial-time algorithms (Gen, Mtac, Vrfy), such that:*

1. *The key-generation algorithm Gen takes as input the security parameter n and outputs a key k with $|k| = n$.*

2. *The code-generation algorithm Mtac takes as input a key k and a message $m \in \{0,1\}^{n_b}$ and outputs a real-valued vector $\mathbf{c} = (c_1, \ldots, c_{n_p})$. Since this algorithm may be randomized, we write this as $\mathbf{c} \leftarrow Mtac_k(m)$.*

3. *The verification algorithm Vrfy takes as input a key k, a real-valued vector $\mathbf{c}'$ of length n_p, and message m'. It outputs a bit b. We assume that Vrfy is deterministic, and so write $b := Vrfy_k(m', \mathbf{c}')$.*

[2] Although MTACs can be constructed so as to be information-theoretically secure, most practical schemes will require that the attacker is computationally restricted.

Figure 3.2: The wireless channel introduces a BER at the receiver and limits the rate at which information can be recovered.

In the above definition, we assume that m may be transmitted separately from $\mathbf{c}$; however $\mathbf{c}$ can also 'carry' m, which case we assume the existence of an efficient algorithm to extract m from $\mathbf{c}$. In this situation, we can also assume that m' can be extracted from $\mathbf{c}'$ and could choose to suppress it as an input to *Vrfy*. The value of b output by *Vrfy* is intended to convey that message time of arrival is correct ($b = 1$) or that it cannot be securely verified ($b = 0$).

An MTAC can be seen as a keyed signal verification scheme that guarantees the integrity of the message time-of-arrival. $\mathbf{c} = (c_1, \ldots, c_{n_p})$ is a vector of samples corresponding to the digital representation of the analog signal after A/D conversion.

We make no assumptions on the confidentiality or authenticity of m. We assume that these can be achieved through other means, e.g., using encryption or message authentication codes.

Before information can be verified, it has to be transmitted over a wireless channel and detected by the receiver. Strictly speaking, *Vrfy* involves not only verification but also time-selective *detection* of physical-layer information. As highlighted in Figure 3.2, detection performance and the resulting security level are fundamentally connected. In general, received samples $\mathbf{c}'$ are affected by channel noise and, in consequence, not identical to $\mathbf{c}$. The detection rate R_{out}, which depends on channel and modulation, is the rate of verifiable information at the receiver. Due to temporal aggregation, it is, in general, smaller than the input data rate, i.e., $R_{out} \leq R_{in}$. Within our assumptions, the ratio R_{in}/R_{out} is given by n_{ppb}. Moreover, detection of this information over a channel is error-prone, which is reflected by a nonzero BER. Consequently, an MTAC will have a non-zero likelihood of false negatives, as well. This we address in a verification criterion that we call *robustness*.

Definition 2 *An MTAC is robust if*

1. *In the absence of an attacker, for any channel, Vrfy applied on* $\mathbf{c}'$ *is falsely negative with probability at most* $1 - (1 - BER)^{n_b}$*, where BER is the error rate in detecting the bits carried by* $\mathbf{c}$*.*

This means that the false negative rate should remain bounded by the frame error rate on the bit level. Note that we will impose robustness only on detection of distance advancement. As mentioned earlier, detection of delay attacks involves a multi-hypothesis test in time and is, therefore, inherently more prone to false positives.

Distance modification can mean either distance reduction or distance enlargement. The former requires the attacker to *advance* the signal in time, the latter to *delay* the signal in time. We define two different MTAC security models, one for each type of attack (a single model would be unwieldy and difficult to use).

MTAC-A: Modelling Advancement Attacks

In what follows, $\alpha \geq 0$ denotes the *observation delay* of the adversary, measured in samples, representing how long it takes for an attacker to observe and react to a given sample.[3] On the other hand, $\delta \geq 1$ denotes the number of samples by which the adversary tries to advance the signal, quantifying its attack goal. Informally, we allow the adversary access to MTAC code values $\mathbf{c}$ for message inputs of its choice in a fully adaptive manner. Then we challenge it to produce an "advanced" signal $\mathbf{c}'$ for a message m of its choice. We model the latter by requiring the adversary to produce component $c'_{i+\delta}$ of its output before being given samples $(c_1, \ldots, c_{i-\alpha})$ of $\mathbf{c} = \mathit{Mtac}_k(m)$. The adversary wins if it eventually produces a vector $\mathbf{c}'$ for which $\mathit{Vrfy}_k(m, \mathbf{c}') = 1$. An MTAC scheme is (informally speaking) secure against advancement attacks if the probability that any efficient adversary wins is small.

We formalise these ideas in terms of a message time-of-arrival forgery experiment $\mathit{Mtac\text{-}A\text{-}forge}_{A,\Pi}(n)$. In this experiment:

1. The experiment sets $k \leftarrow \mathit{Gen}(n)$.
2. The adversary A is given oracle access to $\mathit{Mtac}_k()$; let Q of size q denote the set of queries made by A.
3. A outputs m, and the experiment sets $\mathbf{c} = \mathit{Mtac}_k(m)$.
4. A then sequentially outputs real values $c'_1, \ldots, c'_{n_p}$; however, after outputting $c'_{i+\delta-1}$ (and before outputting $c'_{i+\delta}$), A is given the samples $(c_1, \ldots, c_{i-\alpha})$ of $\mathbf{c}$.

[3] Although in our attacker model, we pose no restriction on the adversary's abilities to reactively record and inject samples, α allows us to model weaker attackers whose reaction speed is bounded.

5. Let $\mathbf{c}'$ denote $(c'_1, \ldots, c'_{n_p})$. Then the output of the experiment is defined to be 1 (and A is said to win) if and only if (1) $Vrfy_k(m, \mathbf{c}') = 1$ and (2) $m \notin Q$. Otherwise, the output of the experiment is defined to be 0.

Note that for schemes in which a message m' (possibly different from m) can be extracted from $\mathbf{c}'$, we can define a different win condition: (1) $Vrfy_k(m', \mathbf{c}') = 1$ and (2) $m' \notin Q$. Here, A still outputs a message m for which she receives a delayed version of $\mathbf{c} = Mtac_k(m)$, but she can win by "forging" a code vector $\mathbf{c}'$ for a different message m' altogether.

Definition 3 *Let $\Pi = \{Gen, Mtac, Vrfy\}$ be an MTAC-A, and let A be an adversary with observation delay α and advancement goal δ that makes at most q queries to its MTAC oracle and that runs in time at most t (across all steps of the Mtac-A-forge$_{A,\Pi}(n)$ experiment). The advantage of A is then defined as:*

$$Adv^{MTAC-A}_{A,\Pi}(n) := \Pr[\textit{Mtac-A-forge}_{A,\Pi}(n) = 1].$$

We associate with Π an insecurity function $\boldsymbol{Adv}^{MTAC-A}_{\Pi}(\cdot,\cdot,\cdot,\cdot,\cdot)$, defined as:

$$\boldsymbol{Adv}^{MTAC-A}_{\Pi}(q, t, \alpha, \delta, n) := \max_{A}\{Adv^{MTAC-A}_{A,\Pi}(n)\}$$

where the maximum is taken over all adversaries with observation delay α, advancement goal δ, making at most q queries to its MTAC oracle and running in time at most t.

It is not hard to see that, with all other parameters fixed, the insecurity function is maximised w.r.t. α and δ when $\alpha = 0$ and $\delta = 1$. This corresponds to the situation where the adversary has no observation delay and is given the next sample c_i from $\mathbf{c}$ immediately after outputting its own guess c'_i. The latter corresponds to an adversary who tries to advance the signal by one pulse.

MTAC-D: Modelling Delay Attacks

In the following, we consider an adversary interested in removing all traces of the legitimate signal to perform a delay attack. Under the condition that all evidence of the legitimate signal is removed, the adversary can trivially achieve any delay goal δ without a risk of detection. As the

value of δ does not help or limit the adversary, we are not using it in the model. However, by limiting the observation delay $\alpha \geq 0$, we constrain the attacker in its ability to observe (and suppress) the samples that are transmitted by the legitimate transmitter. Generally, we assume that the attacker will not be able to detect the legitimate sample, transmit an opposite sample and thus suppress the legitimate sample. Informally, we allow the adversary access to MTAC code values $\mathbf{c}$ for message inputs of its choice in a fully adaptive manner. Then, we challenge it to produce an "advanced" signal $\mathbf{c}'$ for the message m of its choice. We model the latter by requiring the adversary to produce component c'_i of its output before being given samples $(c_1, \ldots, c_{i-\alpha})$ of $\mathbf{c} = \mathit{Mtac}_k(m)$, i.e., the adversary needs to produce at least one sample in advance for $\alpha = 0$. The adversary wins if it eventually produces a vector $\mathbf{c}'$ for which $\mathit{Vrfy}_k(m, \mathbf{c}'') = 0$ for $\mathbf{c}'' := \mathbf{c} + \mathbf{c}'$. $\mathit{Vrfy}_k(m, \mathbf{c}'')$ outputs 0 if it does not find a trace of $\mathbf{c}$ in $\mathbf{c}''$ and is unable to detect the existence of $\mathbf{c}'$.

We formalise these ideas in terms of a message time-of-arrival forgery experiment $\mathit{Mtac\text{-}D\text{-}forge}_{A,\Pi}(n)$. In this experiment:

1. The experiment sets $k \leftarrow \mathit{Gen}(n)$.

2. The adversary A is given oracle access to $\mathit{Mtac}_k()$; let Q of size q denote the set of queries made by A.

3. A outputs m, and the experiment sets $\mathbf{c} = \mathit{Mtac}_k(m)$.

4. A then sequentially outputs real values $c'_1, \ldots, c'_{n_p}$; however, after outputting c'_i (and before outputting c'_{i+1}), A is given the samples $(c_1, \ldots, c_{i-\alpha})$ of $\mathbf{c}$. Samples c_i and c'_i arrive at the receiver at same time, resulting in the superposition $c''_i = c_i + c'_i$.

5. Let c'' denote $(c''_1, \ldots, c''_{n_p})$. Then the output of the experiment is defined to be 1 (and A is said to win) if and only if (1) $\mathit{Vrfy}_k(m, c'') = 0$ and (2) $m \notin Q$. Otherwise, the output of the experiment is defined to be 0.

Definition 4 *Let $\Pi = \{\mathit{Gen}, \mathit{Mtac}, \mathit{Vrfy}\}$ be an MTAC-D, and let A be an adversary with observation delay α that makes at most q queries to its MTAC oracle and that runs in time at most t (across all steps of the $\mathit{Mtac\text{-}D\text{-}forge}_{A,\Pi}(n)$ experiment). The advantage of A is then defined as:*

$$\mathit{Adv}^{MTAC-D}_{A,\Pi}(n) := \Pr[\mathit{Mtac\text{-}D\text{-}forge}_{A,\Pi}(n) = 1].$$

We associate with Π *an insecurity function* $\boldsymbol{Adv}_{\Pi}^{MTAC-D}(\cdot,\cdot,\cdot,\cdot)$*, defined as:*

$$\boldsymbol{Adv}_{\Pi}^{MTAC-D}(q,t,\alpha,n) := \max_{A}\{Adv_{A,\Pi}^{MTAC-D}(n)\}$$

where the maximum is taken over all adversaries with observation delay α*, making at most* q *queries to its MTAC oracle and running in time at most* t*.*

With all parameters fixed, the insecurity function is maximized for $\alpha = 0$. This corresponds to the situation when an attacker's observation delay is limited due to its position or hardware capabilities such that he cannot detect the legitimate sample and suppress them when they are already being transmitted. However, he can observe sample c_i from $\mathbf{c}$ immediately after outputting its own guess c_i'.

Practical MTAC instantiations are likely to rely on a scheme to expand some finite sequence of ideal randomness into a longer one, e.g., using Pseudo-Random Function (PRF)s. We note that, in practice, this is the component vulnerable to higher values of q and t. On the other hand, the security of the verification does not necessarily depend on q and t, i.e., is not affected by those under the assumption of ideal randomness going into signal generation. This is equivalent to stating that verification is not necessarily randomized (beyond the randomness in the signal). However, verification has to be reliable given some, within the computational model bounded, knowledge of the attacker about the PRF output used for signal generation.

3.3 MTAC Design Space

In this section, we shift to a statistical viewpoint on the design space of secure MTAC schemes and explain how this approach relates to the computational model presented earlier. A statistical analysis entails the advantage of summarizing the infinite number of possible attack strategies. This is particularly beneficial because legitimate as well as adversarial signals can assume uncountably many realizations due to their real-valued nature and due to the uncertainty introduced by noise. Moreover, an attacker is free to choose any amplitude level for each sample of the transmitted signal. The resulting complexity does not allow a straightforward evaluation of all possible strategies in a closed-form computational setting. Also, the security of the verification procedure itself is best analyzed in information-theoretic terms, since verification itself does not have to be randomized, i.e., its security is not necessarily limited to a bounded adversary. Therefore, we present a signal theoretic approach to evaluate different designs of MTACs and argue about the distinguishability of legitimate and attack signals in statistical terms. We compare different signals using both, distance on the bit level (Hamming distance) and distance on the sample level (L2-distance), which is motivated by the fact that attack success directly depends on the receiver's inability to distinguish an attacker's guessing error from noise.

Using our statistical model, we identify the symbol-wise mean[4] and (residual) variance as the two main axes of optimization in any attack. We then derive meaningful over-approximations for these two properties that a successful attack signal needs to exhibit and define a strong attacker that will form the basis for the analysis in Section 3.5

3.3.1 Distance-reducing attacker

We ignore for a moment that the attacker has to provide a bit sequence that is accepted by the receiver and assume that the adversarial message passes bit-level verification. In that case, detecting a distance-reducing attacker means distinguishing adversarial guessing errors from benign noise on the sample level.

To formulate such a test, we model noise and attacker error as stochastic processes $\mathbf{N}$ and $\mathbf{A}$. The noise process $\mathbf{N}$ is i.i.d. Gaussian (AWGN

[4]With mean we refer to the accumulated statistics per symbol after inner product with the expected polarity sequence.

channel), an assumption that holds as long as signal modulation places samples/pulses reasonably far apart to avoid inter-pulse interference. The attacker process $\mathbf{A}$, on the other hand, reflects the errors produced by the strategy to guess $\mathbf{c}$. An attacker can freely choose the amplitude of its signal based on any strategy, however, $\mathbf{A}$ is random w.r.t. the polarity of the adversarial samples since the attacker has to guess each sample of $\mathbf{c}$. We can capture this in the following hypothesis test:

$$\mathcal{H}_0 : \mathbf{r} \sim \mathbf{N}$$

$$\mathcal{H}_1 : \mathbf{r} \sim \mathbf{A} + \mathbf{N}$$

For each time j (corresponding to one sample), the noise process is distributed as $\mathbf{N}[j] \sim \mathcal{N}(0, \sigma_n)$, the attacker residual as $\mathbf{A}[j] \sim \mathcal{A}_j(A)$, for an attack strategy A. The best strategy is the one for which the hypothesis test distinguishing $\mathbf{A}$ from $\mathbf{N}$ fails with the highest likelihood.

Together with the bit-level requirement that we have so far ignored, we can now formulate any attacker's universal goals as:

1. **Create the correct bits:** In order to achieve correct detection of each bit, the attacker needs to shift the signal mean $\mu_{b'_i}$ w.r.t. the polarity sequence of each symbol $i \in \{1, \ldots, n_b\}$ beyond the sensitivity of the receiver.

2. **Minimize the error energy:** The attacker aims to minimize the residual energy, i.e., the variance of his error distribution $\mathcal{A}_j$ at any time j.

3. **Make the error as indistinguishable from noise as possible:** The attacker aims to hide in the noise the unavoidable[5] guessing error, i.e., to bring the distribution $\mathcal{A}_j$ close to the legitimate noise distribution $\mathcal{N}(0, \sigma_n)$.

Goal 1 targets correctness on the bit level, whereas Goals 2 and 3 are about indistinguishability of the guessed signal from the expected signal on the physical layer. As we will show, for Goal 2, there exists a clear relation to the hardness of guessing each signal sample of $\mathbf{c}$.

In the presented statistical model, achieving all three goals together represents a sufficient condition for attack success, irrespective of potential countermeasures (i.e., detection techniques). There are different ways an attacker can go about these goals: an attacker can (1) select the subset of

[5]Since being related to the underlying hardness of guessing the pulses correctly.

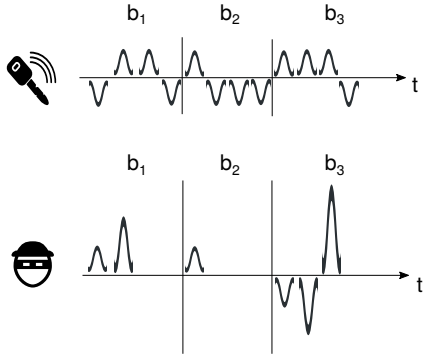

Figure 3.3: Power-increase strategy. Even under a fully randomized pulse sequence holds: If the receiver (i.e., verifier) combines the pulses to symbols in a predictable manner, the attacker has high chances of getting a sufficiently high symbol-wise mean, by increasing the power in reaction to wrong polarity guesses.

samples/pulses she wants to interfere with, (2) choose arbitrary amplitude levels for each targeted pulse, and (3) decide how many samples need to be observed before interfering. A meaningful attack strategy will be concerned with how to make these choices in order to satisfy all three goals jointly.
We now describe two general concepts that guide any attack strategy and lead to the definition of a strong attacker by over-approximating signal mean and residual energy.

Steering the mean: Power-increase strategy

Even if the signal is fully randomized at the pulse level, an attacker can guess symbols by employing a *power-increase strategy* as shown in Figure 3.3. Fundamentally, pulse level randomization under sample-level feedback does not keep an attacker from steering his signal to an arbitrarily high mean under inner product with the hidden polarity sequence with high probability. An attacker starts by sending a pulse containing the entire symbol power. The attacker will keep on doubling the power per pulse until he guesses a pulse of the symbol correctly. This attack succeeds with probability $1 - 0.5^{n_{ppb}}$ per symbol. The core takeaway from this attack is that a sample-level guessing error of the

attacker does not necessarily translate to a bit-level error, due to the dimensionality reduction applied at the receiver. As long as the attacker can hide the error in the null space of this linear transformation, there is no incentive against the attacker using progressively higher energy levels to 'force' the bits. This means, Goal 1, in isolation, is easy to achieve for an attacker. However, achieving the goal with high likelihood, i.e., more attempts, is associated with higher power levels, which puts Goals 2 and 3 in increasing jeopardy.

Minimizing guessing error by learning pulse polarities

Goals 2 and 3 are directly related to the pulse-guessing performance of the attacker. Depending on how the information bits are modulated, the attacker can potentially use bit-level information to infer the signal or rely on knowledge of past pulses to anticipate the pulse polarities ahead. This would reduce the guessing error and make it harder to detect the attack. Our attacker, as introduced in Section 3.2.4 has full knowledge about the transmitted bits. In general, any unmasked signal redundancy in time can potentially help the attacker. An example of this is repetition coding or bit-level Error correction codes (ECCs) as used in the coherent mode of IEEE 802.15.4z HRP [8]. Also, nonidealities in the underlying PRF can help an attacker.

A strong attacker

We abstract away from all possible strategies and only describe the attack signal statistically subject to an over-approximation of its properties that are linked to the attacker's success: signal mean[6] and residual energy (i.e., residual variance).

As will be motivated, residual variance emerges as observable under a maximum entropy assumption on the attacker's strategy. A result from information theory states that the Kullback-Leibler divergence (i.e., relative entropy) determines the exponent of the error in distinguishing two statistical distributions [31]. Consequently, an attacker that brings its residual closest to the legitimate signal is the strongest. Therefore, we can define the strongest attacker $\hat{A}$ as the one that is closest in the KL-sense over all times:

[6] i.e., the inner product with the expected polarity sequence. Correct guesses contribute to it, wrong guesses diminish it.

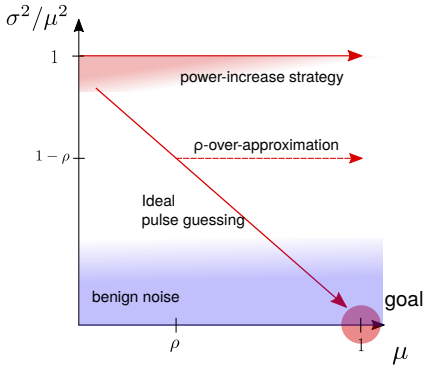

Figure 3.4: Attacker's strategy space. An attacker needs to exceed a certain symbol-wise mean to produce the correct bits a the receiver. This he can achieve with high likelihood using a power-increase strategy. However, there does not exist any reliable strategy for decreasing the normalized error variance. An attacker can only do so by maintaining an edge in guessing pulse polarities. This we model by over-approximating the attacker, e.g., by giving him a pulse-guessing bias ρ.

$$\begin{aligned}\hat{A} &:= \arg\max_{A} Adv(A) \\ &= \arg\min_{A} \sum_{j=1}^{n_p} D_{KL}(\mathcal{A}_j(A) + \mathcal{N} \parallel \mathcal{N}) \\ &= \arg\min_{A} \min_{j} n_p \; D_{KL}(\mathcal{A}_j(A) + \mathcal{N} \parallel \mathcal{N}) \\ &= \arg\min_{A} D_{KL}(\mathcal{A}(A) + \mathcal{N} \parallel \mathcal{N})\end{aligned}$$

The strategy that produces the smallest statistical distance at any j can be converted into the best strategy over the entire signal, by applying the same technique at any other time, since the noise is i.i.d. Therefore, we argue that the attacker that is locally optimal at any time is also optimal over the entire process. The strongest attacker is, therefore, the one that can produce a residual distribution $\mathcal{A} + \mathcal{N}(0, \sigma_n)$ that has smallest relative entropy compared to the legitimate noise distribution $\mathcal{N}(0, \sigma_n)$. Under the condition that the attacker's error has nonzero energy, the process $\mathcal{A}$ that minimizes relative entropy to the AWGN only is also a

Gaussian.

Therefore, as an over-approximation, we can model the attacker residual signal process as normally (i.e., maximum entropy) distributed stochastic process with zero mean and a variance given by the pulse-level guessing performance, which we over-approximate. This is equivalent to assuming maximum ignorance about the attacker's process beyond the existence of some residual energy. Under these conditions, we know that the signal energy is a sufficient statistic for distinguishing two i.i.d. $\mathcal{N}(0, \sigma_1)$, $\mathcal{N}(0, \sigma_2)$-distributed processes.

Observation 1 *The signal residual variance constitutes a sufficient statistic for detection of a guessing attack with a maximum-entropy residual under AWGN noise.*

Basing the classification on the residual energy is optimal if we can extract the attacker's error perfectly and within the assumptions, we can universally impose on the attacker's error process (i.e., being close to satisfying the three goals). A practical attacker will likely deviate from these assumptions, but in ways that *add* distinctive properties (i.e., non-zero higher moments) to the residual distribution. Conversely, an attacker that gets mean and variance right will win.

Observation 2 *The attacker getting the mean per bit right* and *minimizing signal residual variance together constitute a sufficient condition for attack success.*

We have seen that, without countermeasures, a power-increase strategy leads to a guessing bias in the receiver-side security parameter (i.e., the bits). As an over-approximation for the course of a power-increase strategy, we can tilt the guessing performance in the attacker's favor *on the pulse level*. For instance, we can assume that the attacker never makes a wrong guess twice in a row. This means, after at most two interferences (i.e., pulses), the attacker is guaranteed to have made a positive net contribution to the receive statistics. We refer to this attacker as having a *non-ideal bias* of $l = 2$ and illustrate it in Figure 3.5. There, we contrast it to an *ideally-biased* attacker, which knows a given fraction ρ of pulses.

In Figure 3.4, we highlight the two-dimensional nature of the attack strategy. It is easy for an attacker to steer the mean by varying his energy levels, i.e., to move along the x-axis. However, the attacker cannot control the error variance at the same time. So, any practical attacker strategy will be concerned with trading off those two goals. Providing the attacker

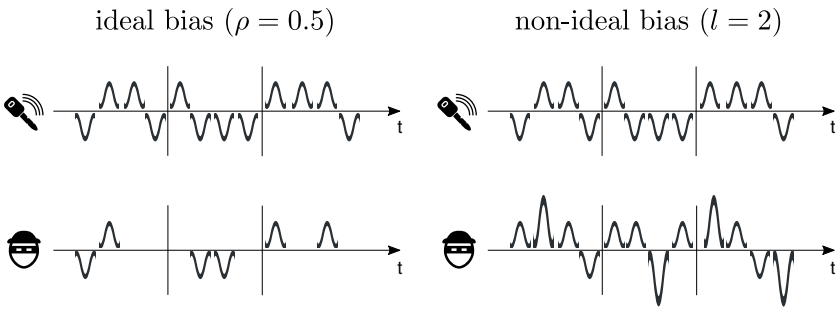

Figure 3.5: Attack modelling. We model two different over-approximations for the attacker's error variance level: An ideal bias, where an attacker knows a fraction ρ of the pulse polarities and a non-ideal bias, where we give the attacker a bound l on the number of power levels for a successful power-increase attack.

an ideal bias results in a diagonal towards the desired spot of high mean and low variance. In addition, as part of any over-approximation, we assume the attacker to be successful regarding the mean (e.g., through a power-increase strategy). This means the attacker can move arbitrarily on the x-axis. In the following, we motivate a specific over-approximation for the error variance, i.e., the attacker's position on the y-axis.

Observation 3 *For an attacker, reducing the signal error variance, while increasing its mean, is 'pulse-guessing-hard'. This means, without a systematic guessing bias, the (normalized) error variance is bound to increase in a guessing attack.*

Persistent interference with a non-ideal bias (i.e., correct symbol polarity after one or two pulses) alone results in an expected normalized variance (i.e., distortion) of more than 0.8. We can estimate the strength of this over-approximation as $0.75^{n_p/2}$. This results from the fact that for every two pulses guessed by an attacker, we omit the possibility of two wrong guesses, an event with probability 0.25. By comparing this value to the bit-equivalent MTAC security level of 2^{-n_b}, we can see that an over-approximation with $\rho = 0.2$ is actually stronger than the bit-equivalent MTAC target security level for modulations with $n_{ppb} > 2\frac{\log(0.5)}{\log(0.75)} \approx 4.82$, i.e., at least five pulses per bit. A decrease of the relative variance to 0.8 or, equivalently, an ideal bias of $\rho = 0.2$ are, therefore, very strong over-approximations, i.e., on the order of the (receiver-side) security

parameter, that become even stronger (less likely) for modulations over longer communication distances.

3.4 Variance-Based MTAC

In the following, we propose the Variance-Based MTAC for direct variance estimation, consisting of rules for signal creation and a receiver-side verification procedure. We then embed this technique into a generic verification algorithm and address side requirements for its practical instantiation.

3.4.1 Tx-side signal generation (*Gen*, *Mtac*)

We assume each sample to follow a binary encoding, achieved either through On-Off Keying (OOK), Frequency-Shift Keying (FSK) or Phase-Shift Keying (PSK), but not Pulse-Position Modulation (PPM). The reason is that, in PPM, the fundamental signal contribution representing each sample becomes vulnerable to ED/LC. Within our assumptions about the modulation, we can represent the transmit signal as a binary pulse sequence of length $n_p = n_{ppb} \cdot n_b$. In particular, we assume that pulses are separated by more than the channel delay spread, i.e., there is no inter-pulse interference. Without this assumption, signal degradation under benign conditions might be hard to distinguish from attacks. The bits are first encoded in a frame $\mathbf{b} = (\mathbf{s}_{b_1} \| \dots \| \mathbf{s}_{b_{n_b}})$, consisting of symbols that each represent message bit under repetition coding, either as $\mathbf{s}_1 = \{1\}^{n_{ppb}}$ or $\mathbf{s}_0 = \{-1\}^{n_{ppb}}$. Preventing an attacker from inferring pulse polarities from either the content of the message m or past samples is achieved by relying on full pulse-level randomization, i.e., by applying a secret sequence $\mathbf{x}$ on the pulses, as in

$$\mathbf{c} = \mathbf{b} \oplus \mathbf{x}.$$

We can either idealize $\mathbf{x}$ being perfectly random, as in

$$\mathbf{x} \leftarrow \{-1, 1\}^{n_p},$$

and shared between transmitter and receiver, or being generated using a pseudorandom function that operates on a previously shared secret.

3.4.2 Rx-side operations (*Vrfy*)

A message time of arrival code has to combine bit detection and verification with an additional signal verification for ensuring the correct signal time of arrival. The bit-level tests are a sequence of binary hypothesis

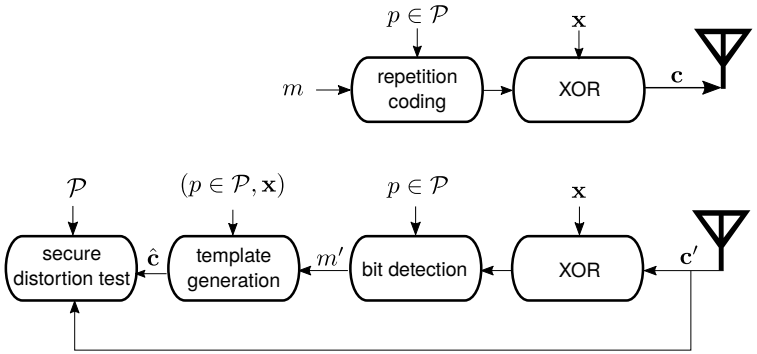

Figure 3.6: Tx/Rx structure of a Variance-Based MTAC. A keyed XOR and a secure distortion test are the central security components. For simplicity, we omit the modulation of each value in $\mathbf{c}$ onto a UWB pulse in the picture. Bit encoding and decoding are parametrized by a performance level p, whereas the secure distortion test applies to an entire performance region $\mathcal{P}$.

tests. The additional check is a single binary test applied to the entire signal, parametrized by the bits received. We illustrate the whole pipeline in Figure 3.6.

Bit detection

Each bit is carried by n_{ppb} pulses. The receiver combines the energy of those pulses subject to the bit-wise hypothesis and the XOR-mask and applies a binary hypothesis-test per bit. The outcome is a received bit sequence $m' = (b'_1, \ldots, b'_{n_b})$.

Signal residual extraction

In order to test the signal integrity on the physical layer, we need to extract the signal-level residual. We exemplify the residual extraction at the receiver in Figure 3.7. Under our stated assumptions about channel and modulation, the received signal $\mathbf{c}'$ consists of the actual signal $\mathbf{c}$, attenuated by path loss, as well as AWGN. At the receiver, the expected pulse polarity sequence (i.e., the template) $\hat{\mathbf{c}}$ is constructed based on the detected bits b'_i and the shared XOR-sequence $\mathbf{x}$, as in $\hat{\mathbf{c}} = (\mathbf{s}_{b'_1} \| \ldots \| \mathbf{s}_{b'_{n_b}}) \oplus \mathbf{x}$. We refer to this step as *template generation* in

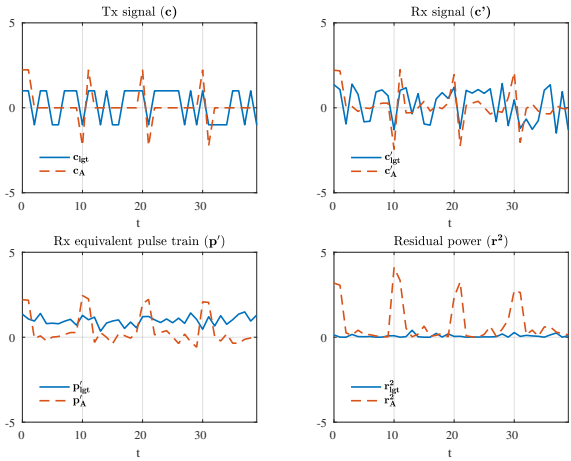

Figure 3.7: Variance-based MTAC. Legitimate (blue) and attack (red) signals in a scenario with four bits and 10 pulses per symbol and repetition coding. The first plot shows the shape of the transmitted signal. Without countermeasure, the attack signal is winning since each bit contains sufficient power, despite the attacker only guessing 2 out of 10 pulses per symbol. The second plot shows the noisy signals at the receiver. The third plot shows the received signal after removing the data modulation. The residual after the expected signal component has been removed is shown in the rightmost plot. It becomes evident that the attack residual can be discerned easily from the legitimate residual, despite the attack on repetition coding (i.e., the bit level) being successful.

Figure 3.6. The receiver-side equivalent pulse train is then given by the element-wise multiplication of the received signal with the expected pulse polarity sequence $\hat{\mathbf{c}}$, as $\mathbf{p}' = \mathbf{c}' \odot \hat{\mathbf{c}}$. The residual is then obtained by subtracting the expected value from the receiver-side equivalent pulse train, as in $\mathbf{r} = \mathbf{p}' - \mu_{\mathbf{p}'}$. A variance-based hypothesis test is concerned with whether the receive signal error consists of model error only or also contains an attacker error. As we argue in Section 3.3, the property to test for is the variance of the signal residual, i.e., if the variance matches the expected noise or is too large, i.e., was caused by attacker errors. However, we require some normalization since the overall receive Signal-to-Noise Ratio (SNR) will vary.

Secure distortion test

We need to normalize the observed signal error to the overall signal energy. This way, we do not need to maintain an explicit noise estimate. The worst-case SNR is found by maximizing over the performance region $\mathcal{P}$, guiding the choice of a threshold for the legitimate distortion. We are then able to check if the observed distortion, i.e., the overall normalized signal error, is within this bound.

This involves a hypothesis test on the normalized variance of the received signal, after being XOR-ed with the expected sequence. As secure distortion function, we propose taking the ratio between the power of the signal residual and the overall received power:

$$\mathcal{D} = \frac{\mathbf{r}^2}{\|\mathbf{c}'\|^2} = \frac{\sigma^2_{\mathbf{p}'}}{\|\mathbf{c}'\|^2} = \frac{\sum_i \left(\mathbf{c}'[i]\hat{\mathbf{c}}[i] - \frac{\sum_j \mathbf{c}'[j]\hat{\mathbf{c}}[j]}{n_p} \right)^2}{\|\mathbf{c}'\|^2}$$

The distortion can be interpreted as the inverse of a receive SNR estimate, based on a hypothesis on the pulse-level structure of the received signal. A random, zero-mean process will, for instance, evaluate to a distortion of $\mathcal{D} = 1$.

Consequently, for a sufficiently large number of pulses, we can write the hypothesis test given by *Vrfy* as a decision between a signal containing the expected structure

$$\mathcal{H}_0 : \ \mathcal{D} < 1,$$

and the signal being only (attacker-induced) random noise:

$$\mathcal{H}_1 : \ \mathcal{D} = 1.$$

Performance region, decision threshold

We assume the transmitter to choose the number of pulses per bit appropriately given a previously selected performance level $p = (d', BER', \Gamma'_{nlos})$, $p \in \mathcal{P}$, i.e., such that

$$Q\left(\sqrt{\frac{n_{ppb} P_{rx}(d', \Gamma'_{nlos})}{\sigma_n^2}} \right) \overset{!}{\leq} BER'.$$

To satisfy our robustness criterion, the maximum legitimate signal distortion needs to be chosen such that the false negative rate does not

exceed the underlying frame error rate, i.e.,

$$T_{\mathcal{D}}(p) = \max\left(T'_{\mathcal{D}} \in [0,1]\right) \text{ , s.t. } P[\mathcal{D}_{lgt} > T'_{\mathcal{D}}] \overset{!}{\leq} FER.$$

The effective threshold is then chosen as the maximum threshold over the entire performance region, i.e., $\hat{T}_{\mathcal{D}} = \max_{p\in\mathcal{P}} T_{\mathcal{D}}(p)$. As a result, $\hat{T}_{\mathcal{D}}$ results in a robust test under any performance tradeoff within the performance region $\mathcal{P}$.

3.4.3 Variance-Based MTAC: Summary

To summarize and illustrate how to embed the Variance-Based MTAC into a distance-measurement system, we highlight the steps involved in the detection of an advancement attack by a receiver (Rx) on a signal originating from a transmitter (Tx).

(**a**) Pre-configuration

1. Rx determines the maximum accepted distortion threshold $\hat{T}_{\mathcal{D}}$ based on the maximum communication distance and maximum tolerated noise level, subject to a performance region $\mathcal{P}$.

(**b**) Key generation (*Gen*)

1. Tx and Rx derive a fresh pseudorandom XOR sequence $\mathbf{x}$ from some shared secret. $\mathbf{x}$ could theoretically also be secretly shared before each round[7].

(**c**) Mtac generation (*Mtac*)

1. Tx encodes the message m using repetition coding according to a chosen configuration $p \in \mathcal{P}$ and applies the XOR sequence.

(**d**) Mtac verification (*Vrfy*)

1. Rx constructs the message m' by multiplying the received pulse sequence $\mathbf{c}'$ with the expected XOR sequence and applying a bit-wise binary hypothesis test on the overall symbol energy.
2. Based on the received message m' and the XOR sequence, Rx constructs the expected pulse-level sequence $\hat{\mathbf{c}}$ (i.e., the template).
3. Rx computes the signal distortion $\mathcal{D}(\mathbf{c}', \hat{\mathbf{c}})$ between received and expected pulse sequence.
4. Rx checks if $\mathcal{D}$ exceeds $\hat{T}_{\mathcal{D}}$. If so, it declares attack.

[7] We don't have any requirements on ToA protection in this step.

3.4.4 Practical concerns

Time reference: Distance commitment

We assume the detection of an advancement attack to be limited to verification of the data relative to some established time frame. This can be achieved by a distance commitment as introduced in [118]. This means the prover is assumed to have already responded in quick fashion to the query by transmitting a deterministic preamble, i.e., is committed to certain temporal reference. Relative to this temporal reference, the prover then has to deliver the secret information (i.e., m, correctly modulated) at a pre-agreed time relative to the preamble. It is realistic to assume a channel to be coherent throughout the frame, as the duration of a UWB frame used for distance measurement is typically less than 1ms. Through a distance commitment, the vulnerabilities of a back-search [84] on the data-bearing part can be avoided.

Ranging precision

Under a distance commitment, the back-search for the acquisition of the first signal path is only necessary on the preamble of the frame. Therefore, the precision of the ranging procedure is not determined by any operation applied to the data-bearing part. Consequently, the precision of our proposal cannot be worse than that of existing schemes relying on a distance commitment. It has been shown that such a system can achieve a precision of 10cm, irrespective of communication distance [85, 111].

Bit-level security

We assume a bit-level procedure to detect if the received bits m' do not match the transmitted message m. This could be achieved by a Message Authentication Code (MAC) appended to the frame or even transmitted on a separate, potentially ToA-agnostic channel.

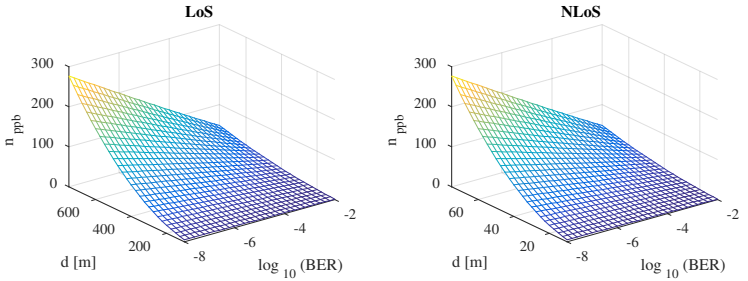

Figure 3.8: Performance Requirement. The number of pulses per symbol as a function of the target performance level, i.e., the target BER and operating distance under FCC/ETSI constraints. These numbers refer to a LoS (left) as well as a NLoS scenario (right) with 20dB attenuation of the direct path. Lower BERs over longer distances require more pulses per symbol.

3.5 Analysis

In the following, we explore the trade-off between security and performance by modeling the effect of the channel and evaluating the classification performance of our Variance-Based MTAC from the previous section. The results are based on simulations, which, however, make assumptions in line with realistic UWB-based distance measurement systems. From these results, we can derive the performance region in which our proposal maintains bit-equivalent security (i.e., $Adv(\hat{A}) < 2^{-n_b}$) and how to scale to longer distances.

3.5.1 Model

Path loss model

To evaluate the impact of distance on a) the modulation required and b) the implications on security, we assume a free-space path loss model. This means the received power degrades inversely to the square of the distance, as in

$$P_{rx} = P_{tx} \left(\frac{\lambda}{4\pi d} \right)^2 \Gamma_{nlos}.$$

We assume the antennas to be operated in each other's far field, as the goal of this analysis is to understand the tension between long distance

and security. As input power, we rely on the constraints put forward by the FCC and ETSI regarding UWB in licensed spectrum. This is, a maximum peak power of 0dBm within the 50 MHz around the peak and an average limitation on signal power spectral density of -41.3dBm/Hz. We assume that our pulses are sufficiently spaced, such that each pulse can be sent at peak power. We assume a signal bandwidth of 620 MHz at a center frequency of 6681.6 MHz, which is a typical UWB channel configuration [8]. For receiver-side noise, we consider the thermal noise figure at room temperature, given by -174dBm/Hz. In a separate NLoS scenario, we assume an additional attenuation of 20 dBm which is roughly the attenuation the signal experiences when traversing the human body. In Figure 3.8, we show the number of pulses per symbol required under both LoS and NLoS conditions. The required number of pulses increases with longer distances and decreases if the requirement on target BER gets relaxed.

Gaussian model for variance distributions

The variance constitutes a sum of n_p independent random variables. Due to the central limit theorem, for a sufficiently high overall number of pulses, the variance distribution converges to a Gaussian, i.e.,

$$\mathcal{D}_{\hat{A}}(d) \sim \mathcal{N}\left(\mu_{\mathcal{D}_{\hat{A}}}(d), \sigma_{\mathcal{D}_{\hat{A}}}(d)\right) \tag{3.2}$$

$$\mathcal{D}_{lgt}(d) \sim \mathcal{N}\left(\mu_{\mathcal{D}_{lgt}}(d), \sigma_{\mathcal{D}_{lgt}}(d)\right). \tag{3.3}$$

In general, these distributions are a function of the communication range as well as the target BER. Through simulations, we can verify that in the area of interest (i.e., where the distributions significantly overlap), these distributions indeed fit a Gaussian hypothesis well, as we show in detail in Appendix A.

3.5.2 Results

We model the bit error rate of the underlying modulation according to Equation 3.1. We simulate this in MATLAB for a frame of 32 bits. For robustness, the choice of the decision threshold should result in the same false negative rate of *Vrfy* as under bit-wise detection, i.e., $FNR_{Vrfy} \stackrel{!}{=} 1 - (1 - BER)^{n_b}$. Under the Gaussian hypothesis for the distortion distribution, we can derive the practical decision threshold by choosing it $Q^{-1}(FNR_{Vrfy})$ normalized standard deviations above the

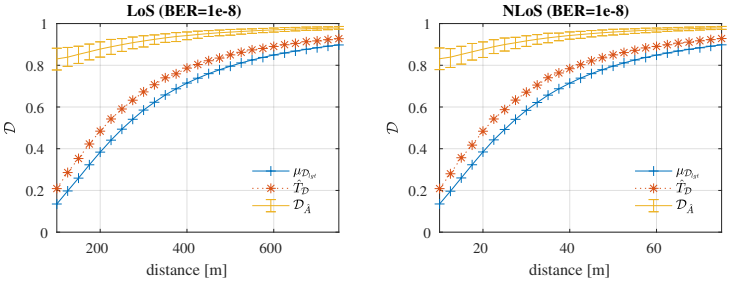

Figure 3.9: Distance dependence. Over longer distances, the legitimate distortion increases. The gap between maximum legitimate distortion and minimum attack distortion becomes smaller for longer distances, eventually vanishing altogether. This means, under our strong attacker model, MTAC security can only be maintained up to some distance.

expected legitimate distortion. The resulting threshold is indicated in Figure 3.9. We evaluate the probability of attacker success for a given maximum communication distance based on the attacker's best case statistics and the legitimate worst-case statistics, over a range of target BER values. This is in line with our attacker model, which does not make any assumptions about the attacker's position. For a given performance region, the upper bound of the attacker's advantage is given by

$$Adv(\hat{A}) = Q\left(\frac{\hat{\mu}_{\mathcal{D}_{\hat{A}}} - (\hat{\mu}_{\mathcal{D}_{lgt}} + Q^{-1}(FNR_{Vrfy}) \cdot \hat{\sigma}_{\mathcal{D}_{lgt}})}{\hat{\sigma}_{\mathcal{D}_{\hat{A}}}}\right),$$

whereas the statistical parameters, i.e., means and variances, are chosen in favor of forger $\hat{A}$. Specifically, we choose the attacker's parameters under minimization of the worst-case distortion and the parameters of the legitimate transmitter under maximization of the distortion, within the defined performance region.

In particular, we choose the attacker's parameters under minimization of the worst-case distortion, i.e., as

$$(\hat{\mu}_{\mathcal{D}_{\hat{A}}}, \hat{\sigma}_{\mathcal{D}_{\hat{A}}}) = (\mu_{\mathcal{D}_{\hat{A}}}(d_{\hat{A},ideal}), \sigma_{\mathcal{D}_{\hat{A}}}(d_{\hat{A},ideal}))$$

$$d_{\hat{A},ideal} = \underset{d\in[0,d_{max}]}{\arg\min}\ \mu_{\mathcal{D}_{\hat{A}}}(d) - \sigma_{\mathcal{D}_{\hat{A}}}(d),$$

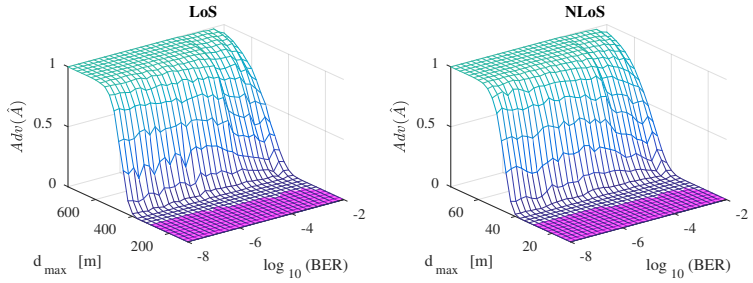

Figure 3.10: Security Level. Attacker's advantage as a function of the performance level. We highlight the performance region within which the MTAC provides bit-equivalent security. The secure distortion test provides us with a bit-equivalently secure MTAC for distances up to 200m and 20m for LoS and NLoS scenarios, respectively.

and the parameters of the legitimate transmitter under maximization of the distortion, within the defined performance region, i.e., as

$$(\hat{\mu}_{\mathcal{D}_{lgt}}, \hat{\sigma}_{\mathcal{D}_{lgt}}) = (\mu_{\mathcal{D}_{lgt}}(d_{lgt,worst}), \sigma_{\mathcal{D}_{lgt}}(d_{lgt,worst}))$$

$$d_{lgt,worst} = \underset{d\in[0,d_{max}]}{\arg\max}\ \mu_{\mathcal{D}_{lgt}}(d) + \sigma_{\mathcal{D}_{lgt}}(d).$$

Unsurprisingly, the worst-case distance for the legitimate transmitter amounts typically to the maximum distance. The numerical values of those statistical parameters (i.e., means and variances) were obtained through simulation. We thereby modeled the attacker as having an ideal pulse-level bias of 20%, as motivated in Section 3.3. In the following, we are interested in the performance region in which the MTAC provides bit-equivalent security, i.e., $Adv(\hat{A}) \leq 2^{-n_b}$.

Performance-equivalent MTAC region

Figure 3.10 shows the attacker's advantage as a function of the performance level. The figure highlights the performance region in which we have bit-equivalent MTAC security.

Observation 4 *Under any tradeoff between symbol length and target bit error rate: For any frame m of at least 32 bits, we can find a distortion threshold $\hat{T}_{\mathcal{D}}$ resulting in an MTAC with bit-equivalent security for distances up to 200m under LoS conditions and up to 20m under NLoS conditions.*

Extending the MTAC region

By comparing the results for LoS and NLoS conditions, we see that the MTAC region seems to degrade proportionally to the attenuation added, i.e., the results are invariant under amplification/attenuation. This means we can extrapolate to any communication range if we allocate a security link margin $\Gamma_{sec} \geq 0$ satisfying

$$\Gamma_{sec} \overset{!}{\geq} 20 \cdot \log_{10}\left(\frac{d_{max}}{200m}\right) + \Gamma_{nlos}.$$

3.5.3 Constant Power Attacker

To validate the model and previous findings, we can evaluate the security level also against an attacker that sends a random pulse sequence at constant power per pulse. If an adversary sends a random pulse sequence at constant power, the resulting distribution for the distortion is given by

$$\mathcal{D} = 1 - \left(\frac{n_p - 2X}{n_p}\right)^2$$

where

$$X \sim \mathcal{B}\left(n_p, 0.5\right),$$

i.e., assumes a Binomial distribution.

To derive the threshold that is secure given a target security level, we can evaluate above expression as a function of the inverse cumulative distribution function of the Binomial. Assuming the configuration with the smallest number of pulses, i.e., 64 pulses consisting of 32 bits with 2 pulses each, up to 55 pulses can be guessed correctly with probability lower bounded by 2^{-32}. I.e., a secure threshold needs to be chosen smaller than the distortion corresponding to 55 correct pulses, which evaluates to $\mathcal{D} = 0.4834$. Importantly, the probability of an incorrect bit is still > 0.99, even under such a bias, which allows relaxation of the threshold to less than the distortion corresponding to 54 correct pulses, i.e., < 0.5273.

3.6 Conclusion

With MTAC, we propose a physical-layer primitive for secure distance measurement. We formally define the security of its underlying algorithms. We then derive design principles for the practical instantiation of an MTAC: A randomized pulse sequence and a secure distortion test over the entire signal. The results indicate that the bit-equivalent security level can be regained over a meaningful performance region, thereby resulting in a fundamental building block preventing any physical-layer, distance-reducing attacks.

Chapter 4

UWB Ranging

4.1 Introduction

Ultra-Wide Band signals consist of a series of short pulses with duration in the order of ca. 2 ns each. Their high bandwidth makes this class of modulations a premier candidate for high-precision ranging in both indoor and outdoor environments, offering a precision of ca. 10 cm. UWB operates in licensed spectrum and is subject to stringent regulatory constraints on output power. While the low transmission power allows long operation on a single battery, facilitating its use for PKES system, its practical range is limited to a few tens of meters.

UWB chips that measure distance are currently being massively deployed in various consumer devices, such as smartphones, cars, and other products [6, 62, 101]. Applications range from car entry and start systems 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 [47], which have been used in practice for car thefts and attacks on contactless payments [39, 66, 114].

In the past, standards for UWB-based ranging, such as IEEE 802.15.4a, did not explicitly consider physical-layer attacks. Hence, the data modulation was found vulnerable to such attacks [82].

The recently adopted IEEE 802.15.4z standard [5] aims to address known distance reduction attacks by introducing cryptographically protected parts of the ranging frame. It introduces two physical layers for

ranging: Low-Rate Pulse Repetition Frequency (LRP) and High-Rate Pulse Repetition Frequency (HRP). In addition, it defines different ranging procedures.

Although both modes are used in automotive applications, primarily for PKES systems [6,21,40,107], HRP has seen adoption in Apple iPhones and AirTags, as well as Samsung phones and SmartTags [16, 100, 106]. Despite its standardization and deployment, no public example implementations or standardized algorithms for security-relevant functionality exist. IEEE 802.15.4z focuses on ranging procedures and message formats without mandating in detail how the integrity of the derived ToA statement is protected by the receiving endpoint.

In this chapter, we demonstrate 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 [113]. 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 [78]. 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. Prior work [113] 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 [43] 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 [28]. 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 [21]. Since UWB as an access system is a new protocol, it might take time until malicious actors can fully understand and bypass security checks [123].

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 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 demonstrate our attack, implemented on inexpensive (USD 65), commercial off-the-shelf components, 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 for different applications and use cases and the applicability of different mitigation techniques in practical settings.

The remainder of this chapter is organized as follows. In Section 4.2, we provide background on the amendments of the IEEE 802.15.4 standard related to ranging. In Section 4.3, we present our attack. We discuss our experimental results in Section 4.4. Finally, we reflect on the security of

HRP UWB in Section 4.5 and compare it to related work in Section 4.6 before concluding in Section 4.7.

4.2 Background on IEEE 802.15.4

We briefly cover the developments of the IEEE 802.15.4 amendments related to ranging that lead up to the most recent 2020 amendment which is the focus of this chapter.

4.2.1 Previous amendments for ranging

In its 2007 amendment for ranging, IEEE 802.15.4a uses a combination of Burst-Position Modulation (BPM) and BPSK to accommodate for both coherent and non-coherent transmitters and receivers [1]. In BPM, pulses are repeated within a short interval many times. In the case of coherent operation, the burst also contains information in its polarity (phase). Due to the high rate of these pulses (499.2 MHz) as well as channel multipath, it is unlikely for a non-rake receiver to resolve individual pulses. More likely, a receiver will integrate the energy over the entire time slot of a burst and obtain the timing and phase as an aggregate over all the pulses of a burst. Individual bursts can, in consequence, become a target for ED/LC attacks due to time-wise integration at the receiver. It has indeed been observed that, in IEEE 802.15.4a, an attacker can always decrease the distance by some value slightly smaller than the distance corresponding to the burst duration [82].

IEEE 802.15.4f [2], in 2012, added the LRP Physical Layer (PHY), where a bit is either represented as a single UWB pulse, or as a sequence of pulses (extended mode). These pulses are separated by more than the typical power delay profile of the channel, thereby avoiding the problem of Inter-Pulse Interference.

On the physical layer, there emerged two configurations that are still part of the most recent amendment: LRP (similar to IEEE 802.15.4f), HRP (similar to IEEE 802.15.4a).

4.2.2 Ranging Procedures in IEEE 802.15.4z

IEEE 802.15.4z [8], released in 2020, adds PHY and MAC improvements to increase the integrity and accuracy of ranging measurements. The amendment proposes different ranging procedures. In all of them, devices establish their relative distance by measuring the RTT of a particular message exchange. In IEEE 802.15.4z, this time-critical exchange can either involve one or two RTT measurements. In the Single-Sided Two-Way Ranging (SS-TWR) configuration, an initiator sends a frame to the

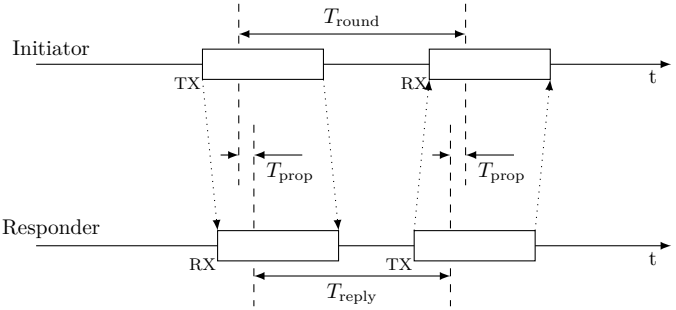

Figure 4.1: Single-sided two-way ranging, as standardized in IEEE 802.15.4z [5].

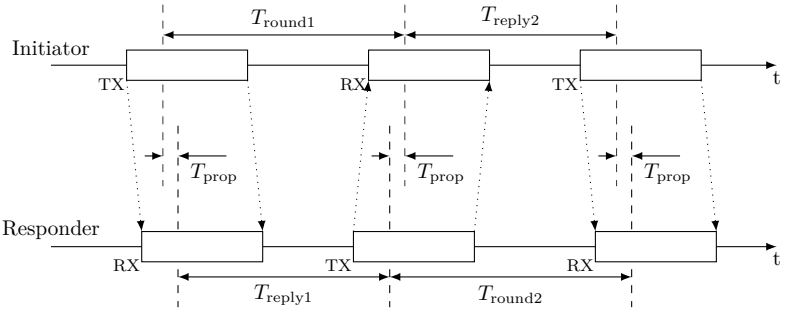

Figure 4.2: Double-sided two-way ranging using three messages, as standardized in IEEE 802.15.4z [5].

responder, which responds after a pre-agreed reply time. The initiator then calculates the RTT based on the difference between the ToA of the response and the time of the initiator's transmission and derives the relative distance by subtracting the (pre-agreed) reply delay of the prover (T_{reply}). This ranging procedure is shown in Figure 4.1. Finally, the initiator sends the calculated distance in an additional data frame to the responder. Both initiator and responder have to agree on which part of the frame the ToA refers to, IEEE 802.15.4z chooses this reference point to be after the preamble and just before the start of the cryptographically generated sequence.

I.e., in SS-TWR, the distance is calculated as the RTT minus the reply time at the responder:

$$\hat{T}_{\text{prop}} = \frac{T_{\text{round}} - T_{\text{reply}}}{2}$$

In general, the clock instability at both the initiator and responder affects the precision of the range estimate. This relative offset can be compensated by measuring the RTT as an average of two symmetrical measurements.

In Double-Sided Two-Way Ranging (DS-TWR), the propagation time is measured based on two RTT measurements. IEEE 802.15.4z proposes to do this either with two consecutive rounds of SS-TWR, involving four frames, or a version that interleaves the two rounds and uses a total of three frames. The latter we illustrate in Figure 4.2. For DS-TWR, the standard proposes a more complicated calculation of the propagation time:

$$\hat{T}_{\text{prop}} = \frac{T_{\text{round1}}T_{\text{round2}} - T_{\text{reply1}}T_{\text{reply2}}}{T_{\text{round1}} + T_{\text{round2}} + T_{\text{reply1}} + T_{\text{reply2}}}$$

A derivation of this formula can be found in [3].

Under appropriate choice of the frame structure, modulation, and cryptographic algorithms, both procedures could be used for implementing the time-critical part of an authenticated ranging protocol.

4.2.3 IEEE 802.15.4z LRP

In IEEE 802.15.4z LRP mode [8], pulses are transmitted at a relatively low rate, i.e., at either 1 MHz or 2 MHz, which helps avoid self-interference due to channel effects, assuming the power delay profile in a short-range scenario is up to ca. 100 ns. Due to the limited effect of the channel on the signal quality, the receiver can recover each symbol (pulse) as a single

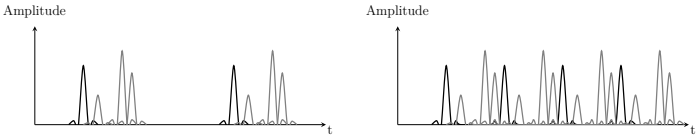

Figure 4.3: IEEE 802.15.4z LRP vs. HRP: LRP mode places the pulses at intervals exceeding the power-delay profile of the channel, whereas HRP is prone to inter-pulse-interference caused by a typical channel.

bit. This allows a receiver to base its decision on the samples that occur at pre-agreed relative delays from the preamble, and, thus, for direct implementation of the distance-commitment concept.

Different techniques can be employed to validate the (cryptographically blinded) pulse sequence. As introduced in the previous chapter, a method that detects bits by combining pulses can be complemented by a test of the signal variance. Alternatively, each pulse could be detected independently, and an adequate threshold on the number of correct bits could be applied. Moreover, the mapping from pulses to bits can be randomized, i.e., the pulses reordered, to thwart a power-increase guessing attack.

4.2.4 Secure ranging in IEEE 802.15.4z HRP

In HRP PHY mode, pulses are sent at a relatively high rate of 64 MHz. The inter-pulse-interval can therefore be as short as 16 ns. In contrast to LRP, this configuration is prone to inter-pulse-interference, as illustrated in Figure 4.3. A typical channel will cause pulses corresponding to different paths to overlap at the receiver. This means that cross-correlation noise caused by a late, strong signal path can distort the cross-correlation profile of earlier, weaker paths. However, for precise ranging, a receiver is required to identify the earliest, i.e., direct signal path.

Frame structure

For initial acquisition the frame contains of a preamble. This preamble consists of repetitions of a publicly known ternary sequence (i.e., consisting of -1, 1, 0) with good correlation properties, i.e., a low cross-correlation noise profile. After the preamble there is a static, pre-negotiated pulse sequence called SFD, marking the end of the preamble. The preamble

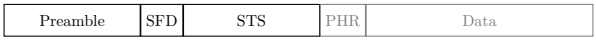

Figure 4.4: HRP Frame Structure: The cryptographically generated Scrambled Timestamp Sequence (STS) occurs after a publicly known preamble and Start-of-Frame Delimiter (SFD).

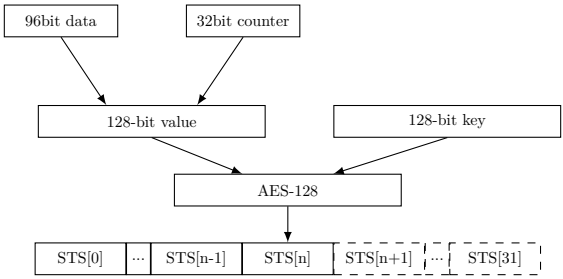

Figure 4.5: STS generation

detection is important for initial acquisition and also recovery of the carrier offset between transmitter and receiver. After that, and as a novel addition in IEEE 802.15.4z, the frame contains a random binary pulse sequence, called STS. Figure 4.4 shows the frame structure for the no-data mode, which is predominantly used for ranging in currently available products, as well as the Mode 1 configuration that places a PHY header and data after the STS. Depending on the configuration, the frame can also contain data, e.g., containing the time-stamp of a previous RTT measurement. This data is modulated similar to IEEE 802.15.4a, using burst-position modulation, optionally combined with BPSK. As this modulation is not secure against physical-layer attacks [82], the data part is not used for secure ToA acquisition. In the context of the attacks that will be outlined in the following, it is worth noting that the overall frame duration as well as the duration of its components significantly exceed the relative ToA reduction required for meaningful distance reduction. Components of the frame span tens of microseconds, a meaningful ToA reduction only tens of nanoseconds.

Scrambled Timestamp Sequence (STS)

The STS consists of pseudo-random BPSK pulses (i.e., consisting of -1, 1) that are generated as shown in Figure 4.5. The devices are assumed to

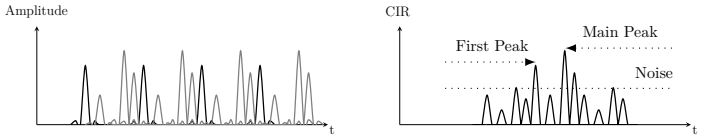

Figure 4.6: Receiver-side correlation for obtaining the CIR.

have a shared key and are synchronized, i.e., generate a fresh STS before each ranging procedure by incrementing the counter. The receiving device knows the expected STS in advance and can create a local template in advance. However, without knowledge of the key material, it is impossible for an adversary to predict the correct STS and send it in advance to reduce the signal ToA. Also a meaningful ED/LC attack is effectively prevented since the pulses only span ca. 2 ns in duration, which translates to a maximum distance reduction of ca. 60 cm. As a result, ranging devices can, in theory, 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. Since the STS consists of random pulses chosen without any regard for cross-correlation properties, there is the problem of cross-correlation noise.

Receiver operation

The receiver constantly scans the medium for the occurrence of a static (pre-agreed) preamble. For this purpose, the received signal is down-converted, digitized, and correlated with a local template. As a consequence of the preamble's good correlation properties, for initial acquisition, only a small part of the preamble is sufficient, and the overall correlation is associated with a low cross-correlation noise and, hence, good ToA precision. The SFD marks the end of the preamble and the start of either data or the STS, depending on the mode. By convention, the receiver denotes the ToA as the reception time of the SFD. Then, for validation of that timestamp, the receiver correlates the received STS with a locally generated template to obtain an estimate of the Channel Impulse Response (CIR), as

$$\mathrm{CIR}[t] = (\hat{c} * c)[t] = \sum_{m=0}^{|\hat{c}|-1} \overline{\hat{c}[m]} \cdot c[m+t],$$

where c is the complex, time-discrete received signal, and $\hat{c}$ is the template of the expected signal.

While the algorithm for peak acquisition is not known and may be different among implementations, we can reason about the underlying challenges such an algorithm faces. The receiver must distinguish a legitimate, potentially attenuated, early NLoS from a backdrop of cross-correlation noise, as illustrated in Figure 4.6. In practice, this process has to be optimized to cope with channel distortions, most notably multi-path fading. Objects in the vicinity reflect the signal, which creates copies of the signal that are slightly delayed in time. The sum of these individually delayed and attenuated copies constitutes the received signal. In general, high correlation values imply a similarity between the received and expected signal at a given offset. However, the peak clearly distinct from cross-correlation noise is only under clean (e.g., LoS) channel conditions. The challenge is to detect with sufficient confidence the signal corresponding to the direct path under constructive and destructive interference of other, potentially stronger, reflections. While the preamble is selected for good correlation properties, the STS is pseudo-random and therefore exhibits a significant cross-correlation noise profile. As a consequence, under challenging channel conditions, the preamble might indicate an early path that is less pronounced in the STS. This creates, in practice, a difficult performance-security trade-off between a receiver's ability to detect weak early peaks and accepting only peaks that adequately represent the expected STS.

Revisiting the steps outlined in Section 2.4.2, we can assume that (fine) signal acquisition can be facilitated by the good correlation of the preamble. However, the STS pulses cannot be directly demodulated in practice due to Inter-Pulse Interference. Therefore, the STS needs to be acquired by a correlation procedure, and the first path detected based on the cross-correlation profile. The last requirement, ToA validation, can then potentially be achieved by detecting a valid cross-correlation peak at a temporal offset consistent with the early peak acquired in the preamble.

Commercial HRP UWB Chips

HRP-based location and tracking tags, as well as mobile-device support have recently entered the consumer market at scale [62] and automotive manufacturers are planning to release cars featuring PKES systems using HRP chips, such as the BMW iX and Genesis GV60 models [21, 101].

The FiRa consortium considers HRP viable for both consumer-grade and security-critical applications [44].

Apple has a diverse UWB software and hardware stack. Different versions of Apple U1 chips have been released in different products, such as the iPhone (since iPhone 11), the HomePod mini, the AppleWatch (since Series 6), and the AirTag. On the iPhone, Apple integrated UWB into AirDrop with iOS 13 [80], 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 [60], exposing a selected set of UWB-based ranging functionality to application developers. A compatibility mode for third-party accessory support has been available since the release of iOS 15.

NXP advertises their Trimension chip series for secure ranging and precise positioning [78]. Development kits exist for the SR150 and SR040 [107]. Our analysis showed that several Samsung products, for example, the SmartTag+ and phones starting from Samsung Note20 Ultra [106], contain NXP chips to enable ranging and improve Point to Share [99] data transfers. Examples for cars that comprise NXP chips are upcoming BMW and VW models [104, 105], whereas VW seems to rely on LRP chips for the PKES use case [6].

Qorvo, also known as Decawave before their acquisition [88], manufactures the DW3000 chip series. These chips are interoperable with the Apple U1 chip [85]. 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 [86], and DWM3001CDK, an integrated board that contains an nRF52833 with Bluetooth 5.2 [87].

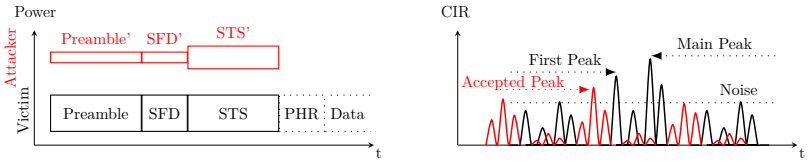

Figure 4.7: Attack Principle: The attacker transmits a carefully crafted packet (red), coarsely synchronized with the legitimate signal (black) and with power low enough to avoid jamming. The STS is secret, and STS' is randomly chosen by the attacker. Therefore, at the receiver, the correlation peaks caused by the attacker (red) are lower than those of the strongest path (black). However, one of them is higher than the threshold for accepted 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.

4.3 A Practical Distance-Reduction Attack

We show a practical attack on widely deployed HRP UWB ranging systems.

4.3.1 Attacker Model and Attack Execution

We consider an attacker that aims to reduce the measured distance between two HRP UWB devices. This corresponds to the Mafia Fraud attacker model introduced in Chapter 2. In practical terms, this means an attacker trying to unlock and start a car by tricking it into believing that the legitimate owner's car keyfob is closer than it is. The attacker has no access to any secrets shared between victim devices and cannot predict the message fields with content that is assumed to be unpredictable in HRP UWB. In particular, the attacker cannot predict the Scrambled Timestamp Sequence (STS). We assume the attacker to be located in proximity to the victim devices; however, unable to physically access or tamper with the devices.

We illustrate the attack principle in Figure 4.7. During an initial observation phase, the attacker device behaves like an HRP UWB packet analyzer and resolves the sequence of packets exchanged by the victim devices. Once the ranging sequence and its timings have been identified, the attacker devices reactively transmit a signal over selected packet

components (*overshadowing*), as shown in Figure 4.8. The receiver mistakenly detects a cross-correlation peak caused by the adversary's contribution as an early copy of the legitimate signal, reducing the measured distance (Figure 4.7).

The adversarial overshadowing signals follow the HRP UWB frame structure but are crafted such that different packet fields are transmitted at different power levels. Specifically, the attacker's signal consists of a preamble, an SFD, and an STS. Moreover, the attacker does not transmit the same preamble as the victim since this would result in denial of service by shifting the entire search window because the attacker is only coarsely synchronized. The attacker adjusting the power level and (deliberately) choosing a wrong sequence both result in the necessary increase in cross-correlation noise without affecting the victim receiver's ability to detect the main correlation peak of the benign transmission. Based on the first packet sent during a ranging procedure, which triggers the attack, the attacker can adjust its timing and then overshadow the following packets belonging to the same ranging procedure. During overshadowing, the attacker's synchronization accuracy only needs to be in the order of µs, despite attacking a procedure that measures timing with ps resolution [33]. The effect of distance on timings (approximately 3.3 ns per meter) being significantly smaller than the required synchronization accuracy, fine-grained tuning is not necessary, even under changing distance. To find an adequate transmission power level, the attacking device can start with a strong signal, then reduce power until the communication is not jammed and distance reductions occur.

This attack falls under the term *overshadowing* since the adversary transmits its signal on top of the legitimate signal but requires the legitimate signal to be still in part detectable. In particular, the attack relies on the fact that the main peak of the correlation profile still corresponds to the legitimate signal, while the surrounding noise profile is modified by the adversary. This is in contrast to a *spoofing* attack, where an attacker introduces an earlier ToA claim based on its signal alone, which is not possible due to the unpredictability of the STS.

The attack can be implemented and executed using a simple and inexpensive off-the-shelf HRP UWB device. No complex laboratory equipment is needed, making the attack practical and easy to implement and mount. Figure 4.7 shows an adversarial packet aligned with a legitimate packet, and the corresponding CIR at the receiver. The adversarial packet consists of a preamble, SFD, and STS, where the preamble different from the legitimate preamble and the STS is randomly

chosen 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 adversarial peaks is high enough to be misclassified as a legitimate early peak, corresponding to an (inexistent) early path. The power of each field is independently adjusted to obtain optimal results, as explained in more detail in Section 4.3.2.

In many practical scenarios, 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 [28, 29], NFC, UHF). However, this pre-negotiation does not impact the attack, which only targets the ToA measurement of packets in the ranging sequence. As shown in Figure 4.8, the attack can be easily generalized to target either side or different frames of the ranging procedure. The attacker can target any desired packet in the sequence by configuring the delay of reaction after the reception of the first packet. In the case of DS-TWR this can be leveraged to select the receiving device to attack. Alternatively, an attacker can also use two devices to attack both ends simultaneously, increasing the success rate.

4.3.2 Attack Principle

We explain why and how the attack works and compare it to existing distance reduction attacks.

Secure Leading Edge Detection

Accurate timestamps require reliably 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, in general, be presented with multiple copies of the transmitted signal, arriving with different power and delay from different paths.

In HRP UWB, pulses are sent at a relatively high rate, meaning the pulse spacing is less than the power delay profile. Specifically, at a pulse repetition rate of 64 MHz, the pulses are only spaced by 16 ns. This can cause inter-pulse-interference, especially in a channel with strong indirect paths (reflections). For Time of Flight measurements used in Two-Way

Ranging (TWR), the receiver must find the earliest copy, corresponding to the shortest path (Line-of-Sight). The correlation peak corresponding to the direct path can, in general, be weaker than the main correlation peak. For adequate performance under NLoS conditions, such weak early paths need to be detected. The noise profile can be manipulated by an attacker in order to increase the likelihood that the receiver accepts a faulty early path as the first path of the signal, resulting in a reduced ToA and, hence, a reduced distance estimate. The attack presented in this chapter exploits the inherent difficulty in separating a potentially weak direct path from benign cross-correlation noise. In particular, the attacker reactively overshadows a subset of the frames of the ranging procedure with a pulsed signal. The adversary's signal is coarsely synchronized to the legitimate frame, which allows for tuning power levels to respective sections of the frame.

In the following, we empirically analyze the behavior of the unknown receiver-side algorithms deployed in real products (Apple U1) under adversarial interference. Because of their closed-source nature, we do not know most of the design choices. For instance, 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-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 power levels of the fields to maximize the chances of reduction while avoiding jamming and other errors.

It is worth noting that the attacker does not have direct control over the amount of distance reduction. A method to gain partial control has been proposed by simulation in [113]. This approach, evaluated in simulation, assumes an attacker to keep track of the correlation profile during the process of overshadowing the STS, and stopping as soon as a correlation peak at the desired offset emerges.

Selective Overshadowing to Avoid Jamming

An attack against leading edge detection can be successful in practice only if the injection of the adversarial signal does not accidentally produce other errors that invalidate the ranging procedure. To achieve this goal, our attacker 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 attacker transmits a wrong preamble 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.

Selecting Victim Packet(s)

Typically, in available HRP UWB implementations, DS-TWR is used because it compensates for relative clock error and asymmetric reply times. We have confirmed this in our analysis of many HRP UWB configurations. As shown in Figure 4.8, 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 not strictly necessary. Moreover, the attacker needs to be in range of the initiator, in order to receive the first ranging packet. The effect of overshadowing the second or third packet is not entirely symmetric. Attacking the third packet has the effect of reducing the Round-Trip Time measured by the responder (T_{round2}), the following relationship between the corresponding maximum reduction at the receiver (d) and the maximum reported reduction ($\hat{d}'$):

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

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

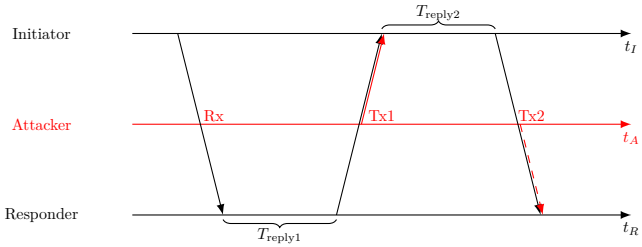

Figure 4.8: Reactive overshadowing attack: The attacker receives the first frame of the ranging procedure and overshadows parts of the second and/or third frame of the procedure.

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

Hence, by attacking the second packet, the attacker can obtain reductions that are twice the amount 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 direct and (stronger) indirect path accepted by the receiver (i.e., the width of the backsearch window). For a practical attack, precise control is not necessarily required. For example, any reduction below 2 m would break a PKES system and unlock a car.

Alternatively, the attacker can use devices to target both the second packet (near the initiator) and the third packet (near the responder). We can consider the two attacks as independent events. Therefore, the attacker will obtain reductions of up to $c \cdot \delta/4$, $2c \cdot \delta/4$, and $3c \cdot \delta/4$, each with decreasing probability.

We confirm these calculations in Section 4.4. We tested our targets in their operating range (ca. 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 for proximity detection (e.g., car keys, item finder) only envision short ranges.

4.3.3 Attack Implementation

The attack works by configuring a third device to reactively overshadow the second and third frame of a DS-TWR ranging procedure. The signal contributed by the adversary increases the noise profile at the receiver. Since the adversarial signal is modulated similarly (also pulses), the resulting noise profile shares similarity with NLoS paths. In principle, the attack 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 [86], controlled by a Nordic Semiconductor nRF52 DK [77], for a total cost of around USD 65 only. These devices can be easily programmed with open-source firmware [50], 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 by the victim so that the attack signal is transmitted on top of one of the following packets (Figure 4.8). The attacker can find this and other reception parameters in an attack preparation phase. The preparation phase is only required once per protocol, since the parameters tend to be stable for a particular device pairing.

4.3.4 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.

Figure 4.9a 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 the responder. A Qorvo DWM3000EVB acts as an attacker placed at around 30 cm distance 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 4.9b 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 message of the DS-TWR procedure, the attacker causes reductions from 10 m to less than 2 m in the raw measurements plotted on the laptop.

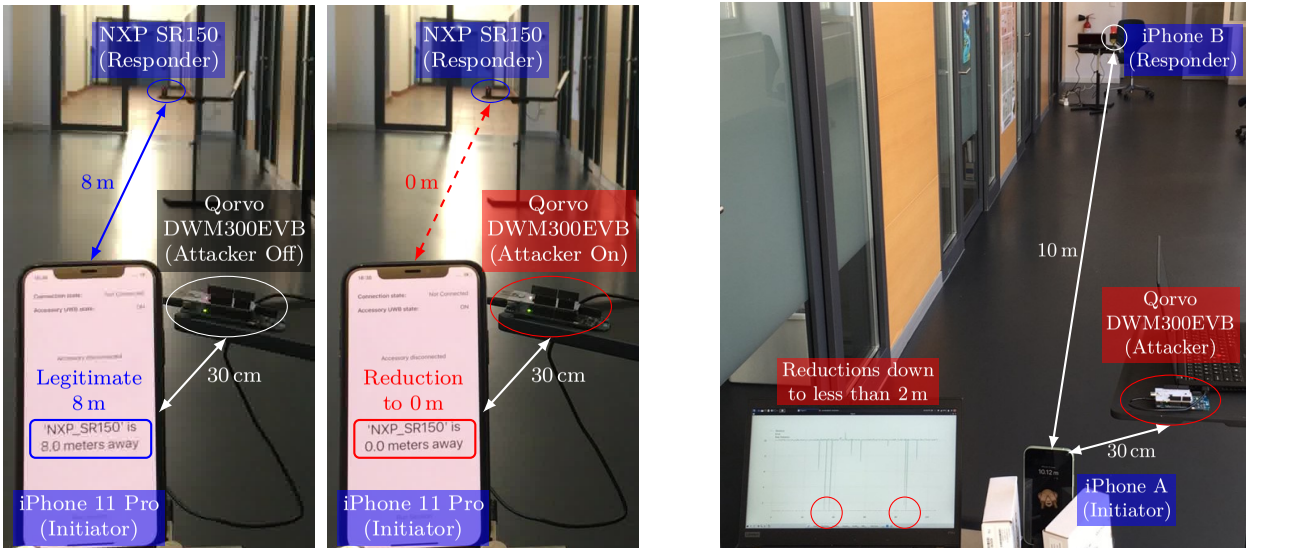

(a) iPhone (initiator, victim) + NXP SR150 (responder): reduction from 8 m to 0 m visible on the screen of the iPhone.

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

Figure 4.9: Two concrete examples of distance reduction attacks.

4.4 Attack Experimental Evaluation

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

4.4.1 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, 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) [60], iPhone-AirTag (FindMy ranging) [10], iPhone-HomePod (Handoff music) [11], iPhone-NXP and iPhone-Qorvo (compatibility mode).

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 4.3, i.e., they are programmed to detect an initial frame of the ranging exchange and then overshadow preamble and STS of one or two subsequent frames. It is important to not 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 victim devices eventually report the same measurement time series.

4.4.2 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. 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 [58]. Then, detailed measurement information appears in the logs, including the distance. These readings were used in the evaluation of our attack.

4.4.3 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 4.1. 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.

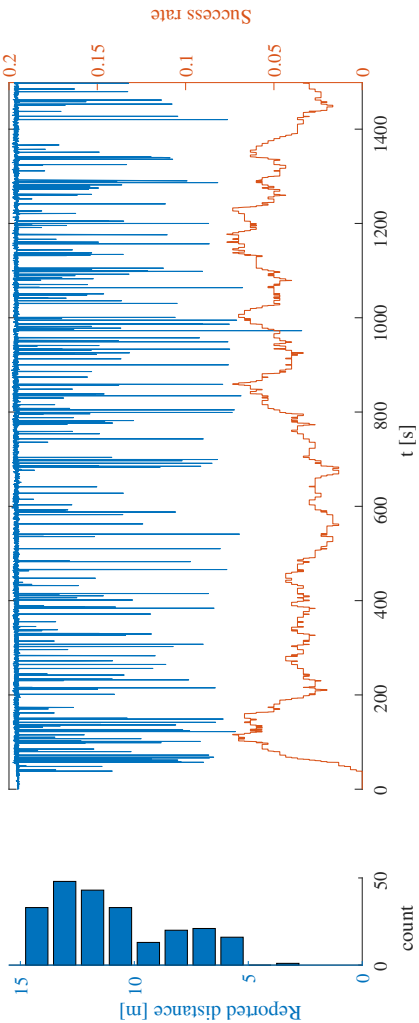

Figure 4.10: This figure shows a 25 min ranging experiment. Two iPhones are placed at a real distance of 15 m between each other, under attack with two devices. The right part shows the distances reported for each measurement in blue, with obvious reductions (i.e., reported distance less that 15 m). 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.

To exemplify the cumulative effect of multiple attacking devices on the overall reduction, Figure 4.10 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 $\leq$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 4.3.2. 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 entrance[1]. 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 4.11 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.

[1] Precisely this attack scenario has become an increasing concern for PKES that do not rely on signal ToF [39, 47, 114, 126].

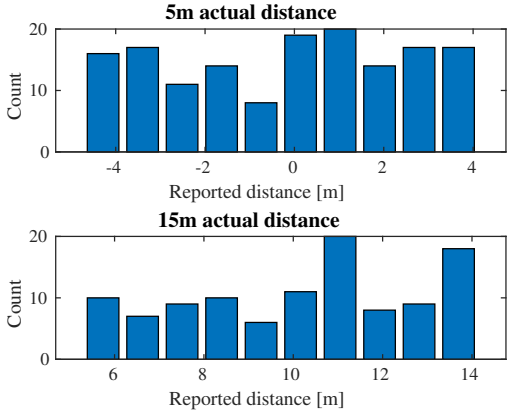

Figure 4.11: Distribution of reduced distance reports for the iPhone–iPhone (5 m and 15 m) setup, attacked with single device over a 15 min observation period. The overall rate of successful distance reductions (i.e., less than 5 meter and 15 meter, respectively) is ca. 2.2 % in both cases.

4.4.4 Device Pairings

In Table 4.1 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.

Table 4.1: Overview of attack scenarios against Apple U1 (primary victim) and results.

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

4.5 Discussion

In the following, we discuss the underlying challenge in practical HRP-based ranging and point towards potential countermeasures. We then address other implications of the difference between LRP and HRP ranging and close with an outlook on potential directions in UWB secure ranging.

4.5.1 What is the fundamental problem?

The fact that the vulnerable receiver accepts wrong (reduced) distances only in a certain time interval before the strongest peaks suggests that the confidence level of that peak influences the likelihood of accepting an earlier (weaker) path. While this makes sense from a performance viewpoint, this violates the principle underpinning the distance commitment, i.e., the fact that the data needs to be validated at a fixed time interval from the preamble for any accepted path. Since the backsearch window is longer than a practically tolerable measurement error, the deviation has security implications.

Orthogonalization strategies, applied to the received signal, are important countermeasures. The idea being to separate the paths as good as possible and only base the decision on signal contributions that belong to that path. In consequence, the statistics of the first are considered in the acceptance decision.

4.5.2 Towards countermeasures

Since the attack exploits the difficulty in distinguishing a weak early path from self-interference, there exists a potential for countermeasures that seek to reduce the impact of self-interference. Temporal correlation can be interpreted as a *statistical* interference cancellation technique that, on average, cancels ill-aligned copies of the template, with random correlation noise being distributed around the average.

This could be improved, i.e., the noise profile lowered, by actual Successive Interference Cancellation (SIC). Such an approach removes the footprint of late by subtracting them from the received signal and, in turn, improves the SNR of weaker early copies. This idea has been proposed in various contexts, e.g., to address the near-far-problem in augmented GNSS [70] or in the context of a narrowband ranging system [124]. The

practical drawback of such an approach is the increased complexity of the receiver design since detection occurs repeatedly.

A SIC approach can rely on channel information learned from a preamble. This channel estimation for consistency is different from filtering for performance enhancement since late copies are subtracted and not accumulated into the decision statistics. The latter would offer an opportunity for distance reduction to an adversary.

4.5.3 LRP vs. HRP

The difference between LRP and HRP is a consequence of different implementation approaches to address the difficult transmit power constraints. LRP sends pulses at a lower rate, thereby operating closer to the peak power constraint. HRP sends pulses at a higher rate, meaning each pulse needs to be sent weaker due to the average power constraint. In theory, HRP allows for coherent demodulation; however, it faces the practical challenge of clock offset estimation and compensation (for accurate phase detection). The drawback of LRP is amplification requirement (stronger pulses) that might be difficult to reconcile with hardware limitations of resource-constrained, highly integrated devices. Both LRP and HRP limit frame duration to a few dozen microseconds, primarily to limit the effect of accumulating clock offset.

4.5.4 Outlook

Longer frames are always useful for an improved security level. However, with longer overall durations, relative clock inaccuracies between initiator and responder become increasingly problematic since the absolute error accumulates over time and can prevent correct demodulation. Especially demodulating coherently modulated signals requires highly accurate clock offset estimation. Solving this challenge could prove beneficial for increasing the security level due to a frame containing more random pulses and the ability to space them wider to avoid IPI.

4.6 Related Work

The UWB IEEE 802.15.4 standard is documented in [4,5]. Chips following the HRP mode of the standard have been implemented by several vendors, such as Apple (U1) [12], NXP (SR040, SR150, SR100T) [78], and Qorvo (DWM3000) [86]. Chips implementing the standard LRP mode have been implemented by Microchip (ATA8352, ATA8350) [74] and Renesas [96]. To the best of our knowledge, these LRP mode chips are not available in consumer electronic devices such as mobile phones.

The first implementation-independent security evaluation of HRP UWB at the physical layer has been conducted in [113]. That work proposed two attacks on HRP, derived from the Cicada attack [83, 84], and shows in simulations that even conservative receiver implementations could be susceptible to distance reductions. In contrast to the work in this chapter, the authors neither conducted experiments with real UWB chips nor proved 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, the significance of different message fields, or their power.

Further research on UWB ranging has been done in [111]. This work proposes improvements to LRP that aim at securely extending the range of IEEE 802.15.4 LRP mode through pulse reordering.

Previously documented attacks against UWB [46], which applied to earlier standards (IEEE 802.15.4a), cannot be used against HRP because it does not combine pulses into symbols and the individual pulses are very short. For example, an attacker cannot acquire the polarity of a 2 ns pulse in time to advance it to conduct an ED/LC [82] attack.

4.7 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 physical-layer security of HRP UWB, we have proposed and discussed several countermeasures.

Chapter 5

Multicarrier domain

5.1 Introduction

Although secure precision ranging can be realized with UWB, its deployment is currently confined to short-range applications (i.e., 10s of meters). Due to its use of wide segments of licensed spectrum, UWB technology is subject to stringent constraints on transmit power. Moreover, the fact that the signal power is compressed in short pulses makes amplification difficult and limits the distance in practical use.

Compared to UWB, Orthogonal Frequency-Division Multiplexing (OFDM) is a modulation technique that is widely used today, especially in wireless systems that offer high throughput, such as in WiFi or cellular (i.e., 4G, 5G). A lot of infrastructure supporting these communication standards has been deployed with an ongoing trend towards high-bandwidth OFDM signals (5G). With OFDM, data is transmitted over many subcarriers in parallel. This provides robustness against frequency-selective channel drops (fading) [75]. However, because the subcarriers are closely spaced in frequency in most OFDM-based systems, an OFDM receiver requires multiple time samples for correct decoding. The transmitted symbols are significantly longer than for most single-carrier systems, which is not ideal for (secure) ranging. Over the last decade, there was a lot of research dedicated to overcoming this challenge and re-purposing OFDM signals in WiFi for ToF-based ranging and positioning [51, 67, 122], achieving ranging precision on the order of meters or less.

Such performance numbers are sufficient for many applications, and OFDM signals are a viable candidate for ranging. However, in the context of distance bounding and ranging, the security of OFDM systems is unclear to date, unlike UWB based systems that are thought to be secure against a powerful, Dolev-Yao-like attacker with idealized reaction times [111]. Given the vast proliferation of OFDM systems today and in the foreseeable future (5G), it is therefore of great importance to also assess the security of OFDM-based ranging implementations. This concern has been identified by the ongoing IEEE 802.11az standardization effort for next-generation positioning based on WiFi signals. Current proposals for secure ranging that have been made by the respective Task Group [14] include different OFDM modulations where random symbol sequences are transmitted to acquire the ToF.

To the best of our knowledge, the Task Group has not yet decided on the final technique that would provide the most resilience against a distance-reducing attacker. The fact that discussions have been ongoing for more than four years [15] clearly indicates the challenging nature of OFDM-based ranging. Undoubtedly, one needs to fully understand the security implications of a ranging scheme before its design is "baked" into billions of hardware chips supporting the upcoming IEEE 802.11az standard.

In light of this development, we aim to identify the pitfalls of OFDM-based ranging and assess whether multicarrier Time of Flight ranging can be secured against physical-layer attacks. We choose a theoretical angle to approach the question and assume an idealized adversary with no hardware constraints. Therefore, our results serve as a guideline for real-world systems that might relax the adversarial model by constraining reaction time, sensitivity, and computational power of a potential attacker. In addition to the theoretical insights, we make our own proposal for secure multicarrier ranging that is based on orthogonal noise and can be used in conjunction with other approaches.

In order to increase positioning accuracy, some OFDM-based ranging systems exploit signal phase and directionality alongside time-of-flight information. Since these features do not contribute to the system's overall security—phase information can easily be subverted, see, e.g., [90]—, the focus of this work will only be on the security guarantees provided by time-of-flight measurements.

In particular, we make the following contributions:

- We provide mathematical proof that robust OFDM constellations,

namely BPSK and 4-QAM, are vulnerable to early-detection. An attacker can identify (almost) any symbol with access to only a quarter (plus one) of the time-domain samples for BPSK and half (plus one) of the time-domain samples for 4-QAM.

- For the highly performant BPSK, we constructively prove the existence of valid late-commit attack sequences for all non-pulsed symbols. Those factors jointly lead to a deterministically achievable, significant distance reduction.
- We identify a possible countermeasure that involves a continuous extension of the constellation grid.

The chapter is organized as follows. The following Section 5.2 introduces secure ranging and summarizes the main results. Section 5.3 introduces the vulnerabilities of highly robust OFDM configurations. In Section 5.4, we address a potential countermeasure. We discuss our findings in a broader sense in Section 5.6 and provide related work in Section 5.5 before concluding in Section 5.7.

5.2 Background and Summary of Results

Over the last two decades, OFDM and its variants have become the predominant modulation techniques for high-throughput wireless communication in both the WiFi and cellular domains (4G, 5G). In the cellular domain, we see a trend towards high signal bandwidths (100 MHz and more), which furthers the adoption of OFDM modulation and increases the utility of those signals for ranging based on Time of Flight measurement. The security of such systems against physical-layer attackers depends on certain time-domain properties of the modulation. However, due to the information being encoded in the frequency domain, the resulting physical-layer security properties against a distance-modifying attacker do not follow trivially. To the best of our knowledge, they have not been studied so far.

5.2.1 OFDM

Orthogonal frequency-division multiplexing encodes message bits in frequency domain and transforms them into time-domain by an inverse Fourier transform, i.e.,

$$\mathbf{c} = \mathcal{F}^{-1}\{\mathbf{C}\},$$

which is defined as

$$c_n = \sum_{k=0}^{n_s-1} C_k e^{\frac{i2\pi k}{n_s} n}.$$

The values of $\mathbf{C}$ are determined by the symbol bit-sequence $\mathbf{b}$ and the constellation mapping $MAP(\cdot)$, e.g., BPSK, Quadrature Phase Shift Keying (QPSK), 16-QAM, etc. This results in the transmitted signal

$$\mathbf{c} = OFDM(\mathbf{b}) = \mathcal{F}^{-1}\{MAP(\mathbf{b})\},$$

which is sent over the wireless channel. The receiver then detects the information bits after performing a Fast Fourier Transform (FFT) on the incoming signal,

$$\mathbf{b}' = OFDM^{-1}(\mathbf{c}) = DEMAP(\mathcal{F}\{\mathbf{c}\}),$$

where the demapping operation is a hypothesis test based on the constellation set. As information bits are transmitted on orthogonal subcarriers (illustrated in Figure 5.1), OFDM provides resilience to frequency-selective fading. The dips in the channel transfer function caused by fading remain

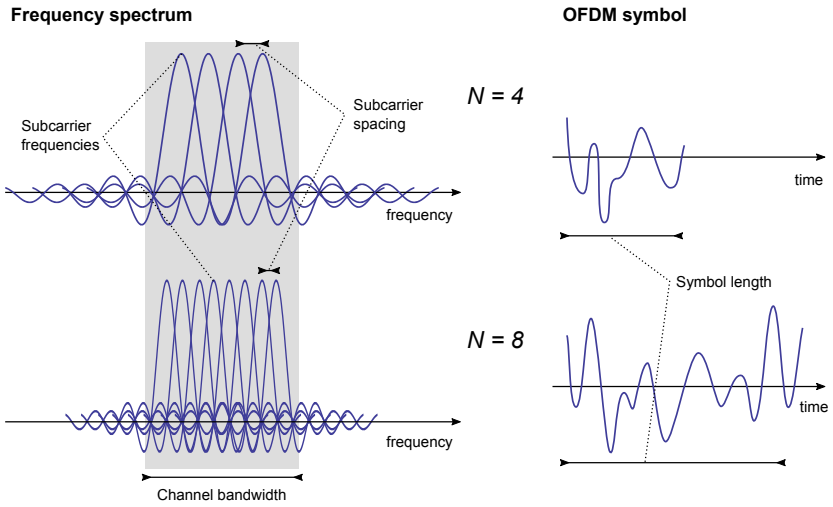

Figure 5.1: OFDM signal in frequency and time domain for different numbers of subcarriers (N). The frequency spectrum shows how the subcarriers share the channel bandwidth. The transmitter modulates message bits on individual subcarriers and applies an Inverse Fourier Transform to arrive at the time-domain samples (on the right).

constrained to a subset of the subcarriers. The receiver can maintain the orthogonality under a channel by adding a Cyclic Prefix (CP), which means to prepend the last few samples of the symbol at the beginning, thus circularizing the symbol. This allows simple equalization on a per-subcarrier level, as orthogonalization ensures an independent impact of the channel on each subcarrier.

To enable reliable communication, an OFDM transceiver has to perform additional tasks, namely synchronization, frequency and sampling offset correction, channel estimation, and equalization. Introducing a Cyclic Prefix can help with those tasks. However, the use of a cyclic prefix has a detrimental effect on ranging security. The cyclic prefix adds redundancy such that an attacker can predict the last part of the symbol with absolute certainty, even after only listening to the first part of the symbol (i.e., the cyclic prefix). For the remainder of this chapter, we are therefore only concerned with "plain" OFDM symbols that neither contain a cyclic prefix nor any guard symbols or bands. This

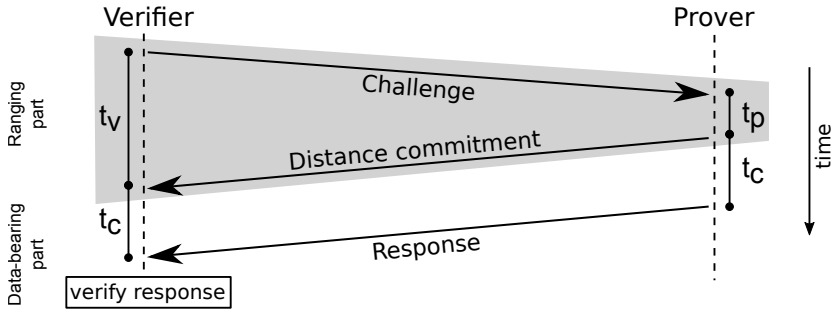

Figure 5.2: ToF ranging with a known static distance commitment. ToF $= (t_v - t_p)/2$ where $t_p \ll t_c$ are fixed parameters to accommodate hardware delay (t_p) and time to compute the response (t_c).

is a realistic assumption, which the IEEE 802.11az standardization group has also made. When in ranging mode, OFDM symbols must not feature (additional) redundancy, such as a cyclic prefix.

5.2.2 Distance bounding and secure ranging

Typically, a distance-bounding or secure ranging protocol allows a prover to convince a verifier to be within a certain distance. Among the different techniques to measure physical distance based on a radio signal, time-of-flight measurement is the only one with the potential of being secure against a physical-layer attacker. This is based on the observation that an unknown signal's arrival time cannot be meaningfully modified (i.e., reduced) by an attacker, as opposed to the signal's absolute strength or phase. We focus on a scenario where two entities, a verifier and a prover, determine their distance by measuring the ToF of a signal exchange, as illustrated in Figure 5.2. We assume the prover to be trusted and, in particular, entrusted with maintaining a time schedule that feeds into ToF estimation. We will henceforth assume the use of a *distance commitment*, as presented in [118]. This allows us to separate the fast reply from the data-bearing part in a challenge-response protocol for secure ranging, removing the need for fast processing of the challenge, i.e., to decouple the time-critical part of the protocol, unlike rapid bit exchange in Brands and Chaum [22]. Alternatively, the reply time could even be communicated by the prover after the ranging exchange.

Irrespective of this protocol design choice, the crucial requirement on the data-bearing part is that an attacker cannot advance the response in time through reactive interference.

Attacks against distance bounding and ranging.

Like the previous part, this chapter also focuses on Mafia Fraud, where both verifier and prover are honest, and an external attacker (a separate entity) attempts to modify the ToF measurement such that the prover appears to be closer to the verifier. This is also the attack scenario the IEEE 802.11az task group is mainly concerned with.

In order to achieve a distance reduction, the attacker has to make sure the challenge message is registered at the prover at an earlier time than the legitimate challenge, and/or, advance both distance commitment and response message in a way that they arrive at the verifier at an earlier time. The attacker can operate either on the protocol/data-layer or on the symbol level to inject and advance the messages. If the adversary cannot predict the content of the messages, it is forced to resort to the symbol level and has to mount an ED/LC or another physical-layer attack. ED/LC attacks are a likely threat to an OFDM PHY since the symbols are relatively long. We explain the ED/LC attack assuming the attacker attempts to advance the challenge message. The same technique can be applied to the response message.

For every symbol the verifier transmits, the attacker also emits a symbol, such that it registers with a time advantage at the prover. Because the attacker does not know the exact symbol a priori, the first part of the adversarial symbol can be random noise, tricking the prover into believing that the wireless channel has distorted the start of the symbol. The adversary starts transmitting early even though the exact symbol is not known yet.

Since wireless transmission is not instant and the symbols have a certain duration, the attacker listens to the verifier's transmission (while interfering with the prover) and tries to detect the verifier's actual symbol based on an initial segment of each symbol. This process is called early detection. Assuming the attacker succeeds and early-detects the verifier's symbol with high probability, it changes its own transmission from noise to a valid symbol—or a signal that is interpreted by the prover as the intended symbol. The adversary superimposes its signal onto the legitimate signal. The adversarial signal has to take such effects into account and has to be transmitted at higher power for it to be decoded

correctly at the receiver. This step of the attack is called late-commit and, if successful, makes the second part of the adversarial symbol appear as a valid symbol to the prover. If this succeeds for a sufficient number of symbols of a ranging frame, the attacker succeeds in reducing the ToF measurement and in accomplishing the Mafia Fraud.

5.2.3 IEEE 802.11az

Within the IEEE 802.11az task group, there is an ongoing standardization effort towards secure OFDM-based ranging [14]. Publicly available, preliminary documents indicate that a physical-layer attacker is considered a threat and part of the ongoing discussion. These documents discuss an attacker with limited reaction times and countermeasures evolving around coarser measures, such as avoiding cyclic prefixes and highly redundant encoding. Some documents treat a similar attacker as introduced in this work, operating on the sub-symbol level, however, without a rigorous study underpinning the presented measures. Our work aims to help bridge this gap with a rigorous physical-layer analysis that can motivate the choice of the modulation for the symbol sequences used for ToA estimation. This allows extending the security argument to a physical-layer attacker that is not constrained in its reaction time.

5.2.4 Known principles clash with implementation and performance constraints

Low Peak-to-Average Power Ratio (PAPR) is an important signal property for performant operation in real-world systems. The reason is that fast changes in the signal (the opposite of low PAPR) are challenging to amplify without encountering non-linearities of the hardware, causing inter-carrier interference and limiting overall performance (i.e., communication distance). Due to power constraints, many end-devices have amplifiers optimized for efficiency, which makes them, in turn, highly nonlinear. Therefore, OFDM uses different techniques to limit the PAPR. One of them is to limit the code set and exclude high-PAPR symbols, of which pulses are the most extreme examples.

On the other hand, existing proposals for modulations enabling secure time-of-arrival measurement all assume pulses that are spaced by more than the channel delay-spread [69]. The existing understanding of secure physical layer design for ranging and the requirements on practical OFDM systems cannot be reconciled without either a heavy performance

(data rate, range) penalty or hardware changes. E.g., [112] makes a proposal to use a multicarrier system like a single-carrier (UWB-like) system, a technique that provides security, however heavily constrains the information content per symbol and relies on time-domain techniques at the receiver. The question addressed in this work is to investigate the security of multicarrier modulations in general and whether we can find a technique that allows for secure ranging within the practical OFDM assumptions, those being parallel transmission on all subcarriers and frequency-domain mapping and demapping.

5.2.5 Summary of results

In this chapter, we show that the OFDM configurations that offer the highest robustness, i.e., BPSK and 4-QAM, are prone to ED/LC attacks. We provide mathematical proof that irrespective of the number of subcarriers, the first quarter and first half of the symbol allow the attacker to learn the full BPSK or 4-QAM symbol, respectively. In the case of BPSK, every symbol can be late-committed with only half the samples. For BPSK, the most robust constellation, the susceptibility to both early-detection and late-commit attack leads to a deterministically achievable distance reduction of more than 200 m for a typical IEEE 802.11 OFDM configuration of 20 MHz split into 64 subcarriers. In the case of 4-QAM, we show that an attacker's late-commit success can be significantly improved with an optimization technique, resulting in a considerable adversarial advantage in a distance-reducing attack. We identify the structure of the frequency-domain constellation grid as the main enabler of strong early-detection strategies and identify a technique that uses orthogonal noise and a random phase shift as a possible countermeasure since those operations limit structural information about the frequency-domain constellation.

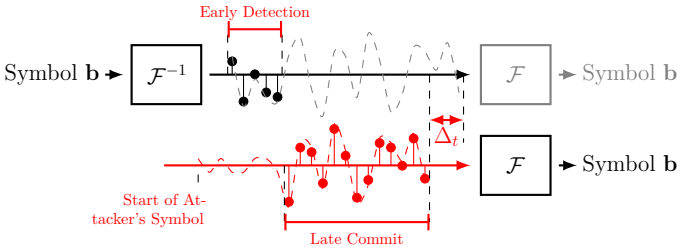

Figure 5.3: OFDM distance-reduction attack: An attacker that can detect the symbol based on a number of initial samples can send a symbol representation consisting of the later samples preemptively and thereby achieve advancement (Δ_t) of the symbol (ED/LC attack). This enables a distance-reduction attack in the context of ToF ranging.

5.3 The OFDM ED/LC attack

In the context of ToF distance measurement, it is well-known that an attacker can exploit the time redundancy of symbols to decrease the measured distance, irrespective of cryptographic primitives [46, 82]. For example, if a modulation uses repetitions of a certain signal shape for improved robustness, an attacker can detect this symbol early by only decoding the first repetition and can late-commit to such a symbol by only transmitting the last repetition. This behavior is illustrated in Figure 5.3. For the outcome of such an attack, it is not important whether the symbols are sent in direct succession or in separate frames. It is, however, not straightforward how such an attacker performs in OFDM since the symbols are encoded in the frequency domain and only before transmission transformed into a time-domain symbol. The physical-layer attacker presents us with a heavy asymmetry in the information-theoretic sense between the attacker's observation and the verifiable information at the receiver (prover and verifier). An attacker can "understand" and interact with the signal at the physical layer. However, the receiver will only be able to assess the validity of the signal after demodulating it into bits. This demodulation must be robust against noise and multi-path channel propagation for reliable operations over long communication ranges, and we do not want to break with this requirement.

5.3.1 Attacker model

As mentioned in Section 5.2.2, we assume a Mafia Fraud attack scenario, where the external attacker is located between the two legitimate entities that measure their relative distance via signal ToF. The attacker's goal is to decrease the measured time of arrival of the communication protocol employing an ED/LC attack or a similar technique. While an attacker could also work on the protocol/data-layer, in the following, we constrain ourselves with an attacker that operates on the symbol level. If not taken care of at the physical layer, such an attacker can be successful irrespective of cryptographic primitives and protocols (such as distance bounding) on higher layers. Moreover, we assume an attacker that can receive and react to signals at the physical layer at arbitrarily high sensitivity and arbitrarily small reaction times. We understand that this is an unattainable attacker model in the real world; however, in order to account for future technological advances, we do not want to limit ourselves to the current state-of-the-art results. As the attacker's aim is distance reduction, we assume the attacker has full control over his signal power, and the legitimate signal is negligible in relative power. This naturally applies in a scenario where the legitimate devices are out of communication range, however, an attacker can relay signals, e.g., by wire (relay attack). This provides the attacker the advantage of amplifying the signal as needed and, in particular, establishing a communication path, whereas in reality, the victim devices might be out of range. This attacker model is in line with the ones chosen in recent proposals for secure ToF estimation [69, 111].

5.3.2 Robust OFDM configurations

Performance-enhancing techniques such as channel compensation, cyclic prefix, and coding can create additional vulnerability to an ED/LC attacker since those techniques create dependencies between parts of the symbol. The absence of such techniques can be compensated by using highly robust constellations for the symbol sequences used for ToF estimation. For this reason, we cover the two most robust constellations in our analysis.

BPSK

BPSK uses a maximally robust symbol constellation. In BPSK, each subcarrier can only assume one of two possible values: +1 or -1. Robustness

is an important characteristic and a key design goal on the bit sequences used for ranging in recent standardization efforts [8, 14]. Unfortunately, as opposed to the pulsed scenario, OFDM BPSK proves a particularly bad choice regarding a distance-reducing attacker, especially an ED/LC attacker. The reason is that a limited set of constellation points in frequency-domain results in strong time-domain *symmetry*. Because all n_s frequency-domain values are real, any BPSK symbol exhibits Hermitian symmetry in time-domain. This means the last $n_s/2-1$ time-domain samples are complex-conjugated versions of the $n_s/2-1$ samples after the initial sample c_0. Indeed, we will prove constructively that strong late-commit sequences exist for all non-pulse BPSK symbols, requiring an attacker to send only half the samples. In addition, we will see that the time-domain samples contain a substantial amount of differential information about the entire symbol, granting a steep learning curve to the early-detecting attacker.

4-QAM

4-QAM is the minimal constellation that transmits bits on both signal-space dimensions in parallel, resulting in four possible constellation points per tone. As a consequence, it provides double the data rate, however at slightly less robustness under equal overall signal strength, compared to BPSK.

5.3.3 Early detection

An early-detecting attacker is looking for the algorithm that will detect the correct message with highest probability, for a given detection delay δ_{ed}. The advantage of an early-detect algorithm *ED* at detection delay δ_{ed} over a symbol set $\mathcal{B}$ is defined as

$$A_{ED}(\mathcal{B}, \delta_{ed}) = \underset{\mathbf{b}\leftarrow\mathcal{B}}{P}\left(ED\left(\mathbf{c}_{ED}||0^{(n_s-\delta_{ed})}\right) = \mathbf{b}\right),$$

where

$$\mathbf{c}_{ED} = c_0||...||c_{\delta_{ed}-1},$$

and

$$\mathbf{c} = OFDM\left(\mathbf{b}\right).$$

In the following, we introduce two different viewpoints on early-detection. The first is standard OFDM demodulation, which simply

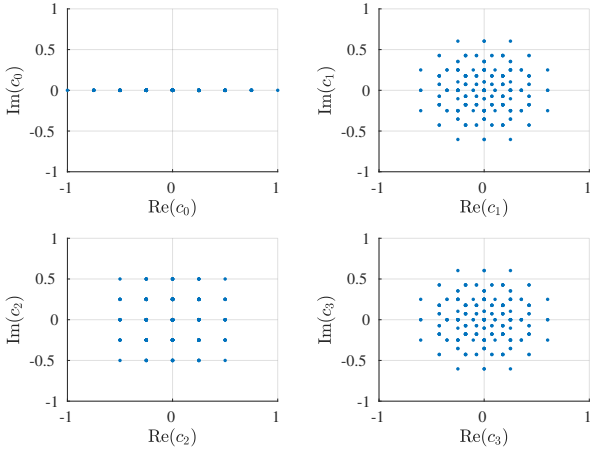

Figure 5.4: All possible values of the first four time-domain samples of an BPSK OFDM symbol with eight subcarriers. Odd samples are numerically diverse, i.e., contain a lot of information about the symbol sequence.

applies an FFT on the zero-padded time-domain signal before testing on the polarity of each tone. Then, we analyze a time-domain sample-by-sample matching strategy, assuming an attacker with optimal sensitivity. This second viewpoint shall grant insights into optimized strategies, e.g., strategies that compensate for Inter-Carrier Interference (ICI) imposed by the fact that later time-domain samples are unknown, an effect that highly impacts standard demodulation.

Direct demodulation

This approach feeds the early-detect signal with trailing zeros into an OFDM demodulator. This is equivalent to applying the FFT on the ideal symbol multiplied with a 1-0 step function. Doing so, the attacker directly maps the time-domain samples to all frequency subcarriers in order to then detect the bits. The shortcoming of this approach is that the ED condition (i.e., the later samples being cut off) is equivalent to applying a sharp filter in time domain, which corresponds to a wide (1/f) dispersion profile in frequency domain. This means every bit is subject

to significant inter-carrier interference, which results in a relatively high bit error rate. The computational complexity of this approach is given by the FFT algorithm, i.e., $\mathcal{O}(n_s \log(n_s))$.

Number-theoretic viewpoint

In this section, we deal with an idealized early-detection attacker that matches the observed time-domain samples against all possible symbols. Security against such an attacker can only be based on numerical ambiguity of the initial samples. However, as we will show, the initial samples of BPSK-modulated OFDM symbols contain a substantial amount of information about the entire sequence—a fact directly related to the FFT size being a power of two.

The set of possible time-domain values of each sample is limited, as we illustrate in Figure 5.4. The figures show the possible values that can be assumed by the first four time-domain samples for all possible bit sequences of a BPSK OFDM symbol with eight subcarriers. We can observe that the odd samples (i.e., samples c_1 and c_3) can assume many different distinct values because those are based on a linear combination of all distinct complex exponentials. Numerical matching exploits the systematic nature of the modulation, i.e., the fact that the limited frequency-domain constellation points, together with distinct complex exponentials, result in distinct numerical time-domain samples. By analyzing the conditions under which the numerical samples represent unique bit combinations, we can arrive at a concrete upper bound of the number of time-domain samples representing the bits of the symbol unambiguously.

Theorem 1 *An attacker with infinite sensitivity operating on a non-pulsed BPSK OFDM symbol*[1] *(with $n_s = 2^M$ for $M \in \mathbb{Z}^{>1}$) requires at most $n_s/4 + 1$ samples to detect the symbol.*

Proof 1 *An OFDM time-domain sample can be represented as the inverse Fourier transform of the frequency-domain modulated symbol samples. In the following, we will show that left-right antivalent bits (i.e., bits that do not repeat after $n_s/2$ samples) are leaked with the first odd*

[1] As a non-pulsed BPSK symbol we define a symbol that under no (time-domain) circular shift has $\mathbf{C} = \pm\{1, 1, ...\}$

(respectively, any odd) sample. For the first odd sample, we have

$$c_1 = \sum_{k=0}^{n_s-1} C_k e^{2\pi ik/n_s}$$

$$= \sum_{k=0}^{n_s/2-1} C_k e^{2\pi ik/n_s} + \sum_{k=n_s/2}^{n_s-1} C_k e^{2\pi ik/n_s}$$

$$= \sum_{k=0}^{n_s/2-1} \left(C_k + e^{\frac{2\pi i n_s/2}{n_s}} C_{k+n_s/2} \right) e^{2\pi ik/n_s}$$

$$= \sum_{k=0}^{n_s/2-1} (C_k - C_{k+n_s/2}) e^{2\pi ilk/n_s},$$

since

$$e^{\frac{2\pi i n_s/2}{n_s}} = e^{\pi i} = -1.$$

Due to the limited constellation set of BPSK modulation (i.e., $C_k \in \{-1, 1\}$), we can express the difference between frequency-domain samples in terms of a logical bit-level operation:

$$c_1 = 2 \sum_{k=0}^{n_s/2-1} C_k (b_k \oplus b_{k+n_s/2}) e^{2\pi ik/n_s}$$

The negative sign is equivalent to a π-phase rotation of the complex exponential, therefore, equivalently:

$$c_1 = 2 \sum_{k=0}^{n_s-1} a_k e^{2\pi ik/n_s}, \quad a_k \in \{0, 1\}$$

In order to understand whether c_1 uniquely represents the sequence a, we consider the difference between two of these polynomials for different sequences $a^{(1)}$ and $a^{(2)}$:

$$\sum_{k=0}^{n_s-1} a_k^{(1)} e^{2\pi ik/n_s} - \sum_{k=0}^{n_s-1} a_k^{(2)} e^{2\pi ik/n_s}, \quad a_k^{(1)}, a_k^{(2)} \in \{0, 1\}$$

$$= \sum_{k=0}^{n_s-1} \varepsilon_k e^{2\pi ik/n_s}, \quad \varepsilon_k \in \{0, 1, 2\}$$

We assume the sequences $a^{(1)}$ and $a^{(2)}$ not to be identical, therefore there exists a $k \in \{0, ..., n_s - 1\}$, for which $\varepsilon_k > 0$. Therefore, this is a sum over up to $n_s/2$ of the n_s-th roots of unity, with $\varepsilon_k \varepsilon_{k+n_s/2} = 0$. For the sake of contradiction, we consider the above expression to be zero, i.e.,

$$\sum_{k=0}^{n_s-1} \varepsilon_k e^{2\pi i l k/n_s} = 0, \quad \varepsilon_k \in \{0, 1, 2\}.$$

A result from algebraic number theory reveals interesting properties of such vanishing sums of roots of unity [68]. Corollary 3.4 in [68] states that if $m = p^a q^b$, where p, q are primes, then, up to a rotation, the only minimal vanishing sums of m-th roots of unity are $1 + \zeta_p + ... + \zeta_p^{p-1}$ and $1 + \zeta_q + ... + \zeta_q^{q-1}$ (where ζ_p denotes a p-th primitive root of unity), and rotations thereof. A minimal vanishing sum is defined as a sum of roots of unity that amounts to zero, yet contains no sub-sum that is zero. In our case, due to the FFT size being a power of two, we have $p = q = 2$, meaning the only minimal vanishing sum is given by one plus the 2nd primitive root of unity (and rotations thereof). This means, $1 - 1$, and rotations thereof, i.e., $e^{\rho i} + e^{(\rho+\pi)i}$ for $\rho \in [0, \pi)$. However, since we have $\varepsilon_k \varepsilon_{k+n_s/2} = 0$, the expression above does a) not contain any minimal vanishing sum nor b) constitute a minimal vanishing sum, which proves the contradiction.

In consequence, every left-right antivalence in the bit sequence results in a unique contribution to every odd time-domain sample. Left-right equivalence, on the other hand, cancels out the contributions. This means, the odd sample does not convey any information on bits that repeat after $n_s/2$ samples, however, conveys all information about bits that are inverted after $n_s/2$ samples. Conversely, the first non-zero even sample (i.e., the sample c_2) is oblivious to information about tones that repeat after $n_s/4$ samples, however conveys information about antivalence of tones $n_s/4$ apart.

In the following, we consider the sequence of samples with indices that are powers of two ($l = 2^L$ and $0 \leq L < M$).

From sample c_1, we learn the sequence $C_k(b_k \oplus b_{k+n_s/2})$, i.e., the values of the left-right antisymmetric bits. We can create a compensation term $\sum_{k=0}^{n_s/2-1} C_k(b_k \oplus b_{k+n_s/2}) e^{2\pi i 2k/n_s}$ and add it to c_2, which recovers the equivalent sample of the inverse Fourier transform on the first half of

the spectrum only, since

$$c_2 = \sum_{k=0}^{n_s-1} C_k e^{2\pi i 2k/n_s}$$

$$= \sum_{k=0}^{n_s/2-1} \left(C_k + e^{\frac{2\pi i 2 n_s/2}{n_s}} C_{k+n_s/2} \right) e^{2\pi i 2k/n_s}$$

$$= \sum_{k=0}^{n_s/2-1} \left(C_k + C_{k+n_s/2} \right) e^{2\pi i 2k/n_s}.$$

Hence,

$$c_2 + 2 \sum_{k=0}^{n_s/2-1} C_k (b_k \oplus b_{k+n_s/2}) e^{2\pi i 2k/n_s} = 2 \sum_{k=0}^{n_s/2-1} C_k e^{2\pi i 2k/n_s}$$

This allows, in turn, to recover the sequence $C_k(b_k \oplus b_{k+n_s/4})$ and so on.

This procedure can be invoked recursively until the sequence consists of four samples only. The remaining uncertainty is only given by the center pulse (i.e., 1,-1,1,-1 vs. -1,1,-1,1) (single equivalence is leaked by DC sample), which we exclude from the proof. Hence, under the last recursion step we have $n_s/2l = 2 \Leftrightarrow l = n_s/4$. This means, we need in total $n_s/4 + 1$ samples for ideal detection.

For BPSK, we show the resulting bit error rate as a function of the early-detection delay δ_{ed} in Figure 5.5, and contrast it to direct demodulation. Our bound indicates full symbol knowledge at sample $c_{n_s/4}$, whereas direct demodulation requires more than half the samples for error-free detection.

A brute-force attacker that matches time-domain samples against pre-computed traces faces a space complexity of $\mathcal{O}\left(3^{n_s/2}\right)$. However, it is expected that polynomial-time maximum likelihood detectors exist. The fact that both the nature of the inter-carrier interference and the possible constellation points are known to the attacker makes a compelling case for the existence of efficient cancellation techniques.

Furthermore, we can reduce the problem of ideal time-domain matching in 4-QAM to the same problem on two interleaved BPSK symbols.

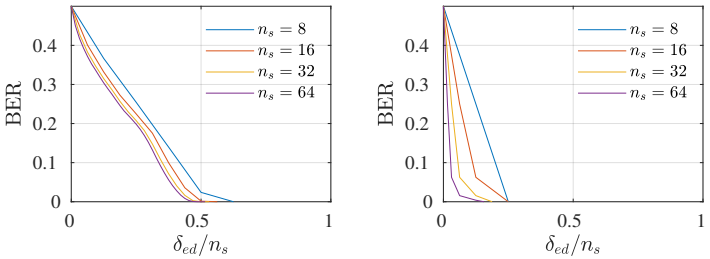

Figure 5.5: BPSK early detection: Early-detection bit error rate under direct demodulation (left) and ideal attacker behavior (right) as a function of the relative detection delay δ_{ed}/n_s.

Corollary 1 *An attacker with infinite sensitivity operating on a non-pulsed 4-QAM OFDM symbol*[2] *requires at most* $n_s/2+1$ *samples to detect the symbol.*

Proof 2 *Without loss of generality, an attacker can run the early-detection on the signal that is circularly shifted by* $n_s/4$ *to the left and start the early-detection procedure* $n_s/4$ *delayed. This is equivalent to a multiplication of the frequency-domain representation by a sequence* $1, -i, -1, i, ...$ *Starting with* c_1*, the attacker can then separate every sample of the shifted representation in its symmetric and antisymmetric components, which correspond to the time-domain representation of the real and imaginary parts of the frequency-domain symbols. These components are individually BPSK-modulated. Theorem 1 states that a non-pulsed BPSK symbol requires at most* $n_s/4 + 1$ *samples for early-detection. Together with the offset required for the shift operation, we arrive at* $n_s/4 + n_s/4 + 1 = n_s/2 + 1$ *samples for ideal early-detection of the non-pulsed 4-QAM symbol.*

The same separation strategy for frequency-domain I and Q components can be applied irrespective of constellation density. Our main insight from the number-theoretic analysis is that, without assumptions on an attacker's sensitivity, even early samples contain a substantial amount of differential information about the entire symbol which, due

[2] As a non-pulsed 4-QAM symbol we define a symbol that under no (time-domain) circular shift has $\mathcal{R}(\mathbf{C}) = \pm\{1, 1, ...\}$ or $\mathcal{I}(\mathbf{C}) = \pm\{1, 1, ...\}$

to the structure of the constellation, directly translates to information about the symbol bits.

5.3.4 Late-commit

The late-commit problem for the attacker consists in finding a sequence of samples that result in the correct symbol at the receiver under a delayed onset of transmission. For a given symbol, the ability of an attacker to late-commit with a certain delay is not probabilistic but an immutable property of this symbol.

The fundamental principle behind late-committing to a symbol is reflected by the fact that the attacker does not have to provide a signal that is actually close on the physical layer (e.g., in the L2-sense), but only one that creates the correct bits at the receiver. In general, finding a valid late-commit sequence for an OFDM symbols is not straightforward. There is room for optimization on a per-symbol basis beyond just sending the late part of the symbol, as we illustrate in Figure 5.6.

Irrespective of the optimization technique, for a given symbol sequence $\mathbf{b}$ and transmission delay δ_{lc}, the goal of the attacker is to find a late-commit signal $\mathbf{c}^{lc}$ consisting of $n_s - \delta_{lc}$ samples that, if prepended with δ_{lc} zeros, minimizes the Hamming Distance $H(\ ,\ )$ between the demodulated late-commit signal and the actual symbol sequence $\mathbf{b}$. The optimal late-commit algorithm LC is defined as

$$LC(\mathbf{b}, \delta_{lc}) = \arg\min_{\mathbf{c}^{lc}} \left\{ H\left(OFDM^{-1}\left(0^{\delta_{lc}} || \mathbf{c}^{lc}\right), \mathbf{b}\right)\right\}.$$

We say LC is a δ_{lc}-LC algorithm under symbol set $\mathcal{B}$ iff

$$H\left(OFDM^{-1}\left(0^{\delta_{lc}} || LC(\mathbf{b}, \delta_{lc})\right), \mathbf{b}\right) = 0, \forall \mathbf{b} \in \mathcal{B},$$

meaning an OFDM receiver will correctly interpret each symbol sequence despite the attacker omitting the first δ_{lc} samples of each time-domain symbol. A $\hat{\delta}_{lc}$-LC algorithm is optimal if there exists no δ_{lc}-LC algorithm for $\delta_{lc} < \hat{\delta}_{lc}$.

In the following, we will constructively prove the existence of a $n_s/2$-LC algorithm for the full BPSK symbol set without c_0-pulses, i.e. for $\mathcal{B}' = \mathcal{B} \setminus \mathcal{P}$, whereas $\mathcal{P} = \{\pm(1, 1, 1, ...)\}$.

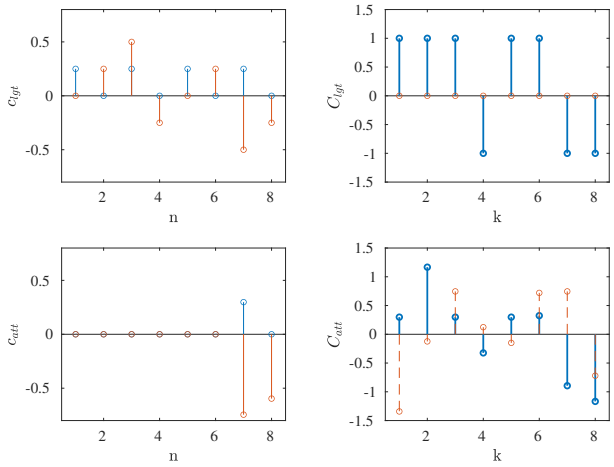

Figure 5.6: Example of a late-commit attack on BPSK OFDM. Ideal signal (top) vs. adversarial signal (bottom), both in time (left) and frequency domain (right). The attacker only provides the last two time-domain samples yet can create the correct BPSK symbol, as only the real part (blue) of the frequency-domain representation is of interest.

Deterministic BPSK late-commit ($n_s/2$-LC)

As a consequence of Hermitian symmetry, a late-committing attacker can generate any non-pulse BPSK symbol using only the samples corresponding to the second half of the symbol.

Theorem 2 *There exists a $n_s/2$-LC algorithm under the set of all non-pulsed BPSK OFDM symbols.*

Proof 3 *Consider a split of the frequency-domain symbol* $\mathbf{C}$ *into its even and odd contributions, i.e.,*

$$C_k^{(E)} \coloneqq C_{2k}, \;\; k = 0, ..., n_s/2 - 1,$$
$$C_k^{(O)} \coloneqq C_{2k+1}, \;\; k = 0, ..., n_s/2 - 1.$$

The corresponding time-domain contributions are given by the inverse Fourier transform:

$$\mathbf{c}^{(E)} = \mathcal{F}^{-1}\left\{\mathbf{C}^{(E)}\right\}$$

$$\mathbf{c}^{(O)} = \mathcal{F}^{-1}\left\{\mathbf{C}^{(O)}\right\}$$

From the definition of the Discrete Fourier Transform (DFT), we know that

$$c_n = c_n^{(E)} + e^{-\frac{2\pi i}{n_s}n} \cdot c_n^{(O)}, \quad n = 0, ..., n_s/2 - 1.$$

Hence, the late-commit condition, i.e.,

$$c_n = 0, \quad for \;\; n = 0, ..., n_s/2 - 1,$$

imposes a clear relationship between even and odd frequency-domain samples (as given by the trigonometric interpolation of every second sample), respectively, its individual time-domain contributions, i.e.,

$$c_n^{(E)} = -e^{-\frac{2\pi i}{n_s}n} \cdot c_n^{(O)} = -g_n \cdot c_n^{(O)}, \tag{5.1}$$

where we define the half-period complex exponential $\mathbf{g}$ *as*

$$g_n = e^{-\frac{2\pi i}{n_s}n}, \quad n = 0, ..., n_s/2 - 1.$$

Taking the Fourier transform of Equation 5.1 yields

$$\mathbf{C}^{(E)} = -\frac{1}{n_s}\mathbf{G} * \mathbf{C}^{(O)}, \tag{5.2}$$

where $\mathbf{G}$ *is defined as*

$$G_k = \sum_{n=0}^{n_s/2-1} e^{-\frac{i2\pi}{n_s}n} e^{-\frac{i2\pi}{n_s/2}nk} = \sum_{n=0}^{n_s/2-1} e^{-\frac{i2\pi}{n_s}n(1-2k)} \tag{5.3}$$

and can be considered a frequency-domain 'filter' that corresponds to said time-domain relationship, representing the resulting dispersion profile through inter-carrier interference. Importantly, the real part of Equation 5.3 constantly evaluates to 1*, due to circular symmetry.*

In the following, we treat the late-commit signal as a sum of the perfect odd and even contributions separately. Without loss of generality, we assume the odd contributions as ideal.

Only the real part of Equation 5.2 matters for BPSK symbols, for which the circular convolution evaluates to

$$\mathbb{R}\left\{\tilde{\mathbf{C}}^{(E)}\right\} = \mathbb{R}\left\{\mathbf{G}\right\} * \mathbf{C}^{(O)} = -\frac{1}{n_s}\sum_{k=0}^{n_s/2-1} C_k^{(O)}.$$

This follows from odd contributions being ideal, i.e. real values +1,-1 only, which means that only the real part of $\mathbf{G}$ *matters.*

Inter-carrier interference terms are given by respective first time-domain samples, for contribution with odd samples ideal:

$$\mathbb{R}\left\{\tilde{C_k}^{(E)}\right\} = -c_0^{(O)},$$

and for the contribution with even samples ideal:

$$\mathbb{R}\left\{\tilde{C_k}^{(O)}\right\} = -c_0^{(E)}$$

If we now assume the two contributions are added, we can imagine the value of every bit to contain an ideal contribution and an interference term. Correct detection is achieved if no bit is flipped due to the interference term. We, therefore, need to limit the inter-carrier-interference to be less than the legitimate signal value. Consider the superposition, where $c_0^{(O)'} = \alpha \cdot c_0^{(O)}$ *and* $c_0^{(E)'} = \alpha \cdot c_0^{(E)}$*, for* $\alpha \in (0, 1)$*. This corresponds to a dampening of the first signal sample sent by the attacker by real-valued constant* α*. The resulting interference term will amount to* $\alpha \cdot c_0^{(O)}$*. The amplitude will be less affected, i.e.,* $1 \pm (1-\alpha) \cdot c_0^{(E)}$*. Without loss of generality (due to symmetry), we assume the bit to be 1. For correct detection of each bit, we need to have*

$$\underbrace{1 - (1-\alpha) \cdot c_0^{(E)}}_{Amplitude} - \underbrace{\alpha \cdot c_0^{(O)}}_{ICI} \overset{!}{>} 0$$

For $\alpha > 0$*, this is equivalent to*

$$c_0^{(E)} - c_0^{(O)} \overset{!}{>} \frac{c_0^{(E)} - 1}{\alpha},$$

which holds iff $\mathbf{c}$ *is not a pulse (since 1 is maximum DC), and the condition is not satisfied iff both even and odd frequency samples have full DC, which corresponds to the spectral profile of a pulse.*

5.3.5 ED/LC attack

We have presented independent strategies for early-detection and late-commit. This section deals with how an attacker can combine these elements into a successful distance-reduction attack. This combination is

characterized by a transition step from early-detection to late-commit. An ED/LC attack consists of independent stages for detection and late-commit, separated by the attacker's reduction target. We propose three fundamental strategies for transitioning from early-detection to late-commit. The first uses the same transition time for all symbols in the symbol set, the second chooses the transmission time adaptively, given the symbol. Thirdly we propose an adversarial strategy that is more general, without a strict transition.

Fixed transition

This attacker uses a fixed portion of each symbol for early-detection and late-commit. This corresponds to an attacker that does not pre-generate all late-commit signals in advance but generates the signal on the fly and transmits it at the earliest time required for any symbol in the symbol set. The latest late-commit time of any symbol in the message set will therefore be considered a strict upper bound for the delay at which the attacker has to guess the symbol under a given target for distance reduction.

The resulting adversarial advantage can be expressed in terms of the early-detection advantage, as

$$A(\mathcal{B}, \delta_{adv}) = A_{ED}\left(\mathcal{B}, \hat{\delta}_{lc}(\mathcal{B}) - \delta_{adv}\right),$$

for an advancement goal δ_{adv}.

Symbol-adaptive transition

This attacker incrementally learns about the symbol and uses this knowledge to optimize the start time of the late-commit attack. For this purpose, the attacker can be thought to maintain an uncertainty-set of symbols at each stage of the early-detection process and chooses the late-commit time to satisfy the lowest late-commit delay within this set. This way, the attacker optimizes the late-commit start time subject to his knowledge gained from early-detection. This behavior requires the attacker to pre-generate a significant fraction of all late-commit symbols in order to generate statistics on the latest possible late-commit delays subject to every symbol. The adversarial advantage under this model is bounded by the attacker's ability to correctly guess the symbol at the latest possible transmission time, given a certain reduction goal:

$$A(\mathcal{B}, \delta_{adv}) \leq \mathbb{E}_{\mathbf{b}\leftarrow\mathcal{B}}\left[A_{ED}\left(\mathbf{b}, \hat{\delta}_{lc}(\mathbf{b}) - \delta_{adv}\right)\right]$$

As an over-approximation of the attacker, we can consider above at equality. This corresponds to an attacker never waiting too long to start transmission of the late-commit symbol, i.e., being optimally informed about $\hat{\delta}_{lc}(\mathbf{b})$.

Interleaved ED/LC

This is the most generalized model with regards to the attacker's transition from ED to LC. It assumes the attacker continuous detects the legitimate symbol, even after starting to transmit late-commit samples. In other terms, this is an attacker that might start transmitting before getting a clear picture from the early-detection, and adjust each transmitted samples to new observations. This corresponds to the attacker model put forward in [69].

5.3.6 Distance reduction attack

We evaluate the vulnerability of BPSK and 4-QAM OFDM to an ED/LC distance-reduction attack by combining our findings for early-detection and late-commit.

BPSK

BPSK OFDM is vulnerable to an ED/LC attack that results in a deterministically successful distance reduction by a physical-layer attacker under the fixed transition model. We have proven that the attacker requires only half the samples for successful late-commit and a quarter of the samples for early-detection. This means the attack succeeds irrespective of asymptotic properties on the bit- and frame level (i.e., independently of the quality of entropy of the message bits and how many messages are exchanged).

Corollary 2 *An (ideal) attacker operating on a non-pulsed BPSK OFDM symbol can achieve a distance reduction corresponding to up to* $n_s/4 - 1$ *samples deterministically.*

Proof 4 *Theorem 2 states that for any non-pulsed BPSK OFDM symbol, there exists an* $n_s/2$*-LC algorithm. Theorem 1 states that an attacker*

	n_s			
	8	**16**	**32**	**64**
$\delta_{lc} - \delta_{ed}$ at $A = 1$	1	3	7	15
$\Delta_t = (\delta_{lc} - \delta_{ed})(n_s/20\,\text{MHz})$	50 ns	150 ns	350 ns	750 ns
$\Delta_d = \Delta_t * c$	15 m	45 m	105 m	225 m

Table 5.1: Maximum time advancement for BPSK OFDM at adversarial advantage $A = 1$, using ideal early-detection and a fixed transition. We assume a total bandwidth of 20 MHz split into varying numbers of subcarriers.

requires up to $n_s/4 + 1$ *samples to detect a non-pulsed BPSK OFDM symbol ideally. This leaves an attacker with* $n_s/2 - (n_s/4 + 1) = n_s/4 - 1$ *samples for distance reduction with* $A = 1$ *under the fixed transition model.*

Table 5.1 exemplifies the impact of the sample-level advancement on time and distance. For the numerical example, we assume a system bandwidth of 20 MHz, split into various numbers of subcarriers. The sample spacing is determined as the inverse of the system bandwidth. We observe a higher impact for systems with more and narrower subcarriers, e.g., the typical configuration for an IEEE 802.11 system consisting of 64 subcarriers is vulnerable to a distance reduction of up to 225 m if BPSK is used. It becomes evident that, under a fixed system bandwidth, higher numbers of subcarriers come at a loss for secure ToF measurement since the symbol duration is increased.

4-QAM

Corollary 1 states that an attacker requires only $n_s/2 + 1$ samples to ideally detect a 4-QAM symbol. Because the learning curve for early detection is steep, i.e., the sample with index $n_s/4 + 3$ already reveals more than 90 % of the symbol information, there is a potential for an interleaved strategy, which helps reduce the bit-error rate under any late-commit strategy.

5.4 Can OFDM be secured?

After identifying the major problems with secure ranging based on OFDM, this section proposes a potential direction for securing OFDM-based ranging. The underlying observation is that the possible set of constellation points can be randomized and extended to cover a continuous disk in the IQ plane, minimizing an adversary's structural knowledge about the modulation.

5.4.1 Continuous extension of constellation

We can increase the constellation density on the transmit-side by limiting the modulation to one dimension and adding a *noise dimension* to each tone. The rationale is to increase the numerical diversity of the resulting time-domain samples.

In addition, we can add a random phase shift to each tone that is pre-shared and inverted by the receiver before demodulation. Phase randomization is a common technique for PAPR reduction, i.e., existing hardware is expected to implement it. This approach leverages the same procedure for security against early detection. The random phase offsets create a dense, concentric constellation pattern if jointly applied with orthogonal noise and a denser than minimal constellation set (e.g., eight constellation points in the information dimension). If we move beyond BPSK for the information dimension, the resulting frequency-domain constellation covers a concentric disk. We can choose orthogonal noise and phase offset at an arbitrarily fine resolution without any impact on performance. This leaves an attacker with minimal a-priori numerical knowledge, only a lower and upper bound on each tone's amplitude.

5.4.2 Evaluation

We analyze our proposal, consisting of eight constellation points in the information dimension, together with orthogonal noise and phase randomization, in terms of its security against early-detection.

Information-theoretic security against ED

Phase randomization together with fine-grained orthogonal noise can provide information-theoretic security against an early-detecting attacker.

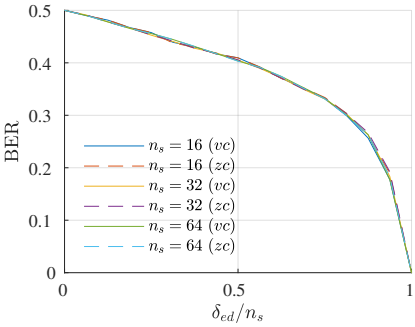

Figure 5.7: Bit-error rate of an early-detecting attacker against full phase-randomization with orthogonal noise as a function of the relative detection delay and for different numbers of subcarriers.

The fundamental reason that any point in a continuous area in the IQ-plane is a valid value for each frequency-domain sample. This means, any partial time-domain sequence can be continued in many ways such that each tone ends up within the valid range. This uncertainty is associated with a certain bit error rate. We verified this in a simulation, where we randomly sampled valid continuation (vc) sequences for many different symbols and evaluated their resulting bit-error rate, as shown in Figure 5.7. We contrast them to zero-extended (zc) symbols and see no difference in the resulting bit error rate.

5.5 Related Work

We compare our analysis of multicarrier-based ranging with existing proposals for secure single-carrier ranging, as well as other physical-layer concepts in wireless communication. In particular, we focus on mechanisms that attempt to protect a wireless signal on the physical layer. Secure ranging achieves a similar goal since it has to guarantee that the arrival time of the signal can not be subverted by external influence, in addition to the protection of physical-layer attributes and data integrity.

5.5.1 Single-carrier Ranging

Research has yielded a handful of protocols for secure single-carrier ranging and distance measurements. The majority of them focus on UWB, a technology that provides non-cooperative communication at bandwidths of up to 500 MHz. Due to their wide spectral use, UWB devices have to operate at limited output power, but the high bandwidth allows them to send short pulses that have high immunity to multi-path fading. If data is encapsulated in nano-second pulses, the surface for ED/LC attacks is very narrow since an attacker is forced to advance or delay single pulses. Different effective proposals that describe how pulses need to be emitted can be found in [69, 111].

The UWB technology has also resulted in few commercial products [7, 85]. However, the main disadvantage of UWB ranging is its limited power output and as a consequence, distances greater than 50 to 100 meters (depending on channel conditions) are difficult to overcome. As a remedy, frame size has to be increased, but this leads to long communication times in an already uncoordinated spectrum. UWB ranging is therefore mainly used for indoor positioning or in two-device configurations, such as key-less entry systems for vehicles.

OFDM, on the other hand, has proven to be an extremely reliable modulation technique. While techniques for improving the performance of OFDM-based ranging have been proposed [53], its security against physical-layer attacks has, to the best of our knowledge, not been studied so far. OFDM-based communication systems can cover distances on the order of kilometers and coordinate many co-existing devices, such as in 4G and the new 5G standard. On the downside, symbol length for OFDM-encoded data is generally longer than UWB pulses—an important reason to study the security of OFDM systems when used for ranging.

5.5.2 Physical-Layer Integrity Protection Schemes and Jamming

There exist many physical-layer schemes that aim to guarantee the integrity of transmitted data. They can roughly be divided into randomness extraction from the channel (key establishment), Multi-Input and Multi-Output (MIMO)-based approaches (orthogonal blinding, zero-forcing), friendly jamming and integrity codes [25].

The concept of friendly jamming is related to the countermeasures for OFDM-based ranging that we propose in this work. The idea behind phase randomization and orthogonal noise is similar to that of friendly jamming [34, 89] where an attacker can not separate the information-bearing message from a jamming signal transmitted by a friendly jammer. The concept of intentional signal interference can be used to establish confidentiality, message authentication or access control [64, 119]. Reactive jamming on the other hand tries to analyze and react to packets in the air [20] in order to annihilate/overwrite certain packets or prevent communication altogether. This is related to the problem statement of the ED/LC attack described in this chapter. In reactive jamming, it is crucial to detect a signal very early on, i.e., only based on parts of it, to have maximum impact when interfering with the remainder of the signal.

Probably most related to our work is the survey in [73] that compares different approaches to physical-layer security in OFDM. Most of the presented methods are concerned with confidentiality either on the data bits or on the symbol level. The main idea is to encrypt or obfuscate the signal and/or provide resiliency against interference [108]. The idea we present in this chapter is similar in the sense that an attacker should not be able to predict the transmitted signal. However, we propose secure ranging schemes that protect the communication on the symbol level, rather than entire messages. Furthermore, we are specifically concerned with the learning/listening time that an attacker requires until the remainder of the symbol can be predicted since this is the crucial factor that facilitates secure ToF ranging.

5.6 Discussion

In the following, we cover the main avenues that can be taken to secure OFDM signals against an ED/LC attacker and highlight a few additional OFDM features that are linked to physical-layer security.

5.6.1 Preventing late commit: ICI sensitivity

Late-commit detection is enabled by a receiver's ability to detect deviations from the expected signal shape.

Utilizing both signal-space dimensions

Utilizing both signal-space dimensions in frequency domain breaks up the symmetry of the time-domain signal around symbol center and is, therefore, a necessity against both early-detection and late-commit.

Denser Constellation

Late-commit attacks become less effective if information is modulated on a denser constellation grid. As a consequence, late-commit needs to start earlier, as inter-carrier-interference has more impact. The denser the constellation, the less dispersion can be tolerated for correct detection. In general, denser constellations and increased throughput come at a tension with robustness, a requirement which is especially important since cyclic prefix and channel compensation cannot be used for security reasons.

Error integration

A possible way to resolve this tension is to map the signal into a denser constellation at the receiver and then post-process the received bits in a way that approximates the L2 distance to the expected spectrum (and selecting an appropriate symbol-wide decision threshold). This way, significant deviations of only a few tones can be weighted accordingly, and late-commit strategies that optimize for low bit-error rate under a coarse constellation mapping lose their utility.

5.6.2 Secure time-domain signals over OFDM

Due to inherent drawbacks of OFDM for secure ranging, i.e., the countermeasure requiring additional power for a noise dimension, and significant shared entropy for phase correction, it might be of use to retrofit OFDM transceivers with time-domain modulation capability. One such proposal is the use of DFT-spread OFDM. Outside of Orthogonal Frequency-Division Multiple Access (OFDMA), this means precoding the IQ values as the spectral representation of a desired time-domain signal. This results in similar properties of the signal to any time-domain pulsed modulation.

A way to create secure time-domain signals without the need for additional DFT blocks is to transmit a single pulse per symbol, as in [112]. This can be achieved by using identical tones (of a certain polarity) and verifying correctness either by a time-domain technique or, alternatively, by evaluating the Hamming weight per symbol at the receiver (serving as an approximation for the polarity of a constrained pulse in time-domain). The drawback of this approach is its data rate, i.e., one symbol can only transmit one bit, and a long sequence of ranging symbols has to be exchanged.

5.6.3 Other aspects

Different mechanisms that are commonly used for enhancing the performance of OFDM systems can have a detrimental impact on physical-layer security.

Channel sensing and equalization

A secure ranging implementation based on OFDM cannot rely on channel sensing and equalization. Channel sensing can be manipulated by an attacker, which brings equalization under adversarial control. Fundamentally, channel compensation is about the compensation of time-dispersion, which causes delayed signal components to be included in the decoding.

Cyclic prefix

The cyclic prefix, commonly applied on OFDM symbols to achieve orthogonal equalization under a channel, should not be used in symbols used for secure ranging, as it provides an additional advantage to the early-detecting adversary. The rationale behind the cyclic prefix is to

prepend the trailing samples of the symbol at its beginning and, in turn, to circularize the Fourier matrix under a time-dispersive channel. This creates symbol redundancy which helps an early-detecting attacker.

PAPR reduction techniques

Orthogonal noise with a random phase shift is compatible with techniques for peak-to-average power reduction, as phase randomization is one of those techniques. Another method for PAPR reduction is to reduce the symbol set to low-PAPR symbols. With BPSK and QPSK, pruning the symbol set of high-PAPR symbols tends to remove symbols with very stringent late-commit constraints, which might add to the overall vulnerability of those configurations.

5.7 Conclusion

We highlighted the vulnerability of highly performant OFDM modulation schemes for ToF distance measurement against an ED/LC attacker operating on the physical layer. Existing proposals for secure ToF distance measurement developed for single-carrier modulation methods require time-domain focusing of bit-information (pulsing) and time-domain padding. This work identified another possible direction suited to OFDM systems, using all subcarriers in parallel with randomized constellations.

Chapter 6

Discussion and Outlook

6.1 Summary

Attacks on real-world ranging and, in particular, PKES systems have been an ongoing concern. In the absence of secure ToF ranging, even simple relay attacks cannot be prevented. Secure ranging aims to protect against physical-layer and other distance modification attacks by combining principles from the design of distance-bounding and authenticated-ranging protocols with modulation and signal processing techniques enabling precise ToF measurement. This reconciliation causes non-trivial challenges that have not been emphasized in the design of the protocols. In practice, ranging performance, security, and signal robustness can be at odds, especially if modulation techniques are used that were originally envisioned for other design goals than secure ranging. A common understanding of how the security of a ranging implementation can be quantified and strengthened against such attacks is therefore important. Meanwhile, the need for quantifiable security in secure ranging also came to the attention of relevant industry consortia and standardization groups. This work studies the main concerns faced by real-world applications of secure ranging under current and foreseeable industry developments.

In Chapter 3, we provided the definitions for the attacker and the security against physical-layer attacks. We proposed the notion of the Message Time of Arrival Code, a physical-layer primitive for secure ToA measurement. We introduced and evaluated a particular instantiation based on the measurement of the signal variance. In Chapter 4, we focused

on IR-UWB implementations that are currently used in the consumer market. This is important for proximity verification, e.g., mobile payment and access systems. We demonstrated and evaluated the first over-the-air distance reduction attack on a real-world IEEE 802.15.4z UWB ToF ranging system. In Chapter 5, we analyzed a physical layer optimized for throughput, in particular OFDM, assuming a robust symbol mapping. The main reason is that these modulations are also increasingly important for localization, whether in WiFi [14] or cellular [35].

6.2 Future Work

In the following, we highlight a few directions for future work.

6.2.1 More security-performance trade-offs

Performance-enhancing techniques that are taken for granted in typical communication systems, such as channel equalization, filtering techniques, and clock offset compensation, can potentially have detrimental security implications in a ranging system. These techniques aim to compensate for channel-induced temporal dispersion or mismatch, thereby weakening the (idealized) notion of the distance commitment. Which of these techniques can be employed without significant deterioration of the security level is an ongoing concern.

6.2.2 MAC design

There are unresolved questions concerned with how to scale a secure localization system efficiently to larger numbers of participating nodes. Since broadcast localization is vulnerable to relay attacks even if authenticated [76], measuring the RTT in a two-way exchange is important for security and, hence, the cornerstone of any distance bounding or authenticated ranging protocol. However, this also means that all localizing nodes need to transmit, which increases the occupation of the medium. Moreover, schemes that allow multiple nodes to reply to a single challenge result in longer reply times. This can be a potential vulnerability if the clocks of the initiator and responder are not perfectly synchronized. This shows that, especially in the multi-user setting, protocol and physical-layer aspects are, to an extent, coupled. Whereas a distance-bounding protocol imposes a challenging requirement of replying

as fast as possible, a multi-user protocol with large reply times and no dedicated synchronization protocol requires the clock to be *as stable as possible*. The resulting implications on security can only be quantified from a joint PHY and MAC perspective.

In addition, MAC control information needs to be integrity protected or is, otherwise, prone to message manipulation or overshadowing attacks that modify the per-user resource allocations and, hence, message timings, as has been demonstrated in the context of 4G [41].

6.2.3 Integration with positioning infrastructure

Node-centric secure positioning can be used to securely establish the relative position between individual nodes. Typically, global location is established by means of broadcast-based positioning that is vulnerable to relay attacks. Moreover, there are also important ongoing developments in the availability of communication infrastructure, potentially enabling infrastructure-based two-way ranging. With 5G, we see a trend towards denser cellular deployments. In addition, there are growing deployments of Low Earth Orbit (LEO) satellite networks providing global coverage [115]. The latter, in theory, provides an opportunity for globally available ground-satellite RTT-based positioning.

6.3 Closing Remarks

Secure proximity verification and distance measurement are rooted in the integrity of a physical measurement (ToA). This creates a need for an attacker model dedicated to ToA acquisition and validation and evaluation criteria that can yield a concrete security level given a particular ranging procedure. Higher-level security mechanisms, such as distance bounding protocols, can only be as secure as the underlying modulation and the signal acquisition and validation processes. Depending on the modulation, establishing an achievable security level can be challenging, especially if the modulation was not designed for secure ranging. Designing modulations with the purpose of secure ToA establishment in mind simplifies the establishment and validation of security claims, reduces ambiguities due to implementation differences between manufacturers, and facilitates inter-operable solutions with quantifiable guarantees.

Appendix A

Validating the Gaussian Variance Model

In the following, we motivate the Gaussian model for the distortion distribution put forward in Equation 3.2 and Equation 3.3.

A.1 Extrapolation vs. fully empirical results

In the following, we compare our extrapolated results from Section 3.5 to a fully empirical (i.e., Monte-Carlo) simulation. The probability of winning as a function of the performance level is shown in Figure A.1 for LoS conditions and Figure A.2 for NLoS conditions. Both results refer to a frame of 20 bits. For both scenarios, we see that the the attacker's advantage evolves almost identically. We see that the fully empirical results indicate a slightly wider MTAC region, which suggest our Gaussian model to be a conservative estimate.

A.2 Variance distribution vs. Gaussian

We provide quantile-quantile (QQ) plots that compare the empirical distributions against normal distributions. This allows to validate the model we use in Section 3.5 which serves extrapolate the empirical classification performance to small likelihoods. We provide those plots

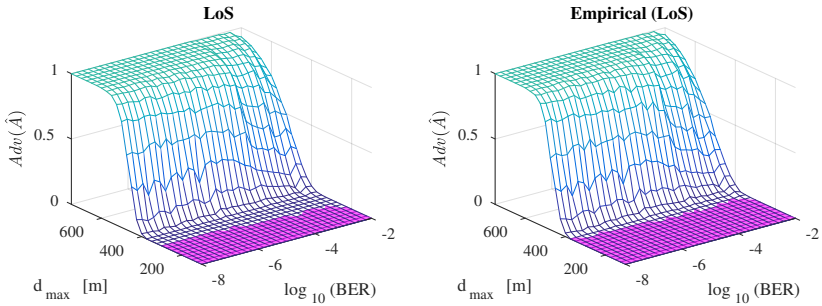

Figure A.1: Attacker's advantage as a function of the performance level for LoS conditions under a Gaussian extrapolation (left) and fully empirical simulation (right). Overall, the empirical result is very similar, in particular its MTAC region is not smaller than the one resulting from the Gaussian model.

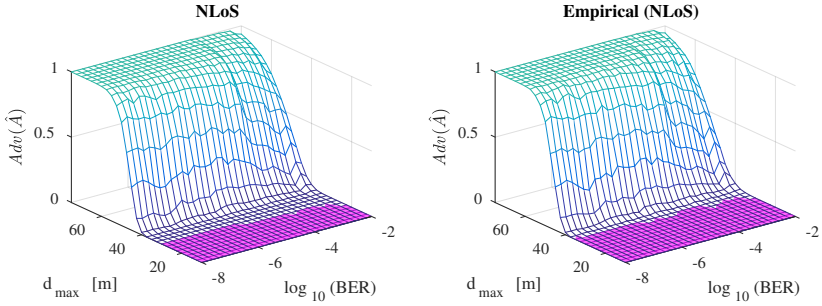

Figure A.2: Attacker's advantage as a function of the performance level for NLoS conditions under a Gaussian extrapolation (left) and fully empirical simulation (right). Overall, the empirical result is very similar, in particular its MTAC region is not smaller than the one resulting from the Gaussian model.

for a frame of 32 bits and a selection of communication distances, both for LoS and NLoS scenarios. Figure A.3 presents those results for the attacker's variance distribution. The relevant distance for the resulting MTAC region boundary is around 100 m for LoS and around 10 m for NLoS. This is the distance at which the distortion for the attacker is minimal, see Figure 3.9. There is a slight downwards bend of the empirical value for higher quantiles. This means, a slightly bit more than expected high-variance outliers compared to the Gaussian hypothesis. This is in line with our requirements, i.e., the normal estimate being conservative regarding distinguishability. The plots for those distances show that the empirical quantiles are well aligned with the straight diagonal. Figure A.4 presents those results for the attacker's distortion distribution. The relevant distance for the resulting MTAC boundary is around 200 m for LoS and around 20 m for NLoS, i.e., mid-range. The plots for those distances show that the empirical quantiles are well aligned with the straight line at those distances relevant for the MTAC region derived in Section 3.5.

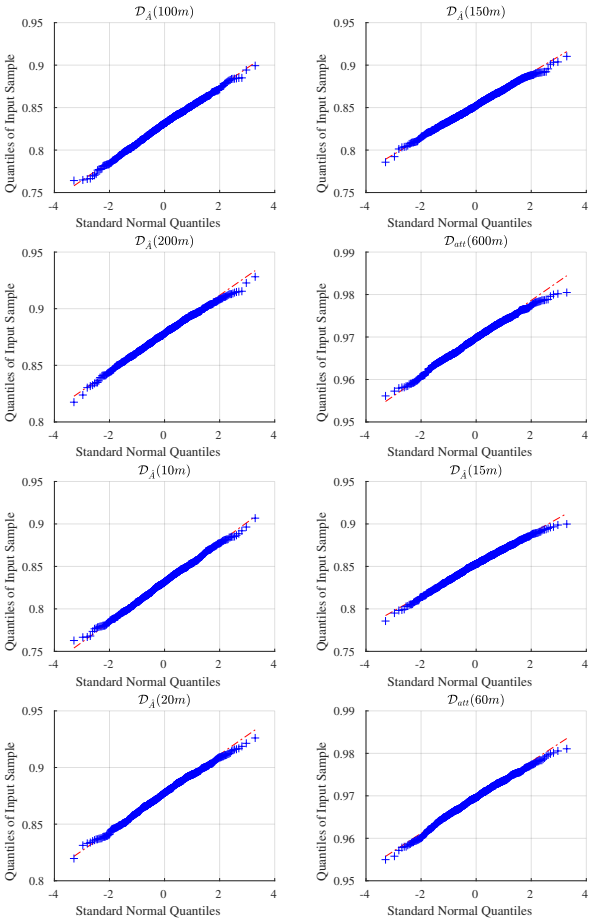

Figure A.3: QQ plots comparing the attacker empirical distortion distribution for LoS (top) and NLoS (bottom) conditions for a frame of 32 bits and different distances against a normal distribution. For validity of results w.r.t. the MTAC region boundary, the attack signal distortion at a distance of 100 m (LoS) and 10 m (NLoS) should be close to a Gaussian. Indeed, the QQ plots of the second column are close to the diagonal.

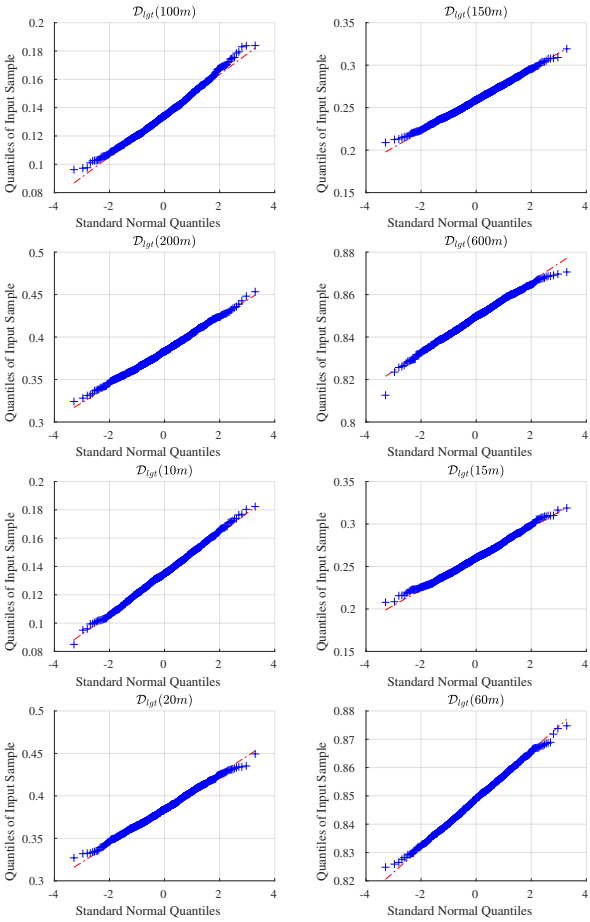

Figure A.4: QQ plots comparing the legitimate empirical distortion distribution for LoS (top) and NLoS (bottom) conditions for a frame of 32 bits and different distances against a normal distribution. For validity of results w.r.t. the MTAC region boundary, the attack signal distortion at a distance of 200 m (LoS) and 20 m (NLoS) should be close to a Gaussian. Indeed, the QQ plots of the third column are close to the diagonal.

Bibliography

[1] Ieee standard for information technology– local and metropolitan area networks– specific requirements– part 15.4: Wireless medium access control (mac) and physical layer (phy) specifications for low-rate wireless personal area networks (wpans): Amendment 1: Add alternate phys. *IEEE Std 802.15.4a-2007 (Amendment to IEEE Std 802.15.4-2006)*, pages 1–210, 2007.

[2] Ieee standard for local and metropolitan area networks– part 15.4: Low-rate wireless personal area networks (lr-wpans) amendment 2: Active radio frequency identification (rfid) system physical layer (phy). *IEEE Std 802.15.4f-2012 (Amendment to IEEE Std 802.15.4-2011)*, pages 1–72, 2012.

[3] Ieee standard for wireless medium access control (mac) and physical layer (phy) specifications for peer aware communications (pac). *IEEE Std 802.15.8-2017*, pages 1–322, 2018.

[4] 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.

[5] 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.

[6] 3db Access. VW Adopts UWB for Secure Car Access. `https://www.3db-access.com/article/18`. Accessed 2022-08-02.

[7] 3dB Access AG. 3db access - technology. `https://www.3db-access.com/technology`. Accessed 2022-08-02.

[8] Task Group 4z. IEEE 802.15 WPAN enhanced impulse radio. http://www.ieee802.org/15/pub/TG4z.html. Accessed 2022-08-08.

[9] Apple. Apple AirTag. https://www.apple.com/airtag/. Accessed 2022-08-02.

[10] Apple Inc. Find your keys, wallet, and more with AirTag. https://support.apple.com/en-us/HT210967. Accessed 2022-08-08.

[11] Apple Inc. Hand off audio to HomePod. https://support.apple.com/en-nz/guide/homepod/apdfb81a72e4/homepod. Accessed 2022-08-08.

[12] Apple Inc. Ultra Wideband security in iOS. https://support.apple.com/guide/security/ultra-wideband-security-sec1e6108efd/web. Accessed 2022-08-08.

[13] 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 Singelée, Aslan Tchamkerten, Rolando Trujillo-Rasua, and Serge Vaudenay. Security of distance-bounding: A survey. *ACM Comput. Surv.*, 51(5):94:1–94:33, 2019.

[14] Task Group az. IEEE 802.11 next generation positioning. http://www.ieee802.org/11/Reports/tgaz_update.htm. Accessed 2022-08-08.

[15] Task Group az. Versioning for phy security. https://mentor.ieee.org/802.11/dcn/20/11-20-1972-01-00az-versioning-of-phy-security.pptx. Accessed 2022-08-08.

[16] Brian Barrett. The biggest iPhone news is a tiny new chip inside it. https://www.wired.com/story/apple-u1-chip/. Accessed 2022-08-08.

[17] David A. Basin, Srdjan Capkun, Patrick Schaller, and Benedikt Schmidt. Formal reasoning about physical properties of security protocols. *ACM Trans. Inf. Syst. Secur.*, 14(2):16:1–16:28, 2011.

[18] BBC. Mercedes 'relay' box thieves caught on CCTV in Solihull. https://www.bbc.com/news/uk-england-birmingham-42132689. Accessed 2022-08-08.

[19] Samy Bengio, Gilles Brassard, Yvo Desmedt, Claude Goutier, and Jean-Jacques Quisquater. Secure implementations of identification systems. *J. Cryptol.*, 4(3):175–183, 1991.

[20] Daniel S. Berger, Francesco Gringoli, Nicolò Facchi, Ivan Martinovic, and Jens B. Schmitt. Gaining insight on friendly jamming in a real-world IEEE 802.11 network. In Gergely Ács, Andrew P. Martin, Ivan Martinovic, Claude Castelluccia, and Patrick Traynor, editors, *7th ACM Conference on Security & Privacy in Wireless and Mobile Networks, WiSec'14, Oxford, United Kingdom, July 23-25, 2014*, pages 105–116. ACM, 2014.

[21] 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 2022-08-02.

[22] Stefan Brands and David Chaum. Distance-bounding protocols (extended abstract). In Tor Helleseth, editor, *Advances in Cryptology - EUROCRYPT '93, Workshop on the Theory and Application of of Cryptographic Techniques, Lofthus, Norway, May 23-27, 1993, Proceedings*, volume 765 of *Lecture Notes in Computer Science*, pages 344–359. Springer, 1993.

[23] Agnès Brelurut, David Gérault, and Pascal Lafourcade. Survey of distance bounding protocols and threats. In Joaquín García-Alfaro, Evangelos Kranakis, and Guillaume Bonfante, editors, *Foundations and Practice of Security - 8th International Symposium, FPS 2015, Clermont-Ferrand, France, October 26-28, 2015, Revised Selected Papers*, volume 9482 of *Lecture Notes in Computer Science*, pages 29–49. Springer, 2015.

[24] Srdjan Capkun, Mario Cagalj, Ghassan Karame, and Nils Ole Tippenhauer. Integrity regions: Authentication through presence in wireless networks. *IEEE Trans. Mob. Comput.*, 9(11):1608–1621, 2010.

[25] Srdjan Capkun, Mario Cagalj, Ram Kumar Rengaswamy, Ilias Tsigkogiannis, Jean-Pierre Hubaux, and Mani B. Srivastava. Integrity codes: Message integrity protection and authentication

over insecure channels. *IEEE Trans. Dependable Secur. Comput.*, 5(4):208–223, 2008.

[26] Srdjan Capkun and Jean-Pierre Hubaux. Secure positioning of wireless devices with application to sensor networks. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies, 13-17 March 2005, Miami, FL, USA*, pages 1917–1928. IEEE, 2005.

[27] Srdjan Capkun and Jean-Pierre Hubaux. Secure positioning in wireless networks. *IEEE J. Sel. Areas Commun.*, 24(2):221–232, 2006.

[28] Car Connectivity Consortium. Digital Key Release. `https://carconnectivity.org/press-release/car-connectivity-consortium-publishes-digital-key-release-3-0/`. Accessed 2022-08-05.

[29] Jiska Classen and Alexander Heinrich. Wibbly Wobbly, Timey Wimey – What's Really Inside Apple's U1 Chip. Presentation at Black Hat USA 2021, August 2021.

[30] 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.

[31] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[32] Cas Cremers, Kasper Bonne Rasmussen, Benedikt Schmidt, and Srdjan Capkun. Distance hijacking attacks on distance bounding protocols. In *IEEE Symposium on Security and Privacy, SP 2012, 21-23 May 2012, San Francisco, California, USA*, pages 113–127. IEEE Computer Society, 2012.

[33] Decawave Ltd. DW3000 Family User Manual. `https://www.decawave.com/wp-content/uploads/2021/05/DW3000-User-Manual-1.pdf`.

[34] Bhaswati Deka, Ryan M. Gerdes, Ming Li, and Kevin Heaslip. Friendly jamming for secure localization in vehicular transportation. In Jing Tian, Jiwu Jing, and Mudhakar Srivatsa, editors, *International Conference on Security and Privacy in Communication Networks - 10th International ICST Conference, SecureComm 2014, Beijing, China, September 24-26, 2014, Revised Selected Papers, Part I*, volume 152 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 212–221. Springer, 2014.

[35] José A. del Peral-Rosado, José A. López-Salcedo, Gonzalo Seco-Granados, Francesca Zanier, and Massimo Crisci. Achievable localization accuracy of the positioning reference signal of 3gpp lte. In *2012 International Conference on Localization and GNSS*, pages 1–6, 2012.

[36] Yvo Desmedt, Claude Goutier, and Samy Bengio. Special uses and abuses of the fiat-shamir passport protocol. In Carl Pomerance, editor, *Advances in Cryptology - CRYPTO '87, A Conference on the Theory and Applications of Cryptographic Techniques, Santa Barbara, California, USA, August 16-20, 1987, Proceedings*, volume 293 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 1987.

[37] Danny Dolev and Andrew Chi-Chih Yao. On the security of public key protocols. *IEEE Trans. Inf. Theory*, 29(2):198–207, 1983.

[38] Saar Drimer and Steven J. Murdoch. Keep your enemies close: Distance bounding against smartcard relay attacks. In Niels Provos, editor, *Proceedings of the 16th USENIX Security Symposium, Boston, MA, USA, August 6-10, 2007*. USENIX Association, 2007.

[39] 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/amp/`. Accessed 2022-08-08.

[40] Embedded. Ultra-wideband (UWB) adoption picks up pace. `https://www.embedded.com/ultra-wideband-uwb-adoption-picks-up-pace/`. Accessed 2022-08-08.

[41] Simon Erni, Martin Kotuliak, Patrick Leu, Marc Roeschlin, and Srdjan Capkun. Adaptover: adaptive overshadowing attacks in

cellular networks. In *ACM MobiCom '22: The 28th Annual International Conference on Mobile Computing and Networking, Sydney, NSW, Australia, October 17 - 21, 2022*, pages 743–755. ACM, 2022.

[42] AE Fawal and J-Y Le Boudec. A robust signal-detection method for ultra-wideband networks with uncontrolled interference. *IEEE transactions on microwave theory and techniques*, 54(4):1769–1781, 2006.

[43] FiRa Consortium Inc. FiRa Consortium Website. `https://www.firaconsortium.org/`. Accessed 2022-08-02.

[44] FiRa Consortium Inc. Technical FAQ - How secure is UWB positioning? `https://www.firaconsortium.org/discover/technical-faq`. Accessed 2022-08-02.

[45] Manuel Flury, Ruben Merz, and Jean-Yves Le Boudec. Robust non-coherent timing acquisition in ieee 802.15. 4a ir-uwb networks. In *2009 IEEE 20th International Symposium on Personal, Indoor and Mobile Radio Communications*, pages 1642–1646. IEEE, 2009.

[46] 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.

[47] 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.

[48] Lishoy Francis, Gerhard P. Hancke, Keith Mayes, and Konstantinos Markantonakis. Practical relay attack on contactless transactions by using NFC mobile phones. In Nai-Wei Lo and Yingjiu Li, editors, *Radio Frequency Identification System Security - RFIDsec'12 Asia Workshop Proceedings, Taipei, Taiwan, November 8-9, 2012*, volume 8 of *Cryptology and Information Security Series*, pages 21–32. IOS Press, 2012.

[49] 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.

[50] GitHub user foldedtoad. Nordic nRF52-series + Decawave DWM3000 on Zephyr v2.5. `https://github.com/foldedtoad/dwm3000`. Accessed 2022-08-08.

[51] Stuart A. Golden and Steve S. Bateman. Sensor measurements for wi-fi location with emphasis on time-of-arrival ranging. *IEEE Trans. Mob. Comput.*, 6(10):1185–1198, 2007.

[52] Andy Greenberg. Radio attack lets hackers steal cars with just $20 worth of gear. `https://www.wired.com/2017/04/just-pair-11-radio-gadgets-can-steal-car/`. Accessed 2022-08-08.

[53] Azadeh Haghparast, Traian Abrudan, and Visa Koivunen. Ofdm ranging in multipath channels using time reversal method. In *2009 IEEE 10th Workshop on Signal Processing Advances in Wireless Communications*, pages 568–572, 2009.

[54] 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.

[55] Gerhard P. Hancke, Keith Mayes, and Konstantinos Markantonakis. Confidence in smart token proximity: Relay attacks revisited. *Comput. Secur.*, 28(7):615–627, 2009.

[56] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Packet leashes: A defense against wormhole attacks in wireless networks. In *Proceedings IEEE INFOCOM 2003, The 22nd Annual Joint Conference of the IEEE Computer and Communications Societies, San Franciso, CA, USA, March 30 - April 3, 2003*, pages 1976–1986. IEEE Computer Society, 2003.

[57] Todd E Humphreys, Brent M Ledvina, Mark L Psiaki, Brady W O'Hanlon, Paul M Kintner, et al. Assessing the spoofing threat: Development of a portable gps civilian spoofer. In *Proceedings of the 21st International Technical Meeting of the Satellite Division*

of The Institute of Navigation (ION GNSS 2008), pages 2314–2325, 2008.

[58] Apple Inc. Apple Profiles and Logs. `https://developer.apple.com/bug-reporting/profiles-and-logs/`. Accessed 2022-08-08.

[59] Apple Inc. Exposure notification. `https://developer.apple.com/exposure-notification/`. Accessed 2022-11-16.

[60] Apple Inc. Nearby Interaction. `https://developer.apple.com/documentation/nearbyinteraction`. Accessed 2022-08-02.

[61] Google Inc. Exposure notification. `https://www.google.com/covid19/exposurenotifications/`. Accessed 2022-11-16.

[62] 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 2022-08-02.

[63] Chong Hee Kim, Gildas Avoine, François Koeune, François-Xavier Standaert, and Olivier Pereira. The swiss-knife RFID distance bounding protocol. In Pil Joong Lee and Jung Hee Cheon, editors, *Information Security and Cryptology - ICISC 2008, 11th International Conference, Seoul, Korea, December 3-5, 2008, Revised Selected Papers*, volume 5461 of *Lecture Notes in Computer Science*, pages 98–115. Springer, 2008.

[64] Jongyeop Kim and Jihwan P. Choi. Cancellation-based friendly jamming for physical layer security. In *2016 IEEE Global Communications Conference, GLOBECOM 2016, Washington, DC, USA, December 4-8, 2016*, pages 1–6. IEEE, 2016.

[65] KJ Kim. Samsung expects uwb to be one of the next big wireless technologies. `https://news.samsung.com/global/samsung-expects-uwb-to-be-one-of-the-next-big-wireless-technologies`. Accessed 2022-08-02.

[66] 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.

[67] Manikanta Kotaru, Kiran Raj Joshi, Dinesh Bharadia, and Sachin Katti. Spotfi: Decimeter level localization using wifi. In Steve Uhlig, Olaf Maennel, Brad Karp, and Jitendra Padhye, editors, *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM 2015, London, United Kingdom, August 17-21, 2015*, pages 269–282. ACM, 2015.

[68] Tsit Yuen Lam and Ka Hin Leung. On vanishing sums of roots of unity. *Journal of algebra*, 224(1):91–109, 2000.

[69] 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.

[70] Premala H Madhani, Penina Axelrad, Kent Krumvieda, and John Thomas. Application of successive interference cancellation to the gps pseudolite near-far problem. *IEEE Transactions on Aerospace and Electronic Systems*, 39(2):481–488, 2003.

[71] Catherine Meadows and Paul Syverson. Range authentication protocols for localization.

[72] Catherine Meadows, Paul Syverson, and LiWu Chang. Towards more efficient distance bounding protocols for use in sensor networks. In *2006 Securecomm and Workshops*, pages 1–5, 2006.

[73] Reem Melki, Hassan N. Noura, Mohammad M. Mansour, and Ali Chehab. A survey on OFDM physical layer security. *Phys. Commun.*, 32:1–30, 2019.

[74] 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 2022-08-08.

[75] Andreas F Molisch. *Wireless communications*, volume 34. John Wiley & Sons, 2012.

[76] Maryam Motallebighomi, Harshad Sathaye, Mridula Singh, and Aanjhan Ranganathan. Cryptography is not enough: Relay attacks on authenticated GNSS signals. *CoRR*, abs/2204.11641, 2022.

[77] Nordic Semiconductor. nRF52 DK - Nordic Semiconductor. https://www.nordicsemi.com/Products/Development-hardware/nRF52-DK. Accessed 2022-08-08.

[78] NXP Semiconductors. Secure Ultra-Wideband (UWB) Positioning and Ranging Optimized for IoT Use Cases. https://www.nxp.com/products/wireless/secure-ultra-wideband-uwb/trimension-sr040-secure-uwb-solution-for-iot-tags:SR040?tab=Documentation_Tab. Accessed 2022-08-02.

[79] 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.

[80] Malcolm 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 2022-08-02.

[81] Panagiotis Papadimitratos and Aleksandar Jovanovic. Gnss-based positioning: Attacks and countermeasures. In *Military Communications Conference, 2008. MILCOM 2008. IEEE*, pages 1–7. IEEE, 2008.

[82] 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.

[83] 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.

[84] 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. IEEE, 2010.

[85] Qorvo Inc. DW3110 - Qorvo. https://www.qorvo.com/products/p/DW3110. Accessed 2022-08-02.

[86] Qorvo Inc. DWM3000EVB - Qorvo. https://www.qorvo.com/products/p/DWM3000EVB. Accessed 2022-08-02.

[87] Qorvo Inc. DWM3001CDK - Qorvo. https://www.qorvo.com/products/p/DWM3001CDK. Accessed 2022-08-02.

[88] Qorvo Inc. Qorvo Completes Acquisition of Decawave. https://www.qorvo.com/newsroom/news/2020/qorvo-completes-acquisition-of-decawave. Accessed 2022-08-02.

[89] Hanif Rahbari and Marwan Krunz. Friendly cryptojam: a mechanism for securing physical-layer attributes. In Gergely Ács, Andrew P. Martin, Ivan Martinovic, Claude Castelluccia, and Patrick Traynor, editors, *7th ACM Conference on Security & Privacy in Wireless and Mobile Networks, WiSec'14, Oxford, United Kingdom, July 23-25, 2014*, pages 129–140. ACM, 2014.

[90] Aanjhan Ranganathan and Srdjan Capkun. Are we really close? verifying proximity in wireless systems. *IEEE Secur. Priv.*, 15(3):52–58, 2017.

[91] Aanjhan Ranganathan, Boris Danev, and Srdjan Capkun. Proximity verification for contactless access control and authentication systems. In *Proceedings of the 31st Annual Computer Security Applications Conference, Los Angeles, CA, USA, December 7-11, 2015*, pages 271–280. ACM, 2015.

[92] 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.

[93] Aanjhan Ranganathan, Nils Ole Tippenhauer, Boris Skoric, Dave Singelée, and Srdjan Capkun. Design and implementation of a terrorist fraud resilient distance bounding system. In Sara Foresti, Moti Yung, and Fabio Martinelli, editors, *Computer Security -*

ESORICS 2012 - 17th European Symposium on Research in Computer Security, Pisa, Italy, September 10-12, 2012. Proceedings, volume 7459 of *Lecture Notes in Computer Science*, pages 415–432. Springer, 2012.

[94] 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.

[95] Kasper Bonne Rasmussen, Srdjan Capkun, and Mario Cagalj. Secnav: secure broadcast localization and time synchronization in wireless networks. In Evangelos Kranakis, Jennifer C. Hou, and Ram Ramanathan, editors, *Proceedings of the 13th Annual International Conference on Mobile Computing and Networking, MOBICOM 2007, Montréal, Québec, Canada, September 9-14, 2007*, pages 310–313. ACM, 2007.

[96] 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 2022-08-08.

[97] Ettus Research. Usrp product selector. `https://www.ettus.com/products/usrp-product-selector/`. Accessed 2022-08-08.

[98] Michael Roland, Josef Langer, and Josef Scharinger. Applying relay attacks to google wallet. In *2013 5th International Workshop on Near Field Communication (NFC)*, pages 1–6. IEEE, 2013.

[99] Samsung Electronics Co., Ltd. . How do I use Point to Share? `https://www.samsung.com/global/galaxy/what-is/uwb/`. Accessed 2022-08-02.

[100] 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 2022-08-02.

[101] 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-experie`

nce-galaxy-users-can-now-use-secure-digital-key-with-the-genesis-gv60. Accessed 2022-08-02.

[102] Naveen Sastry, Umesh Shankar, and David A. Wagner. Secure verification of location claims. In W. Douglas Maughan and Adrian Perrig, editors, *Proceedings of the 2003 ACM Workshop on Wireless Security, San Diego, CA, USA, September 19, 2003*, pages 1–10. ACM, 2003.

[103] Sahar Sedighpour, Srdjan Capkun, Saurabh Ganeriwal, and Mani B. Srivastava. Distance enlargement and reduction attacks on ultrasound ranging. In Jason Redi, Hari Balakrishnan, and Feng Zhao, editors, *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems, SenSys 2005, San Diego, California, USA, November 2-4, 2005*, page 312. ACM, 2005.

[104] 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-uwb-fine-ranging-capabilities:NW-VOLKSWAGEN-SHOWCASES-UWB. Accessed 2022-08-08.

[105] 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 2022-08-02.

[106] 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-uwb-deployed-in-samsung-galaxy-note20-ultra-bringing-the-first-uwb-enabled-android-device-to-market:NW-SECURE-UWB-SAMSUNG-GALAXY. Accessed 2022-08-02.

[107] NXP Semiconductors. Secure Ultra-Wideband Developer Resources. https://www.nxp.com/products/wireless/secure-ultra-wideband-uwb:UWB-TRIMENSION. Accessed 2022-08-02.

[108] Chowdhury Shahriar, Matthew J. La Pan, Marc Lichtman, T. Charles Clancy, Robert W. McGwier, Ravi Tandon, Shabnam

Sodagari, and Jeffrey H. Reed. Phy-layer resiliency in OFDM communications: A tutorial. *IEEE Commun. Surv. Tutorials*, 17(1):292–314, 2015.

[109] Dave Singelée and Bart Preneel. Location verification using secure distance bounding protocols. In *IEEE 2nd International Conference on Mobile Adhoc and Sensor Systems, MASS 2005, November 7-10, 2005, The City Center Hotel, Washington, USA*. IEEE Computer Society, 2005.

[110] Mridula Singh, Patrick Leu, AbdelRahman Abdou, and Srdjan Capkun. UWB-ED: distance enlargement attack detection in ultra-wideband. In Nadia Heninger and Patrick Traynor, editors, *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, pages 73–88. USENIX Association, 2019.

[111] 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.

[112] Mridula Singh, Marc Roeschlin, Aanjhan Ranganathan, and Srdjan Capkun. V-range: Enabling secure ranging in 5g wireless networks. In *NDSS*, 2022.

[113] Mridula Singh, Marc Roeschlin, Ezzat Zalzala, 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.

[114] 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 2022-08-08.

[115] Starlink. Starlink availability. `https://www.starlink.com/map`. Accessed 2022-07-26.

[116] Sultan Khan. Technical Advisory – Tesla BLE Phone-as-a-Key Passive Entry Vulnerable to Relay Attacks. `https://research.nccgroup.com/2022/05/15/technical-advisory-tesla-ble-phone-as-a-key-passive-entry-vulnerable-to-relay-attacks/`. Accessed 2022-08-08.

[117] Nils O Tippenhauer. *Physical-Layer Security Aspects of Wireless Localization*. PhD thesis, ETH Zurich, 2012.

[118] 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.

[119] Nils Ole Tippenhauer, Luka Malisa, Aanjhan Ranganathan, and Srdjan Capkun. On limitations of friendly jamming for confidentiality. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 160–173. IEEE Computer Society, 2013.

[120] Nils Ole Tippenhauer, Christina Pöpper, Kasper Bonne Rasmussen, and Srdjan Capkun. On the requirements for successful GPS spoofing attacks. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS 2011, Chicago, Illinois, USA, October 17-21, 2011*, pages 75–86. ACM, 2011.

[121] Carmela Troncoso, Mathias Payer, Jean-Pierre Hubaux, Marcel Salathé, James R. Larus, Wouter Lueks, Theresa Stadler, Apostolos Pyrgelis, Daniele Antonioli, Ludovic Barman, Sylvain Chatel, Kenneth G. Paterson, Srdjan Capkun, David A. Basin, Jan Beutel, Dennis Jackson, Marc Roeschlin, Patrick Leu, Bart Preneel, Nigel P. Smart, Aysajan Abidin, Seda Gurses, Michael Veale, Cas Cremers, Michael Backes, Nils Ole Tippenhauer, Reuben Binns, Ciro Cattuto, Alain Barrat, Dario Fiore, Manuel Barbosa, Rui Oliveira, and José Pereira. Decentralized privacy-preserving proximity tracing. *IEEE Data Eng. Bull.*, 43(2):36–66, 2020.

[122] Deepak Vasisht, Swarun Kumar, and Dina Katabi. Decimeter-level localization with a single wifi access point. In Katerina J. Argyraki and Rebecca Isaacs, editors, *13th USENIX Symposium on*

Networked Systems Design and Implementation, NSDI 2016, Santa Clara, CA, USA, March 16-18, 2016, pages 165–178. USENIX Association, 2016.

[123] Roel Verdult, Flavio D. Garcia, and Josep Balasch. Gone in 360 seconds: Hijacking with hitag2. In Tadayoshi Kohno, editor, *Proceedings of the 21th USENIX Security Symposium, Bellevue, WA, USA, August 8-10, 2012*, pages 237–252. USENIX Association, 2012.

[124] Tien Dang Vo-Huu, Triet Dang Vo-Huu, and Guevara Noubir. Spectrum-flexible secure broadcast ranging. 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 300–310. ACM, 2021.

[125] Volkswagen. Realtime safety with uwb. `https://www.volkswagen-newsroom.com/en/stories/realtime-safety-with-uwb-5438`. Accessed 2022-08-08.

[126] West Mercia Police. Car theft prevention advice. `https://www.westmercia.police.uk/news/west-mercia/news/2022/january/car-theft-prevention-advice/`. Accessed 2022-08-08.

[127] Liuqing Yang and Georgios B Giannakis. Ultra-wideband communications: an idea whose time has come. *IEEE signal processing magazine*, 21(6):26–54, 2004.

[128] Zebra Technologies Corp. Ultra wideband (uwb) technology. `https://www.zebra.com/us/en/products/location-technologies/ultra-wideband.html`. Accessed 2022-08-08.