Secure Neighbor Discovery and Ranging in Wireless Networks

THÈSE N^{O} 5131 (2011)

PRÉSENTÉE LE 14 OCTOBRE 2011

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
LABORATOIRE POUR LES COMMUNICATIONS INFORMATIQUES ET LEURS APPLICATIONS 1
PROGRAMME DOCTORAL EN INFORMATIQUE, COMMUNICATIONS ET INFORMATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Marcin POTURALSKI

acceptée sur proposition du jury:

Prof. A. Lenstra, président du jury
Prof. J.-P. Hubaux, Prof. P. Papadimitratos, directeurs de thèse
Prof. D. Basin, rapporteur
Prof. S. Capkun, rapporteur
Dr C. Dehollain, rapporteur

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2011

Abstract

This thesis addresses the security of two fundamental elements of wireless networking: neighbor discovery and ranging. Neighbor discovery consists in discovering devices available for direct communication or in physical proximity. Ranging, or distance bounding, consists in measuring the distance between devices, or providing an upper bound on this distance. Both elements serve as building blocks for a variety of services and applications, notably routing, physical access control, tracking and localization. However, the open nature of wireless networks makes it easy to abuse neighbor discovery and ranging, and thereby compromise overlying services and applications. To prevent this, numerous works proposed protocols that secure these building blocks. But two aspects crucial for the security of such protocols have received relatively little attention: formal verification and attacks on the physical-communication-layer. They are precisely the focus of this thesis.

In the first part of the thesis, we contribute a formal analysis of secure communication neighbor discovery protocols. We build a formal model that captures salient characteristics of wireless systems such as node location, message propagation time and link variability, and we provide a specification of secure communication neighbor discovery. Then, we derive an impossibility result for a general class of protocols we term "time-based protocols", stating that no such protocol can provide secure communication neighbor discovery. We also identify the conditions under which the impossibility result is lifted. We then prove that specific protocols in the time-based class (under additional conditions) and specific protocols in a class we term "time- and location-based protocols," satisfy the neighbor discovery specification. We reinforce these results by mechanizing the model and the proofs in the theorem prover Isabelle.

In the second part of the thesis, we explore physical-communication-layer attacks that can seemingly decrease the message arrival time without modifying its content. Thus, they can circumvent time-based neighbor discovery protocols and distance bounding protocols. (Indeed, they violate the assumptions necessary to prove protocol correctness in the first part of the thesis.) We focus on Impulse Radio Ultra-Wideband, a physical layer technology particularly well suited for implementing distance bounding, thanks to its ability to perform accurate indoor ranging. First, we adapt physical layer attacks reported in prior work to IEEE 802.15.4a, the de facto standard for Impulse Radio, and evaluate their performance. We show that an adversary can achieve a distance-decrease of up to hundreds of meters with an arbitrarily high probability of success, with only a minor cost in terms of transmission power (few dB). Next, we demonstrate a new attack vector that disrupts time-of-arrival estimation algorithms, in particular those designed to be precise. The distance-decrease achievable by this attack vector is in the order of the channel spread (order of 10 meters in indoor environments). This attack vector can be used in previously reported physical layer attacks, but it also creates a new type of external attack based on malicious interference. We demonstrate that variants of the malicious interference attack are much easier to mount than the previously reported external attack. We also provide design guidelines for modulation schemes and devise receiver algorithms that mitigate physical layer attacks. These countermeasures allow the system designer to trade off security, ranging precision and cost in terms of transmission power and packet length.

Keywords neighbor discovery, ranging, distance bounding, relay attack, formal analysis, physical layer attacks

Riassunto

La presente tesi affronta il problema della sicurezza di due elementi fondamentali delle reti wireless: neighbor discovery e ranging. Neighbor discovery consiste nello scoprire apparecchi disponibili per una comunicazione diretta oppure in prossimità fisica. Ranging, oppure distance bounding, consiste nel misurare la distanza fra apparecchi oppure nel fornire un limite superiore a questa distanza. Entrambi gli elementi costituiscono due componenti essenziali per diversi servizi e applicazioni, quali routing, controllo di accesso fisico, tracking e localizzazione. Tuttavia, l'accessibilità delle reti wireless rende più facile l'abuso di neighbor discovery e ranging, compromettendo i servizi di rete e applicazioni overlay. Per prevenire ciò, numerosi scritti hanno proposto protocolli che mettano in sicurezza questi due elementi fondamentali. Malgrado ciò, due aspetti cruciali per la sicurezza di tali protocolli hanno ricevuto relativamente poca attenzione: verifica formale e attacchi sul layer di comunicazione fisico. Essi sono infatti il fulcro della presente tesi.

Nella prima parte di questa tesi, il nostro contributo é un'analisi formale di protocolli per la "secure communication neighbor discovery". Il modello formale che vi costruiamo comprende le caratteristiche salienti dei sistemi wireless, come la posizione dei nodi, i tempi di propagazione dei messaggi e la variabilità del link wireless. Inoltre, provvediamo a una specificazione di "secure neighbor discovery". In seguito, deriviamo un risultato d'impossibilità per una classe generale di protocolli che nominiamo "time-based protocols", asserendo che non vi siano tali protocolli che forniscano un secure neighbor discovery. Per tale risultato, offriamo inoltre le condizioni in cui esso si manifesta. Dopodiché, proviamo che protocolli specifici della classe "time-based", sotto condizioni aggiuntive, e protocolli specifici della classe che denominiamo "time- and location-based" soddisfano le specificazioni di neighbor discovery. Rinforziamo questo risultato tramite una meccanizzazione del modello e prove nel theorem prover Isabelle.

Nella seconda parte della tesi, esploriamo attacchi sul layer fisico che possano apparentemente diminuire il tempo di arrivo di un messaggio senza modificarne il contenuto. Di conseguenza, essi possono eludere i protocolli di neighbor discovery basati sul tempo e i protocolli di distance bounding. (In effetti, essi violano le supposizioni necessarie per provare la correttezza dei protocolli nella prima parte della tesi.) Ci focalizziamo sull'Impulse Radio Ultra-Wideband, una tecnologia di layer fisico particolarmente adatta per l'implementazione di distance bounding grazie alla sua abilità di effettuare un'accurato ranging indoor. In primo luogo, adattiamo gli attacchi riportati negli scritti precedenti a IEEE 802.15.4a, lo standard de facto per Impulse Radio, e valutiamo la loro performance. Mostriamo come un avversario possa avere successo in un attacco di riduzione della distanza fino a centinaia di metri, con una probabilità di successo arbitrariamente elevata ed un costo minimo in termini di potenza di trasmissione (qualche dB). Di seguito, dimostriamo un nuovo vettore di attacco che scombussola gli algoritmi di estimazione basati sul tempo di arrivo, in particolare quelli sviluppati per essere precisi. La riduzione della distanza che può essere ottenuta da questo vettore di attacco é nell'ordine del channel spread (nell'ordine di 10 metri in ambiente indoor). Tale vettore di attacco può essere utilizzato in attacchi riportati in scritti precedenti e, inoltre, crea un nuovo tipo di attacco basato sull'interferenza maligna. Dimostriamo che varianti dell'attacco basato sull'interferenza maligna sono di gran lunga più semplici da portare a termine rispetto a attacchi esterni riportati negli scritti precedenti. Inoltre, forniamo linee guida per il design di schemi di modulazione e concepiamo algoritmi per il ricevente che mitighano attacchi al layer fisico. Tali contromisure permettono al designer di sistema di arrivare ad un compromesso fra sicurezza, precisione del ranging e costi in termini di potenza di trasmissione e lunghezza di pacchetto.

Parole Chiave neighbor discovery, ranging, distance bounding, relay attack, analisi formale, attacchi sul layer fisico

Acknowledgments

First and foremost, I am grateful to my advisors: Prof. Jean-Pierre Hubaux and Prof. Panagiotis Papadimitratos for guiding and supporting me throughout the PhD process. I learned a great deal from both of them, and not only about doing research. I would also like to express my appreciation to Jean-Pierre, for creating the stimulating and friendly environment of LCA1.

I would like to thank my thesis committee members: Prof. David Basin, Prof. Srdjan Čapkun, Dr Catherine Dehollain and Prof. Arjen Lenstra for the time and effort they put into reviewing this dissertation.

I am in debt to all my colleagues who contributed to this thesis in many ways: discussing and forging ideas, reviewing my work, and co-authoring papers. In particular, I would like to thank Dr Parisa Haghani, Reza Shokri and Dr Wojciech Galuba for the engaging and fruitful collaborations. Special thanks goes to Dr Manuel Flury, who apprenticed me to Impulse-Radio and who taught me a great deal about Swiss efficiency. Without Manuel, Part II of my thesis would not exist. I also want to express my appreciation to all the motivated students I had the pleasure to supervise. Furthermore, I would like to thank our secretaries and system administrators for making the working environment a well-oiled machine, and to Holly for her struggle to polish my English.

My PhD experience would not have been the same without the many wonderful people I have met along the way: Adriana, Gleb, Maxim, Šarūnas, Ivana, Wojtek, Marta, Philipp, Julien, Jacques, Mark, Naouel, Panos, Reza, Olga, Simas, Parisa, Dan, Irina, Valka, Sam, Audrius, Hossein, Giorgio, Igor, Mathias, Murtuza, Zotlan, Denisa, Iuli, Tanja, Oana, Razvan, George, Berker, and many more. I am grateful for the many great experiences we shared; from countless discussions about research, life and philosophy (and becoming a doctor of thereof), through exploring Switzerland and the rest of the world on foot, wheels, skies and various water-vessels, to creating the Los-Ann PhD movie industry. Most notably, I would like to thank Philippe for encouraging me to pursue a PhD at EPFL (and for teaching me where to find the best fondue); Maciek, Michał and Kasia for making me feel at home upon my arrival in Switzerland; and my officemate Nevena for sharing the daily struggles of the PhD life and bringing cheer into BC200.

Finally, I want to thank my parents, Mirosława and Zenon, my brother Andrzej, and all my family for their unconditional support.

Contents

Introduction

> It is a sunny Saturday afternoon, and Jean-Luc is going shopping. As he approaches his car, the car key in his pocket is automatically detected, and the car door is unlocked.[1] "Computer, plot a course to the local supermarket" announces Jean-Luc as he enters the car. While he is fastening his seatbelt, the car's navigation unit acquires the GPS signal, estimates the location, fetches the latest traffic information and computes an optimal route to the store.[2] "Make it so" confirms Jean-Luc. On route, the navigation unit records the trajectory, based on which Jean-Luc's account is automatically credited with appropriate road tolls.[3]
>
> When Jean-Luc enters the store, his mobile phone estimates its location based on WiFi access points in the neighborhood.[4] This allows the phone to check-in online into the store[5]. As it is the 10th time Jean-Luc has visited the store this month, he obtains a 5% fidelity discount. A few minutes later, his phone informs him that his friend Beverly (or rather her phone) was also detected in the store.[6] Pleased with this fortunate coincidence, Jean-Luc quickly finds Beverly to arrange a few details about the barbecue that evening. On his way out, he passes by the cash register, and waves his phone in front of the cash register to pay for the groceries;[7] his fidelity discount is automatically applied.

This short scene, although not entirely grounded in reality at the time this thesis is written, can hardly be called science-fiction. It illustrates how prominent a role wireless communication plays and will play in daily activities of people and businesses. The scene provides just a few examples, but they are sufficient to illustrate the two factors that motivated us for this thesis. First, many of these applications and services are security-sensitive, meaning that there is a clear incentive (notably monetary) for an *adversary* to meddle with them. Second, these services rely on two fundamental building blocks: *discovering neighbors*, i.e., wireless devices available for direct communication, or in physical proximity; and *ranging*, i.e., computing the distance between wireless devices. The bad news is that the open nature of wireless communications makes it quite easy for an adversary to disrupt neighbor discovery and ranging. And through disrupting these building blocks, the adversary can abuse overlaying applications and services.

[1] Passive Keyless Entry and Start systems [162].

[2] TomTom [10], Google Navigation [7]

[3] The Toll Collect system in Germany, currently for trucks [9].

[4] Skyhook [8].

[5] FourSquare [2], Facebook Places [3].

[6] Nokia Instant Community based on ad-hoc WiFi communications [11]

[7] News and rumors about Google Android [1], Apple IPhone and IPad [6], or Microsoft Windows Phone [5].

Not surprisingly, the research community has proposed a number of protocols for *secure neighbor discovery* (ND) and *secure ranging*, traditionally termed *distance bounding* (DB). However, in most cases the security of these protocols is only argued in an informal fashion. The history of security protocols has shown that an informal analysis is often insufficient, and many design flaws were discovered after protocol publication. This motivated us to pursue the investigation that comprises the first part of this thesis: A formal analysis of secure neighbor discovery. Another aspect that has received relatively little attention is security at the physical-communication-layer (PHY). Distance bounding protocols, as well as many neighbor discovery protocols, are cryptographic protocols that incorporate precise message time-of-flight measurements. It was pointed out in [35] that an adversary can bypass such protocols by mounting physical-communication-layer attacks that decrease the measured distance while preserving the payload of messages. The second part of the thesis is devoted to the study of such PHY attacks, as well as countermeasures.

Contributions

In this thesis, we address the security of two fundamental elements of wireless networking: neighbor discovery and ranging. We focus on two aspects crucial for the security of these protocols: formal analysis and security on the physical-communication-layer. We make the following main contributions:

1. We build a formal model that captures the characteristics of wireless systems crucial for the security on neighbor discovery and ranging, such as message propagation time, device location, and the neighbor relation. This framework allows us to formally reason about two general classes of protocols: "time-based protocols" and "time-and-location-based protocols". We prove an impossibility result for time-based protocols: Essentially, secure discovery of neighbors available for direct communication is not possible if the adversary can relay messages with a delay below a threshold determined by the communication range of the honest devices. We also prove the security of concrete time-based and time-and-location-based neighbor discovery protocols under additional realistic assumptions. We thus demonstrate that time-based protocols can provide secure communication ND if the adversary cannot relay messages with a delay below the threshold. In contrast, time-and-location-based protocols, a construct that we introduce, can provide secure communication ND as long as the adversary introduces a strictly positive relaying delay.

2. We investigate the vulnerability of IEEE 802.15.4a to PHY attacks. IEEE 802.15.4a is a standard for Impulse-Radio Ultra-wideband, a technology particularly well suited for distance bounding, thanks to its ability to perform precise ranging even in indoor environments. We adapt physical layer attacks reported in prior work to IEEE 802.15.4a and evaluate their performance with PHY simulations. We demonstrate that an adversary can achieve a distance-decrease of hundreds of meters with an arbitrarily high probability of success, with only a minor cost in terms of transmission power (few dB). This is in part due to certain features of the standard that are designed to improve performance. We propose simple modifications to IEEE 802.15.4a; they remove the vulnerability while retaining the performance benefits introduced by these features. Combined with simple countermeasures implemented at the receiver side, this limits the distance-decrease that

the adversary can achieve to a value in the order of the channel spread (in the order of 10 meters in indoor environments) at a minor cost in term of packet length.

3. We reveal a new attack vector against IR-UWB ranging that is directed at message time-of-arrival (ToA) estimation. This attack vector can decrease the measured distance by values in the order of the channel spread. Hence, it is a threat for precise distance bounding. It can be be used in previously reported PHY attacks, but it also enables a new type of PHY attack, based on malicious interference. This attack is much easier to mount than alternative external attacks, although it is less precise. We show with simulations and experiments that the attack is effective against a number of receivers. We also identify countermeasures that allow wireless devices to perform distance bounding that is both secure and precise, although at a cost of significantly increasing the packet length.

Thesis Outline

Part I is devoted to formal analysis and focuses on neighbor discovery. In Chapter 1 we provide a general introduction of secure neighbor discovery. In Chapter 2 we define the formal framework and present the results obtained within this framework. Part II is devoted to physical-communication-layer security and focuses on distance bounding. In Chapter 3 we provide an general introduction to distance bounding, physical-communication-layer attacks, and Impulse-Radio Ultra-wideband. In Chapter 4 we cover the security evaluation of IEEE 802.15.4a, whereas in Chapter 5 we address attacks on time-of-arrival estimation and countermeasures.

Publications

Chapter 1 is based on [113]. Chapter 2 combines the results from [130] and [131]. Chapter 3 and Chapter 4 are based on [59] and [128]. Chapter 5 is an extended version of [127]; the Chapter content is currently under submission to IEEE Transactions on Wireless Communications.

Part I

Formal Analysis of Secure Neighbor Discovery

Chapter 1

Introduction to Neighbor Discovery

A major benefit of wireless communications is flexibility, notably in terms of the mobility of devices and their users. Indeed, a device equipped with a wireless interface can start communicating with another device, an access point, or a base station, almost instantly, without setting up a cable connection. As a consequence, wireless connections are frequently established, making *discovering neighbors* an indispensable element of wireless networks. However, due to the open nature of wireless networks, *neighbor discovery* (ND) is easy to abuse: An adversary can convince a device into falsely believing that another device is its neighbor. The adversary can then use these false neighbor links to disrupt the applications and services that use ND as a building block.

Securing ND has therefore attracted considerable attention from the research community, and a number of secure ND protocols have been proposed. However, not much attention has been devoted to the meticulous analysis of ND, and the subtleties of ND are often overlooked. This motivated us to develop a formal framework for analysis of ND (Chapter 2), which is the main contribution of Part I of this thesis.

Chapter Outline In this Chapter, we provide a general introduction to secure ND. In Section 1.1 we clarify the definition of *neighbor*, emphasizing the difference between *communication* neighbor and *physical* neighbor; and we provide a basic taxonomy of ND protocols. In Section 1.2 we give representative examples of using ND as a building block. In Section 1.3 we explain how easy it is to attack ND with a *relay* attack, provide a classification of relay attacks, and explain how, by abusing ND, an adversary can disrupt overlaying applications. Finally, in Section 1.4 we provide an overview of secure ND protocols proposed in the literature.

1.1 Neighborhood and Neighbor Discovery

Devices in existing and upcoming wireless networks are diverse in their characteristics and functionality. To introduce the problem at hand, we abstract away numerous details and consider system entities to be generic *nodes*. Each node has a unique identity, a processing unit, and a wireless transceiver.

Nodes communicate over the wireless medium, based on the state of the medium and the capabilities of their transceivers. We do not dwell on the transceiver characteristics, unless needed. In general, beyond technical characteristics of the transceivers (notably their antennas), parameters and factors that determine the ability to communicate include: (i) the power of the transmitted signal, (ii) the distance between the transmitting and (intended) receiving nodes, (iii) the ratio of the received power over that of noise and interfering signals, and (iv) and propagation environment.

1.1.1 Communication and Physical Neighborhood

We define the *communication neighborhood* of a node U as the set $\mathcal{N}_{\mathcal{C}}(U)$ of nodes able to send information directly to U. In other words, a node V is a *communication neighbor* of U if and only if U is able to receive packets transmitted by node V. Equivalently, we denote the communication neighborhood by stating that the (V, U) (wireless) *link* is *up*. Otherwise, we say that (V, U) is *down*. A graph in which vertices represent nodes, and edges reflect the state of the links is called the *connectivity graph*.

We define the *physical neighborhood* of a node U as the set $\mathcal{N}_{\mathcal{P}}(U, r)$ of nodes that are at a distance smaller than r from U. This notion is intuitively related to communication neighborhood. Indeed, is we set $r = R$, where R is the *nominal communication range* of the wireless technology that nodes U and V use for communication, then node $V \in \mathcal{N}_{\mathcal{C}}(U)$ can be expected to be a physical neighbor of U as well.

However, it is crucial to clarify that communication and physical neighborhood are *not* equivalent in general. On one hand, communication neighborhood does not imply physical neighborhood. For example, a node V that increases its transmission power, exceeds the expected communication range, and thus places itself in $\mathcal{N}_{\mathcal{C}}(U)$ but not in $\mathcal{N}_{\mathcal{P}}(U, R)$. On the other hand, physical neighborhood does not imply communication neighborhood. Consider, for example, $V \in \mathcal{N}_{\mathcal{P}}(U, R)$ that *cannot* send information directly to U because of an obstacle (e.g., a wall); clearly, $V \notin \mathcal{N}_{\mathcal{C}}(U)$. The two types of neighborhood are equivalent only under the *unit disk* communication model, which considers U and V communication neighbors if and only if their (geometric) distance is below R. This model is a useful, but unrealistic approximation: It assumes free space propagation, no interference, and non-existent isotropic antennas.

Note that physical neighborhood is by definition symmetric. However, communication neighborhood, as defined, may be *asymmetric*. Even if $V \in \mathcal{N}_{\mathcal{C}}(U)$, U is not necessarily able to send information directly to V and would therefore not belong to $\mathcal{N}_{\mathcal{C}}(V)$. For communication neighborhood to be symmetric, both links (V, U) and (U, V) must simultaneously be *up*.

1.1.2 Neighbor Discovery Protocols

Neighbor(hood) Discovery (ND) protocols attempt to determine the neighbors (communication or physical) of a given node. Therefore, their main requirement is correctness: to identify only nodes that are actual neighbors, that is, to prevent an adversary from tricking nodes into accepting non-neighbors as neighbors. Verifying that a given node is indeed a neighbor could be viewed as a stand-alone part of secure neighborhood discovery functionality; we term this as *verification*. For example, a node could obtain neighborhood information in an insecure manner, but then perform verification to achieve secure ND.

In practice, ND protocols are only *partial*, as they may fail to discover (and verify) all neighbors. This is because it is difficult to guarantee message delivery in wireless networks. Furthermore, an adversary can jam communication and thereby prevent the discovery of one, many, or even all nodes that would be otherwise part of the neighborhood.

This problem can be avoided in restricted operating environments, where anti-jamming or other measures guarantee the delivery of messages. We call a ND protocol *complete* when it discovers all *honest* (or *correct*) neighbors, that is nodes that abide with the protocol functionality. We restrict ourself to honest, correctly functioning participants because *dishonest* or *faulty* nodes can always refrain from participating in the protocol execution.

1.1.3 Neighbor Discovery versus Distance Bounding

Distance Bounding (DB) protocols attempt to determine the distance between a verifier node and a prover node, or more precisely, an upper-bound on this distance. Hence, they are similar to physical ND protocols. The difference is the output of the protocol: When executed at node U with node V, a physical ND protocol returns *true* if node V is located closer than r (a parameter of the protocol), or *false* otherwise. In contrast, a DB protocol returns $\hat{r}$, which is an upper-bound on the distance to V.

A DB protocol can be easily converted to an physical ND protocol, but typically only a partial one. Indeed, if the DB protocol returns a loose upper-bound, then node V can be declared a non-neighbor even if it is closer than r. Furthermore, a physical ND protocol with $r = R$, the nominal communication range, can be converted into a (very conservative) DB protocol. However, such a protocol would not satisfy an additional requirement that we put on DB protocols in Part II: In a benign setting, we expect a DB protocol to provide an exact range estimate (Section 3.1.1).

1.2 ND as a Building Block

Neighbor discovery enables different types of system functionality, as the following examples illustrate.

Physical Access Control Receiving a packet from an RFID tag with a tag reader can be used to authorize the access of the tag bearer to a building (Figure 1.1(a)), or to trigger the opening of a car door in a Passive Keyless Entry and Start (PKES) system [162]. Packet reception implies the tag is at most within a system-specific predefined distance (e.g., a few centimeters or couple of meters) from the tag reader. Physical access control systems leverage on the range-limited communication capabilities of their hardware (tags), aiming essentially at physical ND.

Network Access Control In general, access to network resources is granted only to registered users or devices. Nonetheless, direct communication with a dedicated system entity can be an important access control criterion in mobile wireless systems. For example, nodes obtain connectivity with the Internet only when they are in range of a WLAN Access Point (AP) or a cellular system base station. Here, access control relies on communication ND.

Routing In multi-hop wireless networks, all types of data communication and dissemination (one-to-one, one-to-(m)any, or broadcast) rely on the notion of neighborhood. The neighbors of each node are always the ones that receive and forward control traffic and data to and from the node, for example, for route discovery and communication with another (destination) node. If a destination is already identified as a neighbor, then no route discovery or calculation

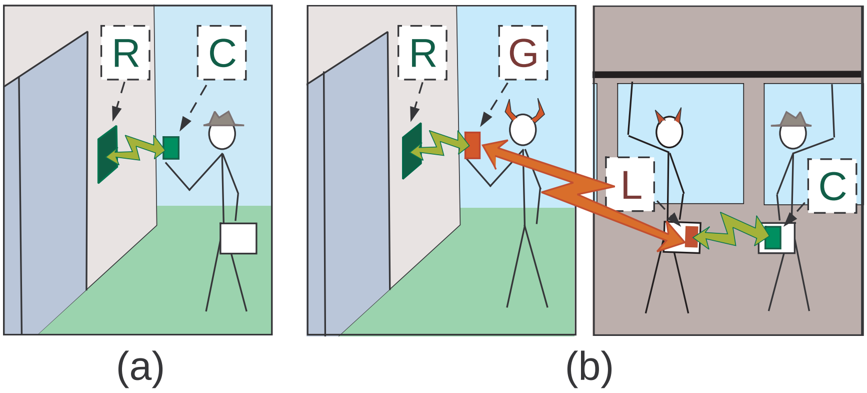

Figure 1.1: RFID Access Control. (a) Normal operation: a legitimate user opens the door controlled by an RFID reader (R) using his RFID card (C). (b) The attack: the *leech* (L), next to the RFID card owner, and the *ghost* (G), next to the RFID reader, use a long-range link to relay transmissions between the card and the reader. As a result, the reader is misled to believe that the legitimate card is in its physical neighborhood and opens the door.

is necessary. In the case a location-based routing protocol is employed, the neighbor closest to the destination's position is selected. In all of these cases, communication ND is necessary. If efficiency or fault-tolerance are sought, a complete ND protocol would be desirable. For example, selecting the appropriate neighbor to forward data to or using alternative paths assume that many or even all neighbors have been discovered. Completeness of a ND protocol would therefore prevent the adversary from disconnecting nodes from their benign neighbors.

Localization Neighbor discovery can be used for localization. In urban areas, with an abundance of infrastructure nodes with publicly known locations (cellular base stations, WLAN access points), this has become a viable alternative to GPS localization. In the simplest case, a node that wants to localize itself performs neighbor discovery, and derives its own location, e.g., as a center of mass of the neighbors' locations. The node can discover neighbor' locations form a service like Wigle [4]. Alternatively, in services like Skyhook [8], the device sends a query with the identifiers of its neighbors (e.g., MAC addresses), and obtains its own location. This localization process relies on physical ND.

1.3 ND Vulnerabilities and Attacks

If security was not a concern, communication ND would be trivial: Node V sends a beacon "Hello, I am V", and upon receipt of this beacon, U adds V to its communication neighborhood. Or, more implicitly, U adds V to $\mathcal{N}_{\mathcal{C}}(U)$ whenever it overhears any message that identifies V as the sender.

Such a protocol is easy to attack: An adversary M can forge and transmit a message "Hello, I am W" and thus convince node U that W is a neighbor, even if it is not the case, violating ND correctness. The countermeasure against this attack is equally straightforward: Apply a cryptographic authentication mechanisms, for example sing the beacon message with the sender's private key. However, this is not sufficient to secure ND. Indeed, M could still receive a beacon from W and simply *relay* it, to mislead U into adding W to $\mathcal{N}_{\mathcal{C}}(U)$. Although cryptography ensures that the received message has been created and transmitted at some point by node W, it gives no guarantees that the message was received directly from W. Such *relay attacks* constitute a fundamental attack vector against ND protocols. Also known in the literature as *wormhole attacks*, they are effective not only against the naive ND protocol described above, but can also harm the more sophisticated ND protocols that we survey in Section 1.4.

1.3.1 Classification of Attacks

We begin by differentiating between *external* and *internal* adversaries. In contrast to an external adversary, an internal adversary is an entity that is a legitimate, but malicious participant of the network, typically possessing cryptographic keys as all honest participants do. We can then distinguish between the following threat models:

i) An *external* adversary misleading two honest nodes that are not neighbors into establishing a neighborhood relation. A relay attack is a special case of this attack.

ii) An *internal* adversary misleading two honest nodes that are not neighbors into establishing a neighborhood relation.

iii) An internal adversary tricking a non-neighboring honest node to believe that he is a neighbor, possibly with the assistance of other adversarial nodes (external or internal) or even the unwilling assistance of an honest node.

Traditionally, communication ND protocols focus on threat models (i) and (ii). Note that the threat models are similar to the threat models considered for distance bounding protocols (Section 3.1): the mafia fraud (threats i and ii) and the distance fraud, the distance hijacking attack and the terrorist fraud (threat iii).

The adversarial nodes are typically assumed to be able to communicate via fast adversarial links not perceivable to the honest nodes. The number of nodes that the adversary controls can play a crucial role: Some secure ND schemes proposed in the literature are secure against a basic relay attack mounted with two adversarial nodes (*2-end wormhole*), but can be circumvented with, e.g., a 3-end wormhole.

In a typical Dolev-Yao fashion [48], an adversarial node is capable of transmitting arbitrary messages that he is able to generate without breaking cryptographic primitives such as encryption, digital signatures, or random nonces. In particular, the adversary can modify a received message before he retransmit it. Furthermore, the adversary can jam node communication in a selective or brute-force manner, possibly adjusting its transmission power and thus its impact. In the context of ND, jamming can obviously prevent the completeness of ND (even thwart the discovery of any neighbors at all), but also allow subtle attacks against some existing ND schemes (Section 1.4).

Relay Attacks

An important characteristic of a relay attack is the delay introduced when relaying messages. This is a critical parameter in the case where timing bounds are used in the defense against relay attacks, as explained in Section 1.4. We can classify attackers as:

i) *Store-and-forward relays*, if they need to receive the entire message before they are able to relay it.

ii) *Fast relays*, if they can start retransmitting the message while it is being received.

iii) *Distance-decreasing relays*, if they can relay a message with seemingly *negative* delay, by mounting appropriate attacks on physical-communication-layer [35]. We deal with such attacks in Part II of the thesis.

The store-and-forward relay is relatively easy to implement. In contrast, faster relay attacks require more sophistication from the adversary (Section 1.4).

We also differentiate between *short-range* and *long-range* relay attacks. The former resulting in fictitious links shorter than the nominal node communication range R, and the latter in links longer than R. This distinction is meaningful because short-range relays, as opposed to long-range ones, do not violate the correctness of physical ND. As a result, they cannot be detected by mechanisms protecting physical ND.

We further differentiate relay attacks according to the adversarial node behavior: they can either always forward packets or do so selectively. Moreover, they can relay messages using omnidirectional or directional antennas.

Tunneling Attacks It is important to point out the difference between relay attacks and *tunneling* attacks [112], which were introduced in the context of routing. Suppose, for example, that two internal adversarial nodes participate correctly in their respective ND protocols, but "tunnel" (i.e., encapsulate and transmit to each other) control traffic, so that they appear as neighbors on routes discovered by the routing protocol. In contrast to relay attacks, tunneling attacks cannot be thwarted by any secure ND protocol.

We note that in the literature the term *wormhole attack* is used differently by different authors, and some authors use it do describe both relay and tunneling attacks (and adding prefixes such as open/closed or hidden/exposed wormholes). In this thesis, we assume that a wormhole attack is synonymous to a relay attack.

1.3.2 Attacks on Overlaying Applications

Next, we discuss the implications of successful relay attacks for the upper-layer protocols and services discussed in Section 1.2.

Physical Access Control We consider an attack against an RFID-based system that controls physical access to a building, illustrated in Figure 1.1(b). For this attack [81], whose practical implementation was reported in [67], the adversary must control two nodes. The first adversarial node, the *leech*, is placed close to the victim's RFID tag. The second, the *ghost*, is placed next to the RFID reader controlling the building's entrance. Independently of the distance from the victim's tag to the reader, the leech and the ghost relay messages between them, thereby misleading the reader into believing that the legitimate RFID tag is close, granting access to the ghost's bearer. A similar attack against PKES systems was implemented in [61].

As pointed out in Section 1.2, such a physical access control system should guarantee physical ND. However, it fails because its design considers communication and physical neighborhood as equivalent and it relies on a naive communication ND approach. As a result, the attack violates the correctness of communication ND and thus the correctness of physical ND.

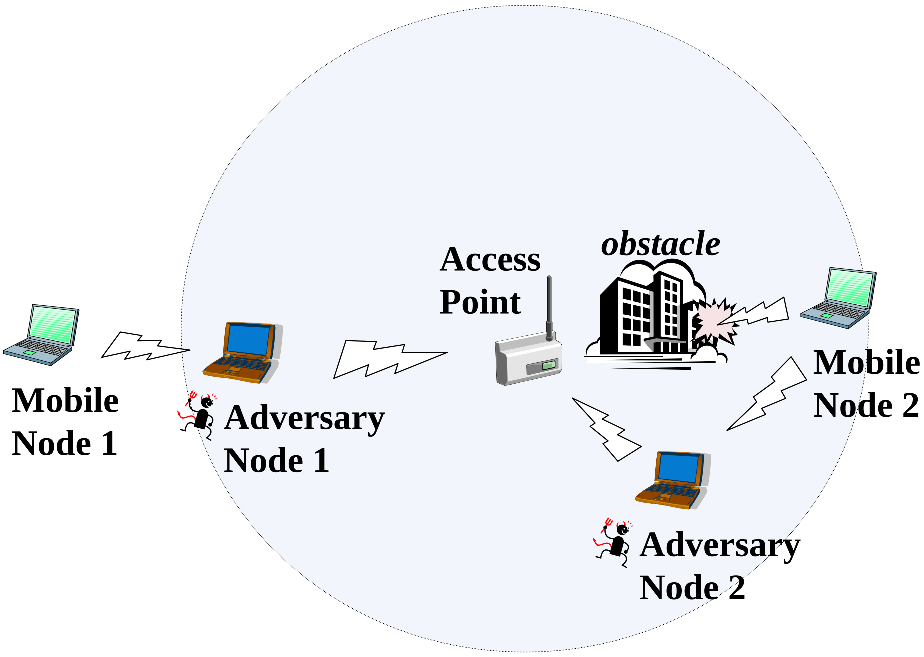

Figure 1.2: Relay Attack on Network Access. Mobile Node 1 cannot directly connect to the Access Point because it is out of range. Similarly, Mobile Node 2 cannot directly connect to the AP due to an obstacle. However, the Attacker Nodes can act as relays between the Mobile Nodes and the Access Point, misleading them they can communicate directly, and thus control the communication.

Network Access Control We consider an attack, illustrated in Figure 1.2, against mobile nodes trying to connect to an Access Point (AP). As Mobile Node 1 (MN1) is out of the AP's range, Adversary Node 1 can easily act as a relay between MN1 and AP. A relay attack is also possible when two nodes are physical neighbors (with r equal to the nominal communication range) but are not communication neighbors; this is the case for Mobile Node 2 and AP in Figure 1.2. In both cases, the correctness of communication ND is violated.

One could argue that the adversarial nodes provide a service to the system, as they essentially extend the AP coverage. But, in doing so, the adversary takes control of the node-to-AP connections. It can then intercept the relayed messages, as well as modify and delete them at will. In wireless networks eavesdropping is easy, yet, without the adversarial relays there would be no communication to eavesdrop on. Moreover, data modification would be more difficult without the relay attack. If MN1 were in range of the AP, the adversary would need at least two strategically positioned and synchronized nodes: one node jamming the AP, to prevent it from receiving the messages of MN1, and the second node recording MN1's transmissions and replaying their modified version.

Eavesdropping can be prevented by encryption, whereas message modification and deletion can only be detected (for example, with the help of digital signatures and message sequence numbers) but not prevented. Nonetheless, the relaying adversaries can delete messages effectively and, more important, stealthily: unlike jamming, the victim nodes (notably, the sending one) can detect the message loss but not its cause. Even worse, the adversary can

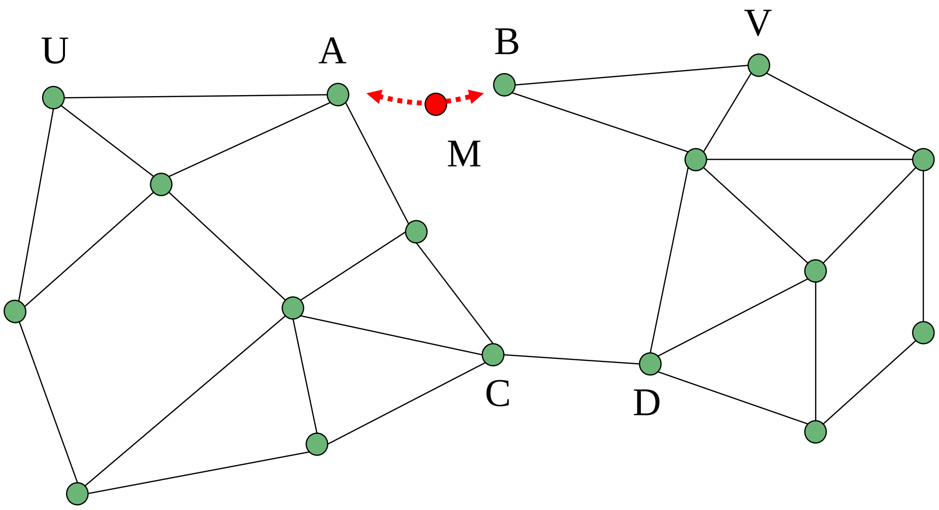

Figure 1.3: Relay Attack on Routing (I). By creating an artificial link (A,B), the adversary M attracts routes, e.g. (U, V)-routes that would otherwise use link (C,D). In this way, acting only locally, M gains control over the communication of remote nodes, e.g. U and V.

choose the point in time to delete messages in order to cause the most harm.

Routing We consider relay attacks against ND in a multihop ad hoc network, such as a sensor network. The significance of such wormhole attacks was first mentioned in [74, 112]. In Figure 1.3, nodes A and B are close to each other but unable to communicate directly due to the terrain and their transceiver limitations. The adversary places a node M within range of A and B, where M acts as a relay, making A and B believe that they are communication neighbors. Then, it is highly likely that U and V will communicate across a route that includes the adversary-controlled link (A, B). Such a $U - V$ route would be shorter than one that includes (C, D), and shorter routes are in general preferable. The result of this attack can be devastating. At first, the adversary-controlled link attracts considerable traffic. In addition, if the network relays a time-critical alarm, the adversary can stealthily cut-off its "link" and prevent the event detection by the network user.

The attacker's control over route establishment can be further enhanced, as shown in Figure 1.4. M_1, M_2, and M_3 are nodes controlled by the attacker, acting as simple relays. At the same time, the adversarial nodes N_1 and N_2 relay messages across a private out-of-band $N_1 \leftrightarrow N_2$ channel. Again, these attacks form short routes for many pairs of nodes, empowering the adversary to control significant amounts of network traffic.

Localization We consider attacks on a ND-based localization system. The attack deceives a node that is trying to localize itself into believing it is at a location chosen by the adversary. If an ND protocol without authentication is used[1], the adversary can simply spoof: The adversary 1) jams the ND messages from the true neighbors of the node, and 2) transmits forged ND messages from infrastructure nodes located at a chosen location. Such an attack against the Skyhook system has been demonstrated in [147]. In contrast, if the ND protocol incorporates authentication, forging of the ND messages is not possible. In that case, step 2

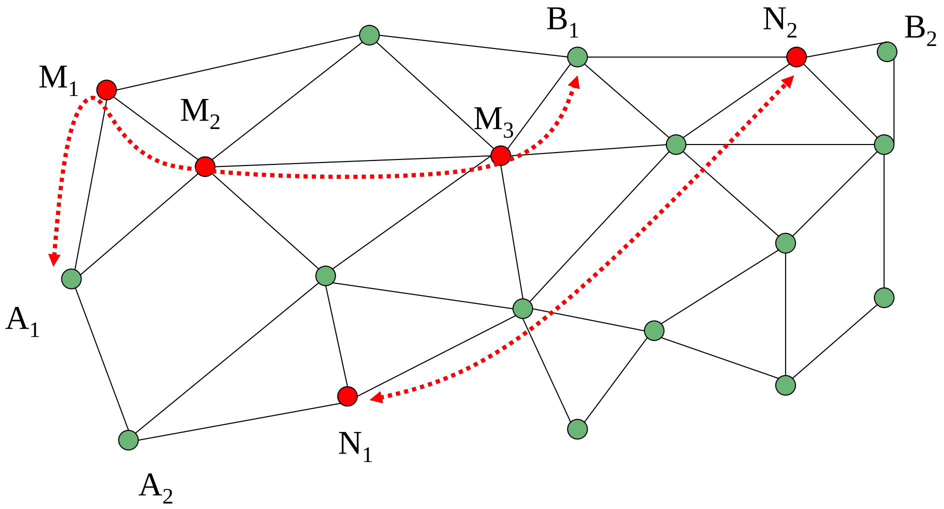

Figure 1.4: Relay Attack on Routing (II). By relaying transmissions between nodes A_1 and B_1, the adversarial nodes M_1, M_2, and M_3 create an artificial, long-range link (A_1,B_1). Similarly, nodes N_1 and N_2 can use an out-of-band channel to relay transmissions between A_2 and B_2. In both cases, the artificial link offers a route much shorter then alternative ones, and thus attracts traffic the adversary has control over.

of the attack can be replaced with a relay attack, as illustrated in Figure 1.5.

1.4 Survey of ND Protocols

We provide an overview of ND (and related) protocols, dividing them into the following basic categories. *Two-party protocols* are protocols that involve only the two nodes discovering if they are neighbors. Such protocols can be applied in any context, including the examples we provided for physical and networks access control. *Multi-party protocols*, in contrast, require the assistance of at least one extra node, and in some cases require a complete multi-hop network, which limits their applicability. Furthermore, schemes differ in hardware requirements, and work under different assumptions; we comment on the practicality of those. By default, we assume that the nodes share some cryptographic materials that allows them to authenticate each other.

We note that, traditionally, most ND protocols focus on the case where both participating nodes are honest. In other words, they only provide protection against threats (i) and (ii), but not (iii).

1.4.1 Two-Party Protocols

Time-based In a time-based approach, node V exchanges one or more messages with a potential neighbor U, measures the message time-of-flight, and multiplies it by the propagation

[1] The most obvious reason for no performing authentication is the lack of shared cryptographic material (e.g., shared symmetric keys, public keys, or certificates) between a node trying to localize itself and the infrastructure nodes. Without such material, authentication is not possible.

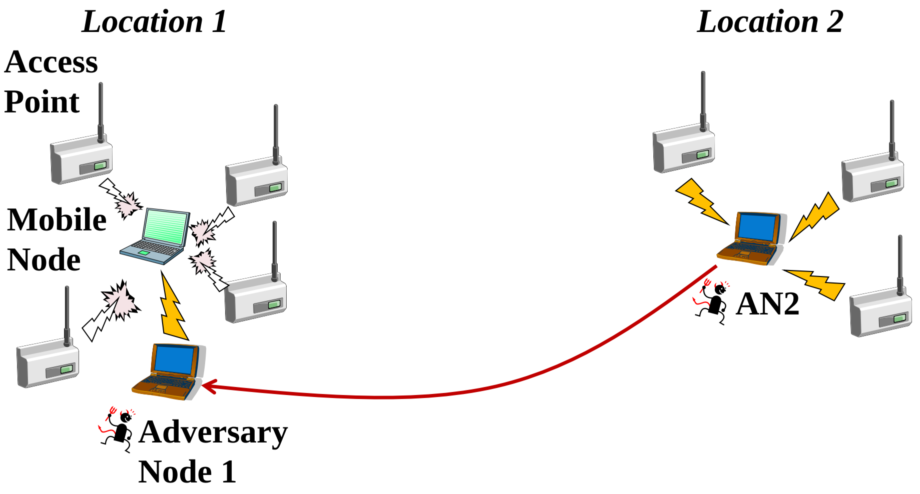

Figure 1.5: Relay Attack on ND-based localization. Adversary Node 1 jams the communication between the Mobile Node and the Access Points in its neighborhood at Location 1. Next, AN1 and AN2 relay the communication between the Mobile Node and Access Points at Location 2. As a result, the Mobile Node believes that the Access Points at Location 2 are its neighbors, and computes its position to be at Location 2.

speed of the wireless medium v to obtain the distance to U. If the distance is below a threshold r, node V concludes that U is a neighbor. The time-of-flight measurement is combined with authenticating U to V. The threshold r is typically the nominal communication range R of the wireless technology used by the nodes. The wireless medium is typically RF, and $v = c$, the speed of light, which has an important security implication: The adversary is not able to make the message travel faster than the speed of light, i.e., he cannot decrease the measured distance. (Unless he resorts to physical-communication-layer attacks [35] which seemingly reduce the propagation time, see Part II of the thesis).

This approach can guarantee physical ND, assuming that the physical-communication-layer is secure against distance-decreasing attacks. On the contrary, a fast short-range relay attack can deceive two nearby nodes that cannot communicate directly into believing they are communication neighbors. In Chapter 2 we prove that no time-based protocol can, in general, provide communication ND. Nevertheless, time-based protocols can be used as an approximation of communication ND: they are secure against long-range relay attacks, and against adversaries that introduce a relatively long delay when relaying. We elaborate on this in Section 2.4.

Time-based ND protocols are related to distance bounding protocols (as mentioned in Section 1.1.3). We provide a survey of distance bounding protocols in Section 3.4. Here, we discuss time-based protocols that were designed to provide (approximate) communication ND, under the assumption that both participating nodes are honest.

Temporal packet leashes [74] assume that nodes have precisely synchronized clocks (down to a nanosecond or microsecond level). A node V broadcast an authenticated message that includes a time-stamp of the time of transmission. Any node U that receives this message can, thanks to clock synchronization, estimate the message time-of-flight, and proclaim V a neighbor if it is below a threshold (typically equal to R/c plus some fixed and known processing time). The big advantage of this scheme is its low communication cost: One message is sufficient to allow all neighbors to securely discover V. It also allows making ND implicit, by attaching an authenticated time-stamp to every message. The obvious drawback of the scheme is the requirement of precise (and secure) clock synchronization, which is non-trivial: It can be achieved with setting the clock time when a node is deployed and using high quality clocks with very small drift. Or, the nodes can run secure clock synchronization algorithms, but this is shown to be as difficult as secure ND ([32], see Section 2.5.1). Or, the nodes can obtain the time from a GPS (or similar) module, which increases the cost and power-consumption of the nodes, and is susceptible to spoofing [76]. Another drawback is that temporal packet leashes do not allow to perform secure ND with an adversarial node, as it can easily forge a time-stamp. In addition, time-stamping the message at the necessary level of precision can be technically challenging on platforms with non-programmable or proprietary medium access control (MAC) layer.

To remove the requirement of clock synchronization, a number of works have proposed to replace a single-message beacon scheme with a challenge-response scheme [174, 55, 105]. In particular, the Truelink protocol [55] is integrating the challenge-response scheme into the RTS/CTS mechanism of IEEE 802.11. As with temporal packet leashes, implementing such a scheme can pose a technical challenge, as the authors of [55] report: They were not able to implement TrueLink on the IEEE 802.11 devices they had available, because of the proprietary firmware. Another drawback of such schemes is the communication overhead: A group of n fully connected nodes needs $O(n^2)$ messages to perform secure ND, versus only $O(n)$ messages requires with packet leashes. This can be a problem notably in highly mobile and highly connected networks such as Vehicular Ad-hoc Networks (VANETs). Nevertheless, in many systems challenge-response time-based ND protocols can be a viable solution, if the MAC layer is sufficiently open. In particular, successful implementations of such schemes for wireless sensor networks have been reported [14, 142].

Location-based If available, trustworthy location information can be utilized for secure ND. Such information can be provided at deployment time (for static networks), or from a GPS (or similar) module. Or, it can be provided by a secure localization scheme that uses the same wireless channel as the ND protocol [155, 89, 90, 156]; however such schemes are often based on ND or distance bounding, creating a vicious dependency circle.

Geographical packet leashes [74] are based on loose clock synchronization. Similar to temporal packet leashes, node V broadcasts an authenticated message that contains a time-stamp and a location-stamp. Any node U that receives such a message, can check if the message is not expired (thanks to loose clock synchronization). If not, U computes the distance to V, and declares V a neighbor if it is closer than R. This provides physical ND. To provide communication ND, the authors of [74] propose that U, knowing V and its own location, can decide if direct communication is possible based on a radio propagation model. This requirement is impractical and hard to satisfy in general in a communication environment that is not known a priori or highly dynamic. The scheme has similar advantages and disadvantages to temporal packet leashes.

Another scheme that provides secure physical ND by using location information is proposed in [173]. The scheme relies on a concept of location-based keys, with the goal of preventing not only relay attacks, but also Sybil, replication and sinkhole attacks.

An alternative way to provide communication ND is to combine time-based protocols with location-based protocols, as we propose in Section 2.3.

Device Fingerprinting Another ND approach relies on device RF fingerprinting, that is, the identification of characteristic signal patterns induced by radio transmitters. These techniques exploit the fact that transmitters, being physically distinct, generate slightly different wireless signals. This, in theory, allows the node receiving a signal to uniquely identify the source of that signal, and thus provide communication ND [134].

A number of works for device fingerprinting have beed presented in the literature, e.g., [34, 153, 134, 154, 24, 78, 42, 40, 139, 170], designed for different technologies and using different techniques. In particular, transient-based techniques are often considered for enhancing the authentication of RFID systems, and prevent RFID tag cloning [42, 139, 170]. These approaches can provide high identification accuracy, but – currently – at the cost complex and expensive hardware (typically an oscilloscope with a few GSamples/s sampling rate). In addition, in such systems the tag is typically placed in a fixed position very close to the reader, such that the wireless channel is known and constant. Generalizing such techniques to variable environments can prove challenging, because the machine learning techniques underlying such approaches would have to distinguish between device-specific factors and the influence of the wireless channel. Furthermore, in [41] it was shown that a sophisticated adversary, by recording the transient and replaying it with a waveform generator can successfully spoof such systems.

Other works propose device fingerprinting based on features that are both more robust to variable propagation environments and easier to measure with commodity hardware. For example, Brik *et al.* [24] propose a "modulation-based" technique based on several PHY features: frequency error, SYNC correlation, I/Q offset error, magnitude error, and phase error. It allows for accurate identification of IEEE 802.11 devices from up to 25m away. Zanetti *et al.* [170] propose another technique based on time interval error that allows for classification of UHF RFID chips from up to 6 meters away. In [78], the authors propose to perform device fingerprinting of IEEE 802.11 access points based on their clock skews. However, such techniques proved relatively easy to spoof: [15] proposes an off-the-shelf-hardware attack against the method proposed in [78]. In [41], the authors use GNU software radios to successfully impersonate devices fingerprinted with the modulation-based method [24]. It is an open question whether it is possible to design device fingerprinting that is robust to environmental factors (distance, antenna orientation, propagation environment), implementable on low- or mid-end devices, and resilient to spoofing.

Channel Fingerprinting In most environments, notably indoor, the wireless channel is much "richer" than wireline channels. On its way from a transmitter to a receiver, the wireless signal traverses multiple paths, and is reflected from various obstacles such as walls, furniture, or people. This makes the design or wireless communication systems challenging. In particular, it is a challenge for device fingerprinting techniques, as they must distinguish between device-specific and channel-induced effects. It also creates vulnerabilities in precise ranging that we explore in Part II of the thesis.

However, the wireless channel has properties that make this diversity useful for security purposes: 1) Channel decorrelation: The channel state, which can be characterized by the *channel impulse response* (CIR), is location-specific. Two devices U and W that receive a message from device V observe CIRs that are not correlated, if they are located more than the RF wavelength apart. This is referred to as *spatial decorrelation*. The channel also decorrelates over time, although this happens less rapidly in static settings. 2) Reciprocity: The CIR that device U observes when communicating with V is the same as the CIR that V observes when communication with U, if measured at the same time instant.

These two properties allow two devices to derive a shared (thanks to reciprocity) and secret (thanks to spatial decorrelation) key. Indeed, most research on using the channel for security purposes is devoted to this topic; a recent overview of such efforts is provided in [91]. Beyond key extraction, spatial decorrelation allows for a form of message authentication [166]. A receiver can record, for every received packet, the CIR estimate, or a similar *channel fingerprint* that its hardware allows it to estimate. Assuming that the channel does not change (significantly) between packet transmissions, two packets transmitted by the same sender U would have similar fingerprints; whereas a transmission from another device W, e.g., trying to impersonate U would have a different fingerprint. It was shown it possible to apply such techniques if the channel varies over time [165] and even if nodes are mobile [164].

Finally, combining such channel-based authentication with channel reciprocity could be used to provide communication ND: Devices V and U need to estimate the channel fingerprints on their ends, and exchange them in a (cryptographically) authenticated fashion. If their estimates are (almost) identical, they can conclude that they are communication neighbors. However, this approach would not be effective against a relay that preserves the CIR of the relayed transmission, e.g., an analog relay. Indeed, consider that node W is such an analog relay, and that the CIR of the $V \leftrightarrow W$ channel is f and the CIR of the $W \leftrightarrow U$ channel is g. Then the CIR of the $V \rightarrow W \rightarrow U$ channel, as well as $U \rightarrow W \rightarrow V$ channel is $f * g$, as convolution ($*$) is commutative.

Beyond the analog relay attack, the security of both authentication and communication ND based on channel fingerprinting relies on the difficulty of spoofing a channel fingerprint by an adversary. This topic requires further investigation.

Directional Antennas-based The use of directional antennas to prevent relay attacks is proposed in [73], under the assumption of the unit disk model, the availability of antennas with an even number n of non-overlapping zones each spanning an angle of $\frac{2\pi}{n}$, and the ability to have zones identically oriented for all nodes (e.g., using a compass). If two nodes are communication neighbors, a message sent over some zone z_i should be received at the opposite zone $\overline{z_i}$. Information (cryptographically protected) on the used zone is included in messages to detect simple relay attacks. However, it is possible that the zones match even under a wormhole attack, if the wormhole end are positioned appropriately with respect to the honest nodes. To detect such cases, two extensions of the scheme with a third *verifier* node is proposed. The extended schemes can prevent 2-end wormholes spanning over more than 2-hops; but a relay attack spanning 2-hops is possible. Furthermore, a k-end wormhole, $k > 2$, can defeat this scheme. In the most extreme case, it is clear that an adversary can place a relay node in every zone of every honest node, and relay messages in a selective fashion to completely defeat the scheme. But even a 3-end wormhole can create a number of false links. A further disadvantage of the extended schemes is the need for an appropriately located verifier node. If one is not available, a valid link cannot be verified. This limits the applicability of the extended schemes in low density networks. In addition, the applicability is further restricted by the requirement of identically oriented directional antennas.

A scheme based on a similar idea, but designed for underwater acoustic networks is proposed in [171].

1.4.2 Multi-Party Protocols

Time- and Distance-based A scheme proposed in [142] extends simple time-based protocols with additional verification based on simple geometric tests performed in a 2-hop neighborhood. This allows the scheme to provably prevent fast, short-range 2-end relay attacks. The geometric tests require the knowledge of distances between the neighbors, which can be measured with, e.g., ultra-sound ranging. This constitutes the biggest drawback of the scheme: It limits the applicability of the scheme to systems where such range measurements are possible.

Location-based If only a subset of nodes is location-aware, these nodes, termed as guards, can help other nodes establish neighbor relations. In the scheme propose in [126], guards broadcast beacon messages reporting their location. Afterwards, other nodes exchange information about received beacon messages and assume they are neighbors if sufficiently many common beacon messages (at least k) were received. Relay attacks are detected based on two principles: 1) any beacon message should be received at most once and 2) all locations in received beacon messages should lie in a circle with a radius two times the guard range. This can prevent relatively simple relay attacks, but not more elaborate attackers. For example, a selective wormhole can avoid detection based on principle 1). Moreover, one end of a wormhole can jam and prevent reception of legitimate beacon messages, relay beacon messages from the other end of the wormhole, and essentially "relocate" the victim node(s), much like the attack in Figure 1.5. The scheme is probabilistic in nature, and the threshold k is calculated in the unit disk model, based on the density of guard and node deployment, resulting in an approximate physical ND. A similar protocol is proposed in [45].

Connectivity-based A number of schemes that use exclusively connectivity information have been designed for multi-hop networks. Technically, for the most part such protocols do not provide secure ND: Rather, they detect the presence of long-range wormhole attacks (2 or more hops) based on the distortion that such wormholes create in the network topology. Only some schemes try to localize the wormhole and/or remove affected links. The most significant advantage of such schemes is that they do not have any hardware requirements. The biggest drawback is that they require a relatively dense multi-hop network, and cannot be applied in services such as physical or communication access control.

A centralized scheme is proposed in [28]. The scheme detects the presence of a wormhole based on two connectivity graph properties: distribution of the node degree, and distribution of the shortest path lengths. The system must known the expected values of these statistics in the benign case, i.e., when there is no wormhole attack, which is somewhat of a drawback. Only wormhole detection is provided.

In [158, 160] a centralized scheme is proposed, based on visualization of the connectivity graph. The (virtual) coordinates of the nodes, necessary for visualization, are computed with multi-dimensional scaling. The scheme requires a human operator that will manually inspect the graph to detect and localize wormholes. Multi-dimensional scaling is also used in [167], which proposes an automatic test for detecting wormhole attacks. The test is based on the

k-hop neighborhood diameter. The scheme is decentralized, and only requires the knowledge of the k-hop neighborhood.

Local network connectivity information is proposed in [93] as the basis of a heuristic to detect wormholes and reject false links. Nodes exchange locally communication neighborhood information, obtained through a non-secure ND mechanism. Afterwards they check for *forbidden structures*, that is, connectivity subgraphs that would exist if a wormhole were present (and would be unlikely otherwise). Forbidden structures depend on node density and the connectivity model. For sufficiently high network density, simulation results show a 100% detection rate with no false alarms, for all connectivity models considered in [93] (unit disk, as well as more realistic models). However, the simulations assume a relatively naive relay, whereas a selective wormhole establishing only one or few fake links would be less likely to create a forbidden structure. Furthermore, although the wormhole detection scheme is evaluated, it is unclear how the link rejection scheme would perform: the authors point out that it might reject valid links. An additional drawback of the scheme is that it requires some knowledge of the connectivity model to determine the parameters defining the forbidden structures.

A similar approach that uses a different feature of the local neighborhood (generalized edge-clustering coefficient) is present in [175]. In [49] a scheme that does not require the knowledge of the connectivity model is proposed. The scheme is based on a topological methodology, assuming that the network is a 2-manifold of genus 0. The authors show that their scheme outperforms [93]. However, their approach is specific to 2D networks, and it is not clear if it would reasonably generalize to 3D networks.

A connectivity-based secure communication ND scheme is proposed in [47]. In this scheme, node V accepts a new node U as a neighbor, if there exists a short U-V path consisting exclusively of previously verified neighbors of U and V. The drawback of the scheme is that is requires a bootstrapping phase free of wormhole attacks.

1.4.3 ND and Routing

In many systems, secure ND is only a step towards secure routing. Hence, there are a number of proposals that integrate detection and prevention of relay attacks with routing protocols. Often, such schemes detect tunneling attacks as well as wormhole attacks.

Two time-based examples of such schemes are [33] and [124]. Essentially, in those schemes the source measures the round-trip-time (RTT) to the destination, and compares it with the number of hops. If the RTT is too large, the source assumes that a wormhole or tunneling attack is present. A similar approach, based on location, is proposed in [159]: Each node on the route must provide its location; a wormhole/tunnel is detected if subsequent nodes are too remote to be communication neighbors. In the scheme proposed in [161], location of the source and destination is used to compute their distance, and compare it against the number of hops. A route it rejected if the number of hops is too small. We should also mention the LiteWorp protocol [82, 83] in which a set of guard nodes monitors the behavior of other nodes, and excludes nodes that do not forward packets correctly. However, as the relay nodes are "invisible" to such a scheme, LiteWorp can only prevent tunneling attacks. An extension of LiteWrop to mobile network, MobiWorp is presented in [84].

The scheme proposed in [133] measures how frequent each link is used in established routes. It detects a wormhole/tunnel attack by identifying link(s) with significantly higher utilization frequency. This is based on the observation that, typically, a wormhole/tunnel creates a shortcut across the network and hence attracts a significant amount of traffic. Once detected, such links can be avoided.

Finally, there are proposals of secure routing protocols that provide reliable data transmission in a failure/attack agnostic fashion. Simply put, in protocols such as SSP, SMT [114] or Sprout [54] the source measures the reliability of the route by counting the number of (cryptographically protected) acknowledgments received from the destination. Routes that offer poor reliability, or links which constitute such routes, gain a "bad reputation" and are eventually avoided. In a similar fashion, in the Castor protocol [63] each node measures the reliability of each neighbor, and avoids neighbors with poor reliability when forwarding packets. Such protocols opportunistically use a wormhole, as long as it is delivering packets, and start avoiding it only when it starts dropping packets.

The security of such protocols benefit from secure ND, as it allows them to instantly avoid wormholes. In particular, the authors of Sprout [54] recommend that the protocol is used with their TrueLink ND protocol [55]. This is because while Sprout can eventually recover from a wormhole attack that starts dropping packets, the number of false links introduced by a wormhole makes this process slow; In contrast, Sprout can recover relatively fast from a tunnel attack.

This is confirmed in the performance comparison shown in Figure 1.6. In the illustrated scenario, no secure ND protocol is used. At time 300s, a wormhole which was previously relaying all packets, starts dropping all data packets. Among all the evaluated protocols, only Castor is able to fully recover from the attack in a reasonable amount of time. (Note that SEAD and AODV do not include a component that would estimate route reliability, hence they do not recover from the attack.)

A drawback of such robust routing schemes is that they cannot prevent an adversary from temporarily disrupting the network operation, which can be sufficient to prevent or delay the delivery of a critical message. Furthermore, an adversary that creates the wormhole to attract traffic and eavesdrop or perform traffic analysis on it, cannot be mitigated by such approaches.

1.4.4 ND and Localization

A number of secure localization schemes have been proposed that are based directly on ND with location-aware guard nodes [89, 88, 30], or – more generally – estimate the distance to guards based on the number of hops [163]. To protect against relay attacks, these schemes apply rules very similar to [126]. However, this makes them vulnerable to the jam-and-relay attack illustrated in Figure 1.5. To prevent against such an attack, the ROPE protocol [90] combines the principle of [89] with distance bounding that provides secure physical ND.

1.4.5 Relay Attacks

From a practical point of view, a store-and-forward relay is relatively easy to implement on off-the-shelf hardware, as access to the physical-communication-layer is not necessary: The adversary simply needs to receive a message on the Medium Access Control (MAC) layer (or higher layer), and retransmit it. Recently, such a relay attack using off-the shelf Nokia NFC-capable mobile phones was implemented in [62] against NFC (*Near-Field Communication*). The implementation relies on APIs public to every application, i.e., it does not require root access. No delay figure for the relay is given.

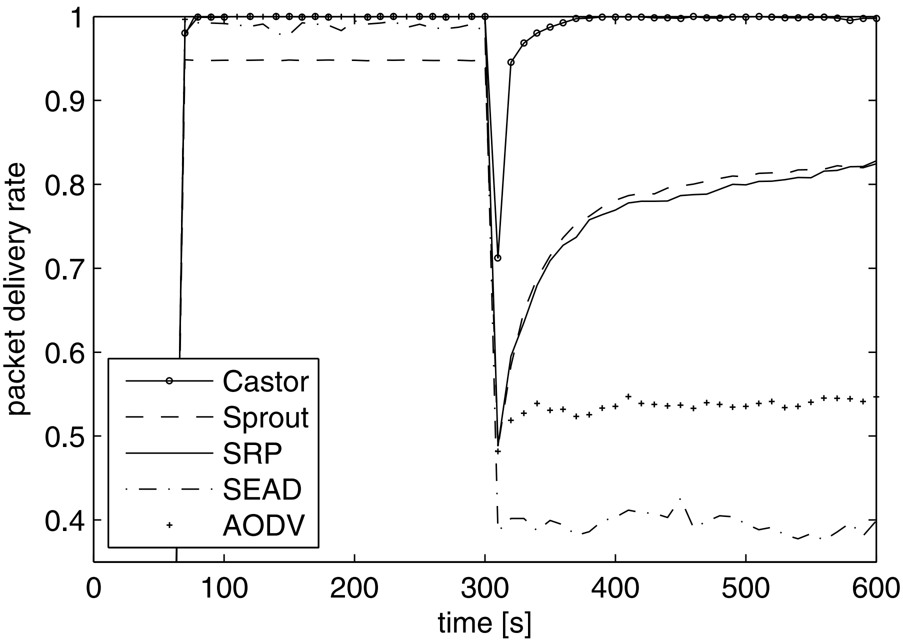

Figure 1.6: Performance (packet delivery rate) of routing protocols under a wormhole attack without secure ND (taken from [63]): Castor [63], Sprout [54], SRP+SSP [112, 114], SEAD [75], AODV [123]. The wormhole relays all packets until the 300s mark, and afterwards drops all data packets and only relays control traffic.

In contrast, faster relay attacks require a more sophisticated adversary, typically able to put together dedicated hardware. In [67], a relay attack against ISO 14443A contactless smart card systems is constructed for under £100. The relay achieves a delay of $15 - 20\mu$s. The delay comes mostly from the need to demodulate the received signal and to modulate it again for transmission.

In [61], the authors construct *analog* relays against a LF wireless link used in Passive Keyless Entry and Start systems. Such relays offers a much lower relaying delay. A relay with a 2.5GHz wireless link connecting two relay ends can achieve a relaying delay as low as 15-20ns (not counting the propagation delays). The delay is even lower for a wireline relay, but the propagation speed of the coaxial cable that connects two relay ends is only 66% of the speed of light. This introduces a non-negligible additional delay if the relay ends are remote.

We discuss the related work on distance-decreasing relay attacks in Section 3.4. For a more elaborate overview of relay attacks, we refer the reader to [70].

1.4.6 Related Attacks

The wormhole attack, in its symptoms, bears similarity to two other fundamental and hard to detect attacks. On one hand, a wormhole end can be perceived as a *Sybil* node, with messages tied to different identities being transmitted by a single node. Hence, seemingly, a Sybil node detection mechanism [108] could be used to thwart relay attacks. However, a wormhole can selectively relay the messages of a single node, and still be effective.

On the other hand, as in the *node replication* attack, messages tied to a single identity are transmitted by more than one node. However, node replication is harder to detect than a wormhole attack: Schemes that address node replication [116, 37] focus on probabilistically detecting replicas located in remote parts of the network and require that nodes are location-aware. Obviously, a long-range wormhole can be easily (and deterministically) prevented using geographical packet leashes.

1.5 Summary

In this Chapter, we have provided a general introduction to neighbor discovery (ND). We have clarified the definitions of neighborhood and the adversary model, provided an overview of the applications of ND protocols and the consequences of attacks against ND. We have also given a survey of the literature on ND. This explains the context and motivation for the rest of the thesis, notably for Chapter 2, where we introduce a formal framework for reasoning about ND protocols.

Chapter 2

Formal Analysis of Secure Neighbor Discovery

Analysis of protocols, notably security protocols, has proven time and again to be notoriously tricky if done in an informal fashion. The canonical example is the Needham-Schroeder protocol [107] published in 1978; only in 1995 did Lowe discover a subtle flaw using formal methods [92]. Informal arguments, in particular those that iterate through a list of possible attack scenarios, are likely to miss the often subtle interaction between participants of the protocol and the adversary, especially if parallel sessions of the protocols are involved. For an example closer to the topic of the thesis, we refer to an informal framework for distance bounding protocol analysis proposed in [16]. This framework does not take into account the distance hijacking attack, recently discovered in [38] (see Section 3.4).

The limitations of informal analysis can be overcome by formal analysis. On one hand, a formal model allows for impossibility results to be proved. Such results establish the limitations of a given class of protocols to solve a particular problem. Hence, they serve as important guidelines in protocol design, highlighting the need to consider a different or broader class of protocols, or modify the problem formulation. On the other hand, with a formal framework it is possible to prove the correctness of concrete protocols, or at least discover their incorrectness. A number of methods for formal verification of security protocol have been proposed, many of which can be computer-verified, or even automatized, thus limiting the fallible "human factor". In particular, a number of approaches deal with authentication protocols. At first glance, ND protocols can be considered a form of authentication protocols. However, as these approaches were developed for Internet-like environments, they abstract away features essential for ND protocols, notably the very notion of being neighbors. In this Chapter, we propose a framework that incorporates such features. More precisely, we make the following contributions:

▶ We build a formal model of wireless networks. The model captures characteristics of wireless communications essential for reasoning about ND such as node location, message propagation time, and the discrepancy between physical ND and communication ND (i.e., nodes that are physically close are not necessarily able to communicate directly). With the model, we propose a specification of *two-party communication ND*, and distinguish two classes of protocols: *time-based protocols* (T-protocols), for which nodes exchange messages and are able to measure time with perfect accuracy, and *time-and-location-based protocols* (TL-protocols), for which nodes in addition are aware of their location.

▶ We prove the following *impossibility result*: No T-protocol can provide secure communication ND if adversarial nodes are able to relay messages with a delay below a certain threshold. This threshold is closely related to the *communication range* of the nodes. On the contrary, if the minimum relaying delay is above the threshold, we show it is possible to achieve secure communication ND.

▶ Within the T-protocol class we distinguish between *beacon* (B) protocols and *challenge-response* (CR) protocols. We define one protocol in each of these classes and prove that they satisfy the secure ND specification, under the assumption that the adversarial relaying delay is above the threshold identified by the impossibility result. Furthermore, we propose two novel TL-protocols (a B/TL-protocol and a CR/TL-protocol) and prove that they can provide secure ND under the assumption that the relaying delay is strictly positive. We reinforce these results by mechanizing the model and some of the proofs in Isabelle/HOL.

Chapter Outline In Section 2.1 we define the framework and the ND specification. In Section 2.2 we derive the impossibility result. We extend the model in Section 2.3, define the protocols and prove that they satisfy a refined ND specification. We also give an overview of the Isabelle/HOL mechanization. We discuss the assumptions and compare the protocols in Section 2.4. In Section 2.5 we discuss the related work and open problem, before concluding in Section 2.6.

2.1 System Model

We build a model of wireless networks that captures features essential for ND protocols and their security. To keep the model tractable, we make a number of simplifying assumptions, which we explain and justify in Section 2.4.1. The model presented in this section is used in Section 2.2 for the impossibility result. To reason about the security of concrete protocols, we extend the model in Section 2.3.

The basic entities in a wireless network, *nodes*, are processes running on computational platforms equipped with transceivers communicating over a wireless channel. We assume that nodes have synchronized clocks (although not all protocols we consider in this paper make use of this assumption) and are static (not mobile). Nodes either follow the implemented system functionality, in which case we denote them as *correct* or *honest*, or they are under the control of an adversary, in which case we denote them as *adversarial* nodes. Adversarial nodes can behave in an arbitrary fashion, also acting as correct nodes or lying dormant for any period of time.

We model communication at the physical layer rather than at higher layers (data link, network, or application), in order to capture the inherent characteristics of ND in wireless networks. For simplicity, correct nodes are assumed to use a single wireless channel and omnidirectional antennas, but we do not require them to have equal transmission power and receiver sensitivity. On the contrary, adversarial nodes use directional antennas to communicate across the wireless channel used by correct nodes, but they can also communicate across a dedicated *adversarial channel* imperceptible to correct nodes.

Our system model comprises: (i) a *setting* $\mathcal{S}$ that describes the type (correct or adversarial) of nodes, their location and the state of the wireless channel; (ii) a *protocol model* $\mathcal{P}$ that determines the behavior of correct nodes; (iii) an *adversary model* $\mathcal{A}$ that establishes the capabilities of adversarial nodes.

We assume that looking at the system at any point in time reveals one or more phenomena. We are interested in those relevant to the wireless communication and the system at hand and thus to our analysis. We denote these phenomena, associated with nodes, as *events* (Definition 3). Then, we model the system evolution over time using the notion of *trace*, i.e., a set of events (Definition 4). More precisely, we use *feasible* traces, which satisfy constraints specified by $\mathcal{S}$ (correspondence between wireless sending and receiving of messages), $\mathcal{P}$ (correct nodes follow the protocol), and $\mathcal{A}$ (adversarial nodes behave according to their capabilities). The constraints are defined by logical formulas we call *rules*.

2.1.1 System Parameters

Our model includes a number of parameters, listed below, which are determined by the technologies used by correct and adversarial nodes.

- $\mathbf{v} \in \mathbb{R}_{>0}$, the *signal propagation speed*, defining how fast messages propagate across the wireless channel, determined by the communication technology,
- $\mathbf{v}_{\text{adv}} \geq \mathbf{v}$, the *information propagation speed* over the *adversarial channel*; as $\mathbf{v}_{\text{adv}} \geq \mathbf{v}$ this is also the maximum speed at which information can propagate,
- $\Lambda \subset 2^{\mathbb{R}^3}$, the set of *antenna patterns* that adversarial nodes can utilize with their directional antennas,
- $\Delta_{\text{relay}} \in \mathbb{R}$, the *minimum relaying delay* introduced by a node when relaying a message; this delay is due to processing exclusively, it does not include propagation time or any other delay.
- $\mathbb{M}$, the *message space*; we keep the message space unspecified for the impossibility result in Section 2.2; we provide a concrete message space when we talk about specific protocols in Section 2.3.
- $|.| : \mathbb{M} \to \mathbb{R}_{>0}$, the *message duration* function.

Further, $\mathbb{V}$ denotes the set of unique *node identifiers*, which for simplicity we will consider equivalent with the nodes themselves.[1]

2.1.2 Settings

A setting describes the type and location of nodes, and how the state of the wireless channel changes over time.

Definition 1. A *setting* $\mathcal{S}$ is a tuple $\langle V, loc, type, link, nlos \rangle$, where:

- $V \subset \mathbb{V}$ is a finite set of nodes. An ordered pair $(A, B) \in V^2$ is called a *link*.
- $loc : V \to \mathbb{R}^3$ is the node *location* function. As we assume nodes are not mobile, this function does not depend on time. We define $dist : V^2 \to \mathbb{R}_{\geq 0}$ as $dist(A, B) = d(loc(A), loc(B))$, where d is the Euclidean distance in $\mathbb{R}^3$. We require the *loc* function to be injective, so that no two nodes share the same location. Thus, $dist(A, B) > 0$ for $A \neq B$.
- $type : V \to \{correct, adversarial\}$ is the *type* function; it defines which nodes are *correct* and which are *adversarial*. This function does not depend on time, as we assume that the adversary does not corrupt new nodes during the system execution. We denote $V_{\text{cor}} = type^{-1}(\{correct\})$ and $V_{\text{adv}} = type^{-1}(\{adversarial\})$.

[1] Although this implies that every node is assigned a single identifier, it does not prevent an adversarial node from using (in the messages in sends) any identifier.

- $\mathit{link} : V^2 \times \mathbb{R}_{\geq 0} \to \{\mathit{up}, \mathit{down}\}$ is the *link state* function. Accordingly, we say that at a given time $t \geq 0$, a link $(A, B) \in V^2$ is *up* (denoted $\mathit{link}(A{\to}B, t)$) or *down* (denoted $\mathit{link}(A{\not\to}B, t)$). We use abbreviations $\mathit{link}(A{\leftrightarrow}B, t) =_{\mathrm{def}} \mathit{link}(A{\to}B, t) \wedge \mathit{link}(B{\to}A, t)$ and $\mathit{link}(A{\not\leftrightarrow}B, t) =_{\mathrm{def}} \mathit{link}(A{\not\to}B, t) \wedge \mathit{link}(B{\not\to}A, t)$. We extend the "$\mathit{link}(A{\to}B, t)$" notation from single time points to sets as follows: $\mathit{link}(A{\to}B, T) =_{\mathrm{def}} \forall t \in T.\ \mathit{link}(A{\to}B, t)$. We establish the convention $\mathit{link}(A{\not\to}A, \mathbb{R}_{\geq 0})$.
- $\mathit{nlos} : V^2 \to \mathbb{R}_{\geq 0}$ is the *non-line-of-sight delay (NLOS)* function. If two nodes A and B can communicate over a line of sight, then $\mathit{nlos}(A, B) = 0$. Otherwise, $\mathit{nlos}(A, B)$ specifies the additional distance that the signal has to propagate compared to line-of-sight propagation $\mathit{dist}(A, B)$. We assume this function is symmetric, because of reciprocity of wireless links.

We denote the set of all settings by $\mathbb{S}$.

The ability to communicate directly, without the intervention or 'assistance' of relays, is expressed in our model by a link being up, thus the following definition:

Definition 2. Node A *is a neighbor* of node B in setting $\mathcal{S}$ at time t, if $\mathit{link}(A{\to}B, t)$. If $\mathit{link}(A{\leftrightarrow}B, t)$ we will say that nodes A and B *are neighbors* at time t.

For simplicity of presentation, we use "$\mathit{link}(A{\to}B, t)$" to denote the neighbor relation and the link relation.

2.1.3 Events and Traces

We use the notion of *trace* to model an execution of the system. A trace is composed of *events*. We model events related to the wireless communication and the ND protocols operation. Each event is primarily associated with (essentially, takes place at) a node we call the *active* node.

Definition 3. An *event* is one of the following terms:

- $\mathsf{Receive}(A; t; m)$
- $\mathsf{Bcast}(A; t; m)$
- $\mathsf{Dcast}(A; t; \alpha; m)$
- $\mathsf{Neighbor}(A; t; B, C, t')$

where $A \in \mathbb{V}$ is the *active node*, $t \in \mathbb{R}_{\geq 0}$ is the event *start time*, denoted by $\mathit{start}(.)$, $m \in \mathbb{M}$ is the transmitted/received message, $\alpha \in \Lambda$ in an antenna pattern, $B, C \in \mathbb{V}$ are node identifiers and $t' \in \mathbb{R}_{\geq 0}$ is a time instant.

The first three events are related to communication on the physical layer. Receive represents message reception. Bcast represents sending a message with an omnidirectional antenna. Dcast represents sending a message with a directional antenna using a pattern $\alpha \in \Lambda$. The pattern α is a subset of $\mathbb{R}^3$ indicating which nodes receive the message, assuming the sending node A is located at $(0, 0, 0)$. We use the notation $B \in \alpha(A)$, meaning that $\mathit{loc}(B) - \mathit{loc}(A) \in \alpha$. The set of allowable antenna patterns, Λ, depends on the antenna used by the adversarial nodes. We do not dwell on the details of the structure of Λ, except for one requirement: $\mathbb{R}^3 \in \Lambda$. This is because, to facilitate proof presentation, we assume that adversarial nodes use Dcast only. Having $\mathbb{R}^3 \in \Lambda$ allows the adversarial nodes to use their antenna in an omnidirectional fashion.

Neighbor can be thought of as an internal outcome of a ND protocol (possibly reported to some higher layer): Node A declares that B is a neighbor of C at time t'. Having t' a single

s1 $\forall A \in V, t \in \mathbb{R}_{\geq 0}, m \in \mathbb{M}.\ \mathsf{Receive}(A;t;m) \in \theta \implies \exists B \in V.\ \mathit{link}(B{\to}A,[t,t+|m|])$
$\wedge\ (\mathsf{Bcast}(B;t-(\mathit{dist}(A,B)+\mathit{nlos}(A,B))\mathbf{v}^{-1};m) \in \theta\ \vee\ (\exists \alpha \in \Lambda.\ A \in \alpha(B)$
$\wedge\ \mathsf{Dcast}(B;t-(\mathit{dist}(A,B)+\mathit{nlos}(A,B))\mathbf{v}^{-1};\alpha;m) \in \theta))$
s2 $\forall A, B \in V, t \in \mathbb{R}_{\geq 0}, m \in \mathbb{M}.\ \mathsf{Bcast}(B;t-(\mathit{dist}(A,B)+\mathit{nlos}(A,B))\mathbf{v}^{-1};m) \in \theta$
$\wedge\ \mathit{link}(B{\to}A,[t,t+|m|])) \implies \mathsf{Receive}(A;t;m) \in \theta$
s3 $\forall A, B \in V, t \in \mathbb{R}_{\geq 0}, m \in \mathbb{M}, \alpha \in \Lambda.\ \ (\mathsf{Dcast}(B;t-(\mathit{dist}(A,B)+\mathit{nlos}(A,B))\mathbf{v}^{-1};\alpha;m) \in \theta$
$\wedge\ A \in \alpha(B)\ \wedge\ \mathit{link}(B{\to}A,[t,t+|m|])) \implies \mathsf{Receive}(A;t;m) \in \theta$
s4 $\forall A \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}, m \in \mathbb{M}, \alpha \in \Lambda.\ \ (\mathsf{Receive}(A;t;m) \in \theta \implies A \in V)$
$\wedge\ (\mathsf{Bcast}(A;t;m) \in \theta \implies A \in V_{\text{cor}})\ \wedge\ (\mathsf{Dcast}(A;t;\alpha;m) \in \theta \implies A \in V_{\text{adv}})$

Figure 2.1: Setting-feasibility rules.

point in time is for simplicity only, and we could easily generalize to arbitrary sets. Next, traces comprising the above events are defined.

Definition 4. A *trace* θ is a set of events.

We denote the set of all traces by Θ. Given a a setting $\mathcal{S}$, a protocol $\mathcal{P}$ and an adversary $\mathcal{A}$, we denote the set of traces feasible with respect to $\mathcal{S}$ by $\Theta_{\mathcal{S}}$, the set of traces feasible with respect to $\mathcal{S}$ and $\mathcal{P}$ by $\Theta_{\mathcal{S},\mathcal{P}}$, and the set of traces feasible with respect to $\mathcal{S}$, $\mathcal{P}$ and $\mathcal{A}$ by $\Theta_{\mathcal{S},\mathcal{P},\mathcal{A}}$.

2.1.4 Setting-Feasible Traces

The feasibility of a trace θ with respect to a setting $\mathcal{S} = \langle V, \mathit{loc}, \mathit{type}, \mathit{link}, \mathit{nlos}\rangle$ ensures a causal and strict time relation between send and receive events; it is formally defined by rules s1 – s4 (Figure 2.1). Rule s1 ensures that every message that is received was previously sent. Dually, rules s2 and s3 ensure that a message broadcasted or sent with a directional antenna is received by all nodes enabled to do so by the link relation and, in the latter case, the antenna pattern used. In other words, communication is causal (a receive is always preceded by a sent), and reliable *as long as the link is up*. Unreliability, expected and common in wireless communications, is modeled by the state of the link being *down*. Furthermore, these rules introduce a strict time relation between events, reflecting the propagation delay from A to B, across the channel, with speed $\mathbf{v}$: $(\mathit{dist}(A,B)+\mathit{nlos}(A,B))\mathbf{v}^{-1}$. Rule s4 is a technical one: It ensures that no communication events are performed by nodes not present in setting $\mathcal{S}$, and that Bcast and Dcast events are used exclusively by correct and adversarial nodes, respectively. Note that this is not a restriction of the adversary: $\mathsf{Bcast}(A;t;m)$ can be emulated (i.e., trigger exactly the same Receive events) by $\mathsf{Dcast}(A;t;\mathbb{R}^3;m)$.

2.1.5 Protocol-Feasible Traces

Intuitively, a trace is feasible with respect to protocol $\mathcal{P}$ if correct nodes behave according to a particular protocol $\mathcal{P}$. To formalize this, we first define the notion of a local view.

A trace is essentially a *global view* of the system execution. To describe what a node observes during a system execution, we use the notion of *local view*, primarily comprising a *local trace* composed of *local events*. We define these next.

Definition 5. A *local event* is one of the terms:

- $\mathsf{Bcast}(t;m)$

P1 $\forall A \in V_{\text{cor}}, t \in \mathbb{R}_{\geq 0}, m \in \mathbb{M}.\ \mathsf{Bcast}(A; t; m) \in \theta \implies \mathsf{Bcast}(m) \in \mathcal{P}(\theta||_{A,t})$
P2 $\forall A \in V_{\text{cor}}, t, t' \in \mathbb{R}_{\geq 0}, B, C \in \mathbb{V}.\ \mathsf{Neighbor}(A; t; B, C, t') \in \theta \implies \mathsf{Neighbor}(B, C, t') \in \mathcal{P}(\theta||_{A,t})$
P3 $\forall A \in V_{\text{cor}}.\ \forall t \in E_A.\ \epsilon \in \mathcal{P}(\theta||_{A,t})$
where $E_A = \mathbb{R}_{\geq 0} \setminus \mathit{start}(\theta|_A \cap I)$
and $I = \{\mathsf{Bcast}(t; m) \mid m \in \mathbb{M}, t \in \mathbb{R}_{\geq 0}\} \cup \{\mathsf{Neighbor}(t; B, C, t') \mid B, C \in \mathbb{V}, t, t' \in \mathbb{R}_{\geq 0}\}$

Figure 2.2: Protocol-feasibility rules for protocol model $\mathcal{P}$.

- $\mathsf{Receive}(t; m)$
- $\mathsf{Neighbor}(t; B, C, t')$

where $B, C \in \mathbb{V}$, $m \in \mathbb{M}$, $t, t' \in \mathbb{R}_{\geq 0}$. For a local event e, $\mathit{start}(e)$ is defined as in Definition 3.

Definition 6. A *local trace* is a set of local events. Given a node identifier $A \in \mathbb{V}$, time $t \geq 0$ and trace $\theta \in \Theta$, we calculate the *local trace of node A at time t in trace θ*, denoted $\theta|_{A,t}$, as follows:

$$\begin{aligned}\theta|_{A,t} = &\{\mathsf{Bcast}(t_1; m) \mid t_1 < t \ \wedge\ \mathsf{Bcast}(A; t_1; m) \in \theta\} \cup \\ &\{\mathsf{Receive}(t_1; m) \mid t_1 + |m| < t \ \wedge\ \mathsf{Receive}(A; t_1; m) \in \theta\} \cup \\ &\{\mathsf{Neighbor}(t_1; B, C, t') \mid t_1 < t \ \wedge\ \mathsf{Neighbor}(A; t_1; B, C, t') \in \theta\}\end{aligned}$$

We call $\theta|_{A,\infty}$ a *complete local trace* of A in θ and denote it shortly $\theta|_A$.

We identify two variants of the local view notion: an *T-local view*, as the basis for defining the class of time-based protocols, and an *TL-local view*, used to define the class of time- and location-based protocols.

Definition 7. Given a trace θ, an *T-local view of node A at time t in θ* is a tuple $\langle A, t, \theta|_{A,t}\rangle$; we denote it $\theta||_{A,t}$.

Definition 8. Given a trace θ and a setting $\mathcal{S}$, an *TL-local view of node A at time t in θ* is a tuple $\langle A, t, \mathit{loc}(A), \theta|_{A,t}\rangle$; we denote it $\theta||_{\mathcal{S},A,t}$, or $\theta||_{A,t}$ is setting $\mathcal{S}$ is clear from the context.

Note that $\mathcal{S}$ is part of Definition 8 as the location of node A is defined only within a specific setting. With the notion of the local view in hand, we can proceed with the definition of a protocol model. This definition captures the property of protocols essential to our investigation: the fact that protocol behavior depends *exclusively* on the local view of the node executing the protocol.

Definition 9. An *T(TL)-protocol model $\mathcal{P}$* is a function which given a T(TL)-local view $\theta||_{A,t}$, determines a finite, non-empty set of *actions*; an *action* is one of the terms: ϵ, $\mathsf{Bcast}(m)$ or $\mathsf{Neighbor}(B, C, t')$, where $m \in \mathbb{M}, B, C \in \mathbb{V}, t' \in \mathbb{R}_{\geq 0}$.

The interpretation of Bcast and Neighbor actions is natural. The ϵ action means that the node does not execute an event, with the exception of possible Receive event(s). Note that modeling the protocol output (i.e., the protocol model codomain) as a family of *sets of actions* allows for non-deterministic protocols.

The feasibility of a trace θ with respect to a protocol model $\mathcal{P}$ ensures that all correct nodes follow the protocol; it is formally defined by rules P1 – P3 (Figure 2.2). Rules P1 and P2 ensure that Bcast of Neighbor actions taken by a node are allowed by the protocol. Rule

A1 $\forall A \in V_{\text{adv}}, t \in \mathbb{R}_{\geq 0}, \alpha \in \Lambda, m \in \mathbb{M}.\ \mathsf{Dcast}(A; t; \alpha; m) \in \theta \implies$
$\exists B \in V_{\text{adv}}, \delta \geq \Delta_{\text{relay}} + \mathit{dist}(B, A)\mathbf{v}_{\text{adv}}^{-1}.\ \mathsf{Receive}(B; t - \delta; m) \in \theta$

Figure 2.3: Adversary-feasibility rule for adversary model $\mathcal{A}_{\Delta_{\text{relay}}}$.

P3, with E_A the set of all time instance in θ at which no event other than Receive happens at node A, ensures that the protocol allows for a node to not perform an action.

Note that our definition of a protocol model only requires that the behavior of the protocol is determined by the local view. This is much broader than a possible alternative approach, in which a protocol is modeled by a Turing machine. But as our definition is an over-approximation, the impossibility result remains valid for more realistic protocol models.

2.1.6 Adversary-Feasible Traces

For the purpose of the impossibility result, we consider first a relatively limited adversary, that is only capable of relaying messages. Note that a weak adversary model strengthens the impossibility result. We denote this model as $\mathcal{A}_{\Delta_{\text{relay}}}$, with the $\Delta_{\text{relay}} > 0$ parameter the minimum relaying delay introduced by an adversarial node; this delay is due to processing exclusively, it does not include propagation or transmission time.

Formally, the feasibility of trace θ with respect to $\mathcal{A}_{\Delta_{\text{relay}}}$ is defined by rule A1 in Figure 2.3. It states that every message sent by an adversarial node is necessarily a replay of a message m that either this or another adversarial node received. In addition, the delay between receiving m and re-sending it, or more precisely the difference between the start times of the corresponding events, needs to be at least Δ_{relay}, plus the propagation delay across the adversary channel (in case another adversarial node received the relayed message). This condition reflects the structure of the adversarial channel: Any two adversarial nodes can establish direct communication.

2.1.7 ND Specification

We consider two classes of properties ND protocols should satisfy. The first class pertains to *correctness* and is expressed through property ND1 (Figure 2.4): If two correct nodes[2] are declared neighbors at some time, then they must indeed be neighbors at that time. More precisely, there are two cases: (i) Node A can declare that B is its neighbor (i.e., A can receive messages from B) or (ii) A can declare that it is a neighbor of C (i.e., C can receive messages from A). In the latter case, property ND1 requires link (C, A) to be up at not exactly time t', but rather $\mathit{dist}(A, C) + \mathit{nlos}(A, C))\mathbf{v}^{-1}$ (propagation delay) after t'. As our model mandates that the link state is determined at the receiving end (node), if A declares that it is a neighbor of C at time t', a message sent by A at t would be indeed received by C. In other words, A is not forced to estimate the propagation delay to make a correct neighbor statement.

The second class of properties pertains to *availability* and is expressed through property ND2 (Figure 2.4), tailored to T-protocols. An additional notion needs to be introduced to formulate satisfiable availability properties: *neighbor discovery (ND) range*, $\mathbf{R} \in \mathbb{R}_{>0}$. Typically, $\mathbf{R}$ is equal to the *nominal communication range* for a given wireless medium and transceiver

[2] The requirement that B and C be correct is explained in Section 2.4.

ND1 $\forall \mathcal{S} \in \mathbb{S}, \theta \in \Theta_{\mathcal{S},\mathcal{P},\mathcal{A}}.\ \forall A, B, C \in V_{\text{cor}}, t, t' \in \mathbb{R}_{\geq 0}.\ \mathsf{Neighbor}(A; t; B, C, t') \in \theta \implies$
$(C = A \ \wedge\ \mathit{link}(B \to A, t')) \ \vee\ (B = A \ \wedge\ \mathit{link}(A \to C, t' + (\mathit{dist}(A, C) + \mathit{nlos}(A, C))\mathbf{v}^{-1}))$

ND2 $\forall d \in (0, \mathbf{R}].\ \forall A, B \in \mathbb{V}, A \neq B.\ \exists \mathcal{S} \in \mathbb{S}.\ V = V_{\mathbf{cor}} = \{A, B\} \ \wedge\ \mathit{dist}(A, B) = d$
$\wedge\ \mathit{link}(A \leftrightarrow B, \mathbb{R}_{\geq 0}) \ \wedge\ \exists \theta \in \Theta_{\mathcal{S},\mathcal{P},\mathcal{A}}.\ \mathsf{Neighbor}(A; t; B, A, t') \in \theta$

Figure 2.4: Basic ND properties.

technology, however, we use $\mathbf{R}$ more freely as the communication range[3] for which ND inferences are drawn. In other words, nodes at a communication range larger than $\mathbf{R}$ will not be required to declare each other neighbors.

Property ND2 requires that for every distance d in the desired ND range $\mathbf{R}$, there should be at least some setting, in which the protocol is able to conclude that a node is a neighbor (in some, not all executions); this setting should contain exactly two nodes at distance d, being neighbors, and both *correct*. The "two-nodes setting" requirement clarifies why we call this *two-party* ND. The ND2 property is the least that can be required from a usable two-party ND protocol: Indeed, a protocol not satisfying this property would be unable to conclude, for some distance(s) in the ND range, that nodes are neighbors. This makes the impossibility result in Section 2.2 more meaningful: impossibility with respect to a weak property implies impossibility for any stronger property.

2.2 Impossibility for T-protocols

We show in this section that no time-based protocol can solve the two-party neighbor discovery problem as specified by properties ND1 and ND2 in Figure 2.4. We base the proof on the fact, captured in Lemma 1, that it is impossible for a correct node to distinguish between different settings based on a T-local view. The impossibility result in Theorem 1 stems from showing two settings which are indistinguishable by a correct node, one in which two nodes are neighbors and one where they are not. We elaborate on the assumptions and implications of this result in Section 2.4.

We emphasize that the non-restricted form of the message space $\mathbb{M}$ encompasses all possible messages including, for example, time-stamps and any type of cryptography, thus contributing to the generality of the impossibility result.

Lemma 1. Let $\mathcal{P}$ be a T-protocol model, $\mathcal{S}$ and $\mathcal{S}'$ be settings such that $V_{\text{cor}} = V'_{\text{cor}}$, and $\theta \in \Theta_{\mathcal{S},\mathcal{P}}$ and $\theta' \in \Theta_{\mathcal{S}'}$ be traces such that local traces $\theta|_A = \theta'|_A$ for all $A \in V_{\text{cor}}$. Then θ' is feasible with respect to T-protocol model $\mathcal{P}$.

Proof. We need to prove that P1, P2, and P3 (Figure 2.2) hold for θ'.

- P1 Take any event $\mathsf{Bcast}(A; t; m) \in \theta'$. Based on Definition 6, we have that $\mathsf{Bcast}(t; m) \in \theta'|_A = \theta|_A$. Using Definition 6 again, we get that $\mathsf{Bcast}(A; t; m) \in \theta$. Since θ is feasible with respect to T-protocol model $\mathcal{P}$, P1 gives us $\mathsf{Bcast}(m) \in \mathcal{P}(\theta|_{A,t})$. Using again the assumption $\theta'|_A = \theta|_A$ we get the desired $\mathsf{Bcast}(m) \in \mathcal{P}(\theta'|_{A,t})$.
- P2 The proof is almost identical as for P1.
- P3 As θ satisfies P3, we have $\forall t \in E_A.\ \ \epsilon \in \mathcal{P}(\theta|_{A,t})$. Since $\theta'|_A = \theta|_A$, we have $E'_A = E_A$ and $\mathcal{P}(\theta'|_{A,t}) = \mathcal{P}(\theta|_{A,t})$, which proves that θ' satisfies P3.

□

[3] By "communication range" we understand the actual distance plus NLOS effects.

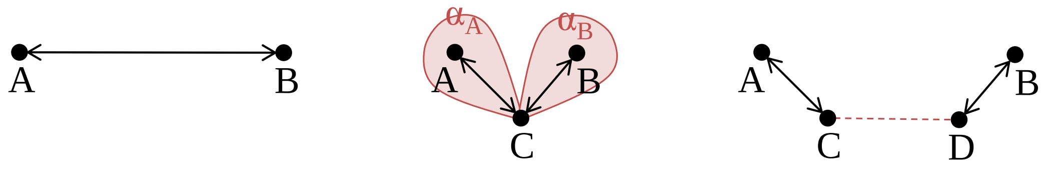

(a) $\mathcal{S}^a$ (b) $\mathcal{S}^b$ (c) $\mathcal{S}^c$

Figure 2.5: Settings used in the impossibility result proof. Settings $\mathcal{S}^a = \langle \{A, B\}, loc^a, type^a, link^a, nlos^a \rangle$, $\mathcal{S}^b = \langle \{A, B, C\}, loc^b, type^b, link^b, nlos^b \rangle$ and $\mathcal{S}^c = \langle \{A, B, C, D\}, loc^c, type^c, link^c, nlos^c \rangle$. In all settings, nodes A and B are correct, nodes C and D are adversarial. The location functions are such that $dist^b(A, C) + dist^b(B, C) + \mathbf{v}\Delta_{\mathrm{relay}} \le dist^a(A, B) \le \mathbf{R}$ and $dist^c(A, C) + dist^c(D, B) + \mathbf{v}\mathbf{v}_{\mathrm{adv}}{}^{-1} dist^c(C, D) + \mathbf{v}\Delta_{\mathrm{relay}} \le dist^a(A, B)$. The state of links does not change over time and is shown in the figure (lack of arrow means that the link is down). For all links and settings, $nlos = 0$. The dashed arrow in (c) denotes the adversarial channel. The shaded areas in (b) are antenna patterns α_A and α_B.

Theorem 1. If $\Delta_{\mathrm{relay}} < \frac{\mathbf{R}}{\mathbf{v}}$ then there exists no T-protocol model which satisfies ND1 and ND2 (Figure 2.4) for the adversary model $\mathcal{A}_{\Delta_{\mathrm{relay}}}$.

Proof. To prove that under the assumptions of the theorem no T-protocol model can satisfy both ND1 and ND2, we show that any T-protocol model that satisfies ND2 cannot satisfy ND1.

Take any T-protocol model $\mathcal{P}$ satisfying ND2. Pick some distance $d \ge \mathbf{v}\Delta_{\mathrm{relay}}$ in the ND range ($d \le \mathbf{R}$). Property ND2 guarantees the existence of a setting such as the one shown in Figure 2.5(a) (we denote it $\mathcal{S}^a$) and the existance of a trace $\theta \in \Theta_{\mathcal{S}^a, \mathcal{P}, \mathcal{A}_{\Delta_{\mathrm{relay}}}}$ such that $\mathsf{Neighbor}(A; t; B, t') \in \theta$. As θ is feasible with respect to setting $\mathcal{S}^a$, this trace has to be of the form:

$$\begin{aligned}\theta =& \big\{\mathsf{Bcast}(A; t_i; m_i) \mid i \in I_A\big\} \ \cup\ \big\{\mathsf{Receive}(B; t_i + \Delta; m_i) \mid i \in I_A\big\} \cup \\ & \big\{\mathsf{Bcast}(B; t_i; m_i) \mid i \in I_B\big\} \ \cup\ \big\{\mathsf{Receive}(A; t_i + \Delta; m_i) \mid i \in I_B\big\} \cup \\ & \big\{\mathsf{Neighbor}(A; t_i; A, B, t_i') \mid i \in J_A^A\big\} \ \cup\ \big\{\mathsf{Neighbor}(A; t_i; B, A, t_i') \mid i \in J_A^B\big\} \cup \\ & \big\{\mathsf{Neighbor}(B; t_i; A, B, t_i') \mid i \in J_B^A\big\} \ \cup\ \big\{\mathsf{Neighbor}(B; t_i; B, A, t_i') \mid i \in J_B^B\big\}\end{aligned}$$

where $\Delta = dist^a(A, B)\mathbf{v}^{-1}$, $t_i, t_i' \in \mathbb{R}_{\ge 0}$ and $I_A, I_B, J_A^A, J_A^B, J_B^A, J_B^B$ are pairwise disjoint index sets with $J_A^B \ne \emptyset$ (all the other index sets can be empty).

In setting $\mathcal{S}^b$, shown in Figure 2.5(b), we have $link(B \not\leftrightarrow A, \mathbb{R}_{\ge 0})$. Consider the following trace θ', which is is essentially the same as θ, but for node C relaying all the communication between nodes A and B:

$$\begin{aligned}\theta' =& \big\{\mathsf{Bcast}(A; t_i; m_i) \mid i \in I_A\big\} \ \cup\ \big\{\mathsf{Receive}(C; t_i + \delta_1; m_i) \mid i \in I_A\big\} \cup \\ & \big\{\mathsf{Dcast}(C; t_i + \delta_2; \alpha_B; m_i) \mid i \in I_A\big\} \ \cup\ \big\{\mathsf{Receive}(B; t_i + \Delta; m_i) \mid i \in I_A\big\} \cup \\ & \big\{\mathsf{Bcast}(B; t_i; m_i) \mid i \in I_B\big\} \ \cup\ \big\{\mathsf{Receive}(C; t_i + \delta_3; m_i) \mid i \in I_B\big\} \cup \\ & \big\{\mathsf{Dcast}(C; t_i + \delta_4; \alpha_A; m_i) \mid i \in I_B\big\} \ \cup\ \big\{\mathsf{Receive}(A; t_i + \Delta; B, m_i) \mid i \in I_B\big\} \cup \\ & \big\{\mathsf{Neighbor}(A; t_i; A, B, t_i') \mid i \in J_A^A\big\} \ \cup\ \big\{\mathsf{Neighbor}(A; t_i; B, A, t_i') \mid i \in J_A^B\big\} \cup \\ & \big\{\mathsf{Neighbor}(B; t_i; A, B, t_i') \mid i \in J_B^A\big\} \ \cup\ \big\{\mathsf{Neighbor}(B; t_i; B, A, t_i') \mid i \in J_B^B\big\}\end{aligned}$$

where $\delta_1 = dist^b(A,C)\mathbf{v}^{-1}$, $\delta_2 = \Delta - dist^b(C,B)\mathbf{v}^{-1}$, $\delta_3 = dist^b(B,C)\mathbf{v}^{-1}$ $\delta_4 = \Delta - dist^b(C,A)\mathbf{v}^{-1}$, and α_A, α_B are the antenna patterns shown in Figure 2.5(b).

It is simple to check that this trace is feasible with respect to setting $\mathcal{S}^b$. It is also feasible with respect to T-protocol model $\mathcal{P}$: This follows from Lemma 1, as $\theta|_{A,t} = \theta'|_{A,t}$ and $\theta|_{B,t} = \theta'|_{B,t}$. Finally, θ' is feasible with respect to the adversary model $\mathcal{A}_{\Delta_{\text{relay}}}$, because $\delta_2 - \delta_1 = \delta_4 - \delta_3 \geq \Delta_{\text{relay}}$. Therefore θ' belongs to $\Theta_{\mathcal{S}^b,\mathcal{P},\mathcal{A}_{\Delta_{\text{relay}}}}$ and together with $\mathcal{S}^b$ forms the counterexample that we were looking for: A concludes B is a neighbor whereas it is not. Thus, T-protocol model $\mathcal{P}$ does not satisfy ND1. As $\mathcal{P}$ was chosen arbitrarily, this concludes the proof. □

Note that in the above proof, the adversary did not need to make use of the adversarial channel. The same proof technique can be used with settings $\mathcal{S}^a$ and $\mathcal{S}^c$ (Figure 2.5) in which the adversary makes use of the adversarial channel, but does not need to resort to directional transmissions (i.e., relies on $\mathsf{Dcast}(.;.;\mathbb{R}^3;.)$). Furthermore, note that the above proof would be easy to extend to a model in which the correct nodes are allowed to use directional antennas, i.e., using directional antennas does not lift the impossibility result.

2.3 ND Protocols

In this section, we consider four types of ND protocols, and one representative protocol per type. We distinguish between (i) beacon-based protocols (*B-protocols*), represented by $\mathcal{P}^{\mathsf{B/T}}$ and $\mathcal{P}^{\mathsf{B/TL}}$, which require the transmission of one message by one of the protocol participants and synchronized clocks, and (ii) challenge-response protocols (*CR-protocols*), represented by $\mathcal{P}^{\mathsf{CR/T}}$ and $\mathcal{P}^{\mathsf{CR/TL}}$, which require a transmission of messages by both participants but *no* synchronized clocks. Within and across these categories, we distinguish protocols according to their capability to perform time measurements (T-protocols) or time measurements and location awareness (TL-protocols).

Fundamentally, beyond authentication mechanisms, all the ND protocols we consider measure the signal time-of-flight (ToF) between two nodes: B-protocols, with tightly synchronized clocks, are able to estimate ToF by transmitting a single beacon message, whereas CR-protocols require two messages, a challenge and a response, for the same purpose. T-protocols accept neighbor relations as valid if the ToF distance is below a threshold, whereas TL-protocols require this distance to be equal to the geographical distance calculated based on nodes locations.

2.3.1 Message Space

We define the message space $\mathbb{M}$ as follows. Any of the following is a message:

- an identifier $A \in \mathbb{V}$,
- a timestamp $t \in \mathbb{R}_{\geq 0}$,
- a location $l \in \mathbb{R}^3$,
- a nonce $n \in \mathsf{Nonces}$.

Moreover, two messages m_1, m_2 can be *concatenated* to form a message $\langle m_1, m_2 \rangle$. Furthermore, an *asymmetric authenticator* $\mathsf{auth}_A(m)$ and a *symmetric authenticator* $\mathsf{auth}_{AB}(m)$ where $A, B \in \mathbb{V}$ and $m \in \mathbb{M}$, are also messages.[4] We assume that symmetric authenticators are symmetric: $\mathsf{auth}_{AB}(m) = \mathsf{auth}_{BA}(m)$. Essentially, messages are terms, with the subterm relation is denoted by $\sqsubseteq$.

Every message m has a *duration* $|m| \in \mathbb{R}_{\geq 0}$, which determines the transmission delay (*not* including the propagation delay), reflecting the bit-rate of the underlying communication technology. We assume that message duration is preserved by concatenation, but not by an authenticator. For $m = \langle m_1, m_2, \ldots, m_k \rangle$, the duration is $|m| = |m_1| + |m_2| + \ldots + |m_k|$ and the *position* of m_i in m is $pos(m_i \sqsubseteq m) = |m_1| + \ldots + |m_{i-1}|$, with $pos(m_1 \sqsubseteq m) = 0$; in the case of multiple occurrences of $m' \sqsubseteq m$, $pos(m' \sqsubseteq m)$ gives the position of the first occurrence. When we use the duration function for any concatenated message, we omit the brackets: $|m_1, m_2, \ldots, m_k|$. Finally, we assume that the duration of identifiers, timestamps, locations, nonces and authenticators in $\mathbb{M}$ is upper-bounded by some constant.

2.3.2 Events

For the purpose of proving protocol correctness, we extend the set of available events defined in Definition 3 as follows:

Definition 10. An *event* is one of the following terms:

- $\mathsf{Receive}(A; t; m)$
- $\mathsf{Bcast}(A; t; m)$
- $\mathsf{Dcast}(A; t; \alpha; m)$
- $\mathsf{Fresh}(A; t; n)$
- $\mathsf{Neighbor}(A; t; B, C, t')$
- $\mathsf{NDstart}(A; t)$
- $\mathsf{NDstart}(A; t; B)$

where $A \in \mathbb{V}$ is the *active node*, $t \in \mathbb{R}_{\geq 0}$ is the event *start time*, denoted by $start(.)$, and $m \in \mathbb{M}$ is the transmitted/received message, $n \in \mathsf{Nonces}$ is a nonce, $\alpha \in \Lambda$ is an antenna pattern, $B, C \in \mathbb{V}$ are nodes, and $t' \in \mathbb{R}_{\geq 0}$ is a time instant.

Assuming that $m_1 \sqsubseteq m_2$, we use $\mathsf{Bcast}(A; t; m_1 \sqsubseteq m_2)$ to denote the event $\mathsf{Bcast}(A; t - pos(m_1 \sqsubseteq m_2); m_2)$; likewise for Dcast and Receive.

The interpretation of the events Receive, Bcast, Dcast and Neighbor remains unchanged. Fresh is used to declare that nonce n is (freshly) generated by A at time t or, in other words, that it was not sent before t. With NDstart, node A declares that an instance of a ND protocol has been initialized: either with a specific node B or with all neighbors.

2.3.3 Protocol-Feasible Traces

In Section 2.1.5 we have defined feasibility rules for arbitrary protocols. However, for reasoning about specific protocols, it is more convenient to define them with rules, rather than specifying a protocol model function, and applying the general rules in Figure 2.2.

The rules that specify this type of feasibility are protocol-dependent and are defined in Section 2.3.6. However, we introduce one general rule that dictates the behavior of correct nodes with respect to nonces. Rule F1 (Figure 2.6) guarantees that if a nonce n is freshly generated at time t then (i) the node that generated n will not broadcast it *before* t, (ii) any other correct node who broadcasts a message containing nonce n must have receive it (possibly in a different message) at least Δ_{relay} before broadcasting; this time difference is measured with respect to the positions of the nonce in the respective messages.

[4] Examples of asymmetric authenticators are digital signatures, and of symmetric authenticators: message authentication codes (MACs).

F1 $\forall A, B \in V_{\mathrm{cor}}, t_1, t_2 \in \mathbb{R}_{\geq 0}, n \in \mathsf{Nonces}, m_1 \in \mathbb{M}.\ \ n \sqsubseteq m_1 \ \wedge\ \mathsf{Fresh}(A; t_1; n) \in \theta$
$\wedge\ \mathsf{Bcast}(B; t_2; n \sqsubseteq m_1) \in \theta \implies (A = B \ \wedge\ t_2 \geq t_1) \ \vee\ (A \neq B \ \wedge\ \exists \delta \geq \Delta_{\mathrm{relay}}, m_2 \in \mathbb{M}.$
$n \sqsubseteq m_2 \ \wedge\ \mathsf{Receive}(B; t_2 - \delta; n \sqsubseteq m_2) \in \theta)$

A1 $\forall A \in V_{\mathrm{cor}}, B \in V_{\mathrm{adv}}, t_1, t_2 \in \mathbb{R}_{\geq 0}, \alpha \in \Lambda, n \in \mathsf{Nonces}, m_1 \in \mathbb{M}.\ \ n \sqsubseteq m_1 \ \wedge\ \mathsf{Fresh}(A; t_1; n) \in \theta$
$\wedge\ \mathsf{Dcast}(B; t_2; \alpha; n \sqsubseteq m_1) \in \theta \implies \exists C \in V_{\mathrm{adv}}, \delta \geq \Delta_{\mathrm{relay}} + \mathit{dist}(C, B)\mathbf{v}_{\mathrm{adv}}^{-1}, m_2 \in \mathbb{M}.\ \ n \sqsubseteq m_2$
$\wedge\ \mathsf{Receive}(C; t_2 - \delta; n \sqsubseteq m_2) \in \theta$

A2 $\forall A \in V_{\mathrm{adv}}, B \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}, m, m_0, m_1 \in \mathbb{M}, \alpha \in \Lambda.\ \ m = \mathsf{auth}_B(m_0) \sqsubseteq m_1$
$\wedge\ \mathsf{Dcast}(A; t; \alpha; m \sqsubseteq m_1) \in \theta \implies (B \in V_{\mathrm{adv}})$
$\vee\ (\exists C \in V_{\mathrm{adv}}, \delta \geq \Delta_{\mathrm{relay}} + \mathit{dist}(C, A)\mathbf{v}_{\mathrm{adv}}^{-1}, m_2 \in \mathbb{M}.\ \ m \sqsubseteq m_2 \ \wedge\ \mathsf{Receive}(C; t - \delta; m \sqsubseteq m_2) \in \theta)$

A3 $\forall A \in V_{\mathrm{adv}}, B, C \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}, m, m_0, m_1 \in \mathbb{M}, \alpha \in \Lambda.\ \ m = \mathsf{auth}_{BC}(m_0) \sqsubseteq m_1$
$\wedge\ \mathsf{Dcast}(A; t; \alpha; m \sqsubseteq m_1) \in \theta \implies (B \in V_{\mathrm{adv}}) \ \vee\ (C \in V_{\mathrm{adv}})$
$\vee\ (\exists D \in V_{\mathrm{adv}}, \delta \geq \Delta_{\mathrm{relay}} + \mathit{dist}(D, A)\mathbf{v}_{\mathrm{adv}}^{-1}, m_2 \in \mathbb{M}.\ \ m \sqsubseteq m_2 \ \wedge\ \mathsf{Receive}(D; t - \delta; m \sqsubseteq m_2) \in \theta)$

Figure 2.6: Adversary- and common protocol-feasibility rules.

$\mathsf{ND2}^{\mathsf{B/T}}$ $\forall \mathcal{S} \in \mathbb{S}, \theta \in \Theta_{\mathcal{S},\mathcal{P},\mathcal{A}}.\ \ \forall A, B \in \mathbb{V}_{\mathrm{cor}}, t \in \mathbb{R}_{\geq 0}.\ \ \mathsf{NDstart}(A; t) \in \theta \ \wedge\ \mathit{link}(A \rightarrow B, [t, t + T_{\mathcal{P}}])$
$\wedge\ \mathit{dist}(A, B) + \mathit{nlos}(A, B) \leq \mathbf{R} \implies \exists t' \in [t, \infty), t'' \in [t, t + T_{\mathcal{P}}].\ \ \mathsf{Neighbor}(B; t'; A, B, t'') \in \theta$

$\mathsf{ND2}^{\mathsf{B/TL}}$ $\forall \mathcal{S} \in \mathbb{S}, \theta \in \Theta_{\mathcal{S},\mathcal{P},\mathcal{A}}.\ \ \forall A, B \in \mathbb{V}_{\mathrm{cor}}, t \in \mathbb{R}_{\geq 0}.\ \ \mathsf{NDstart}(A; t; B) \in \theta \ \wedge\ \mathit{link}(A \rightarrow B, [t, t + T_{\mathcal{P}}])$
$\wedge\ \mathit{nlos}(A, B) = 0 \implies \exists t' \in [t, \infty), t'' \in [t, t + T_{\mathcal{P}}].\ \ \mathsf{Neighbor}(B; t'; A, B, t'') \in \theta$

$\mathsf{ND2}^{\mathsf{CR/T}}$ $\forall \mathcal{S} \in \mathbb{S}, \theta \in \Theta_{\mathcal{S},\mathcal{P},\mathcal{A}}.\ \ \forall A, B \in \mathbb{V}_{\mathrm{cor}}, t \in \mathbb{R}_{\geq 0}.\ \ \mathsf{NDstart}(A; t; B) \in \theta \ \wedge\ \mathit{link}(A \leftrightarrow B, [t, t + T_{\mathcal{P}}])$
$\wedge\ \mathit{dist}(A, B) + \mathit{nlos}(A, B) \leq \mathbf{R} \implies$
$\exists t_1, t_2 \in [t, \infty), t', t'' \in [t, t + T_{\mathcal{P}}].\ \ \mathsf{Neighbor}(A; t_1; A, B, t') \in \theta \wedge \mathsf{Neighbor}(A; t_2; B, A, t'') \in \theta$

$\mathsf{ND2}^{\mathsf{CR/TL}}$ $\forall \mathcal{S} \in \mathbb{S}, \theta \in \Theta_{\mathcal{S},\mathcal{P},\mathcal{A}}.\ \ \forall A, B \in \mathbb{V}_{\mathrm{cor}}, t \in \mathbb{R}_{\geq 0}.\ \ \mathsf{NDstart}(A; t; B) \in \theta \ \wedge\ \mathit{link}(A \leftrightarrow B, [t, t + T_{\mathcal{P}}])$
$\wedge\ \mathit{nlos}(A, B) = 0 \implies \exists t_1, t_2 \in [t, \infty), t', t'' \in [t, t + T_{\mathcal{P}}].\ \ \mathsf{Neighbor}(A; t_1; A, B, t') \in \theta$
$\wedge\ \mathsf{Neighbor}(A; t_2; B, A, t'') \in \theta$

Figure 2.7: ND availability properties.

2.3.4 Adversary-Feasible Traces

For the purpose of reasoning about protocol security, we consider an adversary model $\mathcal{A}^{\mathcal{P}}_{\Delta_{\mathrm{relay}}}$, stronger than the model defined in Section 2.1.6. Intuitively, adversarial nodes are allowed to send arbitrary messages, except for messages which would violate properties of authenticators or freshness; these have to be relayed with the relaying delay at least Δ_{relay}.

A trace θ is feasible with respect to $\mathcal{A}^{\mathcal{P}}_{\Delta_{\mathrm{relay}}}$ if rules A1 - A3 (Figure 2.6) are satisfied. Rules A2 and A3 deal with authenticators: An adversarial node is allowed to send a message containing arbitrary authenticators, as long as these authenticators can be generated by an adversarial node (itself or other). This implies that adversarial nodes can share cryptographic keys or any material used for authentication. Furthermore, rules A2 and A3 reflect that the adversary cannot forge authenticated messages: Any message sent by an adversarial node that contains a correct node authenticator must be relayed. In other words, some (possibly the same) adversarial node must have received a message containing this authenticator earlier, at least Δ_{relay} plus the propagation delay between the two nodes over the adversarial channel. This condition reflects the structure of the adversarial channel: Any two adversarial nodes can establish direct communication. Rule A1 is similar to A2, but it is responsible for freshness: An adversary sending a message with a nonce generated by a correct can only be relaying the message (nonce). In this sense rule A1 is an adversarial equivalent of rule F1.

```
1: on $\mathsf{NDstart}(A; t_1)$
2:     $\mathsf{Bcast}(A; t_1; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle)$
3: on $\mathsf{Receive}(B; t_2; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle)$
4:     if $t_2 - t_1 \le \mathbf{Rv}^{-1}$
5:         $\mathsf{Neighbor}(B; t_2 + |A, t_1, \mathsf{auth}_A(t_1)|; A, B, t_2)$
```

Figure 2.8: Pseudo-code for protocol $\mathcal{P}^{\mathsf{B/T}}$.

P1 $\forall A \in V_{\text{cor}}, t_1 \in \mathbb{R}_{\ge 0}.\ \ \mathsf{NDstart}(A; t_1) \in \theta \implies \mathsf{Bcast}(A; t_1; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle) \in \theta$

P2 $\forall A \in V_{\text{cor}}, B \in \mathbb{V}, t_1, t \in \mathbb{R}_{\ge 0}, m \in \mathbb{M}.\ \ \mathsf{auth}_B(t) \sqsubseteq m \ \wedge\ \mathsf{Bcast}(A; t_1; m) \in \theta$
$\implies m = \langle A, t_1, \mathsf{auth}_A(t_1)\rangle$

P3 $\forall B \in V_{\text{cor}}, A \in \mathbb{V}, t_1, t_2 \in \mathbb{R}_{\ge 0}.\ \ \mathsf{Receive}(B; t_2; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle) \in \theta \ \wedge\ t_2 - t_1 \le \mathbf{Rv}^{-1}$
$\implies \mathsf{Neighbor}(B; t_2 + |A, t_1, \mathsf{auth}_A(t_1)|; A, B, t_2) \in \theta$

P4 $\forall B \in V_{\text{cor}}, A, C \in \mathbb{V}, t_2, t \in \mathbb{R}_{\ge 0}.\ \ \mathsf{Neighbor}(B; t; A, C, t_2) \in \theta \implies C = B$
$\wedge\ \exists t_1 \in \mathbb{R}_{\ge 0}.\ \ \mathsf{Receive}(B; t_2; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle) \in \theta \ \wedge\ t_2 - t_1 \le \mathbf{Rv}^{-1}$
$\wedge\ t = t_2 + |A, t_1, \mathsf{auth}_A(t_1)|$

Figure 2.9: Rules defining protocol $\mathcal{P}^{\mathsf{B/T}}$.

2.3.5 ND Specification

The correctness property ND1 introduced in Figure 2.4 is also part of the ND specification that we use for reasoning about concrete protocol security. However, we provide stronger, and class specific *availability* properties. Informally, if two nodes are neighbors for a long enough, protocol-specific time $T_{\mathcal{P}}$, the protocol must declare them neighbors.

Figure 2.7 displays ND2 properties for all types of protocols we consider. These properties differ in four aspects, one depending on whether the protocol is T or TL, whereas the other three aspects depending on the protocol is beacon or challenge-response. The first aspect is the NDstart event: For CR-protocols, a particular neighbor B with which ND is started is specified, whereas no such specification is necessary for B-protocols. Second, it may be required that link (A, B) be up in only one direction (B-protocols) or both directions (CR-protocols). Third, for T-protocols an upper-bound on propagation distance in enforced ($\mathit{dist}(A, B) + \mathit{nlos}(A, B) \le \mathbf{R}$), whereas for TL-protocols line-of-sight propagation is required ($\mathit{nlos}(A, B) = 0$). Forth, different forms of neighbor declaration are possible. The node making the declaration might be the same as (CR-protocols) or different (B-protocols) from the one initiating the ND protocol. Moreover the declaration might be uni-directional (B-protocols) or bi-directional (CR-protocols).

2.3.6 Protocol Definitions

The protocols are formally defined with rules such at the ones in Figure 2.9. To make the presentation more approachable, we present the protocols informally in the form of pseudo-code, and describe how the rules model the behavior of the protocol. The pseudo-code is divided into *blocks* starting with a *triggering* event (**on** clause). If the triggering event occurs, the body of the block is executed, i.e., other events take place.

We start with a simple B/T-protocol we denote $\mathcal{P}^{\mathsf{B/T}}$, which is essentially the *temporal packet leash* protocol proposed by Hu, Perrig and Johnson in [74]. The pseudo-code is shown in Figure 2.8, the rules defining the protocol are presented in Figure 2.9. Block 1-2 describes the behavior after the ND protocol is started at node A (e.g., by a higher layer protocol);

```
on $\mathsf{NDstart}(A; t_1; B)$
    $\mathsf{Fresh}(A; t_1 + |B|; n_1)$
    $\mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle)$
on $\mathsf{Receive}(B; t; \langle B, n_1 \rangle)$
    $\mathsf{Fresh}(B; t + \Delta; n_2)$
    $\mathsf{Bcast}(B; t + \Delta; \langle n_2 \rangle)$
    let $\tau > \Delta$
    $\mathsf{Bcast}(B; t + \tau; \langle loc(B), \mathsf{auth}_B(n_1, n_2, loc(B)) \rangle)$
on $\mathsf{Receive}(A; t; \langle l, \mathsf{auth}_B(n_1, n_2, l) \rangle)$
    if occurred $\mathsf{Fresh}(A; t_1 + |B|; n_1)$
    if occurred $\mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle)$
    if occurred $\mathsf{Receive}(A; t_2; \langle n_2 \rangle)$
    if $\mathbf{v}(t_2 - t_1 - \Delta) = 2d(loc(A), l)$
        $\mathsf{Neighbor}(A; t + |l, \mathsf{auth}_B(n_1, n_2, l)|; A, B, t_1)$
        $\mathsf{Neighbor}(A; t + |l, \mathsf{auth}_B(n_1, n_2, l)|; B, A, t_2)$
```

Figure 2.10: Pseudo-code for protocol $\mathcal{P}^{\mathsf{CR/TL}}$.

P1 and P2 are the two rules that correspond to this block. Block 3-5 describes the behavior of a node after it receives a beacon message, and it is modeled by rules P3 and P4. Rule P1 is straightforward: if ensures that if the triggering event of block 1-2, $\mathsf{NDstart}(A; t_1)$, occurs in the trace, the event in the body of the block also occurs. In the same fashion, rule P3 is defined for block 3-5, with an additional condition coming from the **if** clause.

These two rules are already sufficient to prove the ND2 property, but they only define half of aspects of the the protocol functionality. Indeed, nothing prevents a node running this protocol from making arbitrary neighbor declarations. Rule P4 addresses this, stating that if a node makes a neighbor declaration, this has to be done according to block 3-5, i.e., the node had to receive a "fresh enough" beacon message. Only one aspect remains: Correct nodes are still allowed to broadcast arbitrary messages, including bogus beacon messages. This is addressed by rule P2. To motivate the definition of P2, let us consider an alternative rule would still be coherent with the pseudo-code: If a correct node broadcasts a message at time t_1, this message is $\langle A, t_1, \mathsf{auth}_A(t_1) \rangle$. We can prove that such a defined protocol satisfies the ND specification. However, this is a weak result, precisely because that rule states that correct nodes cannot send any other messages than beacons. If the ND protocol were used along with or by any other protocol, obviously using other forms of messages, the result would no longer apply. To circumvent this undesired composability restriction, rule P2 is defined as follows. It only requires that if a correct node broadcasts at t_1 a message m of a *particular form*, i.e., containing $\mathsf{auth}_B(t)$ as a subterm, then $m = \langle A, t_1, \mathsf{auth}_A(t_1) \rangle$. Hence, rule P2 gives a much less restrictive condition on protocols that can be securely composed with $\mathcal{P}^{\mathsf{B/T}}$: basically, it mandates that any other protocol does not use authenticated timestamps of this form.[5] Rule P4, in terms of composability, implies that the node cannot run any other ND protocol (i.e., a protocol making neighbor declarations), but we do not see this as a real restriction.

Next, we describe $\mathcal{P}^{\mathsf{CR/TL}}$, a CR/TL-protocol (pseudo-code Figure 2.10, rules Figure 2.11). This protocol has a practical design twist: As authentication of a message can be a time-

[5] If this would pose a problem, the protocol can be modified, by e.g., authenticating a timestamp concatenated with some constant in place of simple the timestamp.

P1 $\forall A \in V_{\text{cor}}, B \in \mathbb{V}, t_1 \in \mathbb{R}_{\geq 0}.\ \ \mathsf{NDstart}(A; t_1; B) \in \theta \implies \exists n_1 \in \mathsf{Nonces}.$
$\mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta \ \wedge\ \mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta$

P2 $\forall B \in V_{\text{cor}}, t \in \mathbb{R}_{\geq 0}, n_1 \in \mathsf{Nonces}.\ \ \mathsf{Receive}(B; t; \langle B, n_1 \rangle) \in \theta \implies \exists n_2 \in \mathsf{Nonces}, \tau > \Delta.$
$\mathsf{Fresh}(B; t + \Delta; n_2) \in \theta \ \wedge\ \mathsf{Bcast}(B; t + \Delta; \langle n_2 \rangle) \in \theta$
$\wedge\ \mathsf{Bcast}(B; t + \tau; \langle loc(B), \mathsf{auth}_B(n_1, n_2, loc(B)) \rangle) \in \theta$

P3 $\forall B \in V_{\text{cor}}, C \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}, n_1, n_2 \in \mathsf{Nonces}, l \in \mathbb{R}^3, m \in \mathbb{M}.\ \ \mathsf{auth}_C(n_1, n_2, l) \sqsubseteq m$
$\wedge\ \mathsf{Bcast}(B; t; m) \in \theta \implies \exists \tau > 0.\ \ m = \langle loc(B), \mathsf{auth}_B(n_1, n_2, loc(B)) \rangle$
$\wedge\ \mathsf{Receive}(B; t - \tau - \Delta; \langle B, n_1 \rangle) \in \theta \ \wedge\ \mathsf{Fresh}(B; t - \tau; n_2) \in \theta \ \wedge\ \mathsf{Bcast}(B; t - \tau; \langle n_2 \rangle) \in \theta$

P4 $\forall A \in V_{\text{cor}}, B \in \mathbb{V}, n_1, n_2 \in \mathsf{Nonces}, t_1, t_2, t \in \mathbb{R}_{\geq 0}, l \in \mathbb{R}^3.$
$\mathsf{Receive}(A; t; \langle l, \mathsf{auth}_B(n_1, n_2, l) \rangle) \in \theta \ \wedge\ \mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta$
$\wedge\ \mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta \wedge \mathsf{Receive}(A; t_2; \langle n_2 \rangle) \in \theta \wedge \mathbf{v}(t_2 - t_1 - \Delta) = 2d(loc(A), l) \implies$
$\mathsf{Neighbor}(A; t + |l, \mathsf{auth}_B(n_1, n_2, l)|; A, B, t_1) \in \theta$
$\wedge\ \mathsf{Neighbor}(A; t + |l, \mathsf{auth}_B(n_1, n_2, l)|; B, A, t_2) \in \theta$

P5 $\forall A \in V_{\text{cor}}, B, C \in \mathbb{V}, t, t_0 \in \mathbb{R}_{\geq 0}.\ \ \mathsf{Neighbor}(A; t; B, C, t_0) \in \theta \implies$
$(C = A \ \wedge\ \exists n_1, n_2 \in \mathsf{Nonces}, t_1 \in \mathbb{R}_{\geq 0}, l \in \mathbb{R}^3.\ \ \mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta$
$\wedge\ \mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta \ \wedge\ \mathsf{Receive}(A; t_0; \langle n_2 \rangle) \in \theta$
$\wedge\ \mathsf{Receive}(A; t - |l, \mathsf{auth}_B(n_1, n_2, l)|; \langle l, \mathsf{auth}_B(n_1, n_2, l) \rangle) \in \theta$
$\wedge\ \mathbf{v}(t_0 - t_1 - \Delta) = 2d(loc(A), l))$
$\vee$
$(B = A \ \wedge\ \exists n_1, n_2 \in \mathsf{Nonces}, t_2 \in \mathbb{R}_{\geq 0}, l \in \mathbb{R}^3.\ \ \mathsf{Fresh}(A; t_0 + |C|; n_1) \in \theta$
$\wedge\ \mathsf{Bcast}(A; t_0; \langle C, n_1 \rangle) \in \theta \ \wedge\ \mathsf{Receive}(A; t_2; \langle n_2 \rangle) \in \theta$
$\wedge\ \mathsf{Receive}(A; t - |l, \mathsf{auth}_C(n_1, n_2, l)|; \langle l, \mathsf{auth}_C(n_1, n_2, l) \rangle) \in \theta$
$\wedge\ \mathbf{v}(t_2 - t_0 - \Delta) = 2d(loc(A), l))$

Figure 2.11: Rules defining protocol $\mathcal{P}^{\mathsf{CR/TL}}$

consuming process, in this protocol we remove it from the time-critical ToF estimation phase. Otherwise, if the response needs too much time to be calculated, the clock of the challenging node can drift beyond an acceptable accuracy level. A protocol parameter $\Delta \in \mathbb{R}_{\geq 0}$ determines exactly how long after the challenge reception a node replies.

Note that we assume that a node keeps track of all the events it observes, and it can always refer to this 'history,' as in 10-12. Note also that there is no explicit block responsible for receiving the $\langle n_2 \rangle$ response sent by B in 06, because in this case node A does not take any action other than recording the event occurrence, for later reference in line 11.

Considering again that "triggering event implies block body events," rule P1 is defined for block 01-03, P2 for block 04-08, and P4 for block 09-15. We do not define rules that restrict the occurrence of Fresh events (in lines 02 and 05) or the form of broadcasted messages (in lines 03 and 06), so that there is no obstacle for composability. For line 08, rule P3 is defined: If a node broadcasts a message m containing an authenticator of the form $\mathsf{auth}_B(n_1, n_2, l)$, then m is precisely the message defined in line 08, and all the other events from block 04-08 occur. Finally, rule P5 is defined based on block 09-15. There is only one rule, despite two Neighbor events in lines 14 and 15, because both events match the universally quantified Neighbor event in P5; The rule uses a disjunction, as there are (small) timing differences in the node behavior depending on which of these two event is considered.

The pseudo-code defining the two remaining protocols ($\mathcal{P}^{\mathsf{B/TL}}$ and $\mathcal{P}^{\mathsf{CR/T}}$) is shown in Figure 2.12 and in Figure 2.14. These protocols are similar to the two previous protocols, hence we omit a detail explanation. We note, however, that opposite to the other protocols, $\mathcal{P}^{\mathsf{B/TL}}$ relies on symmetric authenticators. The purpose of this is to demonstrate that the protocols can be modified to work with symmetric cryptography. There is no specific reason

```
on $\mathsf{NDstart}(A;t_1;B)$
    $\mathsf{Bcast}(A;t_1;\langle A,t_1,loc(A),\mathsf{auth}_{AB}(t_1,loc(A))\rangle)$
on $\mathsf{Receive}(B;t_2;\langle A,t_1,l,\mathsf{auth}_{AB}(t_1,l)\rangle)$
    if $t_2-t_1=d(loc(B),l)\mathbf{v}^{-1}$
        $\mathsf{Neighbor}(B;t_2+|A,t_1,l,\mathsf{auth}_{AB}(t_1,l)|;A,B,t_2)$
```

Figure 2.12: Pseudo-code for protocol $\mathcal{P}^{\mathsf{B/TL}}$.

P1 $\forall A\in V_{\text{cor}},t_1\in\mathbb{R}_{\geq 0}.\ \mathsf{NDstart}(A;t_1;B)\in\theta\implies$ $\mathsf{Bcast}(A;t_1;\langle A,t_1,loc(A),\mathsf{auth}_{AB}(t_1,loc(A))\rangle)\in\theta$

P2 $\forall A\in V_{\text{cor}},B,C\in\mathbb{V},t_1,t\in\mathbb{R}_{\geq 0},l\in\mathbb{R}^3,m\in\mathbb{M}.\ \mathsf{auth}_{CB}(t,l)\sqsubseteq m\ \wedge\ \mathsf{Bcast}(A;t_1;m)\in\theta$ $\implies m=\langle A,t_1,loc(A),\mathsf{auth}_{AB}(t_1,loc(A))\rangle$

P3 $\forall B\in V_{\text{cor}},A\in\mathbb{V},t_1,t_2\in\mathbb{R}_{\geq 0},l\in\mathbb{R}^3.\ \mathsf{Receive}(B;t_2;\langle A,t_1,l,\mathsf{auth}_{AB}(t_1,l)\rangle)\in\theta$ $\wedge\ t_2-t_1=d(loc(B),l)\mathbf{v}^{-1}\implies\mathsf{Neighbor}(B;t_2+|A,t_1,l,\mathsf{auth}_{A}(t_1,l)|;A,B,t_2)\in\theta$

P4 $\forall B\in V_{\text{cor}},A,C\in\mathbb{V},t_2,t\in\mathbb{R}_{\geq 0}.\ \mathsf{Neighbor}(B;t;A,C,t_2)\in\theta\implies C=B$ $\wedge\ \exists t_1\in\mathbb{R}_{\geq 0},l\in\mathbb{R}^3.\ \mathsf{Receive}(B;t_2;\langle A,t_1,l,\mathsf{auth}_{AB}(t_1,l)\rangle)\in\theta\wedge t_2-t_1=d(loc(B),l)\mathbf{v}^{-1}$ $\wedge\ t=t_2+|A,t_1,l,\mathsf{auth}_{AB}(t_1,l)|$

Figure 2.13: Rules defining protocol $\mathcal{P}^{\mathsf{B/TL}}$

why we chose $\mathcal{P}^{\mathsf{B/TL}}$ for this demonstration.

The rules defining the protocols (Figure 2.13, Figure 2.15) are also very similar to the rules for $\mathcal{P}^{\mathsf{B/T}}$ and $\mathcal{P}^{\mathsf{CR/TL}}$. The only noteworthy difference is the rule P2 for $\mathcal{P}^{\mathsf{CR/T}}$, which has no equivalent rule in the definition on $\mathcal{P}^{\mathsf{CR/TL}}$. This rule states that whenever a node sends a the challenge message, the message needs to be fresh (we could also demand that the NDstart event is in the trace, but we omit that for simplicity). This restriction is necessary to prove the correctness of the protocol.

2.3.7 Proofs

We prove that protocols defined in Section 2.3.6 satisfy the ND1 property and the appropriate ND2 properties. First, we present three simple lemmas which facilitate subsequent proofs. Lemma 2 deals with freshness and is an extension of rules S2 and F1, whereas Lemma 3 and Lemma 4 deal with authenticators, extending A2 and A3, respectively. We start with the

```
on $\mathsf{NDstart}(A;t_1;B)$
    $\mathsf{Fresh}(A;t_1+|B|;n_1)$
    $\mathsf{Bcast}(A;t_1;\langle B,n_1\rangle)$
on $\mathsf{Receive}(B;t;\langle B,n_1\rangle)$
    $\mathsf{Bcast}(B;t+\Delta;\langle\mathsf{auth}_B(n_1)\rangle)$
on $\mathsf{Receive}(A;t_2;\langle\mathsf{auth}_B(n_1)\rangle)$
    if occurred $\mathsf{Fresh}(A;t_1+|B|;n_1)$
    if occurred $\mathsf{Bcast}(A;t_1;\langle B,n_1\rangle)$
    if $\mathbf{v}(t_2-t_1-\Delta)\leq 2\mathbf{R}$
        $\mathsf{Neighbor}(A;t_2+|\mathsf{auth}_B(n_1)|;A,B,t_1)$
        $\mathsf{Neighbor}(A;t_2+|\mathsf{auth}_B(n_1)|;B,A,t_2)$
```

Figure 2.14: Pseudo-code for protocol $\mathcal{P}^{\mathsf{CR/T}}$.

P1 $\forall A \in V_{\text{cor}}, B \in \mathbb{V}, t_1 \in \mathbb{R}_{\geq 0}.\ \ \mathsf{NDstart}(A; t_1; B) \in \theta \implies$
$\exists n_1 \in \mathsf{Nonces}.\ \ \mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta \ \wedge\ \mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta$

P2 $\forall A \in V_{\text{cor}}, B \in \mathbb{V}, t_1 \in \mathbb{R}_{\geq 0}, n_1 \in \mathsf{Nonces}.\ \ \mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta \implies$
$\mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta$

P3 $\forall B \in V_{\text{cor}}, A \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}, n_1 \in \mathsf{Nonces}.\ \ \mathsf{Receive}(B; t; \langle B, n_1 \rangle) \in \theta \implies$
$\mathsf{Bcast}(B; t + \Delta; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta$

P4 $\forall B \in V_{\text{cor}}, C \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}, n_1 \in \mathsf{Nonces}, m \in \mathbb{M}.\ \ \mathsf{auth}_C(n_1) \sqsubseteq m$
$\wedge\ \mathsf{Bcast}(B; t; m) \in \theta \implies m = \langle \mathsf{auth}_B(n_1) \rangle \ \wedge\ \mathsf{Receive}(B; t - \Delta; \langle B, n_1 \rangle) \in \theta$

P5 $\forall A \in V_{\text{cor}}, B \in \mathbb{V}, n_1 \in \mathsf{Nonces}, t_1, t_2 \in \mathbb{R}_{\geq 0}.$
$\mathsf{Receive}(A; t_2; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta \ \wedge\ \mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta \ \wedge\ \mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta$
$\wedge\ \mathbf{v}(t_2 - t_1 - \Delta) \leq 2\mathbf{R} \implies$
$\mathsf{Neighbor}(A; t_2 + |\mathsf{auth}_B(n_1)|; A, B, t_1) \in \theta \ \wedge\ \mathsf{Neighbor}(A; t_2 + |\mathsf{auth}_B(n_1)|; B, A, t_2) \in \theta$

P6 $\forall A \in V_{\text{cor}}, B, C \in \mathbb{V}, t, t_0 \in \mathbb{R}_{\geq 0}.\ \ \mathsf{Neighbor}(A; t; B, C, t_0) \in \theta \implies$
$(C = A \ \wedge\ \exists n_1 \in \mathsf{Nonces}, t_1 \in \mathbb{R}_{\geq 0}.\ \ \mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta \ \wedge\ \mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta$
$\wedge\ \mathsf{Receive}(A; t_0; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta \ \wedge\ \mathbf{v}(t_0 - t_1 - \Delta) \leq 2\mathbf{R})$
$\vee$
$(B = A \ \wedge\ \exists n_1 \in \mathsf{Nonces}, t_2 \in \mathbb{R}_{\geq 0}.\ \ \mathsf{Fresh}(A; t_0 + |C|; n_1) \in \theta \ \wedge\ \mathsf{Bcast}(A; t_0; \langle C, n_1 \rangle) \in \theta$
$\wedge\ \mathsf{Receive}(A; t_2; \langle \mathsf{auth}_C(n_1) \rangle) \in \theta \ \wedge\ \mathbf{v}(t_2 - t_0 - \Delta) \leq 2\mathbf{R})$

Figure 2.15: Rules defining protocol $\mathcal{P}^{\mathsf{CR/T}}$

L1 $\forall A \in V_{\text{cor}}, B \in V, t_1, t_2 \in \mathbb{R}_{\geq 0}, \alpha \in \Lambda, n \in \mathsf{Nonces}, m \in \mathbb{M}.\ \ A \neq B \ \wedge\ n \sqsubseteq m \ \wedge\ \mathsf{Fresh}(A; t_1; n) \in \theta$
$\wedge\ (\mathsf{Bcast}(B; t_2; n \sqsubseteq m) \in \theta \ \vee\ \mathsf{Dcast}(B; t_2; \alpha; n \sqsubseteq m) \in \theta) \implies t_2 \geq t_1 + \mathit{dist}(A, B)\mathbf{v}_{\text{adv}}^{-1} + \Delta_{\text{relay}}$

L2 $\forall A \in V_{\text{adv}}, B \in V, t \in \mathbb{R}_{\geq 0}, m, m_0, m_1 \in \mathbb{M}, \alpha \in \Lambda.\ \ m = \mathsf{auth}_B(m_0) \sqsubseteq m_1$
$\wedge\ \mathsf{Dcast}(A; t; \alpha; m \sqsubseteq m_1) \in \theta \implies (B \in V_{\text{adv}})$
$\vee\ (\exists C \in V_{\text{cor}}, \delta \geq \Delta_{\text{relay}} + \mathit{dist}(C, A)\mathbf{v}_{\text{adv}}^{-1}, m_2 \in \mathbb{M}.\ \ m \sqsubseteq m_2 \ \wedge\ \mathsf{Bcast}(C; t - \delta; m \sqsubseteq m_2) \in \theta)$

L3 $\forall A \in V_{\text{adv}}, B, C \in V, t \in \mathbb{R}_{\geq 0}, m, m_0, m_1 \in \mathbb{M}, \alpha \in \Lambda.\ \ m = \mathsf{auth}_{BC}(m_0) \sqsubseteq m_1$
$\wedge\ \mathsf{Dcast}(A; t; \alpha; m \sqsubseteq m_1) \in \theta \implies (B \in V_{\text{adv}}) \ \vee\ (C \in V_{\text{adv}})$
$\vee\ (\exists D \in V_{\text{cor}}, \delta \geq \Delta_{\text{relay}} + \mathit{dist}(D, A)\mathbf{v}_{\text{adv}}^{-1}, m_2 \in \mathbb{M}.\ \ m \sqsubseteq m_2 \ \wedge\ \mathsf{Bcast}(D; t - \delta; m \sqsubseteq m_2) \in \theta)$

Figure 2.16: Rules for Lemmas.

proof of Lemma 3, because the proof of Lemma 2 follows a similar patter, but is slightly more involved. The proof of Lemma 4 is virtually identical to the proof of Lemma 3, and thus omitted.

Lemmas

Lemma 2. Rule L1 (Figure 2.16) holds for every trace θ feasible with respect to the adversary model $\mathcal{A}^{\mathcal{P}}_{\Delta_{\text{relay}}}$ $(\Delta_{\text{relay}} > 0)$, some setting $\mathcal{S}$ and rule F1

Lemma 3. Rule L2 (Figure 2.16) holds for every trace θ feasible with respect to the adversary model $\mathcal{A}^{\mathcal{P}}_{\Delta_{\text{relay}}}$ $(\Delta_{\text{relay}} > 0)$ and some setting $\mathcal{S}$.

Lemma 4. Rule L3 (Figure 2.16) holds for every trace θ feasible with respect to the adversary model $\mathcal{A}^{\mathcal{P}}_{\Delta_{\text{relay}}}$ $(\Delta_{\text{relay}} > 0)$ and some setting $\mathcal{S}$.

Proof. (Lemma 3)
The 1st disjunct of L2, $(B \in V_{\text{adv}})$, follows immediately from A2, so we assume that $B \in V_{\text{cor}}$ and focus on the 2nd disjunct. We we prove it by contradiction. Fix $m = \mathsf{auth}_B(m_0)$. Our goal is to show that by assuming:

(1) $m \sqsubseteq m_1$ and
(2) $\mathsf{Bcast}(C; \tau; m \sqsubseteq m_2) \notin \theta$,
for any correct C, $\tau \leq t - \Delta_{\text{relay}} - \mathit{dist}(C, A)\mathbf{v}_{\text{adv}}^{-1}$ and m_2 st. $m \sqsubseteq m_2$, and
(3) $\mathsf{Dcast}(A; t; \alpha; m \sqsubseteq m_1) \in \theta$

we can derive a contradiction. To achieve this, we show, by induction, that for every N:

$(3)_N$ $\mathsf{Dcast}(A; t_N; \alpha; m \sqsubseteq m_N) \in \theta$ where $t_N \leq t - N\Delta_{\text{relay}}$ and $m \sqsubseteq m_N$

The base case $(3)_0$ follows directly from (3). In the inductive step, we show that $(3)_{N+1}$ follows from $(3)_N$.

Apply A2 to $(3)_N$ and obtain:

(4) $\mathsf{Receive}(D; t_N - \delta; m \sqsubseteq m') \in \theta$,
where $D \in V_{\text{adv}}$, $m \sqsubseteq m'$ and $\delta \geq \Delta_{\text{relay}} + \mathit{dist}(D, A)\mathbf{v}_{\text{adv}}^{-1}$.

Next, apply S1. We can rule out the Bcast disjunct of S1 based on S4 and (2) because:

$$t_N - \delta - (\mathit{dist}(E, D) + \mathit{nlos}(E, D))\mathbf{v}^{-1} \leq t - \Delta_{\text{relay}} - \mathit{dist}(E, A)\mathbf{v}_{\text{adv}}^{-1}$$

This inequality follows from $\mathbf{v}_{\text{adv}} \geq \mathbf{v}$, the inductive assumption $t_N \leq t - N\Delta_{\text{relay}}$ an the condition on δ in (4). Hence, S1 gives us:

(5) $\mathsf{Dcast}(E; t_N - \delta - (\mathit{dist}(E, D) - \mathit{nlos}(E, D))\mathbf{v}^{-1}; \alpha'; m \sqsubseteq m') \in \theta$

Obviously, we can define $m_{N+1} = m'$ and t_{N+1} as:

$$t_{N+1} = t_N - \delta - (\mathit{dist}(E, D) - \mathit{nlos}(E, D))\mathbf{v}^{-1} \leq t - (N+1)\Delta_{\text{relay}}$$

which gives us $(3)_{N+1}$, completing the inductive proof.

The final contradiction follows swiftly: Given $\Delta_{\text{relay}} > 0$, for large enough N, the time t_N is negative. This is in contradiction with event start times being non-negative. □

Proof. (Lemma 2)

The proof is similar to the proof of Lemma 3. We show that we can derive a contradiction by assuming:

(1) $\mathsf{Fresh}(A; t_1; n) \in \theta$, where A is correct, and
(2) for some $B \neq A$, m st. $n \sqsubseteq m$, and $t_2 < t_1 + \mathit{dist}(A, B)\mathbf{v}_{\text{adv}}^{-1} + \Delta_{\text{relay}}$ either:
(a) $\mathsf{Bcast}(B; t_2; n \sqsubseteq m) \in \theta$ or
(b) $\mathsf{Dcast}(B; t_2; \alpha; n \sqsubseteq m) \in \theta$

To this end, we use induction over $N \in \mathbb{N}$ to prove:

$(2)_N$ for some $C_N \neq A$, m_N st. $n \sqsubseteq m_N$, and
$\tau_N < t_1 + \mathit{dist}(A, B)\mathbf{v}_{\text{adv}}^{-1} - (N-1)\Delta_{\text{relay}} - \mathit{dist}(B, C_N)\mathbf{v}_{\text{adv}}^{-1}$ either:
(a) $\mathsf{Bcast}(C_N; \tau_N; n \sqsubseteq m_N) \in \theta$ or
(b) $\mathsf{Dcast}(C_N; \tau_N; \alpha; n \sqsubseteq m_N) \in \theta$

The base case $(2)_0$ follows directly from (2). In the inductive step, we shown that $(2)_{N+1}$ follows from $(2)_N$.

We have two cases of $(2)_N$ to cover. First, consider (a). Given (1), F1 implies:

(3) $\mathsf{Receive}(C_N; \tau_N - \delta_1; n \sqsubseteq m') \in \theta$, for some $\delta_1 \geq \Delta_{\text{relay}}$ and m' st. $n \sqsubseteq m'$

Apply S1 and S4 to obtain:

(4) for some $D \in V$ and $\delta_2 = \delta_1 + (\mathit{dist}(D, C_N) + \mathit{nlos}(D, C_N))\mathbf{v}^{-1}$:
(c) $D \in V_{\text{cor}} \ \wedge\ \mathsf{Bcast}(D; \tau_N - \delta_2; n \sqsubseteq m') \in \theta$ or
(d) $D \in V_{\text{adv}} \ \wedge\ \mathsf{Dcast}(D; \tau_N - \delta_2; \alpha'; n \sqsubseteq m') \in \theta$

We can define $C_{N+1} = D$, $m_{N+1} = m'$, and τ_{N+1} as:

$$\tau_{N+1} = \tau_N - \delta_2 \leq \tau_N - \Delta_{\text{relay}} - (\mathit{dist}(C_{N+1}, C_N) + \mathit{nlos}(C_{N+1}, C_N))\mathbf{v}^{-1} <$$
$$< t_1 + \mathit{dist}(A,B)\mathbf{v}_{\text{adv}}^{-1} - N\Delta_{\text{relay}} - \mathit{dist}(B, C_{N+1})\mathbf{v}_{\text{adv}}^{-1}$$

The inequality follows from the triangle inequality for *dist* and non-negativeness of *nlos*. The last step is to show $C_{N+1} \neq A$. In case (d) this is trivial, as $A \notin V_{\text{adv}}$. In case (c) assume that $C_{N+1} = A$ and observe a contradiction with F1 because:

$$\tau_{N+1} < t_1 + \mathit{dist}(A,B)\mathbf{v}_{\text{adv}}^{-1} - N\Delta_{\text{relay}} - \mathit{dist}(B, C_{N+1})\mathbf{v}_{\text{adv}}^{-1} < t_1$$

This completes the proof of the inductive step in case (a). Now consider case (b) of $(2)_N$. Given (1), A1 implies:

(5) $\mathsf{Receive}(E; \tau_N - \delta_3; n \sqsubseteq m'') \in \theta$,
for some $E \in V_{\text{adv}}$, $\delta_3 \geq \Delta_{\text{relay}} + \mathit{dist}(E, C_N)\mathbf{v}_{\text{adv}}^{-1}$ and m'' st. $n \sqsubseteq m''$

Apply S1 and S4 to obtain:

(6) for some $F \in V$ and $\delta_4 = \delta_3 + (\mathit{dist}(F,E) + \mathit{nlos}(F,E))\mathbf{v}^{-1}$:
- (e) $F \in V_{\text{cor}} \ \wedge\ \mathsf{Bcast}(F; \tau_N - \delta_4; n \sqsubseteq m'') \in \theta$ or
- (f) $F \in V_{\text{adv}} \ \wedge\ \mathsf{Dcast}(F; \tau_N - \delta_4; \alpha''; n \sqsubseteq m'') \in \theta$

We can define $C_{N+1} = F$, $m_{N+1} = m''$, and τ_{N+1} as:

$$\tau_{N+1} = \tau_N - \delta_4 \leq \tau_N < t_1 + \mathit{dist}(A,B)\mathbf{v}_{\text{adv}}^{-1} - N\Delta_{\text{relay}} - \mathit{dist}(B, C_{N+1})\mathbf{v}_{\text{adv}}^{-1}$$

The inequality follows from the triangle inequality for *dist* and non-negativeness of *nlos*. The last step is to show $C_{N+1} \neq A$. In case (f) this is trivial, as $A \notin V_{\text{adv}}$. In case (e) we get a contradiction with F1 if we assume that $C_{N+1} = A$. This completes the proof of the inductive step.

The final contradiction follows as in Lemma 3: Given $\Delta_{\text{relay}} > 0$, for large enough N, the time t_N is negative. This is in contradiction with event start times being non-negative.. □

Protocol $\mathcal{P}^{\mathsf{B/T}}$

Theorem 2. Protocol $\mathcal{P}^{\mathsf{B/T}}$ satisfies ND1 and ND2$^{\mathsf{B/T}}$ under the following assumptions:

(A) $\Delta_{\text{relay}} \geq \mathbf{Rv}^{-1}$

(B) $T_{\mathcal{P}^{\mathsf{B/T}}} = \sup\{|A, t, \mathsf{auth}_A(t)|\ A \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}\} + \mathbf{Rv}^{-1}$

Proof. **Property ND1** (Figure 2.4)

Consider a setting $\mathcal{S}$ and a trace $\theta \in \Theta_{\mathcal{S},\mathcal{P}^{\mathsf{B/T}},\mathcal{A}}$ such that:

(1) $\mathsf{Neighbor}(B; t; A, C, t_2) \in \theta$ for some $A, B, C \in V_{\text{cor}}$

As B is correct, apply P4 to get:

(2) $C = B$ and

(3) $\mathsf{Receive}(B; t_2; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle) \in \theta$, where $t = t_2 + |A, t_1, \mathsf{auth}_A(t_1)|$ and

(4) $t_2 \leq t_1 + \mathbf{Rv}^{-1}$

We need to show:

($\star$) $\mathit{link}(A \to B, t_2)$

Apply S1 to obtain:

(5) $\mathit{link}(D \to B, [t_2, t_2 + |A, t_1, \mathsf{auth}_A(t_1)|])$ and

(6) for $\delta_1 = (\mathit{dist}(D,B) + \mathit{nlos}(D,B))\mathbf{v}^{-1}$
- (a) $\mathsf{Bcast}(D; t_2 - \delta_1; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle) \in \theta$ or
- (b) $\mathsf{Dcast}(D; t_2 - \delta_1; \alpha; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle) \in \theta$.

Consider case (a). From S4 we get $D \in V_{\text{cor}}$ and then from P2 we get $D = A$. Thus, given (5), we have shown $(\star)$.
Consider case (b). Let $\tau = pos(\mathsf{auth}_A(t_1) \sqsubseteq \langle A, t_1, \mathsf{auth}_A(t_1)\rangle)$. Apply L2 to obtain:

(7) $\mathsf{Bcast}(E; t_2 + \tau - \delta_1 - \delta_2; \mathsf{auth}_A(t_1) \sqsubseteq m) \in \theta$,
where $\delta_2 > \Delta_{\text{relay}}$ and $m \in \mathbb{M}$ is st. $\mathsf{auth}_A(t_1) \sqsubseteq m$

S4 gives $E \in V_{\text{cor}}$. Then, apply P2 to get:

(8) $E = A$ and
(9) $m = \langle A, t_1, \mathsf{auth}_A(t_1)\rangle$ and
(10) $t_1 = t_2 - \delta_1 - \delta_2 < t_2 - \Delta_{\text{relay}} \leq t_2 - \mathbf{R}\mathbf{v}^{-1}$, given (A)

From (10) derive $t_2 > t_1 + \mathbf{R}\mathbf{v}^{-1}$. This is a contradiction with (4), thus (b) cannot be true. Consequently, (a) is the only valid option, and ND1 is satisfied.

Property $\text{ND2}^{\text{B/T}}$ (Figure 2.7)
Consider a setting $\mathcal{S}$, where:

(1) $A, B \in V_{\text{cor}}$ and
(2) $dist(A, B) + nlos(A, B) \leq \mathbf{R}$ and
(3) $link(A \leftrightarrow B, [t_1, t_1 + T_{\mathcal{P}^{\text{B/T}}}])$

Next, take any trace $\theta \in \Theta_{\mathcal{S}, \mathcal{P}^{\text{B/T}}, \mathcal{A}}$ such that:

(4) $\mathsf{NDstart}(A; t_1) \in \theta$

We need to show:

$(\star)$ $\mathsf{Neighbor}(B; t'; A, B, t'')$ for some $t' \geq t_1$ and $t'' \in [t_1, t_1 + T_{\mathcal{P}^{\text{B/T}}}]$

Start by applying P1 to obtain:

(5) $\mathsf{Bcast}(A; t_1; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle) \in \theta$

Given (2), (3) and (B), S2 implies:

(6) $\mathsf{Receive}(B; t_2; \langle A, t_1, \mathsf{auth}_A(t_1)\rangle)$, where $t_2 = t_1 + (dist(A, B) + nlos(A, B))\mathbf{v}^{-1}$

Given (2) we obtain $t_2 - t_1 \leq \mathbf{R}\mathbf{v}^{-1}$. Then P3 implies:

(7) $\mathsf{Neighbor}(B; t_2 + |A, t_1, \mathsf{auth}_A(t_1)|; A, B, t_2)$

As $t' = t_2 + |A, t_1, \mathsf{auth}_A(t_1)| \geq t_1$ and $t'' = t_2 \in [t_1, t_1 + T_{\mathcal{P}^{\text{B/T}}}]$ we have shown $(\star)$. □

Protocol $\mathcal{P}^{\text{CR/TL}}$

Theorem 3. Protocol $\mathcal{P}^{\text{CR/TL}}$ satisfies ND1 and $\text{ND2}^{\text{CR/TL}}$ under the following assumptions:

(A) $\Delta_{\text{relay}} > 0$
(B) $\mathbf{v}_{\text{adv}} = \mathbf{v}$
(C) $T_{\mathcal{P}^{\text{CR/TL}}} = \infty$[6]

Proof. **Property ND1** (Figure 2.4)
Consider a setting $\mathcal{S}$ and a trace $\theta \in \Theta_{\mathcal{S}, \mathcal{P}^{\text{CR/TL}}, \mathcal{A}}$ such that:

(1) $\mathsf{Neighbor}(A; t; B, C, t_0) \in \theta$, where $A, B, C \in V_{\text{cor}}$, and $t, t_0 \in \mathbb{R}_{\geq 0}$.

Applying P5 gives two cases:

⟨I⟩ $C = A$: according to ND1, we need to prove $link(B \to A, t_0)$
⟨II⟩ $B = A$: according to ND1, we need to prove $link(A \to C, (t_0 + (dist(A, C) + nlos(A, C))\mathbf{v}^{-1}))$

[6] We set $T_{\mathcal{P}^{\text{CR/TL}}} = \infty$ for simplicity: Otherwise, we would need to assume a maximum distance between A and B to have an upper-bound on the protocol execution time.

We will consider both cases simultaneously. In both cases, if we rename C to B, P5 gives, for some $n_1, n_2 \in \mathsf{Nonces}, t_1, t_2, t_3 \in \mathbb{R}_{\geq 0}, l \in \mathbb{R}^3$:

(2) $\mathsf{Bcast}(A; t_1; \langle B, n_1\rangle) \in \theta$ and
(3) $\mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta$ and
(4) $\mathsf{Receive}(A; t_2; \langle n_2\rangle) \in \theta$ and
(5) $\mathsf{Receive}(A; t_3; \langle l, \mathsf{auth}_B(n_1, n_2, l)\rangle) \in \theta$ and
(6) $\mathbf{v}(t_2 - t_1 - \Delta) = 2d(loc(A), l)$

Further, in case ⟨I⟩:

(7/I) $t_2 = t_0$

whereas in case ⟨II⟩:

(7/II) $t_1 = t_0$

Given (5), we apply S1 to obtain for some $D \in V$:

(a) $\mathsf{Bcast}(D; .; \langle l, \mathsf{auth}_B(n_1, n_2, l)\rangle) \in \theta$ or
(b) $\mathsf{Dcast}(D; .; .; \langle l, \mathsf{auth}_B(n_1, n_2, l)\rangle) \in \theta$

(“.” means that we are not concerned with the value.) Assuming (b), S4 implies $D \in V_{\text{adv}}$. Apply L2 to obtain $\mathsf{Bcast}(E; .; m)$ for some $E \in V_{\text{cor}}$ and m st. $\mathsf{auth}_B(n_1, n_2, l) \sqsubseteq m$. Then P3 gives for some $t_4 \in \mathbb{R}_{\geq 0}$:

(8) $\mathsf{Bcast}(B; .; \langle l, \mathsf{auth}_B(n_1, n_2, l)\rangle) \in \theta$ and
(9) $l = loc(B)$ and
(10) $\mathsf{Receive}(B; t_4 - \Delta; \langle B, n_1\rangle) \in \theta$ and
(11) $\mathsf{Bcast}(B; t_4; \langle n_2\rangle) \in \theta$ and
(12) $\mathsf{Fresh}(B; t_4; n_2) \in \theta$

The same is obtained under (a) via S4 and P3. Apply S1 to (4) to get for some $F \in V$:

(13) $link(F \to A, [t_2, t_2 + |n_2|]) \ \wedge \ (\mathsf{Bcast}(F; t''; \langle n_2\rangle) \in \theta \ \vee \ \mathsf{Dcast}(F; t''; .; \langle n_2\rangle) \in \theta)$

where $t'' = t_2 - (dist(F, A) + nlos(F, A))\mathbf{v}^{-1}$. We have two cases: (c) $F = B$ and (d) $F \neq B$. For case (c), given (12), F1 implies:

(c) $F = B \ \wedge \ t_4 \leq t_2 - (dist(A, B) + nlos(A, B))\mathbf{v}^{-1}$

In case (d), under (12), L1 implies $t_4 + dist(F, A)\mathbf{v}_{\text{adv}}^{-1} + \Delta_{\text{relay}} \leq t'' \leq t_2 - dist(F, A)\mathbf{v}^{-1}$. Using (B) and the triangle inequality we derive:

(d) $F \neq B \ \wedge \ t_4 \leq t_2 - dist(A, B)\mathbf{v}^{-1} - \Delta_{\text{relay}}$

Given (10), apply S1 to get for some $G \in V$:

(14) $link(G \to B, [t_4 - \Delta, t_4 - \Delta + |\langle B, n_1\rangle|]) \ \wedge$
$(\mathsf{Bcast}(G; t'''; \langle B, n_1\rangle) \in \theta \ \vee \ \mathsf{Dcast}(G; t'''; .; \langle B, n_1\rangle) \in \theta)$,
where $t''' = t_4 - \Delta - (dist(G, B) + nlos(G, B))\mathbf{v}^{-1}$.

Again, there are two cases: (e) $G = A$ and (f) $G \neq A$. In case (e), given (3), F1 implies:

(e) $G = A \ \wedge \ t_4 \geq t_1 + (dist(A, B) + nlos(A, B))\mathbf{v}^{-1} + \Delta$

In case (f), given (3), L1 implies $t_1 + |B| + dist(A, G)\mathbf{v}_{\text{adv}}^{-1} + \Delta_{\text{relay}} \leq t''' + |B| = t_4 - \Delta - (dist(G, B) + nlos(G, B))\mathbf{v}^{-1} + |B|$. After simple transformations using the triangle inequality, (B), and omitting the non-negative $nlos$:

(f) $G \neq A \ \wedge \ t_4 \geq t_1 + dist(A, B)\mathbf{v}^{-1} + \Delta + \Delta_{\text{relay}}$

Given (6) and (9) obtain:

(15) $t_2 - t_1 - \Delta = 2dist(A, B)\mathbf{v}^{-1}$

There are now four possible cases to consider: (c)+(e), (c)+(f), (d)+(e) and (d)+(f). Consider case (c)+(e):

(16) $t_1 + (dist(A,B) + nlos(A,B))\mathbf{v}^{-1} + \Delta \leq t_4 \leq t_2 - (dist(A,B) + nlos(A,B))\mathbf{v}^{-1}$

Given (15), both inequalities in (16) need to be equalities and $nlos(A,B) = 0$. As $F = B$, (13) implies $link(B \rightarrow A, t_2)$, which is what we needed to prove in case ⟨I⟩ given (7/I). Furthermore, $G = A$ and (14) implies $link(A \rightarrow B, (t_4 - \Delta))$. In case (e), given (15) and (16), $t_4 - \Delta = t_1 + (dist(A,B) + nlos(A,B))\mathbf{v}^{-1}$, which given (7/II) means that property ND1 is also satisfied in case ⟨II⟩. Finally, given (A), it is easy to see that the remaining three cases are in contradiction with (15), which concludes the proof of ND1.

Property ND2$^{\mathsf{CR/TL}}$ (Figure 2.7)

Consider a setting $\mathcal{S}$ and a trace $\theta \in \Theta_{\mathcal{S},\mathcal{P}^{\mathsf{CR/TL}},\mathcal{A}}$.

Given (C), we assume that:

(1) $\mathsf{NDstart}(A; t_1; B) \in \theta$

(2) $link(A \leftrightarrow B, [t_1, \infty))$

(3) $nlos(A,B) = 0$

We need to prove:

(⋆) $\mathsf{Neighbor}(A; t_1'; A, B, t') \in \theta \wedge \mathsf{Neighbor}(A; t_2'; B, A, t'') \in \theta$ for some $t_1', t_2' \in [t_1, \infty), t', t'' \in [t_1, t_1 + T_{\mathcal{P}^{\mathsf{CR/TL}}}]$.

First apply P1 to obtain:

(4) $\mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta$ and

(5) $\mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta$

Next, given (2) and (3), S2 implies:

(6) $\mathsf{Receive}(B; t_2; \langle B, n_1 \rangle) \in \theta$, where $t_2 = t_1 + dist(A,B)\mathbf{v}^{-1}$

Apply P2 to get:

(7) $\mathsf{Bcast}(B; t_2 + \Delta; \langle n_2 \rangle) \in \theta$ and

(8) $\mathsf{Bcast}(B; t_2 + \tau; \langle loc(B), \mathsf{auth}_B(n_1, n_2, loc(B)) \rangle) \in \theta$, where $\tau > 0$.

Given (2) and (3), S2 implies:

(9) $\mathsf{Receive}(A; t_4; \langle n_2 \rangle) \in \theta$, where $t_4 = t_2 + \Delta + dist(A,B)\mathbf{v}^{-1} = t_1 + \Delta + 2dist(A,B)\mathbf{v}^{-1}$ and

(10) $\mathsf{Receive}(A; t_5; \langle loc(B), \mathsf{auth}_B(n_1, n_2, loc(B)) \rangle) \in \theta$

Given (10), (4), (5), (9), and $\mathbf{v}(t_4 - t_1 - \Delta) = 2dist(A,B)$ we conclude the proof by P4. □

Protocol $\mathcal{P}^{\mathsf{B/TL}}$

Theorem 4. Protocol $\mathcal{P}^{\mathsf{B/TL}}$ satisfies ND1 and ND2$^{\mathsf{B/TL}}$ under the following assumptions:

(A) $\Delta_{\text{relay}} > 0$

(B) $\mathbf{v}_{\text{adv}} = \mathbf{v}$

(C) $T_{\mathcal{P}^{\mathsf{B/TL}}} = \infty$[7]

Proof. **Property ND1** (Figure 2.4)

Consider a setting $\mathcal{S}$ and a trace $\theta \in \Theta_{\mathcal{S},\mathcal{P}^{\mathsf{B/TL}},\mathcal{A}}$ such that:

(1) $\mathsf{Neighbor}(B; t; A, C, t_2) \in \theta$ for some $A, B, C \in V_{\text{cor}}$

[7] We assume $T_{\mathcal{P}^{\mathsf{B/TL}}} = \infty$ for simplicity: : Otherwise, we would need to assume a maximum distance between A and B to have an upper-bound on the protocol execution time.

As B is correct, apply P4 to get:

(2) $C = B$ and

(3) $\mathsf{Receive}(B; t_2; \langle A, t_1, l, \mathsf{auth}_{AB}(t_1, l)\rangle) \in \theta$, where $t = t_2 + |A, t_1, l, \mathsf{auth}_{AB}(t_1, l)|$ and

(4) $t_2 - t_1 = d(loc(B), l)\mathbf{v}^{-1}$

We need to show:

($\star$) $link(A \rightarrow B, t_2)$

Given (3), apply S1 to obtain:

(5) $link(D \rightarrow B, [t_2, t_2 + |A, t_1, l, \mathsf{auth}_{AB}(t_1, l)|])$ and

(6) for $\delta_1 = (dist(D, B) + nlos(D, B))\mathbf{v}^{-1}$

(a) $\mathsf{Bcast}(D; t_2 - \delta_1; \langle A, t_1, l, \mathsf{auth}_{AB}(t_1, l)\rangle) \in \theta$ or

(b) $\mathsf{Dcast}(D; t_2 - \delta_1; \alpha; \langle A, t_1, l, \mathsf{auth}_{AB}(t_1, l)\rangle) \in \theta$.

Consider case (a). From S4 we get $D \in V_{\text{cor}}$ and then from P2 $D = A$. Thus, given (5), we have shown ($\star$).

Consider case (b). Let $\tau = pos(\mathsf{auth}_{AB}(t_1, l) \sqsubseteq \langle A, t_1, l, \mathsf{auth}_{AB}(t_1, l)\rangle)$. Apply L3, to obtain:

(7) $\mathsf{Bcast}(E; t_2 + \tau - \delta_1 - \delta_2; \mathsf{auth}_{AB}(t_1, l) \sqsubseteq m) \in \theta$,
where $\delta_2 \geq \Delta_{\text{relay}} + dist(E, D)\mathbf{v}_{\text{adv}}^{-1}$ and $m \in \mathbb{M}$ is st. $\mathsf{auth}_{AB}(t_1, l) \sqsubseteq m$

S4 gives $E \in V_{\text{cor}}$. Then, apply P2 to get one of the two cases:

(c) $E = A \;\wedge\; m = \langle A, t_1, loc(A), \mathsf{auth}_{AB}(t_1, loc(A))\rangle \;\wedge\; t_1 = t_2 - \delta_1 - \delta_2$ or

(d) $E = B \;\wedge\; m = \langle B, t_1, loc(B), \mathsf{auth}_{AB}(t_1, loc(B))\rangle \;\wedge\; t_1 = t_2 - \delta_1 - \delta_2$

First consider (c). Using the triangle inequality, (B) and (A), we derive $t_1 = t_2 - \delta_1 - \delta_2 \leq t_2 - (dist(D, B) + nlos(D, B))\mathbf{v}^{-1} - \Delta_{\text{relay}} - dist(A, D)\mathbf{v}_{\text{adv}}^{-1} \leq t_2 - dist(A, B)\mathbf{v}^{-1} - \Delta_{\text{relay}} < t_2 - dist(A, B)\mathbf{v}^{-1}$. As $l = loc(A)$, this is a contradiction with (4), thus (c) cannot be true.

Consider case (d). In this case $l = loc(B)$, and (4) implies $t_1 = t_2$. This is in contradiction with $t_1 = t_2 - \delta_1 - \delta_2 \leq t_2 - \Delta_{\text{relay}} < t_2$. Hence (d) cannot be true, and thus (b) cannot be true. Consequently, (a) is the only valid option, and ND1 is satisfied.

Property $\mathsf{ND2}^{\mathsf{B/TL}}$ (Figure 2.7)

Consider a setting $\mathcal{S}$, where:

(1) $A, B \in V_{\text{cor}}$ and

(2) $nlos(A, B) = 0$ and

(3) $link(A \leftrightarrow B, [t_1, \infty))$

Next, take any trace $\theta \in \Theta_{\mathcal{S}, \mathcal{P}^{\mathsf{B/TL}}, \mathcal{A}}$ such that:

(4) $\mathsf{NDstart}(A; t_1; B) \in \theta$

We need to show:

($\star$) $\mathsf{Neighbor}(B; t'; A, B, t'')$ for some $t' \geq t_1$ and $t'' \geq t_1$

Start by applying P1 to obtain:

(5) $\mathsf{Bcast}(A; t_1; \langle A, t_1, loc(A), \mathsf{auth}_{AB}(t_1, loc(A))\rangle) \in \theta$

Given (3), S2 implies:

(6) $\mathsf{Receive}(B; t_2; \langle A, t_1, loc(A), \mathsf{auth}_{AB}(t_1, loc(A))\rangle)$, where $t_2 = t_1 + (dist(A, B) + nlos(A, B))\mathbf{v}^{-1}$

Given (2) we obtain $t_2 - t_1 = dist(A, B)\mathbf{v}^{-1}$. Then P3 implies:

(7) $\mathsf{Neighbor}(B; t_2 + |A, t_1, loc(A), \mathsf{auth}_{AB}(t_1, loc(A))|; A, B, t_2)$

Obviously, $t' = t_2 + |A, t_1, loc(A), \mathsf{auth}_{AB}(t_1, loc(A))| \geq t_1$ and $t'' = t_2 \geq t_1$, and we have shown ($\star$). □

Protocol $\mathcal{P}^{\mathsf{CR/T}}$

Theorem 5. Protocol $\mathcal{P}^{\mathsf{CR/T}}$ satisfies $\textsf{ND1}$ and $\textsf{ND2}^{\mathsf{CR/T}}$ under the following assumptions:

(A) $\Delta_{\text{relay}} > 2\mathbf{R}\mathbf{v}^{-1}$

(B) $T_{\mathcal{P}^{\mathsf{CR/TL}}} = \sup\{|B,n| + |\mathsf{auth}_B(n)|\; B \in \mathbb{V}, t \in \mathbb{R}_{\geq 0}, n \in \mathsf{Nonces}\} + 2\mathbf{R}\mathbf{v}^{-1}$

Proof. **Property ND1** (Figure 2.4)

Consider a setting $\mathcal{S}$ and a trace $\theta \in \Theta_{\mathcal{S},\mathcal{P}^{\mathsf{CR/T}},\mathcal{A}}$ such that:

(1) $\mathsf{Neighbor}(A; t; B, C, t_0) \in \theta$, where $A, B, C \in V_{\text{cor}}$, and $t, t_0 \in \mathbb{R}_{\geq 0}$.

Applying $\textsf{P6}$ gives two cases:

⟨I⟩ $C = A$: according to $\textsf{ND1}$, we need to prove $\mathit{link}(B \to A, t_0)$

⟨II⟩ $B = A$: according to $\textsf{ND1}$, we need to prove $\mathit{link}(A \to C, (t_0 + (\mathit{dist}(A, C) + \mathit{nlos}(A, C))\mathbf{v}^{-1}))$

We will consider both cases simultaneously. In both cases, if we rename C to B, $\textsf{P6}$ gives, for some $n_1 \in \mathsf{Nonces}, t_1, t_2 \in \mathbb{R}_{\geq 0}$:

(2) $\mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta$ and

(3) $\mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta$ and

(4) $\mathsf{Receive}(A; t_2; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta$ and

(5) $t_2 - t_1 - \Delta \leq 2\mathbf{R}\mathbf{v}^{-1}$

Further, in case ⟨I⟩:

(6/I) $t_2 = t_0$

whereas in case ⟨II⟩:

(6/II) $t_1 = t_0$

Given (4), we apply $\textsf{S1}$ to obtain for some $D \in V$:

(7) $\mathit{link}(D \to A, [t_2, t_2 + |\mathsf{auth}_B(n_1)|])$ and for $t_3 = t_2 - (\mathit{dist}(A, D) + \mathit{nlos}(A, D))\mathbf{v}^{-1}$:

(a) $\mathsf{Bcast}(D; t_3; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta$ or

(b) $\mathsf{Dcast}(D; t_3; .; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta$

First consider (a). Apply $\textsf{S4}$ to get $D \in V_{\text{cor}}$ and then $\textsf{P4}$ to get:

(8/a) $B = D \;\wedge\; \mathsf{Bcast}(B; t_3; \mathsf{auth}_B(n_1)) \in \theta$ and

(9/a) $\mathsf{Receive}(B; t_3 - \Delta; \langle A, B, n_1 \rangle) \in \theta$

Given (9/a), apply $\textsf{S1}$ to obtain for some $F \in V$, $t' = t_3 - \Delta$, and $t'' = t' - (\mathit{dist}(F, B) + \mathit{nlos}(F, B))\mathbf{v}^{-1}$:

(10/a) $\mathit{link}(F \to B, [t', t' + |B, n_1|]) \wedge (\mathsf{Bcast}(F; t''; \langle B, n_1 \rangle) \in \theta \vee \mathsf{Dcast}(F; t''; .; \langle A, B, n_1 \rangle) \in \theta$

Two cases arise: (c) $F = A$ and (d) $F \neq A$. Given (3), in case (c) apply $\textsf{F1}$, and in case (d) apply $\textsf{L1}$ to obtain:

(c/a) $F = A \;\wedge\; t'' \geq t_1$ or

(d/a) $F \neq A \;\wedge\; t'' \geq t_1 + \Delta_{\text{relay}} + \mathit{dist}(F, A)\mathbf{v}_{\text{adv}}^{-1}$

Consider case (c/a). As $F = A$, given $A \in V_{\text{cor}}$ and $\textsf{S4}$, (10/a) states that $\mathsf{Bcast}(A; t''; \langle A, B, n_1 \rangle) \in \theta$. Apply $\textsf{P2}$ to get:

(11/ac) $\mathsf{Fresh}(A; t'' + |B|; n_1) \in \theta$

Given (11/ac) and (2), $\textsf{F1}$ implies $t'' \leq t_1$, which under (c/a) gives $t'' = t_1$. Thus, $t_1 = t_2 - 2(\mathit{dist}(A, B) + \mathit{nlos}(A, B))\mathbf{v}^{-1} - \Delta$. Given $B = D$, (7) implies $\mathit{link}(B \to A, t_2)$, which proves ⟨I⟩ given (6/I). Further, $F = A$ and (10/a) give $\mathit{link}(A \to B, t_1 + (\mathit{dist}(A, B) + \mathit{nlos}(A, B))\mathbf{v}^{-1})$, which given (6/II) proves ⟨II⟩. All that remains to show is that cases (d/a) and (b) are not possible.

Consider case (d/a). It is straightforward to derive $t_2 - t_1 - \Delta \geq \Delta_{\text{relay}}$. Thus, given (A), $t_2 - t_1 - \Delta > 2\mathbf{R}\mathbf{v}^{-1}$, which is in contradiction with (5). Hence (d/a) is not possible.
Next, consider (b). Apply S4 to get $D \in V_{\text{adv}}$ and then L2 followed by P4 to obtain:
(8/b) $\mathsf{Bcast}(B; t_3 - \delta_1; \mathsf{auth}_B(n_1)) \in \theta$, where $\delta_1 \geq \Delta_{\text{relay}} + \mathit{dist}(D, E)\mathbf{v}_{\text{adv}}^{-1}$ and
(9/b) $\mathsf{Receive}(B; t_3 - \delta_1 - \Delta; \langle A, B, n_1 \rangle) \in \theta$

The only difference between (9/a) and (9/b) is a different timestamp, which in case (b) is $t' = t_3 - \delta_1 - \Delta$. We can thus repeat a nearly identical reasoning, deriving (10/b), (c/b) and (d/b). However, as the timestamp t' in case (b) "includes" Δ_{relay}, we can easily show that (c/b) and (d/b) are in contradiction with (5). This concludes the proof of ND1.

Property ND2$^{\mathsf{CR/T}}$ (Figure 2.7)
Consider a setting $\mathcal{S}$ and a trace $\theta \in \Theta_{\mathcal{S}, \mathcal{P}^{\mathsf{CR/T}}, \mathcal{A}}$.
Given (C), we assume that:
(1) $\mathsf{NDstart}(A; t_1; B) \in \theta$
(2) $\mathit{link}(A {\leftrightarrow} B, [t_1, t_1 + T_{\mathcal{P}^{\mathsf{CR/T}}}])$
(3) $\mathit{dist}(A, B) + \mathit{nlos}(A, B) \leq \mathbf{R}$

We need to prove:
($\star$) $\mathsf{Neighbor}(A; t'_1; A, B, t') \in \theta \;\wedge\; \mathsf{Neighbor}(A; t'_2; B, A, t'') \in \theta$
for some $t'_1, t'_2 \in [t_1, \infty), t', t'' \in [t_1, t_1 + T_{\mathcal{P}^{\mathsf{CR/T}}}]$.

First apply P1 to obtain:
(4) $\mathsf{Fresh}(A; t_1 + |B|; n_1) \in \theta$ and
(5) $\mathsf{Bcast}(A; t_1; \langle B, n_1 \rangle) \in \theta$

Next, given (2) and (3), S2 implies:
(6) $\mathsf{Receive}(B; t_2; \langle B, n_1 \rangle) \in \theta$, where $t_2 = t_1 + (\mathit{dist}(A, B) + \mathit{nlos}(A, B))\mathbf{v}^{-1}$

Apply P3 to get:
(7) $\mathsf{Bcast}(B; t_2 + \Delta; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta$ and

Given (2), S2 implies:
(8) $\mathsf{Receive}(A; t_3; \langle \mathsf{auth}_B(n_1) \rangle) \in \theta$,
where $t_3 = t_2 + \Delta + (\mathit{dist}(A, B) + \mathit{nlos}(A, B))\mathbf{v}^{-1} = t_1 + \Delta + 2(\mathit{dist}(A, B) + \mathit{nlos}(A, B))\mathbf{v}^{-1}$

Given (4), (5), (8), (3), and $t_3 - t_1 - \Delta \leq 2\mathbf{R}\mathbf{v}^{-1}$ conclude the proof by P5. □

2.3.8 Isabelle/HOL Mechanization

We formalize the proposed model, with a few minor modifications, in the theorem prover Isabelle [110] and higher-order logic (HOL). This allows us to mechanically verify the proofs presented in Section 2.3.7, greatly increasing the confidence in the results. The source code is available is available in [129].

In the formalization process, we make two noteworthy modifications:

- We remove the Dcast event, making both the correct nodes and the adversarial nodes use the Bcast event. This simplifies the model and the proofs, and it does not actually restrict the ability of the adversary to use directional transmissions. Indeed, a directional transmission can be modeled by the state of the *link* relation (between the transmitting adversarial node and receiving nodes) – in the same fashion as we model jamming (Section 2.4.1). In the pen-and-paper model and proofs, we found that using the Dcast event is a more straightforward way of illustrating to the reader that the nodes can use directional transmission, and that it is worth the overhead (minor in pen-and-paper proofs).

- We model concatenated messages as lists of simple messages (identifiers, timestamps, locations, nonces, authenticators). With this representation, we have a one-to-one mapping between the messages in the model, and "real-world" messages. Whereas one "real-world" message concatenated from more than 2 simple messages has multiple term representation, depending on the order of concatenation.

In the Isabelle formalization, we use an extension of HOL, the HOL-Complex logic that defines complex and real numbers, because our model requires the latter. Types, such as messages or events, are defined using `datatype` ([110], Section 2.5). Settings, for convenience, are defined not as tuples, but with `record` ([110], Section 8.2). Recursive functions, such as message duration, or the sub-term function are defined with `primrec` and `fun` ([110], Section 3.5). Finally, the simplest constructs such as non-recursive functions, constants and feasibility rules, are defined with `definition` ([110], Section 2.7.2).

We mechanized the most essential proofs: Lemma 2 and Lemma 3, availability and correctness of the $\mathcal{P}^{\mathsf{B/T}}$ protocol and correctness of the (most involved) $\mathcal{P}^{\mathsf{CR/TL}}$ protocol. The proofs follow the pen-and-paper proofs in Section 2.3.7 very closely. Each step of the pen-and-paper proof (i.e., each application of a feasibility rule) translates into an application of a number of Isabelle methods. The Isabelle source code for the model and proofs is roughly 2500 lines long.

2.4 Discussion

In this section, we discuss the abstractions and simplifications introduced in our framework for the sake of modeling and reasoning about secure ND. We also outline the differences between protocols in terms of requirements and satisfied properties.

2.4.1 Abstractions and Simplifications

Mobility and NLOS Delay We assume nodes are static and non-line-of-sight (NLOS) delay is constant over time. This simplifies the model quite significantly, because otherwise propagation delay would vary during the transmission of a message. This is a reasonable assumption, because mobility and NLOS delay changes are very minor at the ND protocol execution time scale. For example, during $100\mu s$, nodes moving at 100kmph traverse $2.7mm$, which is below the accuracy of most RF ranging systems. However, in general, mobility can have security implications. To see why, consider the $\mathcal{P}^{\mathsf{CR/TL}}$ protocol. If nodes move during the protocol execution, it is important when they estimate their location. At the very least, A should estimate its location once when it sends the challenge, and again when it receives the response; whereas the responding node B should estimate its location when it sends the response. But even this might be insufficient under high mobility: If A measures its location at the beginning of the message, while B measures the ToF at the end of the message, there may be space for a stealthy relay attack. Introducing mobility and a dynamically changing NLOS delay in our model is a possible direction for future work.

Medium Access Control and Jamming For simplicity, we do not introduce any MAC restrictions into the model. Hence, a node is able to simultaneously receive any finite number of messages, even though in reality it is limited (to one message, or more for CDMA-like technologies). We could introduce additional rules that model radio interference, e.g., set links *down* if two (or more, depending on the node transceiver capabilities) simultaneous transmissions take place. However, this would *not* affect any of our results. Notably, the availability properties require links to be *up*, but they are agnostic as to why links are *up* or *down*. Similarly, jamming would not affect our results either: we capture jamming with links being *down*, thus availability implies, among other things, no jamming.

Inaccuracies We assume correct nodes have accurate time and location information. However, in reality, inaccuracies are inevitable. Regarding time, clocks may be coarse-grained and they can drift. Furthermore, there is always some error in estimating the message reception time (time-of-arrival) over a noisy channel. Regarding location, infrastructure (e.g., the Global Positioning System (GPS), or base stations) providing location information may be temporarily unavailable, and the location provided also includes some measurement error. Some of the inaccuracies can be decreased: For example, the error in message time-of-arrival can be decreased by averaging over many messages or introducing long physical-communication-layer preambles (see Part II). But some inaccuracy in time and location is unavoidable.

Inaccuracies do not diminish the impossibility result; Rather, they make it stronger. Indeed, we prove the impossibility result holds even in an idealized environment, in which nodes have access to information more accurate than in reality. In contrast, as secure ND protocols rely on distance estimates, their effectiveness can be negatively affected by such inaccuracies. For T-protocols, and even more so for TL-protocols, inaccuracies hinder availability: they can lead to ToF estimates seemingly above the threshold for T-protocols, and make the two distance estimates diverge for TL-protocols. The only way to cope with these is to introduce some tolerance margins for measurements. Nonetheless, this affects correctness: The higher the tolerance margin, the more space is left for fast relay attacks. This manifests the unsurprising tension between correctness and availability. Introducing inaccuracies explicitly into the framework is a possible direction for future work.

Physical Layer Attacks The messages considered in our framework, albeit at the physical layer, are composed of “atomic” components, such as nonces and identifiers, typically assumed in formal security frameworks. In [35], Clulow et al. pointed out a number of physical layer attacks, working at the symbol (or bit) level, and able to decrease message reception time. In the case of *external* adversaries, as considered in our ND specification, the attacks proposed in [35] can result in a (perceivably) negative Δ_{relay}. This can be expressed in our model. Furthermore, the external malicious interference attack we introduce in Chapter 5 can be modeled with minor extensions, e.g., making the *nlos* depend on time and allow negative values. Hence our framework (notably the “atomicity” assumptions) is not limited with respect to those attacks. However, this is not the case for physical-communication-layer attacks introduced by *internal* adversaries. We discuss this in Section 2.5.

Fingerprinting In Section 1.4 we described device fingerprinting and channel fingerprinting, and how it can be used to secure ND. The framework proposed in this section does not capture the capabilities of wireless receivers to perform such fingerprinting. This has implications notably for the impossibility result. If the model would include some form of fingerprinting, then the indistinguishability of local views (Lemma 1) would not hold, and the impossibility result would be lifted. It is not clear how such capabilities could be introduces into a formal framework, notably because these techniques are relatively recent, and there is no clear understanding of the hardness of fingerprint spoofing. Integrating these into a formal framework is a possible direction of future work.

2.4.2 Protocol Comparison

T-protocols versus TL-protocols On one hand, TL-protocols provide stronger security than T-protocols in term of correctness. First, they do not need the notion of ND range, $\mathbf{R}$, needed by T-protocols. More importantly, they are secure as long as $\Delta_{\text{relay}} > 0$. In contrast, T-protocols require that $\Delta_{\text{relay}} \geq \mathbf{R}\mathbf{v}^{-1}$. On the other hand, TL-protocols suffer in terms of availability: (i) they require location-aware nodes with secure and precise location information, a far from trivial requirement, and (ii) they do not work for links with substantial NLOS delay.

In light of these shortcomings, notably (i), T-protocols can be a viable solution to provide communication ND, depending on the environment, the communication technology, and the sophistication of the adversary. First, T-protocol provide a good approximation of communication ND in environments where two nodes in nominal communication range are able to directly communicate, for example in outer-space. However, many environments (notably indoor) do not display such characteristics. Second, if the ND range $\mathbf{R}$ is low, than the adversary needs to be able to relay with a small Δ_{relay}. For example, if we consider relatively short-range IEEE 802.11 radios, with $\mathbf{R}$ in the order of 100 meters, $\Delta_{\text{relay}} \approx 100mc^{-1} \approx 333ns$. This is significantly below the $15-20\mu$s achievable by the non-trivial relay constructed by Hancke in [67]. Simple store-and-forward relays are also thwarted easily. In contrast, for WiMAX, with a range up to $50km$, the lower-bound on Δ_{relay} is around $166\mu s$ leaving much more space for attacks. In fact, as $\mathbf{R} \to \infty$, T-protocols become useless for securing ND.

Consider more powerful relay attacks (covered in Section 1.4.6). One example is the analog relay of Francillon et al [61], with $\Delta_{\text{relay}} \approx 20$ns. A ND range $\mathbf{R}$ secure against such a relay is only a few meters. Furthermore, for many wireless technologies, 20ns falls below the accuracy of the message time-of-arrival estimation. This implies that for practical purposes this relay can be assumed to achieve $\Delta_{\text{relay}} \approx 0$, which defeats not only T-protocols, but also TL-protocols. Furthermore, using physical-communication-layer attacks [35] it is possible to construct a relay with a (seemingly) negative Δ_{relay}, which also defeats both classes of protocols.

Hence, it might appear that for a sophisticated enough adversary, communication ND is impossible not only for T-protocols, but also for TL-protocols. However, there is a significant difference between these two "impossibility results". In case of TL-protocols, the difficulty stems from the inaccuracy of time- and location-measurements. These can be decreased by, e.g., increasing the signal-to-noise ratio through making the message preambles (on which the time-of-arrival is estimated) longer. Furthermore, as we show in Part II of the thesis, physical-communication-layer attacks can be mitigated with appropriate countermeasures deployed on the receivers. In contrast, the T-protocol impossibility is fundamental, and holds even in an idealized model with no inaccuracies and no physical layer attacks.

B-protocols versus CR-protocols We provide a comparison of time-based B-protocols and CR-protocols in Section 1.4; much of this applies to TL-protocols as well. In addition, we note that B-protocols have less stringent requirements for availability, requiring that links be *up* for shorter periods than those needed by CR-protocols. In terms of correct (secure) operation, CR/T-protocols require Δ_{relay}, the minimum relaying delay, to be twice as large as that required by B/T-protocols (for the same $\mathbf{R}$).

Symmetric Authenticators Contrary to other protocols, the $\mathcal{P}^{\mathsf{B/TL}}$ protocol uses a symmetric authenticator. For this reason, this protocol might seem at the first glance susceptible to a reflection attack: An adversarial node could after receiving the beacon message from node A relay it back to A. Yet, as proven in Section 2.3.7, the protocol is actually secure. Furthermore, we could modify all the other protocols by simply replacing asymmetric authenticators with symmetric ones, and they would still be secure under the same assumptions as their asymmetric counterparts. However, if we would remove the time and location information from the symmetric versions of the protocols, in an attempt to use them as regular authentication protocols, they would be all be vulnerable to the reflection attack. This demonstrates an interesting interplay between authentication and time/location features of ND protocols.

2.4.3 Physical ND and Distance Bounding

In this Chapter, we focus on communication ND, but it is straightforward to express a *specification* for physical ND, as well as distance bounding (DB)[8] in the proposed framework. However, as we explain in Section 2.5, most DB protocols presented in the literature cannot be directly modeled in our framework in a meaningful fashion. The fundamental reason for this is that DB protocols attempt to provide a secure distance bound even if the DB protocol is executed with an adversarial node. This opens a new space of internal attacks which are not represented in our model. Communication ND protocols, in contrast, traditionally assume that participating nodes are correct. The framework we propose is sufficient for capturing this case.

2.5 Related Work and Open Challenges

We have provided a broad overview of the literature on ND schemes, relay attacks, and related issues in Section 1.4. Here we review works that focus on analysis of neighbor discovery and distance bounding protocols. The importance of formalizing the analysis of security protocols in wireless networks has been recognized by a number of authors, e.g., in the context of security of routing [106, 12, 111, 168], local area networking [72], or broadcast authentication [71].

2.5.1 Impossibility Results

In [32] the authors study the problem of secure clock synchronization under relay attacks, which turns out to be closely related to communication ND. Compared to our model, the model in [32] includes clock skews and an adversary model with the distinction of half-duplex, full-duplex and double full-duplex transceivers, rather than the relaying delay. The authors obtain impossibility and possibility results for the considered transceiver types, which are complementary to the results obtained in this Chapter. In contrast to our work, the authors abstract away the cryptographic aspect of the protocols – hence their framework cannot be directly use to prove the correctness of concrete protocols in the same way as our framework.

[8] The difference between physical ND and DB is discussed in Section 1.1.3.

This work is further extended in [138] to propose protocols for network-wide clock synchronization and topology discovery, but again abstracting away the cryptographic aspects.

2.5.2 Formal Verification of Protocols

A number of formal frameworks designed for protocol verification have been proposed. However, contrary to our work, they tend to focus on distance bounding (DB) rather than communication ND. We provide an introduction to DB in Section 3.1, and assume here that the reader is familiar with the information provided in that section.

One of the first works where DB has been treated formally is [95] by Meadows et al. The authors build on top of existing formal approaches [29, 118] tailored to "classical" security protocols, and augments it with a notion of distance based on time-stamps. It is not clear how communication neighborhood would be defined in this framework, nor how to model a protocol that uses location information. Beyond this, an interesting characteristic of this approach is that there is no explicit definition of an adversary. On the contrary, our approach starts with a model of a wireless environment, including node location, state of wireless links, and an explicit adversary, controlling a number of nodes in the network. A potential advantage of this is that attack scenarios can be expressed in our model, whereas in [95] a collusion attack is described in an informal manner.

In [141], Schaller et al propose a framework based on the inductive active approach of Paulson [117]. The framework is formalized in the theorem prover Isabelle using Higher Order Logic. Similar to our approach, the authors model directly the link relation, node location and propagation time. The authors use their framework to verify an authenticated ranging protocol (i.e., DB with a trusted prover), and ultrasound DB protocol, and in addition a delayed key disclosure protocol. The authors further extend this approach in [20], where they propose an elegant way of dealing with message spaces based on equational term theories. This, in particular, allows them to include the *exclusive or* (XOR) operation used in the Brands-Chaum style DB protocols [23] into the message space.

In [94], the authors extend the stand space model with timing information. The approach is automatized using the constrained solving techniques proposed in [97] for bounded-process analysis. The authors analyze 4 distance bounding protocols, including a simplified version of the Brands-Chaum protocol, and a protocols proposed in [95]. They replace the XOR operation with symmetric encryption, and the commitment is removed. The authors report that the constraint solver is able to efficiently find attacks in flawed versions of the protocols. In particular, they discover the distance hijacking attack [38] against these protocols.

The authors of [146] extend the strand space formalism [145] with notions of message propagation time and device location to be able to reason about the security of simple DB and related protocols. No mechanization of the security proofs is provided.

2.5.3 Probabilistic Approach and Open Challenges

Interestingly, in none of the above frameworks is it possible to prove the correctness of the (non-simplified) Brands-Chaum protocol, or the Hancke-Kuhn protocol [66]. This is for a number of reasons, most of which are related to modeling of internal attacks that DB protocols should cope with. First, although the XOR operation is modeled in [95] and in [20], both frameworks deem the way the XOR is used in the Brands-Chaum protocol as insecure. In reality, it is secure because it is combined with the commitment. In the same vain, the look-up operation used in the Hancke-Kuhn protocol is not modeled in any of the frameworks. Second, all frameworks assume that all messages are instantaneous (i.e., have no duration), which is not the case in reality, and leads to overlooking of some attacks (such as the attack in Figure 3.3). In our model, we incorporate message duration. However, we assume that the message components such as nonces are atomic, not composed from bits, which leads to abstracting away crucial attacks (such as the attack in Figure 3.2). Third, all frameworks are non-probabilistic, which means that they only capture attacks that can occur with probability one. In contrast, malicious prover (distance fraud) attacks on DB protocols work with probability that is non-negligible, but lower than 1 (e.g., Figure 3.2). We believe these types of attacks should be possible to capture in any framework to prove the security of ND or DB protocols against internal adversaries. This would require a shift from a non-deterministic, message-based models to a probabilistic symbol/bit-based model.

One formal approach takes this leap: the framework of Pavlovic and Meadows introduced in [120]. It includes probabilistic derivation based on the notion of guards, it models messages on the bit level, and it incorporates the Hancke-Kuhn look-up operation. In fact, [120] provides the first formal analysis of the Hancke-Kuhn protocol. [120] is an extension of [95], thus propagation time, distances, and node locations are not modeled directly. Furthermore, the analysis in [120] has a few limitations. Foremost, it is limited to two specific guessing attack scenarios: one distance fraud and one mafia fraud; nothing is mentioned about the terrorist fraud, to which the Hancke-Kuhn protocol is vulnerable. Such approaches that are based on iterating over a list of high-level attack scenarios offer limited security, as they miss undiscovered attack scenarios. An example of this limitation is the recently discovered distance hijacking attack scenario [38], which is different from the traditional distance fraud, mafia fraud, and terrorist fraud. In contrast, the strength of most formal approaches lies in considering a much broader scope of adversarial actions, and only restricting the adversary from violating causality, physical time and location constraints, and the security of cryptographic primitives. In [94], such an approach allowed the authors to detect the distance hijacking attack.

Second, the framework in [120] uses the notion of negligible probability, abstracting away the exact probability of attack. However, many DB protocols are designed specifically to lower an already negligible probability, with the goal of making the protocols more efficient. It would prudent to allow formal approaches to capture such quantitative results. Third, the approach is not mechanized, and experience has shown that a lack of mechanization can overlook some flaws even in a formal analysis (in [20], an attack against one of the protocols proven correct in [95] is demonstrated). Nevertheless, we believe that this work takes an important step in the right direction.

Finally, we would like to note that beyond the challenges mentioned above, the years of work on distance bounding protocols have lead to the discovery of a number of interesting attacks, surveyed in Section 3.4.3. Ideally, a formal framework should be capable of capturing all these attacks.

2.6 Conclusions

In this Chapter, we contribute a formal analysis of secure communication ND. We build a formal framework, and provide a specification of communication ND or, more precisely, its most basic variant, two-party ND. We consider two general classes of protocols: time-based protocols (T-protocols) and time-and location-based protocols (TL-protocols). For the T-protocol class, we identify a fundamental limitation governed by a threshold value depending on the ND range: We prove that no T-protocol can provide secure communication ND if and only if adversarial nodes can relay messages faster that this threshold. This result is a useful measure of the security achieved by ND T-protocols, which we believe are one of the more practical and universal solutions to ND.

We use the proposed framework to analyze and design concrete provably secure communication ND protocols. We prove that two T-protocols – one beacon (B) and one challenge-response (CR) – can provide secure communication ND if the adversary relaying delay is lower-bounded by the threshold. Further, we show that location information, if available, can improve the security. We propose novel TL-protocols (a B/TL protocol and a CR/TL protocol) and prove that they can provide secure communication ND as long as the adversarial relaying delay is strictly positive. In practice, this implies that the protocol is secure as long as the adversarial relaying delay exceeds the timing and location inaccuracies.

We argue that the proposed framework is adequate for reasoning about ND (and distance bounding) under the assumption that both participating nodes are honest. However, if one of the nodes is allowed to be adversarial – a common assumption for distance bounding protocols – then this opens a whole new range of attacks, in particular attacks on the physical-communication-layer. This mandates, in our opinion, a shift from (non-)deterministic message-oriented models to probabilistic models that explicitly consider bits or even symbols at the physical-communication-layer. This is an interesting direction for future work, but it requires a good understanding of the possible physical layer attacks and countermeasures. This motivated us to pursue the investigation that constitutes Part II of the thesis.

Acknowledgments

I would like to extend my thanks to Patrick Schaller, Benedikt Schmidt, David Basin, Srdjan Capkun, Pascal Lafourcade and Paul Hankes Drielsma. We shared numerous inspiring discussions that helped forge this Chapter. Special thanks go to Benedikt Schmidt for giving me an introduction to Isabelle.

Part II

Physical Layer Attacks against Distance Bounding

Chapter 3

Introduction to Distance Bounding, PHY Attacks and Impulse Radio

In this part of the thesis, we focus on physical-communication-layer (PHY) attacks that cause a (seeming) shift of the message reception time, without changing the content of the message. These attacks were first discovered in [35]. Clearly, they are a threat to the time-based (and time-and-location-based) ND protocols[1] that we consider in Part I, but they are equally relevant to *ranging* and *distance bounding* (DB) protocols. DB protocols allow a (wireless) device to estimate, in a *secure* manner, the distance from itself to another device, or more precisely, the *upper-bound* on this distance. DB protocols can provide physical ND in applications such as physical access control. Furthermore, they can be used in applications such a secure tracking [100] or secure localization [155].

We concentrate on DB protocols in this part of the thesis because traditionally DB protocols consider the case where one of the participants is adversarial. In contrast, ND protocols typically assume that both participants are honest. The former leads to a broader class of PHY attacks. However, as PHY attacks work on the packet level, not the protocol level, the external attacks that we investigate apply to ND protocols, as well as DB protocols. Furthermore, we focus on DB protocols that aim at providing a *strict* upper-bound (i.e., the *precise distance*) whenever possible. For such protocols, distance-decreasing attacks are especially relevant. We focus on a particular PHY: *Impulse-Radio Ultra-Wideband* (IR-UWB or simply IR), because it is an ideal candidate for implementing DB protocols [66], thanks to its ability to perform precise ranging.

Distance-decreasing PHY attacks can be classified into *malicious prover attacks* (internal), *distance-decreasing relay attacks* (external) and *malicious interference attacks* (external). In Chapter 4, we adapt the first two attacks, originally introduced in [35], to the IEEE 802.15.4a standard [77] and evaluate their effectiveness. Beyond the security analysis of IEEE 802.15.4a, this investigation also gives us insight into how different performance-enhancing features (such as convolutional coding and time-hopping) affect security of a PHY against distance-decreasing attacks. In Chapter 5 we identify a new attack vector against IR ranging, based on disrupting the ToA estimation. This allows us to introduce the last class of PHY attacks (malicious interference), and in particular a low-complexity variant: the *cicada attack*. We also show how the ToA attack vector can be used in malicious-prover and relay attacks. We evaluate the attacks against a broad class of modulation schemes and a number

[1] For the relation between ND and DB see Section 1.1.3

of receivers. Furthermore, in both Chapters, we examine countermeasures that can mitigate PHY attacks.

Chapter Outline In this Chapter, we provide an introduction to ranging and DB protocols (Section 3.1). In Section 3.2, we give an overview of PHY attacks. In Section 3.3 we introduce Impulse Radio, we define the multipath channel model, we give an overview of wireless transceivers and the simulation environment. Finally, we survey the related work in Section 3.4.

3.1 Distance Bounding

Distance bounding (DB) protocols allow one device, the *verifier* **V**, to securely compute an upper-bound on the distance to another device, the *prover* **P**. Figure 3.1 shows an example of three DB protocols: the Brands-Chaum protocol [23], the Hancke-Kuhn protocol [66], and the Čapkun-Hubaux protocol [155]. Like regular ranging protocols, DB protocols perform distance estimation based on time-of-flight measurements of *ranging messages*: the *challenge(s)* sent by **V** and the *response(s)* of **P**. Additional messages are employed to guarantee authentication, which is typically based on a common shared secret.

Providing a secure upper-bound implies that DB protocols only attempt to prevent distance-decrease attacks. It is easy to see that no protocol can prevent an adversary from increasing[2] the measured distance: First, a *malicious prover* can delay the response. Second, the adversary can place itself in between the two honest devices, and *relay* the ranging messages without any modification, but adding some delay.

Three threat scenarios are traditionally considered [46, 25]. In the *mafia fraud*, the adversary interferes with a DB session between an honest **V** and an honest **P**, and decreases the measured distance below the actual distance. In the other scenarios, a malicious **P** convinces **V** that it is closer than it actually is, working alone (*distance fraud*), or in collusion with other malicious devices (*terrorist fraud*). More recently, a new threat scenario termed *distance hijacking* was added to this list [38]. In this scenario, a malicious prover tricks a honest prover **P**' to perform the ranging, and then "hijacks" the authentication part of the protocol, making **V** believe that the distance to the malicious prover is that to **P**'.

Setting PHY attacks aside, all DB protocols are secure against the mafia fraud. In particular, it is easy to see that a relay attack cannot decrease the measured distance, because that would imply introducing a negative delay when relaying the messages, or transmitting messages at a speed faster than the speed of light. However, some protocols (referred to as *authenticated ranging* or *DB with a trusted prover*) are by design only secure against the mafia fraud. The Čapkun-Hubaux protocol is an example of such a protocol, where the prover sends the response to the challenge after a delay. The verifier takes this delay into account when computing the distance. A malicious prover can respond faster than the declared delay, thus decreasing the measured distance. In contrast, the two other protocols are secure against this attack, because the response is assumed to be send instantly after the challenge is received. This leaves a malicious prover no room for an early response.

In fact, both the Brands-Chaum protocol and the Hancke-Kuhn protocol incorporate a more sophisticated scheme, called the *rapid bit exchange* (RBE). In the RBE, the verifier sends a number of single-bit challenges, one after the other. The prover replies instantly to

[2] Unless some additional assumptions are made [148].

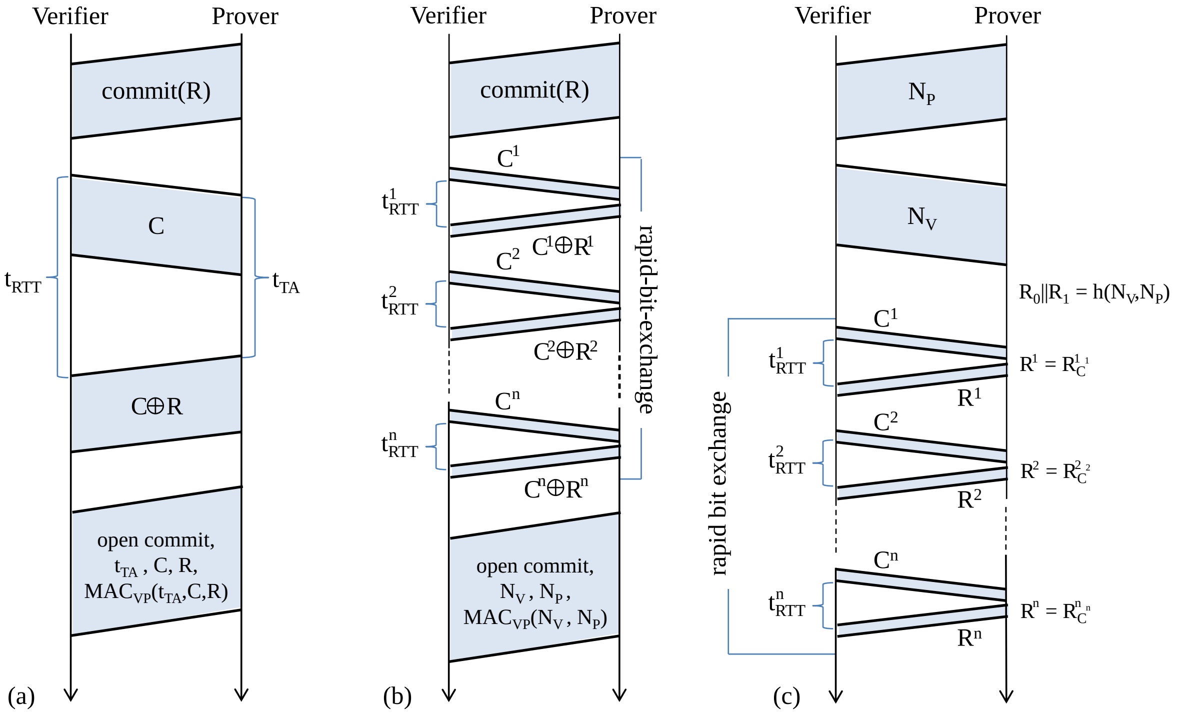

Figure 3.1: Examples of distance bounding protocol: (a) The Čapkun-Hubaux protocol, secure only against the mafia fraud, (b) The Brands-Chaum protocol and (c) The Hancke-Kuhn protocol, both secure against the distance fraud and the mafia fraud. **V** estimates the distance to **P** with the formula $d_{VP} = c(t_{\mathrm{RTT}} - t_{\mathrm{TA}})/2$, where c is the channel propagation speed. $\mathrm{MAC}_{\mathrm{VP}}$ stands for Message Authentication Code with a symmetric key shared between **V** and **P**, C, R, N_V and N_P are freshly generated nonces, h is a cryptographic one-way function, $||$ denotes concatenation, t_{TA} is a constant turn-around time that **V** and **P** know, and which is assumed 0 for protocols (b) and (c), and t_{RTT} is the round-trip-time measured by **V**.

each challenge with a single-bit response. To see why the RBE is needed, consider a simple modification of the Čapkun-Hubaux protocol with no delay between the multiple-bit challenge and response. Then, as shown in Figure 3.2, a malicious prover can reply prematurely by guessing only a few bits (challenge or response) [35]. In contrast, with the RBE, any attempt at an early response implies that the adversary needs to guess *all* the bits.

3.1.1 Distance Bounding and Secure Ranging

Distance bounding protocols, as defined so far, are only required to provide a secure upper-bound on the distance. Hence, a physical ND protocol with proximity threshold $r = R$, the nominal communication range, can be trivially converted to a DB protocol. Indeed, a protocol that returns R when a node is at a distance closer than R, and fails otherwise, is a valid DB protocol.

However, we can also add an additional requirement: If there is no adversarial activity, a DB protocol should return an exact distance to the prover (or at least the best-effort upper-bound of such distance). With this extended definition, DB protocols can be considered a form of *secure ranging*. By this we mean that the integrity of the computed range is preserved

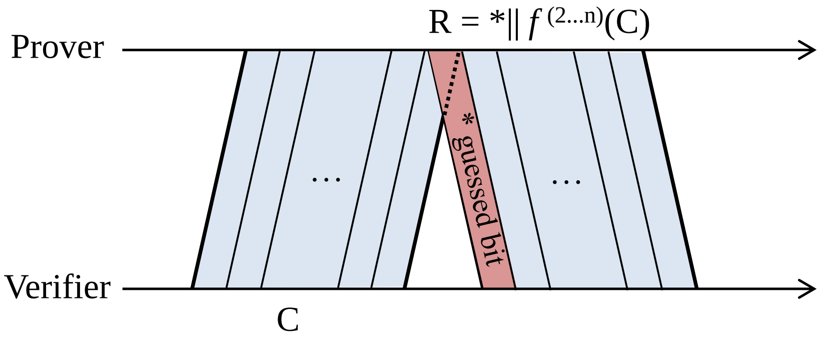

Figure 3.2: Malicious prover attack against a DB protocol without RBE. The response R that the verifier expects is an arbitrary function of the challenge $f(C)$. The malicious prover sends the response one bit early, guessing the first bit of R (denoted by $*$). The remaining bits of R, i.e., bits (2...n) of $f(C)$, can be computed with the challenge C that is fully received after the first response bit is sent. The probability of success of this attack is $\frac{1}{2}$. The adversary can decrease the distance by the duration of k bits with probability $(\frac{1}{2})^k$.

(to the possible extent, keep in mind that the distance can be increased by an adversary); However, the ranging is not confidential (private). In this Part, we concentrate on such DB protocols.

3.1.2 One-way Ranging and Pseudo-Ranging

DB protocol are build on top of two-way ranging, that uses a two message, challenge-response scheme. However, ranging between two devices can also be performed with only a single message, assuming that the clocks of the two devices are tightly synchronized. In terms of security, such a scheme cannot be secure against an internal adversary (who can send the message prematurely or with a delay); but it can be secured against external adversaries.

Furthermore, *time-difference of arrival* localization algorithms also rely on a single message *pseudo-ranging*. In these schemes, a transmitter sends a single message, and a number of synchronized receivers estimate the differences in *time-of-arrival* (ToA) of this message, which then allows them to estimate the location of the transmitter. Or the other way around - a number of synchronized transmitters sends messages to a single receiver, which estimates the differences in ToA of the messages and computes its location accordingly, exactly like in GPS and other GNSS systems. Both TDOA schemes can be secured against external attackers, or even internal attackers, if some additional assumptions are made ([156, 31], or military GPS).

The PHY attacks presented in this Part work by decreasing the time-of-arrival of a single message, hence they apply to such single-message schemes, as well as DB protocols.

3.1.3 Time-of-Arrival Estimation

The description above abstracts away the PHY issues, which have implications for security. In wireless communication, a packet, and in particular a ranging packet, is typically prefixed with a *preamble*, which the receiver knows in advance. The preamble allows the receiver to perform packet detection, synchronization and time-of-arrival (ToA) estimation in an efficient fashion in the harsh wireless environment, and at relatively low SNR (signal-to-noise ratio). In fact, to improve the performance of packet detection and the ToA estimation precision, the

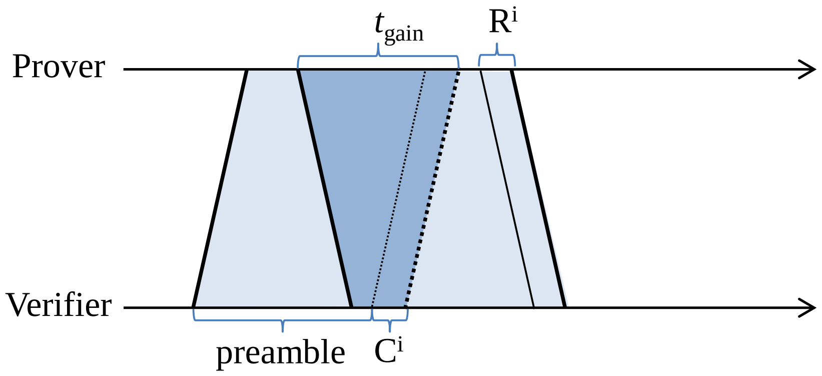

Figure 3.3: Malicious prover attack on the naive implementation of RBE with preamble. The malicious prover replies t_{gain} earlier than an honest prover would, but still receives the challenge bit C^i early enough to compute the response R^i. This decreases the measured distance by $c \cdot t_{\text{gain}}/2$.

preamble is often relatively long. For example, in the IEEE 802.15.4a standard, by "default" (mandatory LPRF mode, Section 4.1.1) the preamble is as long as 248 data symbols, which allows the receiver to combine hundreds of preamble frames for processing gain.

A naive implementation of a DB protocol would put a (long) preamble in front of every single-bit packet in the RBE. However, this would not only be inefficient, but also insecure. Normally, an honest prover would receive the entire challenge (preamble + data bit) before starting to send the response (preamble + data bit). As shown in Figure 3.3, a malicious prover could start to transmit its response prematurely, while the challenge is being received, and still obtain the challenge bit in time to compute the response properly. Essentially, as was observed in [35], an adversary can exploit the presence of any fixed part at the beginning of the ranging packet.

We propose to divide the RBE into two phases (Figure 3.4(a)): In the *ranging* phase, **V** and **P** exchange (long) ranging preambles to estimate the ToA precisely (but insecurely). Second, the verifier *verifies* that the ToA estimates are correct, by sending a number of single-bit challenges without preamble to which the Prover instantly replies with single-bit responses (also without preamble). The verification phase is performed assuming the ToA values estimated in the ranging phase. This allows the challenge and response to be short and free of any fixed parts, preventing the attack in Figure 3.3. Note that it can be possible to perform fine-tuning of the ToA estimate on the challenges and responses.

Note that the range-and-verify scheme can also be applied if the devices use traditional ranging packet structures (Figure 3.4(b)), although this requires full-duplex capabilities from the device [135]. In addition, a number of "buffer" bits should be added at the beginning of the challenge's data part, of total duration equal to the duration of the preamble. This is necessary because the prover cannot start sending the response preamble before it detects the challenge preamble and performs ToA estimation. This is a way to make RBE compatible with the IEEE 802.15.4a, and we are going to assume such an implementation in Chapter 4.

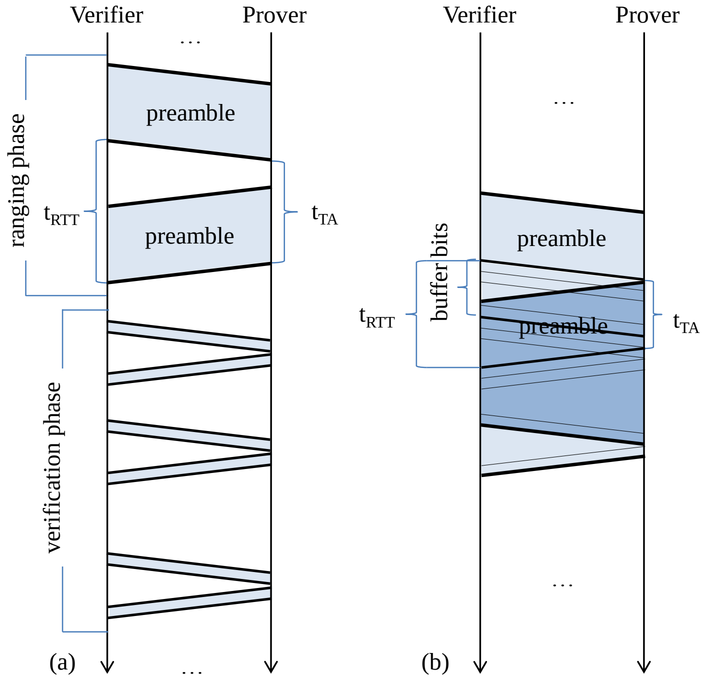

Figure 3.4: Two phase rapid bit exchange implementation (a) compatible with half-duplex transceivers (b) compatible with traditional packet format (in particular, IEEE 802.15.4a), requires full-duplex transceivers.

3.1.4 Security Level and Performance

The challenge and response messages in a DB protocol (Figure 3.1) are essentially random nonces (of length N_{nonce} each), at least from the perspective of the adversary. At the end of the protocol execution, the verifier learns both the true and the received values of these nonces. The verifier accepts a distance measurement only if the received challenge and response messages contain less than N_{err} erroneous bits each. The parameters $N_{\text{nonce}}, N_{\text{err}}$ jointly determine 1) the maximum bit error rate (BER) that the protocol tolerates and 2) the security level, i.e., the the probability that the adversary will succeed in decreasing the measured distance below the actual distance. We derive simple formulas that captures the relation between these two parameters (used in Chapter 4 and Chapter 5). For more elaborate results on this relation, we refer the reader to [144, 122, 98].

We set aside PHY attacks, and assume that the DB protocol is cryptographically secure and that the authenticator is too long to be guessable. Then, the adversary is limited to guessing attacks on the nonces: 1) a malicious **P** guesses the challenge or response and replies early to **V**'s challenge; 2) an external adversary guesses the response and replies early in place of the honest **P**; 3) an external adversary guesses the challenge, sends it early to **P** to extract the correct response, and sends this response to **V**. For the Brands-Chaum protocol the success probability of such guessing attacks is:

$$P_{\text{guess}}^{\text{BC}} = F_{\text{BIN}}(N_{\text{err}}|N_{\text{nonce}}, \frac{1}{2}) \tag{3.1}$$

where $F_{\text{BIN}}(x|n,p)$ is the CDF of a binomial distribution with parameters n and p. For the Hancke-Kuhn protocol the adversary's probability of successfully guessing one bit is $\frac{3}{4}$. This is because for half of the challenges the response is the same, and hence in case 1) no guessing is required and in case 3) an incorrect guess of the challenge will go unnoticed. Hence the security level is:

$$P_{\text{guess}}^{\text{H-K}} = F_{\text{BIN}}(N_{\text{err}}|N_{\text{nonce}}, \frac{1}{4}) \tag{3.2}$$

We focus on the Brand-Chaum protocol. Inverting the CDF yields the maximum N_{err} achieving a desired security level P_{guess}:

$$N_{\text{err}} = F_{\text{BIN}}^{-1}(P_{\text{guess}}|N_{\text{nonce}}, \frac{1}{2}) \tag{3.3}$$

Interestingly, for a fixed security level, this allows a Brands-Chaum-like DB protocol to operate at virtually any bit error rate by simply increasing N_{nonce} and N_{err}. (We will see in Section 4.4.1 why this is an important property for potential countermeasures). Indeed, we have that:

$$N_{\text{err}} = F_{\text{BIN}}^{-1}(P_{\text{guess}}|N_{\text{nonce}}, 1/2) = F_{\mathcal{N}}^{-1}(P_{\text{guess}}|N_{\text{nonce}}/2, N_{\text{nonce}}/4) \tag{3.4}$$

where $F_{\mathcal{N}}^{-1}(x|\mu, \sigma^2)$ is the inverse of the CDF of a normal distribution with mean μ and variance σ^2 and the second equality follows from a normal approximation of the the binomial distribution. The protocol succeeds as long as there are no more than N_{err} errors in a nonce of length N_{nonce}, resulting in a bit error rate of $\text{BER}^{\max} = N_{\text{err}}/N_{\text{nonce}}$. From (3.4) it then follows that, as the length of the nonce increases, the maximum sustainable bit error rate $\text{BER}^{\max}$ tends to the worst case of 1/2, i.e.,

$$\lim_{N_{\text{nonce}}\to\infty} \text{BER}^{\max} = \lim_{N_{\text{nonce}}\to\infty} \frac{N_{\text{nonce}}/2 + \sqrt{N_{\text{nonce}}/4}\Phi^{-1}(P_{\text{guess}})}{N_{\text{nonce}}} = \frac{1}{2} \tag{3.5}$$

where $\Phi^{-1}(x)$ denotes the inverse CDF of a standard normal distribution. In contrast, for Hancke-Kuhn style protocols the maximum sustainable bit error rate tends to 1/4. Unless stated otherwise, we assume that a Brands-Chaum style protocol is used throughout the thesis.

3.1.5 Choosing the Nonce Length

DB protocols make use of both ranging and communication packets. We have seen in Section 3.1.4 that ranging packets carrying nonces can support very high bit error rates and still achieve the desired security level P_{guess}, provided that the coding rate is properly adjusted through the parameters N_{nonce} and N_{err}. In contrast, communication packets do not necessarily offer the same flexibility because the coding rate can be fixed (as in the case of IEEE 802.15.4a, see Section 4.1.1). Consequently, if we define a performance goal in terms of the maximum tolerable packet error rate for communication packets of a given length PER_{comm}, there is a minimum signal-to-noise ratio (SNR) $\text{SNR}_{\min}$ required to reach this goal. Given PER_{comm}, $\text{SNR}_{\min}$ can be established analytically or with simulations.

We can define a similar performance goal for ranging packets by fixing their maximum tolerable packet error rate PER_{db}. Now, if ranging and communication ought to have the same operating range, ranging packets should achieve PER_{db} at $\text{SNR}_{\min}$. To guarantee this, we can first establish (again analytically or with simulations) the bit error rate BER_{db} that ranging messages experience at $\text{SNR}_{\min}$. A ranging packet of length N_{nonce} subject to BER_{db}

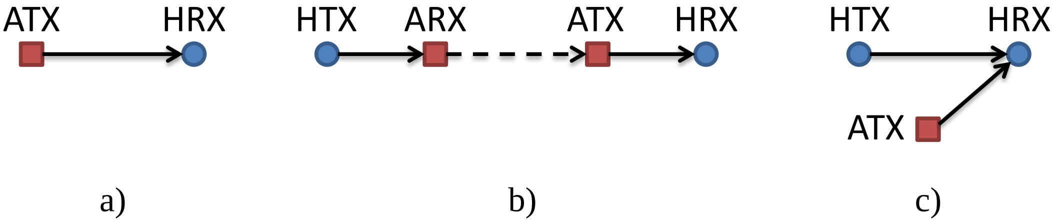

Figure 3.5: Attack types: a) malicious prover attack b) distance-decreasing relay attack c) malicious interference attack. A/H TX/RX stands for adversarial/honest transmitter/receiver. Arrows denote transmissions, the dashed line denotes the adversarial channel.

is considered to be in error if it contains more than N_{err} errors. We can thus derive N_{nonce} and N_{err} by solving the system of equations formed by (3.3), that ensures the desired security level, and

$$N_{\text{err}} = F_{\text{BIN}}^{-1}(1 - \text{PER}_{\text{db}} | N_{\text{nonce}}, \text{BER}_{\text{db}}) \tag{3.6}$$

that ensures the required PER_{db}. Note that we assumed here that the packet reception errors due to factors other than the failure of data demodulation are negligible.

3.2 PHY Distance-Decreasing Attacks

PHY attacks aim at decreasing the estimated packet time-of-arrival while keeping the data payload of the packet unchanged. Hence, in our investigation we focus on the transmission and reception of a single ranging packet. The extension of such an attack to the entire protocol session is straightforward.

Distance-decreasing PHY attacks can be classified as follows (Figure 3.5):

- *Malicious prover attack*: an *internal* attack mounted by a malicious prover against an honest verifier. The malicious prover acts as a *adversarial transmitter* (ATX), and the honest prover acts as a honest receiver (HRX).
- *Distance-decreasing relay attack*: an *external* attack in which the adversary relays a packet between two honest devices, an honest transmitter (HTX) and an HRX. In the relaying process, the packet is received by an adversarial receiver (ARX), forwarded to ATX, which then retransmits the packet to HRX.
- *Malicious interference attack*: an *external* attack that consists in introducing interference. This interference, generated by ATX, interferes with an honest packet transmitted by HTX at HRX.

In Chapter 4 we deal with the first two attack, whereas Chapter 5 is mostly devoted to the latter attack.

Intuitively, we define the *distance-decrease* of a PHY attack as the difference between the actual distance between the prover and the verifier, and the distance that the verifier computes under the PHY attack. A more precise definition is given in Section 4.2.

3.2.1 Attacks on Overlaying Applications

Physical Access Control Consider the RFID physical access control example discussed in Section 1.2 and depicted in Figure 1.1(a). Assume that the normal relay attack (Section 1.3,

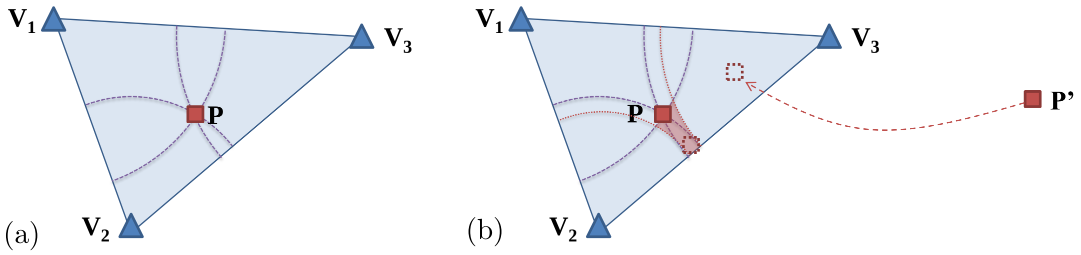

Figure 3.6: Secure localization with verifiable multilateration. (a) Normal operation: The small triangles are the actual locations of the verifiers $\mathbf{V}_i$, the square is the actual location of the malicious prover **P**, the large shaded triangle (*verifier triangle*) is the area for which the verifiers can verify location claims, and the dashed circles centered at the verifiers (*verifier circles*) represent the distance from the respective verifier to **P**. Observe that any location within the verifier triangle (other than the location of **P**) lies strictly inside at least one of the verifier circles. This implies that to claim such a location, **P** needs to decrease at least one of the distances measured by a verifier with a DB protocol. (b) Effect of distance-decreasing malicious prover attacks. Assume that **P** can decrease the distance to $\mathbf{V}_2$ and $\mathbf{V}_3$ by a small amount d (denoted by the dotted circles). Then **P** can claim any location within the red shaded area. In particular, this can be the location marked by the dashed square, which is remote from **P**'s true location compared to d. Furthermore, a prover **P**' that can decrease the measured distance by a large amount can claim any location within the verifier triangle, even if it is itself located outside this triangle.

Figure 1.1(b)) is prevented by a physical ND protocol or a DB protocol. However, if the PHY is susceptible to a *distance-decrease relay attack* in the order of, e.g., 100 meters, then an adversary can use this attack to gain access to the physical resource protected by the reader (the building in Figure 1.1). The adversary simply needs to find a card at distance $d < 100$m form the reader, and use the distance-decreasing relay attack, setting the distance-decrease to d.

Next, consider a car PKES system that uses IR for communication between the car key and the car. Imagine that the car owner, with the key in her pocket, parks the car in front of a small store, and goes inside. An adversary can then mount a *malicious interference attack* with distance-decrease of 19 meters. This will make the car believe that the car key is 1 meter away, whereas it is actually 20 meters away. As a result, the car door will open, giving the adversary access to anything the owner might have left inside.

Secure Localization Consider a secure localization system based on *verifiable multilateration* [155]. In such a system, three (or more) honest anchor nodes that know their location (verifiers), verify the location claim of a (malicious) prover (Figure 3.6(a)). The location claim needs to be positioned with the verifier triangle. Each verifier performs DB with the prover, and checks if the measured distance corresponds to the distance between the verifier own location and the prover's claimed location. The claimed location is accepted only if all checks are successful. The security of this schemes stems from properties of DB: To successfully claim any false location within the verifier triangle, the malicious prover has to decrease at least one of the prover-verifier distances.

Assume that the PHY is susceptible to *malicious prover* distance-decreasing attacks.

Then, even if the maximum achievable distance-decrease d is small, the malicious prover **P** can spoof its location quite considerably, as depicted in Figure 3.6(b). If the achievable distance-decrease is large, than even a prover **P**' located outside the verifier triangle can claim any location within this triangle. Also a prover outside of the nominal communication range can achieve this, assuming that it uses a high gain antenna to increase its communication range.

3.3 IR System Model and Assumptions

Ultra-wideband (UWB) is a physical-communication-layer technology that is characterized by a very large bandwidth. The FCC defines a signal as UWB is its absolute bandwidth is above 500MHz, or it its fractional bandwidth is above 20%. Because of this, UWB transmissions are by regulation required to have relatively low power, not to interfere with existing narrow-band systems. These restrictions result in a relatively low communication range (20-30 m).

Impulse-Radio Ultra-Wideband (IR-UWB or simply IR) is one flavor of UWB[3], characterized by the use of nanosecond pulses. Such narrow (in the time-domain) pulses give IR the capability of high (sub-meter) precision indoor ranging, even in dense multi-path environments, and with relatively low complexity [64]. IR-UWB is envisioned to be used for tracking and access control with a new generation of RFID [100], as well as high-precision localization [64], including secure localization [172]. It has been also advocated as a PHY for distance bounding protocols [66].

3.3.1 Multipath Channel

We adopt the tapped-delay-line channel model, which is commonly used for UWB. The channel impulse response is:

$$h(t) = \sum_{l=0}^{L-1} \alpha_l \delta(t - \tau_l) \tag{3.7}$$

where L is the number of paths (taps), τ_l and α_l are the delay and attenuation of the lth path, and δ is the Dirac delta function. We assume that the channel is invariant for the duration of one packet. We note that the first path is not necessarily the strongest path; we term such a channel *non line of sight.*

A channel parameter most relevant for our investigation is the *channel delay spread* (or simply *channel spread*), that determines how much the wireless channel spreads a single pulse in time. We can define it as the time difference between the last the the first path. However, in practice the last path is hard to determine, because it depends on acceptable attenuation. For this reason, the channel spread is defined such that it captures the majority of the received energy.

The signal observed at a receiver is:

$$r(t) = s(t) * h(t - \nu_0) + n(t) \tag{3.8}$$

where $s(t)$ is the transmitted signal, ν_0 is the unknown (line-of-sight) propagation delay, $*$ denotes convolution, and $n(t)$ accounts for thermal noise assumed to be a zero-mean AWGN process with power spectral density $N_0/2$.

[3] The other flavor being *Multi-Band Orthogonal Frequency-Division Multiplexing* (MB-OFDM).

3.3.2 Wireless Transceivers

We assume that honest devices engaging in distance bounding are equipped with appropriate IR-UWB transmitters and receivers; in particular, in Chapter 4 they are IEEE 802.15.4a compliant. Adversarial devices are equipped with similar transmitters and/or receivers (depending on the attack), but we assume that the adversary can transmit pulse sequences not compliant with the modulation scheme used by honest devices, and can ignore regulatory transmission power limits. The adversary may further equip his devices with high gain antennas, thus allowing him to increase the *signal-to-noise ratio* (SNR) observed by both adversarial and honest devices. Such an increase in SNR can also be achieved by the adversary moving his devices closer to the honest devices. The receivers used by the adversary are modified versions of the receivers uses by honest devices.

The choice of transmitter is of little consequence to our investigation We consider two basic classes of receivers: a low-cost and low-complexity non-coherent *energy-detection receiver* (we also use the term *energy detector*), and a more sophisticated *rake receiver*. An energy-detection receiver is composed of an antenna, a bandpass filter (500MHz), followed by a squaring device and an integrator that outputs a discrete time sample every T_{int}. Unless state otherwise, we assume $T_{\text{int}} = 2\text{ns}$, which gives a sampling rate high enough to allow for precise ranging. A rake reciver is composed of an antenna, a bandpass filter (500MHz) and a filter matched to the pulse shape $p(t)$.

We assume that the receiver is designed to receive packets composed of a *preamble* and a *data* part. The preamble is terminated with a sequence called *start frame delimiter* (SFD), indicating that data is about to start. A receivers always operate in the following basic stages:

- *Coarse synchronization* – the receiver detects a packet (preamble), and achieves a rough synchronization (typically around the strongest path).
- *Fine synchronization / ToA estimation* – following coarse synchronization that provides a rough ToA estimate, the receiver finds a more precise ToA (typically around the first path).
- *Channel estimation* (optional) – the receiver estimates the channel delay profile to improve the performance of data demodulation.
- *SFD detection* – the receiver detects the SFD sequence at the end of the preamble.
- *Data demodulation* – the receiver demodulates the data symbols using the appropriate demodulation method (OOK, BPSK, or BPPM).

In Chapter 4 and Chapter 5 we describe in more detail the modulation scheme and the stages most relevant to the respective investigation.

3.3.3 Simulations

We evaluate the effectiveness of the distance-decreasing attacks and countermeasure with a packet-based system simulator developed in MATLAB. We simulate a full IR-UWB system including all the receiver stages necessary to receive a packet (listed above). The physical layer is simulated with an accuracy of 100 ps. The signal to noise ratio (SNR) is defined as $\text{SNR} = \frac{E_p}{N_0}$ where E_p is the received energy *per pulse* (after the convolution of the pulse with the impulse response of the channel) and $N_0/2$ is the noise power spectral density.

3.4 Related work

The concept of distance bounding, and the first DB protocols (notably the most referenced protocol shown in Figure 3.1(b)), were proposed by Brands and Chaum [23] back in 1993. It was a response to the mafia fraud introduced in [46]. It relies on the idea of accurate time measurements fist proposed in [22]. The main contribution of [23] is the introduction of the rapid bit exchange (RBE). The RBE involves only very simple computational operations (XOR) that can be executed within nanoseconds. This is required to make the RBE secure against malicious provers (the distance fraud in particular). However, it was not until the rise of wireless communications, and notably RFID, that the research on distance bounding gained momentum.

The Hancke-Kuhn DB protocol was proposed in [66]. The structure of this protocol is simpler than the Brand-Chaum protocol (Figure 3.1(c)), with no commitment and no authentication phase after the RBE. This results in a higher success probability of distance fraud and mafia fraud guessing attacks: $(\frac{3}{4})^n$ versus $(\frac{1}{2})^n$ for the Brands-Chaum protocol, where n is the number of bits exchanges (rounds) in the RBE. The Hancke-Kuhn protocol was also the first to address noise during the RBE, by tolerating a certain number of errors. This is a crucial feature, without which DB protocols would not be usable for wireless communications.[4] We note that although the protocol structure used in [66] came to be referred as Hancke-Kuhn style DB in the literature, the first DB protocol that used this structure was proposed in [26] by Bussard and Roudier.

In [155] Capkun and Hubaux propose an authenticated ranging protocol that allows for DB only with an honest prover, i.e., the protocol is secure only against the mafia fraud, not the distance fraud. This relaxation of security guarantees allowed the authors to get rid of the RBE, making the protocol substantially easier to implement, as no rapid response is required on the prover side. Almost all subsequent works on DB protocols focuses on DB protocols that are secure to both distance fraud and the mafia fraud, and in some cases also against the terrorist fraud. The only exception are the protocols proposed in [95]. However, in many application security against the mafia fraud is sufficient: e.g., in physical access control, an adversary that corrupts the prover (i.e., the card) has no need for a distance fraud – he can simply move the card close to the honest reader.

3.4.1 Towards More Efficient DB

Much effort has been devoted to making DB protocols more efficient. For the most part this means decreasing the attack success probability for a given number of RBE rounds n, but other approaches are also used. For example, in [96], the authors propose a Brands-Chaum style DB protocol, in which the there is no commitment, reducing the computational and communication cost.

A number of works attempt to decrease the mafia fraud success probability without resorting to an authentication phase. In [19], the authors propose a modification of the Hancke-Kuhn protocol that uses a set of binary decision trees in place of linear registers for computing the response. Be changing the high of the trees, the system designer can gradually reduce the success probability of the mafia fraud down to $O(n\frac{1}{2^n})$. The success probability of the

[4] Note that error tolerance increases the adversaries probability of success. For simplicity we quote security levels (probability of attack success) for the case where error tolerance is not required.

distance fraud depends on distance-decrease desired by the malicious prover. If the distance-decrease divided by the signal propagation speed is lower than the delay between consecutive challenges, then this probability is $(\frac{3}{4})^n$. Otherwise, the probability gradually drops as the distance-decrease increases. This is because knowing more of the previous challenges makes guessing the response easier. The reduction of the mafia fraud success probability is traded off for memory consumption: For optimal security, $O(2^n)$ memory is required, which is only possible for low n.

Another Hancke-Kuhn style approach is proposed in [85]. This DB protocol mixes random challenges normally used in the RBE, with fixed challenges to which the prover knows the response in advance. Changing the ratio between these challenges allows the protocol to trade off the probability of a successful mafia fraud with the probability of a successful distance fraud. Achieving the optimal mafia fraud security $\frac{1}{2^n}$ comes at a cost of no security against the distance fraud. This trade-off is investigated further in [80]. The authors define *Current Challenge-Dependant* (CCD) protocols in which a response bit depends only on the current challenge bit. They show that for *any* Hancke-Kuhn style CCD protocol, for a single bit, the probability of a successful mafia fraud attack P_{maf} and the probability of a successful distance fraud attack P_{dist} must satisfy $P_{\text{dist}} \geq \frac{3}{4}$ and $P_{\text{dist}} + P_{\text{maf}} \geq \frac{3}{2}$. Further, they define the class of *k-Previous Challenge-Dependent* (k-PCD) protocols, in which the response depends on the current and k previous challenges. They show that for any Hancke-Kuhn 1-PCD protocol $P_{\text{dist}} \geq \frac{5}{8}$ and $P_{\text{dist}} + P_{\text{maf}} \geq \frac{5}{4}$. They also propose two 1-PCD protocols with optimal distance fraud and mafia fraud security, respectively. Note that the protocol in [19] belongs to the n-PCD class. In [151], another n-PCD protocol is proposed, the Poulidor protocol. It exploits a similar idea to [19] but replaces the decision tree with a particular graph, resulting in linear memory consumption. It achieves a better trade-off of distance fraud success probability, mafia fraud success probability and memory consumption than than the protocols in [66, 19, 85, 80].

In [102], the authors propose a protocol that improves the mafia fraud security of the Hancke-Kuhn protocol, decreasing the attack success probability from $(\frac{3}{4})^n$ to $(\frac{3}{5})^n$. However, this is achieved at the cost of introducing void challenges and responses, essentially moving from a binary modulation scheme to a ternary one. Obviously, ternary modulation exhibits a higher bit error rate than binary one (at the same SNR). In [17], the authors extend this idea to using n-ary modulation schemes, and show how to use such schemes to improve the security of Hancke-Kuhn style DB protocols.

In [109], the authors propose a DB protocol in which individual bits in the RBE are replaced with short sequences of bits. This is motivated by the ease implementation. However it makes the protocol vulnerable to packet-level attacks, as described in [35].

3.4.2 Terrorist Fraud Resilience

The first protocol resilient to the terrorist fraud was proposed by Bussard and Bagga [25, 27], but it relies on a relatively computationally expensive proof-of-knowledge protocol. In [137] Reid et al extended the Hancke-Kuhn protocol to be resilient against the terrorist fraud, using much efficient symmetric cryptography primitives.

In [86], the authors propose a DB protocol secure against the terrorist fraud, that offers an optimal security $(\frac{1}{2^n})$ against the mafia fraud, in contrast to the protocol in [137]. This protocol also has na optional privacy feature that prevents an external adversary from learning the identifier of the prover, albeit it is computationally expensive for the verifier.

In [18], the authors investigate using threshold cryptography (secret sharing) to provide security against the terrorist fraud and show that previous protocols resilient to the terrorist fraud all involve a simple form of secret sharing. They also show that protocols in which the threshold is 2 (this includes all previous protocols) are vulnerable to an attack in which an adversary discovers the long-term key. However, to be able to mount such an attack the adversary needs to be able to learn if a given round of the RBE succeeded or failed.

Other protocols secure against the terrorist fraud have been proposed, but they have all been broken, as we discuss next.

3.4.3 Attacks and Design Flaws

A number of minor and major attacks against published DB protocols, as well as errors in their informal analysis have been discovered over the year. This confirms that designing DB protocols is a challenging task, and makes an argument for formal analysis.

In [152], the authors proposed a terrorist fraud resilient DB protocol. It was later extended in [79] to to the case of multiple provers. First, in [101] the authors discovered some flaws in the protocol design. Then both protocols were essentially broken by a severe attack discovered in [86]. This attack allows to extract bits of the *long-term* secret key. It relies on introducing an error into the RBE, and observing if the verifier accepts the prover despite the error. This attack also applies to the variant of the Reid et al protocols published in [137], but not to an earlier, slightly more computationally expensive version of this protocol published in a technical report.

A *full disclosure attack* is discovered in [122]. The attack allows an eavesdropper to extract the long-term secret key of the honest devices by eavesdropping on multiple (in the order of thousands) DB sessions. The authors also propose the Hitomi DB protocol that offers a better security against the full disclosure attack.

In terms of error in the analysis, in [98], the authors show that the security of the Reid et al protocol [137] against the mafia fraud is $(\frac{3}{4})^n$, as was claimed by the authors or the protocol, and not $(\frac{7}{8})^n$ as was reported in [125]. In [103], the authors show that the security of the noise-tolerant mutual DB protocol in [143] is considerably lower than claimed by the authors.

In [16], the authors discover an interesting variant of the distance fraud guessing attack against a version of the Hancke-Kuhn protocol (and other similar protocols). This attack applies if in the phase before RBE the verifier transmits his nonce before the prover (reverse of the nonce sending order in Figure 3.1(c)). Under normal circumstances there is a 50% probability that in the RBE a response is identical for both possible challenge values, in which case the prover does not need to guess the response. However, after receiving N_V, the prover can test a number of candidates for the nonce N_P and choose the one that maximizes the number of identical responses, increasing his probability of guessing attack success.

In [38], a new type of *distance hijacking* attack is revealed. In this attack, a malicious prover $\mathbf{M}$ "hijacks" a DB protocol execution between the verifier $\mathbf{V}$ and an honest prover $\mathbf{P}$. As a result, $\mathbf{V}$ believes that $\mathbf{M}$ is as close as $\mathbf{P}$, while in reality $\mathbf{M}$ can be further away. Such an attack applies to most protocols in the Brands-Chaum style. Instances of this attack against specific protocols were discovered by other authors, e.g., in [94, 146], but it was not identified as a general attack class overlooked in the traditional threat model list. The authors propose a relatively simple fix. However, they also reveal a more fundamental problem: For any DB protocols it is possible to construct another DB protocol such that if both protocols are combined, the original protocol becomes vulnerable to the distance hijacking attack. The authors recommend that different DB protocols are implemented on incompatible PHYs.

3.4.4 Analysis

The analysis of DB protocols is typically performed in an informal fashion, and the assumptions and terminology tend to differ from paper to paper. To partially remedy this situation, Avoine et al propose a unified framework for analysis of DB protocols [16]. They provide a list of threat models: the impersonation fraud (an adversary manages to impersonate an honest prover), the distance fraud, the mafia fraud and the terrorist fraud. They define strategies for querying the prover: pre-ask and post-ask. Furthermore, they distinguish between the black-box and the white-box model. In the former, the malicious prover cannot modify of observe or modify the execution of the algorithms on the card, he can only query it. Establishing the security of the protocol boils down to evaluating all the threat models and strategies, and computing the probability of the resulting attack scenarios. Two DB protocols proposed in [102] are evaluated as an example, and a flaw is found in one.

Although this framework improves on ad hoc arguments for security, it has a drawback. Notably, it only evaluates the security against known attack scenarios, and there is no guarantee that this list is exhaustive. As noted above, recently a new threat model, the distance hijacking attack, has been discovered in [38]. This attack scenario is not included in [16]. Formal verification of protocols can provide stronger proofs of security, although existing approaches [95, 141, 20, 94, 146, 120] have other limitations. We survey them in in Section 2.5.

3.4.5 DB with Additional Features

In [157], a Brands-Chaum style DB protocol that allows for mutual distance bounding (MAD) is proposed, that allows both the prover and the verifier to obtain a distance bound. In [143], the authors extend the MAD protocol to cope with noise during the RBE. The protocol proposed in [86] includes an option for mutual authentication, but not mutual distance bounding (i.e., only the verifier obtains a distance bound). In [169], the authors propose a mutual DB protocol based on the Hancke-Kuhn style. The protocol provides a low success probability for a mafia fraud, $(\frac{3}{8})^n$. However, it requires the ability to detect a lack of challenge in a round (much like the void challenge in [102]).

One aspect of privacy has been addressed in [86], where the identifier of the prover remains secret from anyone but the verifier. In [135], the authors explore a different privacy threat introduced by DB protocols. They show that an eavesdropper can learn the distances between and/or locations of devices performing a DB protocol. They then propose a DB protocol that is secure against such attacks. One of the features of this protocol is performing the RBE using the full-duplex approach (as in Figure 3.4(b)).

In [148], the authors propose a method to make DB protocols resistant to external distance-increasing attacks. To achieve this, they assume that the honest prover is in power range of the verifier. Then, they propose to use an on-off keying modulation scheme, which would prevent an external adversary from "hiding" the provers message from the verifier, and replaying it with a delay. We note that this assumption is rather hard to guarantee in practice.

In [121], the authors combine Brands-Chaum style DB with cryptographic puzzles. The goal is to mitigate attacks from a malicious reader, by forcing the reader to solve a cryptographic puzzle before the card reveals its identity to the reader. Hence, the usual DB roles of the reader and the tags are reversed, the card being the verifier.

3.4.6 PHY Attacks and PHY Design

Physical layer attacks against distance bounding were discovered in [35]. In particular, the authors introduced the early detection and late commit PHY attack primitives and shown how to use them to mount malicious prover and distance-decreasing relay attacks. In addition, attacks exploiting clocks skews were discovered, as well as packet level attacks (Figure 3.2 and Figure 3.3 show representative examples of these attacks). The effectiveness of the late commit attack against concrete PHYs is studied in [68] (ISO 14443 RFID and wireless sensor networks).

An IR-UWB architecture for implementing DB protocols is proposed in [87]. The maximum distance-decrease an adversary can gain against this PHY is $3-6$m. This is achieved with the short symbol duration of 20ns, which limits the applicability of this PHY in dense multipath environments for which IEEE 802.15.4a was designed. An ID-based distance bounding protocol is implemented on proprietary IR radios in [149].

In [21], the authors propose a IR-UWB PHY based on time-hopping. They propose to encode one symbol with a number of pulses distributed according to a secret time-hopping sequence. This approach offers a good level of security against the mafia fraud, but it is not secure against malicious prover attacks, in which the adversary would known secret time-hopping codes, allowing him to mount effective early detection and late commit attacks.

Beyond IR-UWB, a number of DB PHYs tailored to narrow-band HF RFID systems have been proposed. In such systems, the distances between the prover and the verifier are below 10cm, much shorter than the IR-UWB range. At these ranges, the multipath channel spread is practically negligible, hence the ToA estimation vulnerabilities underlaying Chapter 5 are not relevant. In [137], the authors propose a PHY that based on experimental results achieves a timing resolution between 37ns (at 1cm distance) and 300ns (at 5cm distance). The PHY proposed in [104] achieves a timing resolution of 0.5μs (at 5cm distance). In [69], a solution that integrates the HF RFID PHY with wideband pulse detection is proposed, limiting the possible distance-decrease to 1m in case of honest prover and 11m in case of a malicious prover, at similarly short distances.

In [51], the authors first demonstrate the credit card *Chip&Pin* system deployed in the UK is vulnerable to relay attack, and then implement a DB protocol with modest modification to the existing hardware and software. This implementation limits the achievable distance-decrease to at most 6m. It should be noted that this system is wire-line, which eliminates the challenges that wireless implementations have to face.

In [136] Rasmussen and Capkun propose a new, fast method for implementing the rapid response at the prover, which they term *Challenge Reflection with Channel Selection* (CRCS). This method is applicable to Brands-Chaum style DB protocols. Traditionally, it is assumed that the challenge symbol is demodulated by the prover, which performs a simple operations on it (like XORing it with a response), and then sends the response symbol, which incurs some processing delay (at least a few nanoseconds). In CRCS, the challenge is reflect by the prover after shifting it in frequency to one of the two response channels, chosen according to the response bit. CRCS can be implemented in an analog fashion, incurring negligible processing delay (below 1ns). In addition, the challenge is demodulated, as it must be included in the after-RBE authentication phase, but this process is no longer time-critical. Because of the negligible processing delay of the CRCS process, and because reflecting starts when the challenge symbol is starting, this method practically removes the threat of a malicious prover using early detection for distance-decrease. To mitigate the threat of late commit attacks, the authors propose to use auto-correlation for ToA estimation. However, this results in strongest path detection. If DB with higher precision in NLOS environments is required, the security issues of ToA estimation apply (Chapter 5). Note that the CRCS method cannot be applied to a large number of DB protocols, notably the Hancke-Kuhn protocol and the protocols which stem from it. The methods is generic, and can by applied to any modulation scheme, including IR ones. But, it requires dedicated hardware for both the prover, and the verifier, which needs a receiver that can simultaneously receive signals in two frequency bands.

3.5 Summary

In this Chapter, we have provided the context for Chapter 4 and Chapter 5. We have given an introduction to distance bounding (DB). We have covered the basics of PHY attacks and introduced Impulse Radio Ultra-wideband (IR). We have also surveyed the related work.

Chapter 4

Physical Layer Attacks against IEEE 802.15.4a

physical-communication-layer (PHY) attacks against distance bounding were first described in [35]. In particular, two attack primitives are introduced: *early detection* (ED) and *late commit* (LC). The adversary can use these primitives to mount a *malicious prover* attack, or a *distance-decreasing relay* attack. The specifics of the attacks, as well as their effectiveness depends on the PHY. In [35], amplitude shift keying is used as a running example to illustrate and argue the feasibility of the attacks, but no quantitative analysis is given. In a follow-up work [68], Hancke and Kuhn demonstrated with an implementation that LC attacks are feasible for the ISO 14443 PHY and a compliant RFID receiver, as well as 433MHz ASK/FSK modulation and a super-heterodyne receiver, popular in wireless sensor networks. However, none of these technologies were designed for ranging.

In this Chapter, we undertake such an investigation for IR-UWB, and more precisely the IEEE 802.15.4a standard [77] that is particularly well suited for ranging. We adapt the ED and LC attacks to IEEE 802.15.4a and evaluate their effectiveness with detailed PHY simulations. We also examine countermeasures that can mitigate PHY attacks, while minimally degrading the benign-case performance. We make the following contributions:

▶ We show that if honest devices use energy-detection receivers (popular due to their low cost and low complexity), an adversary can mount PHY attacks that decrease the measured distance in the order of hundreds of meters. The adversary can achieve this with energy-detection receivers. Furthermore, by increasing the signal-to-noise-ratio (SNR), the attack success rate can be made arbitrarily large. In particular, to achieve a success rate of 99%, the adversary needs to operate at an SNR only slightly higher (a few dB) than needed for normal system operation. This is easily achievable by using a high-gain antenna, by transmitting at a power exceeding the regulatory limit, or by simply moving the adversarial devices closer to the victim devices.

▶ We observe that, in order to mount a distance-decreasing relay attack, it is not enough to attack the data part of a packet. The attack also has to be extended to the preamble. We then develop ED and LC attacks for the preamble. The cost of these attacks (in terms of SNR) is only a few dB compared to normal system operation.

▶ We unveil a security issue with the convolutional code employed in IEEE 802.15.4a: it can be exploited by an adversary equipped with a rake receiver that attacks energy-detection receivers. It allows the adversary to decrease the distance more substantially than the attacks

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

Table 4.1: IEEE 802.15.4a mandatory modes parameter values.

mentioned above, or to mount an *undetectable* relay attack with a more modest, but still substantial, distance-decrease. This security issue can be patched with a small modification to IEEE 802.15.4a.

▶ We show that time-hopping used in IEEE 802.15.4a allows a malicious prover to augment the distance-decrease of PHY attacks at the cost of decreasing the probability of attack success. A rake-equipped malicious prover can also similarly exploit the combination of BBPM and BPSK employed in IEEE 802.15.4a. The time-hopping vulnerability can be patched with a small modification to IEEE 802.15.4a.

▶ We observe that by increasing the length of nonces used in a DB protocol, we can improve its performance (packet error rate) and preserve the security level. We use this observation to advocate a simple, efficient, and standard-compliant countermeasure to PHY attacks. Employed in an energy-detection receiver, along with the convolutional code patch, this countermeasure effectively limits the distance-decrease of distance-decreasing relay attacks to a value in the order of the channel spread (around 10m). Employed in an energy-detection receiver or in a rake receiver, along with the convolutional code and time-hopping patches, the countermeasure limits the effectiveness of malicious prover attacks to a similar value. In both cases, the cost of the countermeasure in terms of required packet length is modest. We emphasize that these countermeasures preserve the features of IEEE 802.15.4a that improve benign case performance, in particular time-hopping.

Chapter Outline Our basic assumptions about the wireless channel and receivers are presented in Section 3.3. In Section 4.1 we provide the IEEE 802.15.4a specification and additional assumptions on receivers and the DB implementation. We show a range of PHY attacks in Section 4.2, and evaluate their performance in Section 4.3. In Section 4.4 we discuss countermeasures, before concluding in Section 4.5.

4.1 System Model

4.1.1 IEEE 802.15.4a Modulation Scheme

IEEE 802.15.4 is a standard for low-rate wireless personal area networks (WPAN). The IEEE 802.15.4a amendment [77] defines an IR-UWB PHY that allows for low-rate communication and high precision ranging. The IEEE 802.15.4a standard is very flexible, allowing for many combinations of parameter values. However, only two of these combinations need to be implemented by a standard compliant device: one LPRF and one HPRF mode (*high/low pulse repetition frequency*). We focus on the LPRF mode in our investigation. However, our results are easily transferable to other modes (which use different parameter values). The publicly known parameter values most important for our investigation are summarized in Table 4.1. An IEEE 802.15.4a packet is composed of a preamble followed by data.

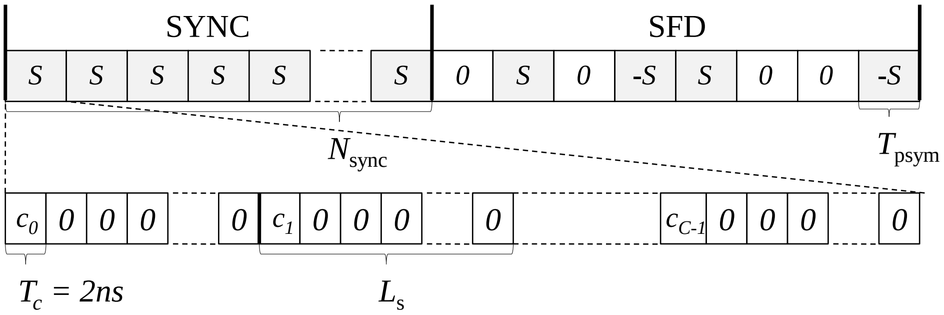

Figure 4.1: IEEE 802.15.4a preamble structure

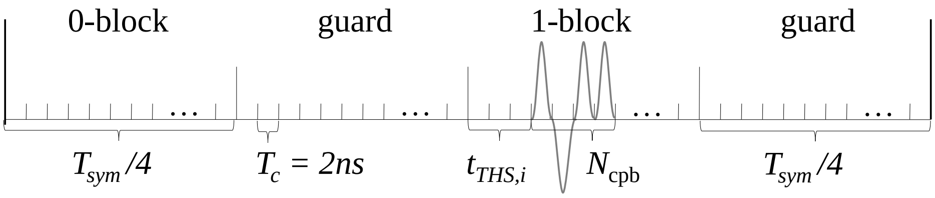

Figure 4.2: IEEE 802.15.4a data symbol structure

Preamble

The preamble (Figure 4.1) consists of a SYNC part and the start frame delimiter (SFD). The SYNC part is composed of $N_{\text{sync}} = 64$ identical preamble symbols of duration T_{psym}. The SFD is composed of a particular sequence of $N_{\text{sfd}} = 8$ preamble symbols. Each preamble symbol is formed by C preamble frames that consist of L_{s} chips of duration $T_{\text{c}} = 2\text{ns}$ (frame duration $T_f^{\text{P}} = L_{\text{s}} \cdot T_{\text{c}}$). Pulses are sent in the first chip of every frame and modulated according to a ternary preamble code of length C. The transmitted signal is given by:

$$s^{\text{P}}(t) = \sum_{i=1}^{N_{\text{sync}}+N_{\text{sfd}}} s_i \sum_{j=1}^{C} c_j \cdot p(t - jL_{\text{s}}T_{\text{c}} - iT_{\text{psym}}) \tag{4.1}$$

where $p(t)$ is the pulse shape, $c_j \in \{-1, 0, +1\}$ are the elements of the preamble code and ν_0 is the propagation delay. Each preamble symbol is modulated by the sequence $s_i = [1, \dots, 1, 0, 1, 0, -1, 1, 0, 0, -1]$ whose last eight elements denote the SFD.

Private Ranging Mode The standard includes a private ranging mode whose main purpose is to prevent an adversary from learning sensitive ranging information. To this end the private ranging mode allows for the encryption of timestamp information that is exchanged during the ranging process. More importantly, the preamble codes used in the ranging packets are secretly agreed on by the legitimate participants. However, this provides a minor increase in security, because the nodes are only allowed to choose from a set of 8 predefined preamble codes. We explain this in detail in Section 4.4.2.

Data

The data (Figure 4.2) is modulated using a combination of *binary pulse-position modulation* (BPPM) and *binary phase-shift keying* (BPSK). In addition, time-hopping is used to allow for multiple-access and every data symbol (synonymous with data frame) is signalled through the transmission of a burst of N_{cpb} pulses. The signal corresponding to the i-th data symbol is:

$$s_i^{\text{D}}(t) = (2a_i - 1) \sum_{j=1}^{N_{\text{cpb}}} b_{ij} \cdot p(t - iT_{\text{sym}} - d_i T_{\text{sym}}/2 - t_{\text{THS},i} - jT_{\text{c}}) \tag{4.2}$$

where $p(t)$ is the pulse shape, a_i is the polarity bit (BPSK), d_i the position bit (BPPM), $T_{\text{sym}} = T_f^{\text{D}} = 1024ns$ is the symbol duration, $t_{\text{THS},i} \in [0, t_{\text{THS}}^{\max}]$, where $t_{\text{THS}}^{\max} = T_{\text{sym}}/4 - N_{\text{cpb}} \cdot T_{\text{c}}$, defines the pseudo-random time-hopping offset, and the scrambling sequence b_{ij} defines the polarity of the j-th pulse of the i-th burst. Both the time-hopping and the scrambling sequences are derived from a fixed and publicly known linear feedback shift register that is initialized to a publicly known state at the beginning of every packet. Both sequences are thus the same for every packet.

Channel Coding A systematic rate 1/2 convolutional code with generator polynomials $g_1 = (0, 1, 0)$ and $g_2 = (1, 0, 1)$ is used. Denote the bits to be transmitted by x_i. Then the position bit is $d_i = x_i$ and the polarity bit is $a_i = x_{i-1} \oplus x_{i+1}$, where $\oplus$ denotes modulo two addition. With this construction, an energy-detection receiver, which cannot recover the polarity bit, can still decode the transmitted bit sequence x_i; whereas a coherent receiver, which can recover both bits, can apply convolutional decoding to improve performance. IEEE 802.15.4a also applies a systematic (55,63) Reed-Solomon (RS) code before modulation.

4.1.2 Receivers

Energy-Detection Receiver

The receiver follows the stages sketched in Section 3.3.2: It employs a traditional synchronization algorithm based on a correlation with the known preamble sequence (details can be found in Section 5.1). The receiver then performs a period of channel estimation where it estimates the energy-delay profile of the channel by averaging a number of preamble symbols.

The default method for SFD detection is based on on-off keying demodulating (maximum likelihood rule) the preamble symbols, and soft-decoding a sequence of length $N_{\text{sfd}} = 8$ until the SFD sequence 01011001 is found. An alternative method is based on correlation with with SFD template. The latter method has inferior performance: In the benign case it requires around 2dB more in SNR to achieve the same performance as the method based on decoding. However, we evaluate both methods for completeness.

By default, to demodulate the i-th BPPM data bit d_i, the receiver uses the optimum decision rule from [39], [57], comparing the (weighted) energies in the first and second half of the symbol:

$$\sum_{m=0}^{M-1} y_{m,i} \cdot p_m \underset{d_i=1}{\overset{d_i=0}{\gtrless}} \sum_{m=0}^{M-1} y_{m+\frac{T_f}{2T_c},i} \cdot p_m \tag{4.3}$$

where $y_{m,i}$ denotes the m-th discrete sample of the i-th symbol. The weighting coefficients p_m are derived from the energy-delay profile of the channel. The number of samples to combine is $M = t_{\text{det}}/T$, where t_{det} defines the *detection time*: the length of the received signal (per half-symbol) that the receiver uses to demodulate the bits; t_{det} is chosen to be large enough to account for the channel delay spread. In our simulations we use the non-line-of-sight residential channel model from [99] and set $t_{\text{det}} = 60\text{ns}$ accordingly.

We also consider an alternative method for data demodulation, in which the receiver does not apply the weights p_m. (Note: This is equivalent to sampling at much lower rate with integration time t_{det}.) The alternative method requires around 2dB more in SNR to achieve the same performance as the default method.

Rake Receiver

The receiver with optimal performance (in a benign setting), but also with the highest complexity, is an all-rake receiver using maximum ratio combining (MRC) [132]. The convolutional code is decoded with the optimal symbol-wise branch metric for BPPM/BPSK given in [13]. For this section, the crucial difference between an energy-detection receiver and a rake receiver is that the latter can recover the polarity bits a_i during data demodulation. This diminishes the effectiveness of data PHY attacks against rake receivers (Section 4.2.4), but also opens a new space for data attacks if a rake receiver is used against energy-detection receivers. In our analysis of the rake receiver, we therefore focus on the data, assuming perfect synchronization and channel estimation.

Similar to the energy-detection receiver, an important parameter for our analysis is the detection time t_{det}, denoting the portion of the received signal that the rake receiver uses to demodulate a bit. We chose t_{det} large enough to account for the channel delay spread.

4.1.3 Distance Bounding Protocol Assumptions

All the assumptions stated in Section 3.1 hold. In addition, we define t_{res} as the time interval at the prover between the start of the reception of a challenge bit and the start of the transmission of the corresponding response bit. To make the response as "instant" as possible with IEEE 802.15.4a, we assume that $t_{\text{res}} = T_{\text{sym}}/2 + t_{\text{THS}}^{\max} + t_{\text{det}}$. This is the smallest t_{res} sufficient to demodulate symbols with maximal time-hopping offsets.

We assume that the DB protocol has access to the received bit sequence before it is decoded with the error correcting codes (Reed-Solomon, convolutional). In a receiver implementation, this assumption can easily be met as these bits have to be received from the channel in any case. If this assumption is violated, the success probability of guessing attacks increases, because coding can mask some of the erroneously guessed bits.

4.2 Distance-Decreasing Attacks

Physical layer (PHY) attacks considered in this Chapter rely on two primitives: 1) In *early detection* (*ED*), an adversarial receiver (ARX) detects a PHY symbol (e.g., data symbol) of duration t_{sym} based only on the beginning part of this symbol of duration $t_{\text{ED}} < t_{\text{sym}}$, where t_{ED} is the *ED delay*. This leads to a detection which is less reliable, but also faster than that of a normal receiver (which takes $t_{\text{res}} > t_{\text{ED}}$ of the PHY symbol into account for detection). 2) In *late commit* (*LC*), only the $(t_{\text{sym}} - t_{\text{LC}})$-long end-part of the PHY symbol is modulated based on the intended value of the symbol (e.g., whether it encodes a 0 bit or a 1 bit), and the beginning part is modulated independently of this value. This allows an adversarial transmitter (ATX) to delay the decision about which symbol it transmits by t_{LC}, where t_{LC} is the *LC delay*. The PHY symbols generated with LC typically differ from regular

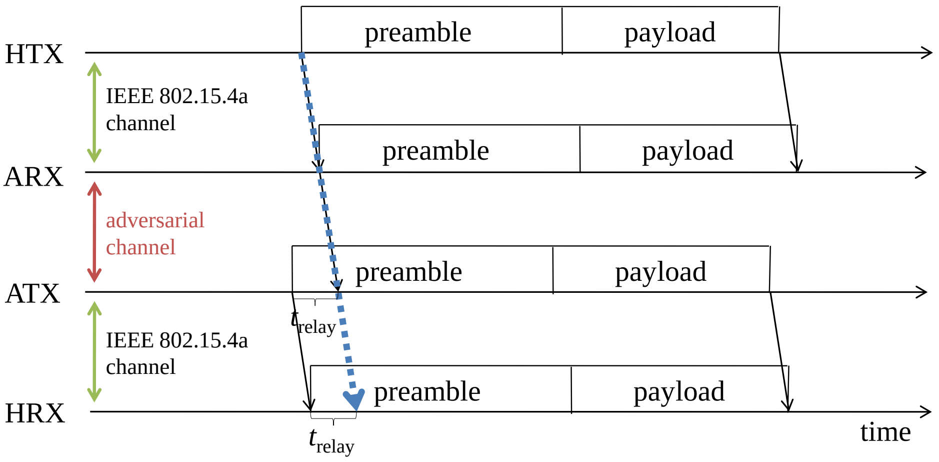

Figure 4.3: Overview of the distance-decreasing relay attack. ARX and ATX are assumed to be located on a line between HRX and HTX. The thick dotted arrow indicates time-of-arrival corresponding to the actual distance between HTX and HRX.

symbols, but, if appropriately chosen, they can be demodulated by an honest receiver, albeit with some performance loss.

The first type of PHY attack we consider is an internal attack mounted by a *malicious prover* (consisting of an ARX and an ATX), which can be used in a distance/terrorist fraud. In this attack, a malicious prover uses ED, LC or their combination to respond prematurely to the verifier's **V** challenge. This decreases the propagation time measured by **V** by an offset t_{gain} that we call the *time-gain*. The time-gain is equal to:

$$t_{\text{gain}} = (t_{\text{C}} + t_{\text{res}} - t_{\text{ED}})/2, \quad \text{for an ED-only attack} \tag{4.4}$$

$$t_{\text{gain}} = (t_{\text{LC}} + t_{\text{res}} - t_{\text{D}})/2, \quad \text{for an LC-only attack} \tag{4.5}$$

$$t_{\text{gain}} = (t_{\text{LC}} + t_{\text{res}} - t_{\text{ED}})/2, \quad \text{for an ED+LC attack} \tag{4.6}$$

where t_{D} and t_{C} are, respectively, the detection delay and commit delay when the adversary chooses not to perform ED and LC, i.e., the adversary performs detection without performance loss and transmits standard-compliant symbols. Note that it is possible that $t_{\text{D}} < t_{\text{res}}$ and $t_{\text{C}} > 0$ if time-hopping is involved. The time-gain translates into a *distance-decrease* of $c \cdot t_{\text{gain}}$, where c is the speed of light.

The second type of attack we consider is a *distance-decreasing relay attack* between two honest devices. This attack is mounted by an external adversary using a combination of ED and LC and it can be classified as a mafia fraud. The general setup for the relay attack is shown in Figure 4.3. The adversary should mount the distance-decreasing relay attack on all ranging messages (challenge, response); other messages, being not time-critical, can be relayed in an arbitrary fashion. Without loss of generality, we focus on the exchange of a single ranging message. In this case one of the honest devices acts as a transmitter (HTX) and the other one as a receiver (HRX), whereas one adversarial device acts as an early detection receiver (ARX), and another as a late commit transmitter (ATX). The channel from HTX to ARX, and from ATX to HRX is the IEEE 802.15.4a channel. ARX and ATX communicate using a dedicated, out-of-band adversarial channel. By default, we assume that HRX does not receive the signal transmitted by HTX, but we also explain what happens if this is not

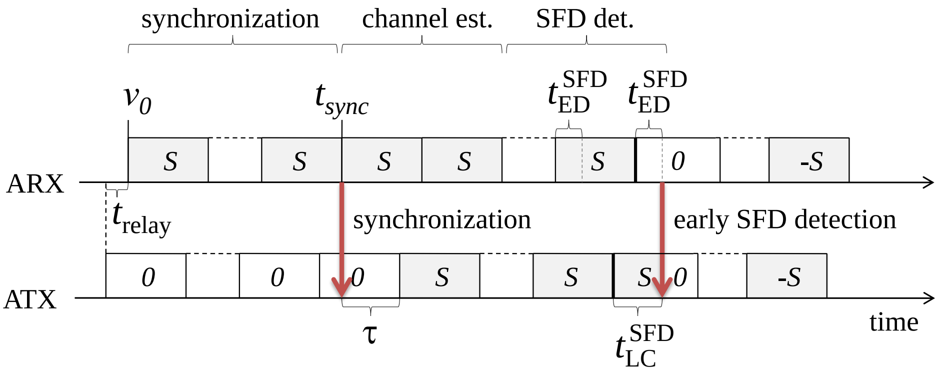

Figure 4.4: Distance-decreasing relay attack on the preamble.

true. The propagation speed of both channels is c, the speed of light.

In a distance-decreasing relay attack the adversary relays messages between HTX and HRX in such a way that to HRX they seem “shifted back in time” by a positive offset

$$t_{\text{relay}} = t_{\text{LC}} - t_{\text{ED}} \tag{4.7}$$

that we call the *relay time-gain* (Figure 4.3). The distance measured by $\mathbf{V}$ is then reduced by the *relay distance-decrease* $c \cdot t_{\text{relay}}$. (Assuming that ARX and ATX are located on a line between HTX and HRX. In other configurations the distance decrease will be smaller. Note, however, that the choice of the configuration rests with the adversary.) In the relay attack, ATX needs to begin the transmission of the preamble at time $t_0 - t_{\text{relay}}$, *before* ARX begins receiving the preamble from HTX a time t_0. However, the adversary learns t_0 only *after* synchronizing to the preamble of HTX. (Guessing t_0 is not practical at the nanosecond precision required.) To escape this vicious circle, the adversary needs to mount some form of ED and LC on the preamble, in addition to ED and LC on the data. In contrast, the preamble attacks are not necessary in the case of malicious prover attacks.

Note that the distance-decrease that the adversary might wish to obtain is not limited by the low communication range of IEEE 802.15.4a. Indeed, a malicious prover can increase his communication range by using a high gain antenna and transmitting with non-regulatory power to reach a remote prover. Further, in a relay attack, the adversary can “connect” remote HRX and HTX by placing ARX close to HTX and ATX close to HRX, and using a long-range ATX–ARX link to which the range limitations of IEEE 802.15.4a do not apply.

We consider three scenarios for data attack, in which the adversary uses different types of receivers against different types of receivers used by the honest devices: Energy Detector against Energy Detector, Rake against Energy Detector, and Rake against Rake. For each scenario we first analyze the delay of the ED and LC primitives. Then, we elaborate on the use of these primitives for malicious prover attacks and relay attacks. Table 4.2 summarizes the upper-bounds on the time-gain and distance-decrease of various variants of PHY attacks. We start by presenting the preamble attack in Section 4.2.1 because this attack applies in all scenarios.

4.2.1 Preamble

The attack, which is part of the distance-decreasing relay attack, is depicted in Figure 4.4; for clarity of presentation, we assume the distance between ARX and ATX to be 0. ARX performs synchronization in the same fashion as an honest receiver. ARX then signals the fact that it has synchronized to ATX. Deviating from honest receivers, ARX performs *early SFD detection*: It chooses an early SFD detection delay $t_{\mathrm{ED}}^{\mathrm{SFD}}$ and tries to detect the presence of the SFD by deliberately considering only the first $t_{\mathrm{ED}}^{\mathrm{SFD}}$ seconds of every received preamble symbol. As the SFD starts with a 0 modulated preamble symbol, as opposed to a 1 modulated symbol used during the SYNC part, early SFD detection boils down to on-off keying (OOK) demodulation.

At the other end of the relay, ATX chooses a late SFD commit delay $t_{\mathrm{LC}}^{\mathrm{SFD}}$ and remains silent until ARX signals that synchronization has been successful. Then, after an appropriately chosen (we explain how shortly) delay $\tau < T_{\mathrm{psym}}$, ATX begins transmitting a sequence of preamble symbols S. This is repeated until ARX signals that the SFD was detected. Immediately afterwards, ATX switches to the transmission of a standard compliant SFD, beginning from $t_{\mathrm{LC}}^{\mathrm{SFD}}$ into the SFD. This concludes the distance-decreasing attack on the preamble.

In contrast to a standard-compliant preamble, the SYNC part of the preamble generated by ATX begins with a number of 0 modulated preamble symbols; The beginning of the SFD corresponds to a 1 modulated preamble symbol for a duration of $t_{\mathrm{LC}}^{\mathrm{SFD}}$, instead of having no signal contribution. The relay time-gain achieved by this attack is $t_{\mathrm{relay}} = t_{\mathrm{LC}}^{\mathrm{SFD}} - t_{\mathrm{ED}}^{\mathrm{SFD}}$. This determines the choice of τ, as $T_{\mathrm{psym}} - \tau = t_{\mathrm{relay}} \mod T_{\mathrm{psym}}$.

4.2.2 Data: Energy Detector against Energy Detector

Data attacks are performed on a symbol basis. As we are considering energy-detection receivers, which are blind to the signal polarity, only the position bit d_i is relevant.

ED ARX performs ED by deciding on the value of d_i^{RX} after an *early detection delay*:

$$t_{\mathrm{ED}}[d_i^{\mathrm{RX}}] = t_{\mathrm{THS},i}^{\mathrm{RX}} + t_{\mathrm{det}}^{\mathrm{A}} < T_{\mathrm{sym}}/2 \tag{4.8}$$

where $t_{\mathrm{det}}^{\mathrm{A}}$ denotes the detection time of ARX and $t_{\mathrm{THS}}^{\mathrm{RX}}$ is the time-hopping offset sequence of the received message. This implies that ARX replaces BPPM demodulation with on-off keying (OOK) demodulation. The time $t_{\mathrm{det}}^{\mathrm{A}}$ can be made arbitrarily short, it determines the attack's performance. If ARX chooses not to perform ED, the detection delay is

$$t_{\mathrm{D}}[d_i^{\mathrm{RX}}] = T_{\mathrm{sym}}/2 + t_{\mathrm{THS},i}^{\mathrm{RX}} + t_{\mathrm{det}} \tag{4.9}$$

LC In the LC attack, ATX always transmits a burst of pulses with energy E_0 (shifted by the appropriate time-hopping offset). In the second half of the symbol, ATX acts according to the value of d_j^{TX}: If $d_j^{\mathrm{TX}} = 0$, ATX transmits nothing in the second part of the symbol; if $d_j^{\mathrm{TX}} = 1$, ATX transmits a burst of pulses with energy $E_1 > E_0$. This attack exploits the fact that HRX performs a simple energy comparison to demodulate. The *late commit delay* for d_j^{TX} is:

$$t_{\mathrm{LC}}[d_j^{\mathrm{TX}}] = T_{\mathrm{sym}}/2 + t_{\mathrm{THS},j}^{\mathrm{TX}} + t_{\mathrm{PLC}} \tag{4.10}$$

where $t_{\mathrm{THS}}^{\mathrm{TX}}$ is the time-hopping offset sequence of the transmitted message, and $t_{\mathrm{PLC}} < t_{\mathrm{det}}$ is the *pulse LC delay*, by which the transmission of the pulse can be additionally delayed, similar

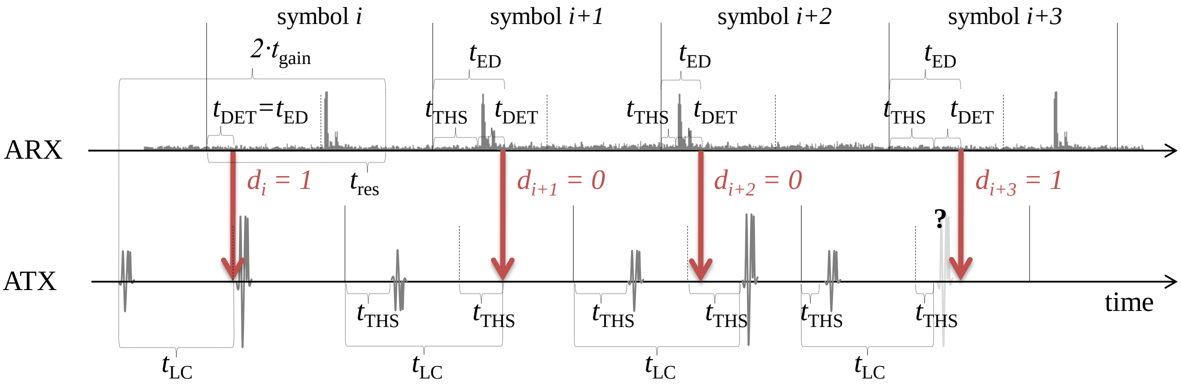

Figure 4.5: Example of a malicious prover ED+LC attack on the data. The first two symbols use equal time-hopping offsets for corresponding symbols, hence their time-gain is identical. The maximal time-gain for the 3rd symbol is larger than the packet-wide time-gain chosen of the adversary. The time-gain for the 4th symbol is smaller than the packet-wide time-gain, and the adversary is forced to guess.

to the LC attacks discussed in [35, 68]. Throughout most of the paper, notably Section 4.3, we assume $t_{\text{PLC}} = 0$. If ATX chooses to send standard-compliant symbols, it can still delay committing to the transmitted symbol by:

$$t_{\text{C}}[d_j^{\text{TX}}] = t_{\text{THS},j}^{\text{TX}} \tag{4.11}$$

Malicious Prover Based on (4.4), the time-gain of the ED-only malicious prover attack for corresponding challenge and response symbols i and j is:

$$\begin{aligned} t_{\text{gain}}[i,j] &= (t_{\text{C}}[d_j^{\text{TX}}] + t_{\text{res}} - t_{\text{ED}}[d_i^{\text{RX}}])/2 = \\ &= (T_{\text{sym}}/2 + t_{\text{THS}}^{\max} + t_{\text{det}} - t_{\text{det}}^{\text{A}})/2 + (t_{\text{THS},j}^{\text{TX}} - t_{\text{THS},i}^{\text{RX}})/2 = C + (t_{\text{THS},j}^{\text{TX}} - t_{\text{THS},i}^{\text{RX}})/2 \end{aligned} \tag{4.12}$$

where C is a constant not dependent on i and j. The time-gain of other malicious prover attacks can also be expressed as $C+(t_{\text{THS},j}^{\text{TX}}-t_{\text{THS},i}^{\text{RX}})/2$. The latter term varies from $-t_{\text{THS}}^{\max}/2$ to $t_{\text{THS}}^{\max}/2$, because $t_{\text{THS}}^{\text{TX}} \neq t_{\text{THS}}^{\text{RX}}$ if different channels are used for for RX and TX, but also because $i \neq j$ due to the "buffer" bits (Figure 3.4(b)). However, the structure of the attack demands that the adversary chooses a constant time-gain t_{gain} for all symbols. This leaves the adversary with a strategic decision: The adversary can set the time-gain conservatively, to make sure there is enough time to perform ED and/or LC on every symbol (i.e., choose $t_{\text{gain}} \leq t_{\text{gain}}[i,j]$ for all corresponding i, j index pairs). Alternatively, the adversary can set the time-gain more aggressively, which will force him to guess the bits with unfavorable time-hopping offsets (i.e., i, j pairs for which $t_{\text{gain}}[i,j] < t_{\text{gain}}$). This is illustrated in Figure 4.5. In this way, the adversary can trade-off a larger time-gain (up to $2 \cdot t_{\text{THS}}^{\max}/2$) for a lower attack success probability. Figure 4.7 shows this trade-off for one particular case ($N_{\text{nonce}} = 42, N_{\text{err}} = 2$, mandatory LPRF mode).

Relay For the relay attack (Figure 4.6), the time-gain, based on (4.7), is:

$$t_{\text{relay}}[i,j] = t_{\text{LC}}[d_j^{\text{TX}}] - t_{\text{ED}}[d_i^{\text{RX}}] = T_{\text{sym}}/2 + t_{\text{PLC}} + t_{\text{THS},j}^{\text{TX}} - t_{\text{THS},i}^{\text{RX}} - t_{\text{det}}^{\text{A}} \tag{4.13}$$

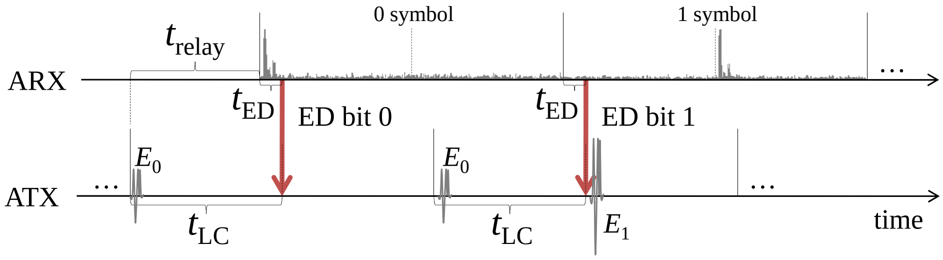

Figure 4.6: Example of a malicious prover ED+LC attack on the data. The first two symbols use equal time-hopping offsets for corresponding symbols, hence their time-gain is identical. The maximal time-gain for the 3rd symbol is larger than the packet-wide time-gain chosen of the adversary. The time-gain for the 4th symbol is smaller than the packet-wide time-gain, and the adversary is forced to guess.

However, in the case of the relay attack $t^{\text{RX}}_{\text{THS}} = t^{\text{TX}}_{\text{THS}}$ and $i = j$. Hence, the time gain is:

$$t_{\text{relay}} = T_{\text{sym}}/2 + t_{\text{PLC}} - t^{\text{A}}_{\text{det}} \tag{4.14}$$

for every symbol. This is also the upper-bound on the overall time-gain of the relay attack, as the time-gains achievable for the preamble are larger.

4.2.3 Data: Rake against Energy Detector

ED and LC If honest devices use energy detectors, using a rake receiver allows the adversary to perform an ED attack with *negative* delay t_{ED} by extracting d^{RX}_i from the (i-*1*)-th symbol. This attack exploits the structure of the convolutional code: The (i-*1*)-th data symbol carries the position bit d^{RX}_{i-1} and the polarity bit $a^{\text{RX}}_{i-1} = d^{\text{RX}}_{i-2} \oplus d^{\text{RX}}_i$. With a rake receiver, ARX can decode both bits, and obtain d^{RX}_i by computing $a^{\text{RX}}_{i-1} \oplus d^{\text{RX}}_{i-2}$. This is all that is necessary to transmit the corresponding j-th symbol: The adversary can compute d^{TX}_j from d^{RX}_i, and a^{TX}_j can be set arbitrarily, as polarity is lost in the squaring operation that the energy-detection receiver performs. The delay of the rake ED attack extracting d_i from the (i-*1*)-th symbol is:

$$t_{\text{ED}}[d^{\text{RX}}_i] = -(1 - d^{\text{RX}}_{i-1}) \cdot T_{\text{sym}}/2 - T_{\text{sym}}/2 + t^{\text{RX}}_{\text{THS},i-1} + t^{\text{A}}_{\text{det}} \tag{4.15}$$

HRX is an energy-detector as in Section 4.2.2, hence the same LC attack applies.

Malicious Prover and Relay With the rake ED attack, the malicious prover attacks, but also the relay attack (as $i = j - 1$ in this case) are subject to per-symbol variability of the time-gain due to time-hopping offsets. An additional time-gain variability is due to BPPM, i.e., the term "$-(1 - d^{\text{RX}}_{i-1}) \cdot T_{\text{sym}}/2$" of the ED delay. As in Section 4.2.2, this presents the adversary with a trade-off between the distance-decrease, and the probability of a successful attack. For example, the additional time-gain of the relay attack is at most $T_{\text{sym}}/2 + 2 \cdot t^{\max}_{\text{THS}}$, and Figure 4.7 shows the trade-off for a particular set of parameters. See Table 4.2 for the malicious prover attack.

Furthermore, a negative ED delay allows for an ED-only distance-decreasing relay attack to be mounted. Although it has a lower time-gain than an ED+LC relay attack, the ED-only attack circumvents any countermeasures that prevent LC attacks, e.g., the countermeasure advocated in Section 4.4.1.

		No guessing		Max. guessing gain	
		(relay) time-gain	distance-decrease	(relay) time-gain	distance-decrease
En.D. against En.D.					
Malicious Prover	ED-only	$T_{\text{sym}}/4 + (t_{\text{det}} - t^{\text{A}}_{\text{det}})/2$	86m	$+\,t^{\max}_{\text{THS}}$	+ 74m
	LC-only	$T_{\text{sym}}/4 + t_{\text{PLC}}/2$	86m	$+\,t^{\max}_{\text{THS}}$	+ 74m
	ED+LC	$T_{\text{sym}}/2 + (t_{\text{PLC}} + t_{\text{det}} - t^{\text{A}}_{\text{det}})/2$	171m	$+\,t^{\max}_{\text{THS}}$	+ 74m
Relay Attack	ED+LC	$T_{\text{sym}}/2 + t_{\text{PLC}} - t^{\text{A}}_{\text{det}}$	171m	$+\,0$	+ 0m
Rake against En.D.					
Malicious Prover	ED-only	$T_{\text{sym}}/2 + (t_{\text{det}} - t^{\text{A}}_{\text{det}})/2$	162m	$+\,T_{\text{sym}}/4 + t^{\max}_{\text{THS}}$	+ 151m
	ED+LC	$3/4 \cdot T_{\text{sym}} + (t_{\text{PLC}} + t_{\text{det}} - t^{\text{A}}_{\text{det}})/2$	248m	$+\,T_{\text{sym}}/4 + t^{\max}_{\text{THS}}$	+ 151m
Relay Attack	ED+LC	$T_{\text{sym}} - t^{\max}_{\text{THS}} + t_{\text{PLC}} - t^{\text{A}}_{\text{det}}$	251m	$+\,T_{\text{sym}}/2 + 2 \cdot t^{\max}_{\text{THS}}$	+ 302m
	ED-only	$T_{\text{sym}}/2 - t^{\max}_{\text{THS}} - t^{\text{A}}_{\text{det}}$	79m	$+\,T_{\text{sym}}/2 + 2 \cdot t^{\max}_{\text{THS}}$	+ 302m
Rake against Rake					
Malicious Prover	ED-only	$(t_{\text{det}} - t^{\text{A}}_{\text{det}})/2$	5m	$+\,T_{\text{sym}}/4 + t^{\max}_{\text{THS}}$	+ 151m
	LC-only	$t_{\text{PLC}}/2$	5m	$+\,T_{\text{sym}}/4 + t^{\max}_{\text{THS}}$	+ 151m
	ED+LC	$(t_{\text{det}} + t_{\text{PLC}} - t^{\text{A}}_{\text{det}})/2$	10m	$+\,T_{\text{sym}}/2 + t^{\max}_{\text{THS}}$	+ 228m
Relay Attack	ED+LC	$t_{\text{PLC}} - t^{\text{A}}_{\text{det}}$	10m	$+\,0$	+ 0m

Table 4.2: Upper-bound on (relay) time-gain and (relay) distance-decrease of various PHY attacks in various “adversarial receiver against honest receiver” configurations. The left column presents conservative attacks, that work with 100% success probability. The right column presents the maximal additional time-gain/distance-decrease that can be achieved by combining PHY attacks and guessing attacks (when time guessing probability approaches the guessing probability of pure guessing attacks). Time-gain is expressed in terms of T_{sym} – data symbol duration, $t_{\text{det}} = 48\text{-}60ns$ – detection time of honest receivers without ED-countermeasure, $t^{\text{A}}_{\text{det}}$ – detection time of the adversary, $t_{\text{PLC}} < t_{\text{det}}$ – pulse LC delay, $t^{\max}_{\text{THS}}$ – maximum time-hopping offset. The distance-decrease is shown for the IEEE 802.15.4a mandatory modes and delay values that maximize the distance-decrease.

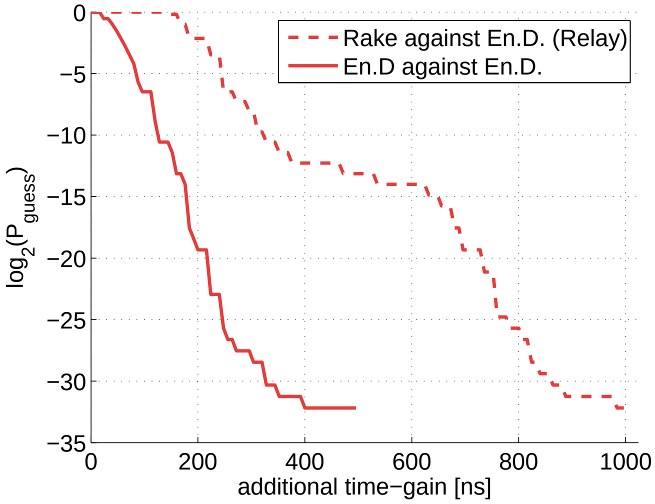

Figure 4.7: Example trade-off between guessing probability and additional time-gain achievable with 1) Malicious prover attacks in the scenario "Energy Detector against Energy Detector" 2) Distance-decreasing relay attack in the scenario "Rake against Energy Detector". The guessing probability $P_{\text{guess}}(t^+_{\text{gain}}) = F_{\text{BIN}}(N_{\text{err}}|B(N_{\text{nonce}}, t^+_{\text{gain}}), \frac{1}{2})$, where $B(N_{\text{nonce}}, t^+_{\text{gain}})$ is the number of bits out of N_{nonce} that the adversary must guess to obtain an additional time-gain t^+_{gain}. The timing parameter values (notably the time-hopping sequence) correspond to the mandatory LPRF mode of IEEE 802.15.4a, and $N_{\text{nonce}} = 42$, $N_{\text{err}} = 2$ (corresponds to security level 2^{-32}).

4.2.4 Data: Rake against Rake

To mitigate the effects of the rake ED attack, an honest rake receiver must demodulate and check the correctness of both positions bits d_i and polarity bits a_i (without applying convolutional decoding, see coding assumption in Section 4.1.3). With this precaution in place, ATX cannot transmit symbol j (with or without LC) without knowing a^{TX}_j.

ED and LC The ED delay of demodulating the polarity bit a^{RX}_i is:

$$t_{\text{ED}}[a^{\text{RX}}_i] = d^{\text{RX}}_i \cdot T_{\text{sym}}/2 + t^{\text{RX}}_{\text{THS},i} + t^{\text{A}}_{\text{det}} \tag{4.16}$$

The LC delay of committing to the polarity bit a^{TX}_j is:

$$t_{\text{LC}}[a^{\text{TX}}_j] = t^{\text{TX}}_{\text{THS},j} + d^{\text{TX}}_j \cdot T_{\text{sym}}/2 + t_{\text{PLC}} \tag{4.17}$$

Note that both delays depend on the value of the position bit. For position bits, the delays for ED of d^{RX}_i and LC of d^{TX}_j are as in Section 4.2.3.

Malicious Prover and Relay The time-gain is computed as the minimum of the time-gain for the position bits d and the time-gain for the polarity bits a. In the relay attack, $i = j$, $d^{\text{RX}} = d^{\text{TX}}$, $a^{\text{RX}} = a^{\text{TX}}$, and $t^{\text{RX}}_{\text{THS}} = t^{\text{TX}}_{\text{THS}}$, hence the relay time-gain is:

$$t_{\text{relay}} = t_{\text{PLC}} - t^{\text{A}}_{\text{det}} < t_{\text{det}} \tag{4.18}$$

which translates to at most 10m assuming $t_{\text{det}} = 32$ns (see Section 4.3.5). This is an order of magnitude lower than the attacks presented so far. The time-gain of malicious prover attacks, without guessing, is of the same order of magnitude, see Table 4.2. However, a malicious prover (but not a relaying adversary) can increase the distance-decrease by as much as $T_{\text{sym}}/2 + t_{\text{THS}}^{\text{max}}$, by lowering the probability of success.

Furthermore, in the case of the relay attack, if only one of the honest devices uses an energy-detection receiver, notably without any countermeasures deployed, (and even though the other one uses a rake receiver) the adversary can achieve a significant time-gain. Even assuming that the distance-decrease against the rake receiver is negligible, the overall distance-decrease of the relay attack is $c \cdot t_{\text{relay}}^{\text{En.D.}}/2$, where $t_{\text{relay}}^{\text{En.D.}}$ is the time-gain of the attack against an energy detector.

4.2.5 Processing Delays

In a distance-decreasing relay attack, an additional factor that reduces the relay time-gain, and hence the amount by which the distance can be decreased, are the ARX's and ATX's processing delays for the IEEE 802.15.4a channel and for the adversarial channel. We discuss these delays here, and argue that it is feasible to keep them in the order of nanoseconds (or a few meters). We focus on the data, as it is the bottleneck in terms of the achieved delay (the adversary has much more time flexibility during the preamble). We distinguish two cases: (i) ARX and ATX integrated into one device, with appropriate shielding and directional antennas, and (ii) remote ARX and ATX. The latter case can lead to a broader scope of attacks, as the adversary has the flexibility of placing its devices close to the corresponding victim devices. On the downside, remote ARX and ATX are subject to an additional processing delay, due to communication over the adversarial channel. Note that in case of malicious prover attacks, the malicious prover is subject to the same processing delay as an honest device, i.e., there are no additional processing delays.

We first consider the processing delay related to the communication with the honest devices, which applies in both (i) and (ii). At ARX the delay consists of the processing due to demodulation, *after* the necessary signal has been received. With an approximate, linearized maximum likelihood test[1] the processing delay would be in the order of a few nanoseconds. At ATX, the delay is of the same order: after the bit value is received from ARX, the transmitter only needs to proceed with or abort the transmission of a previously known burst of pulses (Figure 4.6). Note that these two delays are essentially identical to the delays that an honest prover incurs during the rapid bit exchange. In [87], the authors propose an IR implementation of distance bounding, and they estimate the group delay of the receiver at approximately 4ns. No explicit delay estimate of the transmitter delay is given, but the authors assume the total processing delay to be 4ns in their performance analysis.

In case (ii), there is an additional delay due to communication over the adversarial channel: more precisely, the delay of putting the bit value on the adversarial channel at ARX, and demodulating it at ATX. The exact numbers depend heavily on the technology ARX and ATX use to communicate. The adversary is most likely to choose a wireless communication medium, due to its faster propagation speed, but even more so because of the ease of attack deployment compared to a wired channel.

We emphasize that the adversarial channel has unusual requirements. It does not require

[1] **Details:** The approximate decision consists in comparing $\sum y_{m,n} p_{m,n}$ to a pre-computed threshold [115].

a high bit-rate, as the adversary only needs to transmit a single bit every $1\mu s$. However, the bit has to be transmitted as fast as possible. Many wireless technologies, even those with very high bit-rates, such as 802.11n, are not suitable: They achieve these high bit-rates through large modulation constellation sizes, rather than a short symbol duration. One valid option is IR-UWB with on-off keying and a receiver similar to the ED receiver described in Section 4.2.2, or with BPSK and a rake receiver. Naturally, the adversary will ignore the regulations and transmit with a power high enough to achieve a negligible error rate. To mitigate the multipath delay spread, a highly directive antenna can be used, as proposed for a narrow-band communication system in [50]. The coherent two-level PSK scheme proposed in [50] can also be used as the adversarial channel: It reports bit duration of only 1.6ns. Overall, in case (ii), a processing delay in the order of 10ns (3.5m) seems feasible.

4.3 Performance Evaluation

The basic simulation setup is described in Section 3.3.3. In this section, we use the residential non-line-of-sight channel model [99] with a channel delay spread $T_{\text{spread}} \approx 60$ ns. We use the IEEE 802.15.4a ternary preamble code number 5 of length $N_{\text{pcode}} = 31$ given by the standard. The values chosen for $t_{\text{ED}}^{\text{SFD}}$ and $t_{\text{LC}}^{\text{SFD}}$ are chosen with respect to the structure of this code. We focus on the LPRF mode. Results for HPRF can be found in [59].

Our main performance metrics are the *packet error rate* (PER) and the *synchronization error rate* (SER). We consider a packet to be in error if it was not acquired during synchronization or if the number of bit errors exceeds N_{err}. We assume a desired security level of $P_{\text{guess}} = 2^{-32}$, and performance goals $\text{PER}_{\text{comm}} = \text{PER}_{\text{db}} = 10^{-2}$. According to Section 3.1.5, this results in ranging packets of length $N_{\text{nonce}} = 42$ with a maximum of $N_{\text{err}} = 2$ tolerable bit errors. We consider synchronization to be in error if the packet is not detected (missed detection) or if the synchronization is off by too much for data decoding to be performed correctly (false alarm). Confidence intervals shown are at the 95% level.

To evaluate the cost of the attack, we compare the benign case performance (honest receiver and transmitter) with the performance under attack. We then express the cost as the difference in SNR (between the two cases) necessary for the same performance (SER, PER), which we denote by ΔSNR. This tells us by what factor the adversary needs to improve the received signal level to obtain the same performance as in the case of an honest execution of the protocol. He can achieve this by using a high-gain antenna, by transmitting with a higher power, or by moving closer to the victim transceivers.

4.3.1 Energy Detector against Energy Detector

We first determine the individual performance of the preamble attack primitives and data attack primitives (ED and LC). The latter translate directly into the performance of malicious prover attack. We then derive the performance of the distance-decreasing relay attack, which combines all four primitives.

4.3.2 Preamble Attacks

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

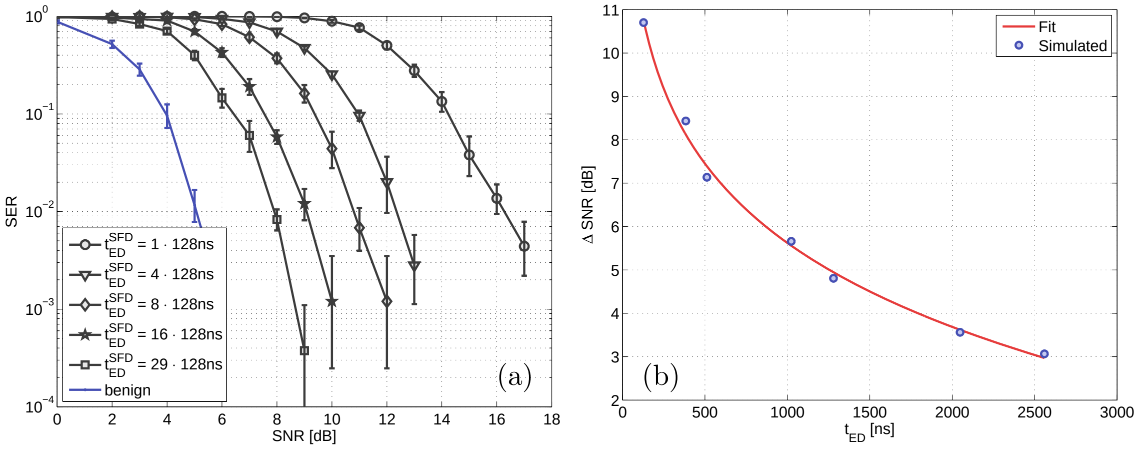

Figure 4.8: SER versus SNR comparing benign performance to ED with varying ED delays $t_{\mathrm{ED}}^{\mathrm{SFD}}$. Figure (b) is a more compact representation of the data in (a), showing the loss ΔSNR with respect to the benign case versus $t_{\mathrm{ED}}^{\mathrm{SFD}}$ for a fixed SER of 10^{-2}.

Figure 4.8 shows the SER for an honest receiver, as well as for an adversary that performs early SFD detection with different early SFD detection delays $t_{\mathrm{ED}}^{\mathrm{SFD}}$. Not surprisingly, the earlier an adversary performs SFD detection, the more additional received power with respect to an honest receiver it is going to cost him to reach a given level of SER. If we fix SER $= 10^{-2}$, detecting the SFD at $t_{\mathrm{ED}}^{\mathrm{SFD}} = 3.712\,\mu$s costs the adversary ΔSNR $=$ 2.8dB in additional received power, detecting at $t_{\mathrm{ED}}^{\mathrm{SFD}} = 0.128\,\mu$s entails a cost of ΔSNR $=$ 11.2dB.

For $t_{\mathrm{ED}}^{\mathrm{SFD}}$, we only consider values shorter than the length of the first SFD symbol. Larger values for $t_{\mathrm{ED}}^{\mathrm{SFD}}$ do not make much sense for the adversary because they also force him to commit after the first SFD symbol, which is only possible at a considerable additional cost.

This can be seen in Figure 4.9, which shows the SER of an adversary that commits late, at time $t_{\mathrm{LC}}^{\mathrm{SFD}}$ into the SFD. Committing at $t_{\mathrm{LC}}^{\mathrm{SFD}} = 8 \cdot 128\text{ns} = 1.02\,\mu$s, or earlier is within 0.6dB of the benign case and thus comes at practically no additional cost at a target SER of 10^{-2}. Committing later comes at an ever increasing cost: Committing at $t_{\mathrm{LC}}^{\mathrm{SFD}} = 29 \cdot 128\text{ns} = 3.712\,\mu$s, already costs ΔSNR $=$ 7.5dB. (Note: According to the preamble and SFD codes, no pulse is sent between the 29th frame of the first SFD symbol and the first frame of the third SFD symbol. So committing anywhere between $t_{\mathrm{LC}}^{\mathrm{SFD}} = 3.712\,\mu$s and $t_{\mathrm{LC}}^{\mathrm{SFD}} = 63 \cdot 128\text{ns} = 8.064\,\mu$s is equivalent to committing at $t_{\mathrm{LC}}^{\mathrm{SFD}} = 8.064\,\mu$s, which costs more than ΔSNR $=$ 9dB.)

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

Alternative Receiver We also evaluated a receiver that uses the correlation-based SFD detection. This receiver is also vulnerable to the attack, and the attack's cost in terms of ΔSNR is close to the cost for the baseline receiver: more precisely, up to 1dB greater (for $t_{\mathrm{LC}}^{\mathrm{SFD}}$ in the order of T_{psym}).

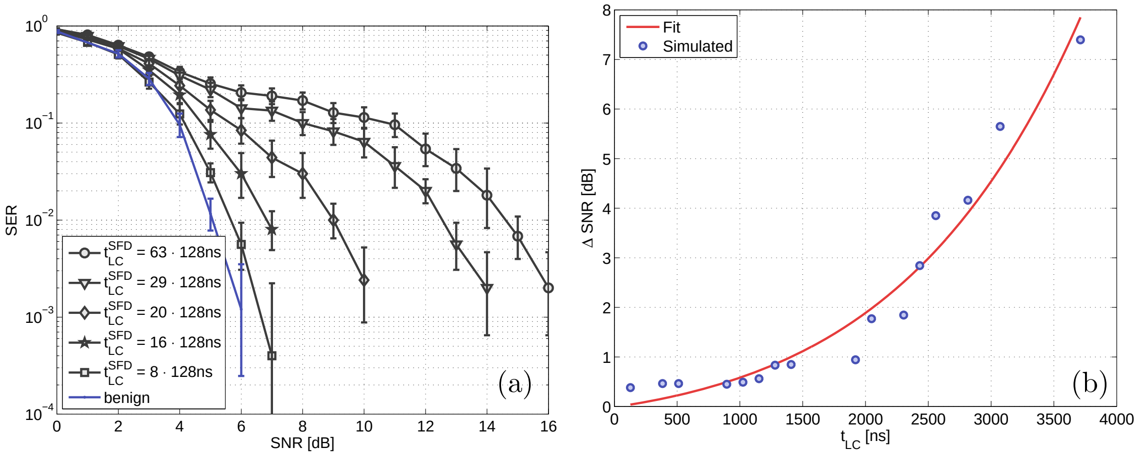

Figure 4.9: SER versus SNR comparing benign performance to LC with varying LC delays $t_{\mathrm{LC}}^{\mathrm{SFD}}$. Figure (b) is a more compact representation of the data in (a), showing the loss ΔSNR with respect to the benign case versus $t_{\mathrm{LC}}^{\mathrm{SFD}}$ for a fixed SER of 10^{-2}.

4.3.3 Data Attacks

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

Figure 4.10(a) shows the PER at different SNRs for the LPRF mode. We show the performance curves for a benign receiver and an adversary performing ED at different ED delays t_{ED}. The optimal ED delay for the adversary is found to be $t_{\mathrm{ED}}^{\mathrm{OPT}} = 68$ns, which is very close to the channel delay spread plus the duration of a burst of pulses. Deciding on the symbol at $t_{\mathrm{ED}}^{\mathrm{OPT}}$ introduces a loss of about 1.6dB with respect to the benign curve at a packet error rate of PER $= 10^{-2}$. This can also be seen in Figure 4.10(b). Here we show the loss in SNR, ΔSNR, with respect to the benign case versus the ED delay t_{ED} for a target packet error rate of PER $= 10^{-2}$. The curve has been obtained from curves such as those shown in Figure 4.10(a) via interpolation. Detecting after $t_{\mathrm{ED}}^{\mathrm{OPT}}$ gives a slightly worse performance because the adversary then merely integrates more noise instead of useful signal. Performing ED much earlier than $t_{\mathrm{ED}}^{\mathrm{OPT}}$ results in substantially larger loss because a large part of the useful signal energy is lost: Deciding at $t_{\mathrm{ED}} = 32$ns, for example, introduces a loss of 4.8dB.

Figure 4.11 shows the performance of LC on the data in the case of LPRF. As explained in Section 4.2.2, the LC delay t_{LC} is fixed to $t_{\mathrm{LC}} = T_{sym}/2 = 512$ns. We show the PER for different ratios γ of the energies $E0$ and $E1$ corresponding to the signal energies transmitted by the adversary during the 0-block and 1-block, respectively. $E1$ here corresponds to the energy a benign receiver would transmit and $E0$ is typically smaller. A ratio of $\gamma^{\mathrm{OPT}} = 0.35$ gives optimal performance throughout the whole operating range, thus this is the energy ratio we will use in all subsequent simulations. The optimal ratio gives a loss of about 3.8dB with

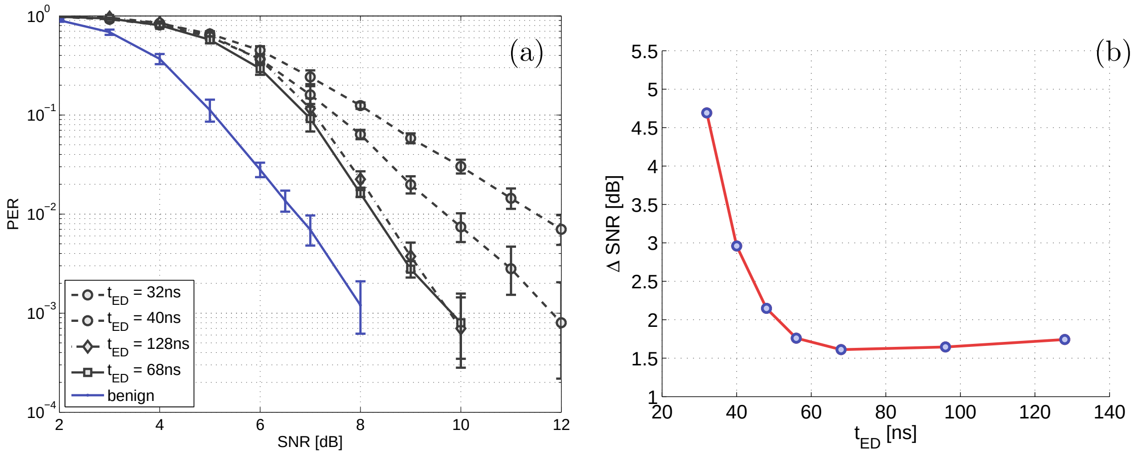

Figure 4.10: PER versus SNR for the data comparing benign performance to ED with varying ED delays t_{ED}. The optimal t_{ED} is in the order of the channel delay spread and gives a loss of about 1.7dB. Figure (b) is a more compact representation of the data in (a), showing the loss ΔSNR with respect to the benign case versus t_{ED} for a fixed PER of 10^{-2}.

respect to the benign case.

Alternative Receiver We also evaluated the receiver that demodulates without weighting with the estimated energy-delay profile. This receiver is vulnerable to the attack as well, and the attack's cost in terms of ΔSNR is within 0.5dB of the cost for the baseline receiver.

4.3.4 Distance-Decreasing Relay Attack

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

Figure 4.12 shows the probability of success of an attack that tries to gain 480 ns when relaying a 42bit packet between HTX and HRX. This relay time-gain is equivalent to a 144 m distance decrease between HTX and HRX.[2] The results shown are for different combinations of $\mathrm{SNR}_{\mathrm{ED}}$ and $\mathrm{SNR}_{\mathrm{LC}}$. For every SNR combination, the probability of success that is reported corresponds to the triple[3] of $(t_{\mathrm{ED}}^{\mathrm{SFD}}, t_{\mathrm{LC}}^{\mathrm{SFD}}, t_{\mathrm{ED}})$ yielding best performance among all the tuples

[2] Here we assume for simplicity that the processing delays at the adversarial transceivers are zero. Processing delays are discussed in Section 4.2.5.

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

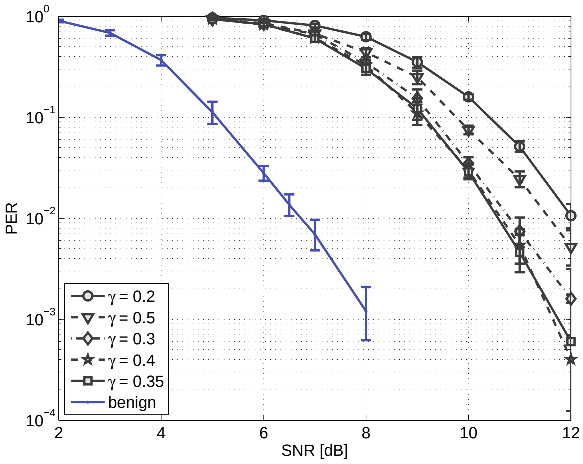

Figure 4.11: PER for LC on the data with varying energy ratios γ. The optimal ratio at $\gamma^{\mathrm{OPT}} = 0.35$ gives a loss of about 4dB with respect to the benign setting.

that achieve the given relay time-gain of 480 ns. In the benign case we achieve a PER of approximately 10^{-2} at an SNR of around 8 dB. In Figure 4.12, a probability of success of $P_s = 0.9896$ is achieved for the pair $(\mathrm{SNR}_{\mathrm{ED}} = 14\mathrm{dB}, \mathrm{SNR}_{\mathrm{LC}} = 12\mathrm{dB})$. For all pairs above $(14\mathrm{dB}, 12\mathrm{dB})$ the probability of success is above 99%. With respect to an honest transmitter-receiver pair, an adversary thus needs an additional 6 dB in SNR for ED and an additional 4 dB for LC, in order to reduce the distance by 144 m with a probability of success in the order of 99%. Attaining SNR values in this range would not pose much of a challenge to the adversary.

HRX Not Isolated From HTX

In the threat model for the distance-decreasing relay attack, we assume that the honest receiver, HRX, cannot receive signals sent by the honest transmitter, HTX. This is inherent in some scenarios, e.g., picking virtual pockets [81], but there are other scenarios where HTX will be in range of HRX. In this case, the adversary can prevent communication between the honest devices through shielding, by placing one of the honest devices in a Faraday cage (such as a "booster bag" coated with aluminium foil). One adversarial device would then be connected via a wired link to the second adversarial device placed outside the Faraday cage.

However, in some scenarios HTX will be in range of HRX, and it might not be feasible for the adversary to shield HRX from HTX. We show here that the attack is still possible, but the cost of the attack (in terms of SNR) increases.

To make sure that HRX locks on the adversarial preamble, and not the preamble of HTX, ATX needs to start transmitting the preamble before HRX acquires HTX's signal. For the parameters of the synchronization algorithm we assume, this happens no sooner than 18 preamble symbols into the preamble. The sooner ATX starts the transmission, the lower the cost (as usual, the cost is measured in terms of SNR necessary to achieve a PER of 10^{-2}).

Figure 4.13 shows the relative cost of the preamble and data LC attack (in comparison with the LC attack where HRX is shielded from HTX) as a function of the SNR obtained by HTX at HRX, and when ARX starts the preamble transmission with a delay of 8, 12 and 16 preamble symbols. For the former two, the cost is in the order of HTX's SNR, but for 16

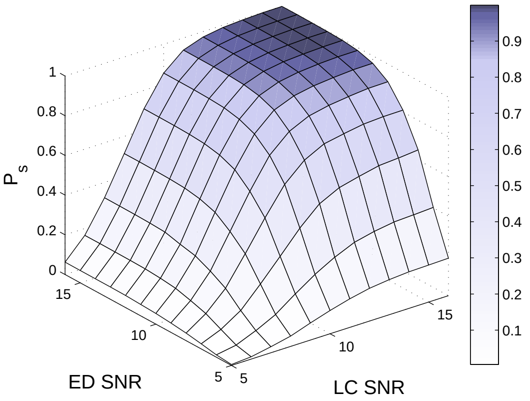

Figure 4.12: Probability of success, P_s, for a distance-decreasing relay attack trying to achieve a distance-decrease of 144 m. $P_s > 99\%$ is reached at a relative cost of $(\Delta\mathrm{SNR}_{\mathrm{ED}}, \Delta\mathrm{SNR}_{\mathrm{LC}}) > (6\mathrm{dB}, 3.8\mathrm{dB})$.

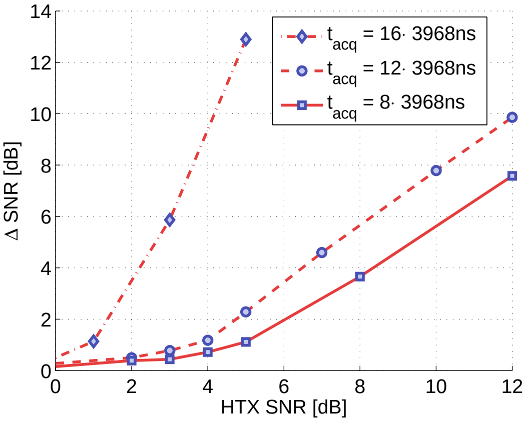

Figure 4.13: Cost of LC when HRX is not shielded from HTX. We show the relative cost ΔSNR of the LC attack (both preamble and data) versus SNR of HTX, for different timing acquisition delays t_{acq}.

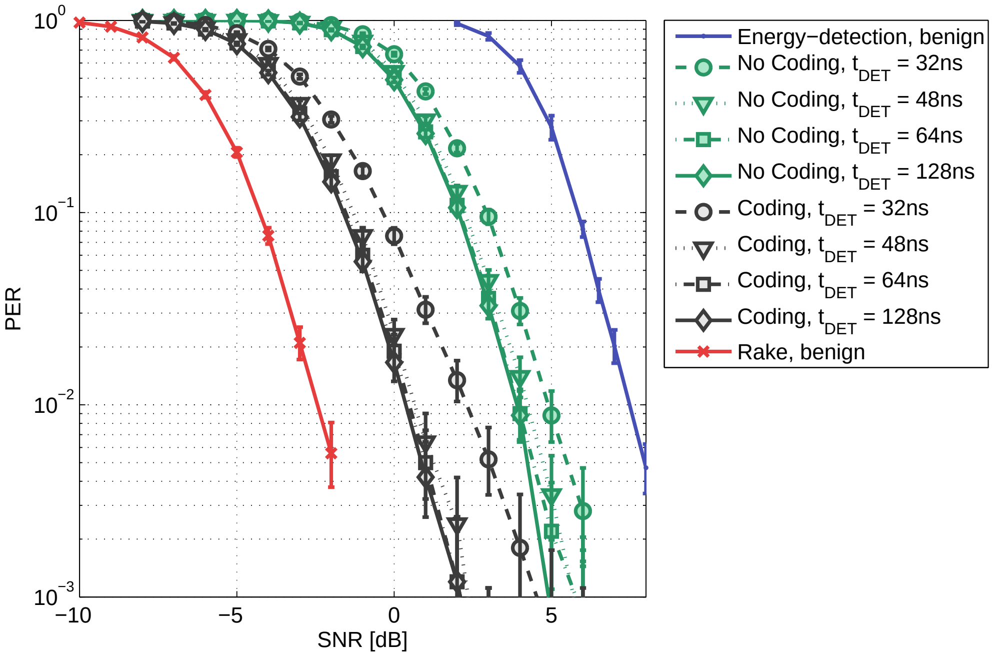

Figure 4.14: Performance of ED on the data if the adversary uses an rake receiver. ED can be performed by decoding the convolutional code using partial information ("Coding") or by neglecting the convolutional code completely ("No Coding").

the cost grows much faster. Furthermore, ARX needs to perform early timing acquisition to make this attack possible. The cost, in comparison with regular acquisition after 18 symbols, is 6dB, 3dB and 2dB, for acquisition before 8, 12 and 16 preamble symbols, respectively.

4.3.5 Rake against Energy Detector

We evaluate the rake ED attack for an optimal all-rake receiver with perfect synchronization and channel estimation. Figure 4.14 shows the PER for an adversary performing the ED attack. When mounting the ED attack, ARX has the option to ignore ("No Coding") or take advantage of the convolutional code ("Coding"). For reference, the performance of a benign energy-detection receiver and a benign rake receiver are shown as well. The benign rake receiver decodes the convolutional code at the end of the packet as in [13], when the full decoding trellis is available. With ED and taking the code into account, only a partial trellis containing information about the symbols received so far is available at the time of decoding. This contributes to the higher cost (in terms of required SNR) of the attack with respect to the benign rake receiver operation. Ignoring the convolutional code is simpler and less computationally expensive, but results in an additional 3.5dB increase of the attack cost.

The adversary also has the choice of t^A_{det}. Optimal performance is experienced for $t^A_{\text{det}} =$ 64ns, in the order of channel spread. $t^A_{\text{det}} =$ 48ns results in a very minor performance loss, $t^A_{\text{det}} =$ 32ns results in noticeable performance loss (around $1-2$dB). Assuming $t^A_{\text{det}} =$ 48ns, the attack costs 2.8dB (at a PER of 10^{-2}, corresponding to an attack with a success rate of 99%) if coding is taken into account and the relay attack achieves a time-gain of $t_{\text{relay}} =$ 728ns (distance-decrease of 218 meters). At the same cost, the alternative, ED-only attack achieves a time-gain of $t_{\text{relay}}^{\text{ED-only}} = 216\text{ns}$ (distance-decrease of 65 meters).

4.4 Countermeasures

When investigating countermeasures and patches, we consider their effectiveness (the maximum relay time-gain the adversary can achieve with the countermeasure in place), the effects they have on benign-case performance, and their compatibility with IEEE 802.15.4a. We discuss first the countermeasures that, in our opinion, provide the best trade-off between these factors. Then, we discuss alternative countermeasures.

4.4.1 Recommended Countermeasures

Convolutional Code Patch

The rake ED attack is possible due to the specific combination of BPPM/BPSK with the convolutional code that IEEE 802.15.4a uses. The attack can be prevented by changing the convolutional code to a code in which the polarity bit a_i does not reveal information about future position bits d_j, where $j > i$. In addition, the code should not allow for decoding a_{i+1} from the i-th symbol i, as this would enable an effective LC attack against a rake receiver. We refer to this solution as the *convolutional code patch*. Alternatively, the convolutional code can be removed altogether. The former is not compatible with IEEE 802.15.4a, whereas the latter is compatible to a very limited extent: Although IEEE 802.15.4a provides a few optional modes that do not use the convolutional code, these modes cannot be used with energy-detection receivers, because, for these modes, polarity carries data information.

Time-Hopping Patch

Time-hopping allows a malicious prover to trade-off attack success probability for additional time-gain. A simple way to remove this vulnerability is to modify the time-hopping sequence such that the corresponding challenge and response symbols have identical time-hopping offsets. Removing time-hopping is also a solution, but an inferior one, as it significantly degrades IEEE 802.15.4a multi-user access properties.

Early Detection at Honest Receiver

The honest energy-detection receiver can choose to only take into account the beginning of the symbol [68], essentially performing early detection with OOK demodulation at an offset $t_{\text{det}}^{\text{C}}$ from the beginning of the symbol. Then, t_{LC} is reduced from $T_{\text{sym}}/2$ to $t_{\text{PLC}}^{C} < t_{\text{det}}^{\text{C}}$, and the (malicious prover) time-gain due to ED is limited to $(t_{\text{det}}^{\text{C}} - t_{\text{det}}^{\text{A}})/2$ (assuming that the prover sends the response symbol immediately after the early detection is done). This countermeasure does not induce inter-symbol interference and is compliant with the mandatory modes of the standard.

Moreover, the performance loss that this countermeasure entails due to the ignoring of half of the symbol, can be compensated for by increasing the length of the nonces N_{nonce}. We can derive the required nonce length by using the method introduced in Section 3.1.5. Figure 4.15 plots the resultant N_{nonce} as a function of $t_{\text{det}}^{\text{C}}$ for performance goals $\text{PER}_{\text{comm}} = \text{PER}_{\text{db}} = 10^{-2}$ and 3 security levels. For example, by employing the countermeasure with $t_{\text{det}}^{\text{C}} = 40ns$

		No guessing		Max. guessing gain	
		(relay) time-gain	distance-decrease*	(relay) time-gain	distance-decrease
En.D./Rake against En.D./Rake with ED-countermeasure and convolutional code patch					
Malicious Prover	ED-only	$(t^{\mathrm{C}}_{\mathrm{det}} - t^{\mathrm{A}}_{\mathrm{det}})/2$	5-6m	$+\,t^{\max}_{\mathrm{THS}}$	+ 74m
	LC-only	$t^{C}_{\mathrm{PLC}}/2$	5-6m	$+\,t^{\max}_{\mathrm{THS}}$	+ 74m
	ED+LC	$(t^{C}_{\mathrm{PLC}} + t^{\mathrm{C}}_{\mathrm{det}} - t^{\mathrm{A}}_{\mathrm{det}})/2$	10-12m	$+\,t^{\max}_{\mathrm{THS}}$	+ 74m
Relay Attack	ED+LC	$t_{\mathrm{PLC}} - t^{\mathrm{A}}_{\mathrm{det}}$	10-12m	$+\,0$	+ 0m
En.D./Rake against En.D./Rake with ED-countermeasure and convolutional code and time-hopping patches					
Malicious Prover	ED-only	$(t^{\mathrm{C}}_{\mathrm{det}} - t^{\mathrm{A}}_{\mathrm{det}})/2$	5-6m	$+\,0$	+ 0m
	LC-only	$t^{C}_{\mathrm{PLC}}/2$	5-6m	$+\,0$	+ 0m
	ED+LC	$(t^{C}_{\mathrm{PLC}} + t^{\mathrm{C}}_{\mathrm{det}} - t^{\mathrm{A}}_{\mathrm{det}})/2$	10-12m	$+\,0$	+ 0m
Relay Attack	ED+LC	$t_{\mathrm{PLC}} - t^{\mathrm{A}}_{\mathrm{det}}$	10-12m	$+\,0$	+ 0m

Table 4.3: Upper-bound on (relay) time-gain and (relay) distance-decrease of various PHY attacks in various "adversarial receiver against honest receiver" configurations. The left column presents conservative attacks, that work with 100% success probability. The right column presents the maximal additional time-gain/distance-decrease that can be achieved by combining PHY attacks and guessing attacks (when time guessing probability approaches the guessing probability of pure guessing attacks). Time-gain is expressed in terms of T_{sym} – data symbol duration, $t^{\mathrm{A}}_{\mathrm{det}}$ – detection time of the adversary, $t^{\max}_{\mathrm{THS}}$ – maximum time-hopping offset, $t^{C}_{\mathrm{det}} = 32\text{-}40ns$ – detection time of honest receiver with ED-countermeasure, $t^{C}_{\mathrm{PLC}} < t^{C}_{\mathrm{det}}$ – pulse LC delay if countermeasure is deployed. The distance-decrease is shown for the IEEE 802.15.4a mandatory modes and delay values that maximize the distance-decrease. The value of $t^{\mathrm{C}}_{\mathrm{det}}$ is chosen to allow operation without significantly increasing the packet length.

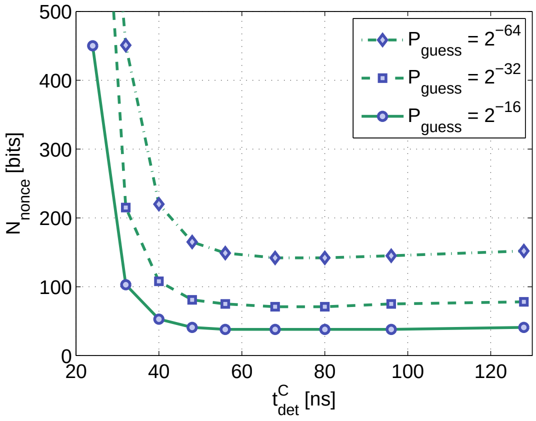

Figure 4.15: Cost of ED countermeasure when the energy-detection receiver decides on the bit value using OOK demodulation at time t^{C}_{det}. $\text{SNR}_{\text{H}} = 6.5\text{dB}$.

and increasing the number of bits per nonce from 42 to 108 we can bring the maximum theoretically achievable time-gain to 40ns (distance decrease of about 12 m) maintaining security level $P_{\text{guess}} = 2^{-32}$. At the same time, this countermeasure does not reduce the performance in terms of PER and we also keep the same security level against guessing attacks. The drawback is generating, sending and receiving of the additional bits required for the longer nonces. As every IEEE 802.15.4a packet carrying a nonce also includes a preamble of considerable length, and as a good deal of receiver complexity during reception stems from synchronization, we argue that the cost of adding a moderate number of bits to the data is in most cases acceptable.

Furthermore, this countermeasure can be employed by both energy-detection and rake receivers to prevent the adversary from exploiting the BPPM variability (Section 4.2.3). See Table 4.3 for upper-bounds on the attack time-gains with the countermeasure and patches deployed.

The countermeasure has a relatively mild cost as long as t^{C}_{det} is not much lower than the channel spread, allowing the receiver to capture a significant part of the symbol energy. However, reducing t^{C}_{det} further implies a more considerable cost in terms os packet length. We explore this further in Section 5.3.

4.4.2 Alternative Countermeasures

Decrease Data Symbol Duration

A straightforward countermeasure is to decrease data symbol duration T_{sym} [35], as the time-gain of any PHY attack is at most T_{sym}. This countermeasure can be even implemented within the IEEE 802.15.4a standard, as some non-mandatory modes have symbols as short as 32 ns. However, significantly reducing T_{sym} (e.g., to make the maximum achievable distance-decrease a few meters), is detrimental to benign performance. Inter-symbol interference (ISI) manifests itself if the symbol duration is close to or below the channel delay spread. Low-complexity non-coherent receivers cannot cope well with ISI and even if some solutions exist, they entail a loss of $5 - 10$ dB in the benign-case [150]. Furthermore, shorter symbols have less resilience to multi-user interference.

Secret Spreading Codes

To make preamble ED harder, if not infeasible within the constrained time budget available to the adversary, the honest devices could generate preamble codes from a shared secret. Alternatively, secret time-hopping sequences could be used to make ED of data symbols more difficult. Naturally, both approaches can only prevent relay attacks, as a malicious prover would know the secret spreading codes. Furthermore, the *external* cicada attack that we investigate in Chapter 5 is also immune to these countermeasures.

Furthermore, neither countermeasure is directly compliant with the current IEEE 802.15.4a standard. The standard includes an optional private ranging mode, in which the ranging devices can secretly agree on a preamble code, but there exist only eight publicly known preamble codes to choose from. This offers little security: The adversary can guess both codes with decent probability, or perform detection using, in parallel, all eight allowable codes. (This can be done entirely in the digital domain by correlating the received signal with each of the 8 codes and choosing the one with the highest correlation output.)

Detect Data LC

We investigated a countermeasure that detects the non-standard signal sent by the adversary during the data LC attack. With this countermeasure, the receiver records, for every bit, the energy in the first half of the symbol, and compares the distribution of these energies for the 0 bits (bits that were decoded as a 0) with the 1 bits (decoded as a 1). In the benign case, the first halves of the 0 bits carry more energy, whereas under attack these energies are the same. To distinguish these cases, one can use a robust statistical test, such as the Mann-Whitney-Wilcoxon test. This countermeasure prevents the attack presented in Section 4.2.2 with virtually no degradation of benign case performance. However, an adversary can modify the attack (vary the energy levels between symbols) to severely degrade the performance of this countermeasure.

Detect Preamble LC

We experimented with countermeasures that attempt to detect the preamble under a LC attack. For example, a countermeasure could check if the first SFD symbol is entirely 0 (as it should be). However, this countermeasure can only reliably detect an attack with relatively high $t_{\mathrm{LC}}^{\mathrm{SFD}}$ (more S than 0 in the preamble symbol) at high SNR, but not attacks with low $t_{\mathrm{LC}}^{\mathrm{SFD}}$ (more 0 than S), especially in the lower SNR regions. A countermeasure could also detect the high number of 0 symbol at the beginning of the preamble – but this can be countered by the adversary by early time acquisition (which comes at some additional cost in terms of SNR, see Section 4.3.4). Finally, countermeasures to preamble LC attacks cannot prevent malicious prover attacks.

4.5 Conclusion

In this Chapter, we have investigated the vulnerability of the IEEE 802.15.4a standard to physical layer distance-decreasing attacks. We have demonstrated that if honest devices use energy-detection receivers without appropriate countermeasures, an adversary can decrease the measured distance by hundreds of meters, with a success rate arbitrarily close to 100%.

However, minor modifications to IEEE 802.15.4a and implementing a simple countermeasure on energy-detection receivers used by honest devices, allow honest devices to reduce the effectiveness of distance-decreasing relay attacks to at most 10m. Alternatively, this can be achieved (even without IEEE 802.15.4a modifications) if honest devices use the more sophisticated rake receivers. Furthermore, to reduce the effectiveness of malicious prover attacks to around 10m, the honest receivers (energy detector and rake alike) should implement the same simple countermeasure, and a time-hopping patch should be applied to IEEE 802.15.4a. In conclusion, with appropriate countermeasures on the receivers and patches to the IEEE 802.15.4a standard, the standard can be used as a DB PHY.

More generally, our investigation has identified PHY features that, although they improve system performance in the benign case, can create vulnerabilities against distance-decreasing PHY attacks if used carelessly. One such potential point of failure is the interaction between the modulation and the coding scheme. Another, perhaps more fundamental one, is data time-hopping, which allows the adversary to additionally decrease the distance by lowering the attack's probability of success. Such features should be approached with caution, or not used at all, in any DB PHY.

The countermeasures that we consider in this Chapter do not prevent PHY attacks completely. Rather, they limit it to a value in the order of the channel spread, roughly 10 meters. In Chapter 5, we show how the adversary can mount attacks that exploit this gap, even without knowing the preamble code. We also investigate countermeasures that can limit the effectiveness of PHY attacks down to a value in the order of the accuracy of the receiver, in particular early detection with very low t_{det}.

Acknowledgments

I would like to thank Yannick Do and Florence Le Goff for their help with the simulations.

Chapter 5

Physical Layer Attacks against ToA Estimation

In this Chapter, we explore a new physical-communication-layer (PHY) attack vector that is targeted at time-of-arrival (ToA) estimation algorithms. The distance-decrease achievable by this attack vector is in the order of the channel spread. Hence, it is relevant if *precise* ranging or DB is required. The *ToA attack vector* exploits the uncertainty inherent to ToA estimation in multipath channels that spread a transmitted signal in time (Section 3.3.1). In typical indoor environments, this *channel spread* is in the order of tens of nanoseconds [99], which translate to at least a few meters. For ranging to be precise, the receiver needs to accurately identify the ToA of the first path, which is not necessarily the strongest path – notably under *weak non-line-of-sight* (weak NLOS) conditions.[1] Impulse-Radio Ultra-wideband is particularly well suited for this task, because the very narrow (in time) pulses it transmits make the task of resolving the channel relatively easy.

The ToA attack vector is most effective against *precise* ToA estimation algorithms that attempt to detect the first arriving path. However, is can also achieve a distance-decrease against ToA estimation algorithms that only detect the strongest path. More specifically, we show the following:

▶ If a receiver implements a precise ToA esitmation algorithm, the adversary can use the ToA attack vector to mount a type of malicious interference attack that we term the *cicada* attack: an *external* attack that can decrease the measured distance by a value in the order of the channel spread (10-20 meters). Compared to the previously known external PHY attack – the distance-decreasing relay (Section 4.2) – the cicada attack has two substantial advantages: It is significantly easier to mount, as it only requires a IR transmitter, and it works even if a secret preamble code is used. The disadvantage of the cicada attack is that the achieved distance decrease is random. We show with simulations and experiments on an IR test-bed that this attack is effective against a wide class of modulation schemes and receivers.

▶ The ToA attack vector can be used in malicious prover attacks and distance-decreasing relay attacks. This is true even if the ToA estimation algorithm is designed to detect only the strongest path. Assuming that the countermeasures proposed in Section 4.4.1 are in place, the ToA attack vector allows the adversary to decrease the measured distance by the amount that these countermeasures permit. Moreover, in a malicious interference attack, if the adversary

[1] In weak NLOS conditions, the direct line-of-sight path is obscured by an obstacle that attenuates the LOS path but does not block it completely.

can synchronize its transmission with with the honest transmitter and knows the preamble code, it can achieve a distance-decrease against strongest-path ToA estimation algorithms.
▶ In terms of countermeasures, we identify a *secure ToA estimation* algorithm for rake receivers [52], and we propose a new secure ToA estimation algorithm for energy-detection receivers; both are precise and secure against the ToA attack vector, if the preamble code is secret and long enough (increasing the length of the code improves the security and/or precision). An alternative countermeasure is to decrease the receiver data detection window (t_{det}, Section 4.4.1) to values in the order of the receiver ranging precision. The latter countermeasure provides higher ranging precision but has a higher failure rate than the secure ToA estimation algorithms. Finally, we propose a hybrid countermeasure that achieves the best of both worlds. These countermeasures allow for implementing DB that is both secure against PHY attacks (in particular the attacks presented in Chapter 4) and is close in precision to non-secure ranging. However, this comes at a cost: To achieve good ranging precision, the ranging packet needs to be orders of magnitude longer than for the countermeasures in Chapter 4 that limit the distance-decrease to values in the order of the channel spread.

Chapter Outline Our basic assumptions about the modulation scheme, channel and receivers are presented in Section 3.3. In Section 5.1 we introduce additional assumptions on the receivers, and the adversary model. We explore the ToA attack space in Section 5.2. Section 5.3 is devoted to an investigation of countermeasures. In Section 5.4 we describe experiments performed on an IR test-bed to confirm the feasibility of the cicada attack. We conclude in Section 5.5.

IEEE 802.15.4a In this Chapter, we assume a generic IR modulation scheme, which gives us more flexibility than IEEE 802.15.4a. Naturally, the attacks and countermeasures apply to IEEE 802.15.4a. We evaluate the performance of the cicada attack against IEEE 802.15.4a in [127].

5.1 System Model

5.1.1 Modulation Scheme

We assume that an IR-UWB *packet* is composed of two parts: 1) the *preamble*, 2) the *data*. More precisely, the transmitted signal is:

$$s(t) = \sum_{i=1}^{N_{\mathrm{P}}} a_i^{\mathrm{P}} p(t - iT_f^{\mathrm{P}} - t_i^{\mathrm{P}}) + \sum_{i=1}^{N_{\mathrm{D}}} a_i^{\mathrm{D}} p(t - iT_f^{\mathrm{D}} - t_i^{\mathrm{D}} - T_{\mathrm{data}}) \tag{5.1}$$

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

The difference between the data and the preamble is that the former is modulated according to the sequence a_i^{D}, which is known only to the *transmitting party*. In contrast, for the preamble the sequence a_i^{P} is known to both the verifier and the prover. The sequences t_i^{P} and

t_i^{D} are known to both parties; they can also be publicly known. Note that a sequence which is only known to the verifier and the prover is derived from a secret shared between them.

Note on Data Modulation For the generic modulation scheme, we assume that one bit is encoded into one data frame, with either *on/off keying* (OOK) or *binary phase-shift keying* (BPSK). In contrast, IEEE 802.15.4a uses a combination of *binary pulse-position modulation* (BPPM) and BPSK. This assumption is dictated by security considerations: As we show in Chapter 4, BPPM leaves room for effective PHY attacks.

5.1.2 Receivers

We distinguish between a number of different algorithms implemented by the receivers: vanilla (basic algorithms), PID (Power Independent Detection [53, 58]), and MINF (min-filter [140, 44]) for the energy-detection receiver, and vanilla and PID for the rake receiver. The PID and MINF receivers are designed to be robust to interference. They are relevant because the malicious interference attack we consider is based on creating interference. Such algorithms could prevent the attack; we show this is not the case.

The receiver operation stages (listed in Section 3.3.2) most relevant for this Chapter are coarse and fine synchronization. The latter three stages (channel estimation, SFD detection and data demodulation) are less crucial. In short, for these stages, the receiver uses classical maximum likelihood algorithms for the vanilla energy detector and the rake (maximal ratio combining), or variants of these algorithms robust to multi-user interference [57] for the PID and the MINF energy detector.

Coarse Synchronization

This stage is performed using a traditional synchronization algorithm, based on correlating the received signal with the known preamble *template*. More precisely, the baseline method from [58] is implemented by the vanilla and MINF energy detectors, and the conventional method from [53] is implemented by the vanilla rake. Given a sequence of samples y_i, the correlator output is computed as:

$$z_j = \sum_{i=1}^{n_{\mathrm{temp}}} b_i \cdot y_{i+j} \tag{5.2}$$

for a block $j = 1, \ldots, m$, where n_{temp} is the template length in samples, and b_i the the binary (ternary for the rake) template sequence corresponding to the preamble. The packet is *detected* if the output from the correlator exceeds a noise-based threshold. Then, the index

$$j_{\mathrm{toa}} = arg\ max_j |z_j| \tag{5.3}$$

becomes the candidate for the coarse ToA estimate. Subsequently, the receiver *verifies* the estimate by checking whether it corresponds to the maximum of the correlator output in M subsequent sample blocks.

The PID receivers rely on the Power-Independent Detection (PID) method [58, 53]. In this method the received signal is first compared to a noise-based threshold, and converted into a binary (energy detector) or ternary (rake) sequence (we call this the *PID filter*). The output of the PID filter is then correlated with the preamble template. Detection and verification is performed as described above. In all cases, coarse synchronization locks on the strongest path component of the received signal.

Fine Synchronization

ToA estimation performs a *back-search* [43, 65] in a window of duration T_{BS} preceding the rough synchronization point found by coarse synchronization. The back-search window duration T_{BS} should be in the order of the channel spread. The back-search identifies the first time-offset in the window at which a noise-based threshold is exceeded – this is considered to be the first arriving path, i.e., the ToA. All receivers except MINF perform the back-search on the output of the correlator z_i (5.2). MINF [140, 44] uses an average of a number of preamble frames, filtered with a moving min filter before averaging. (We set the min-window length $W_{\min} = 8$.) The min filter removes interference based on the assumption that the interference is present in at most $W_{\min} - 1$ consecutive frames (typical for MUI).

Noise-based Thresholds

The synchronization algorithms use three noise-based thresholds: coarse synchronization threshold, fine synchronization threshold, and PID threshold. These thresholds are computed based on the noise statistics. Under AWGN, the noise-only samples are distributed according to the central χ^2 distribution for energy-detection receivers [60], and according to the gaussian distribution for rake receivers. The noise variance can be estimated in a robust fashion [60]. We fixing a *probability of false detection*, i.e., the probability that a noise-only sample is above the threshold. Then, we compute the threshold by inverting the cdf of the respective noise distribution.

We fix the false detection probability for the coarse synchronization and fine synchronization thresholds to be $P_{\mathrm{FD}} = 10^{-4}$. For the PID filter, we fix $P_{\mathrm{FD}} = 0.2$. The latter was derived experimentally; we have optimized it to give a low *missed detection probability*.

In case of the MINF receiver, the fine synchronization threshold is determined experimentally, assuming the presence of a benign interferer (multi-user interference). This results in a threshold that is quite conservative compared to the vanilla receiver.

5.1.3 Adversary Model

Recall that we consider three classes of distance-decreasing attacks: An internal *malicious prover attack* and external *distance-decreasing relay attack* and an external *malicious interference attack*. In this Chapter, we mainly focus is on the latter attack. In this attack, the honest receiver (HRX) receives a ranging packet (5.1) from an honest transmitter (HTX). The honest signal interferes at HRX with the adversarial signal generated by an adversarial transmitter (ATX):

$$s^{\mathrm{A}}(t) = \sum_{i=-\infty}^{\infty} a_i^{\mathrm{A}} p(t - iT_f^{\mathrm{A}} - t_i^{\mathrm{A}}) \tag{5.4}$$

where T_f^{A} is the duration of a frame, a_i^{A} is the amplitude, $t_i^{\mathrm{A}} < T_f^{\mathrm{A}}$ is the time-hopping offset of the ith frame, and $p(t)$ is the same pulse shape as used by the honest transmitter.

We further distinguish between different types of malicious interferers. The simplest *blind* adversary is not equipped with a receiver. This adversary does not attempt to synchronize

frame length	T_f^{P},T_f^{D}	256ns
preamble length	N_{P}	$64 \cdot 31$
data length	N_{D}	32
preamble code	a_i^{P}	IEEE 802.15.4a code 5
back-search window	T_{BS}	64ns
energy detector sampling time	T_{int}	2ns

Table 5.1: Default parameter values used in simulations. The default IEEE 802.15.4a preamble length is used. The back-search window is chosen to match the channel spread.

its transmission with the honest transmitter – i.e., it transmits the adversarial signal blindly. A *reactive* adversary is equipped with a receiver, and synchronizes its transmission to the transmission of HTX. The level of synchronization can be *rough* (in the order of microseconds) or *precise* (in the order of a few nanoseconds).

Furthermore, we consider adversaries with different levels of knowledge. First, an adversary can be oblivious to the codes used by the honest devices. Next, the adversary can know *in advance* the preamble codes (amplitude and time-hopping). Finally, it can in addition know the data time-hopping code. We assume the adversary never knows the amplitude data code (the payload of the ranging packet) and we assume the adversary always knows the frame lengths in the honest signal (T_f^{P}, T_f^{D}).

Finally, the target of the adversary can be all devices in range; this takes advantage of the broadcast nature of the wireless channel. Or, the target can be a specific device, with a location known to the adversary, or a specific location, where an honest device might be located. We term those *broadcast* or *targeted* attacks respectively.

The most sophisticated malicious interference attack we consider is targeted, reactive and precisely synchronized with the HTX, and mounted by an adversary that knows the preamble codes and the data time-hopping code. Such an attack is – from the perspective of HRX – essentially equivalent to a distance-decreasing relay attack or a malicious prover attack. We elaborate on this in Section 5.2.4.

5.1.4 Simulation Setup

As explained in Section 5.1.3, the malicious prover attack and the relay attack are essentially equivalent to a sophisticated malicious interference attack. Hence, we simulate malicious interference attacks only. We assume that an honest receiver (HRX) is exposed to the adversarial signal transmitted by an adversarial transmitter (ATX) at signal-to-noise ratio $\mathrm{SNR_A}$. HRX receives, at random times, ranging packets transmitted by an honest transmitter (HTX) with $\mathrm{SNR_H}$. (In both cases, the SNR is defined as $\frac{E_p}{N_0}$, where E_p is the energy of a single pulse, and the power spectral density is $N_0/2$.) We use the residential (weak) NLOS channel model [99] with a channel spread of roughly 60ns. Unless otherwise stated, we assume the parameter values summarized in Table 5.1. In particular, the back-search window $T_{\mathrm{BS}} = 64$ns is chosen long enough to account for the channel spread.

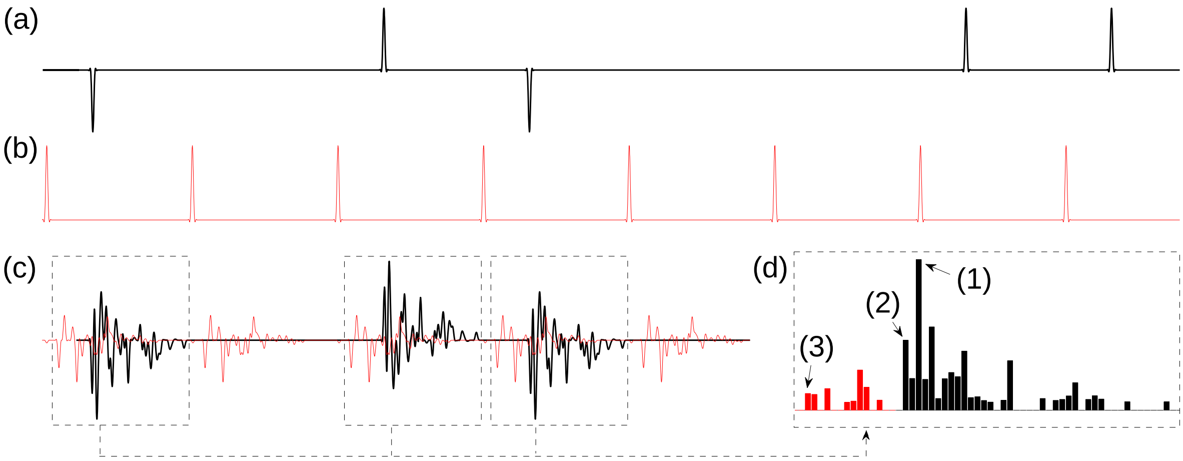

Figure 5.1: The cicada attack (blind constant 1-attack) mounted against a vanilla energy-detection receiver. (a) Benign transmitter T sends a ranging preamble modulated with a preamble code $[-1, 0, 1, -1, 0, 0, 1, 1, \ldots]$. (b) Adversary transmits the adversarial signal. (c) Both signals propagate through the multipath environment before they are received by R. (d) R aggregates the received signal over a number of pulses, and finds the strongest path (1). It then searches back for the first path (2), but instead selects the bogus path introduced by the adversary (3).

Metrics

We consider that *distance-decrease* occurs if a packet is received and the data is recovered without errors, but the estimate ToA is at least $T_{\text{dd}} = 4\text{ns}$ below the actual ToA. (T_{dd} is chosen such that the probability of obtaining such a ToA in benign conditions is negligible.) We consider that *denial* occurs, if the packet is not correctly received – either due to failure to detect the packet in coarse synchronization or due to failure of subsequent reception stages. In our performance evaluation, we measure the percentage of packets subject to distance-decrease or denial. As an additional metric, we measure the amount ot distance-decrease (*ToA error*) for the packets for which distance-decrease was successful.

5.2 ToA Attack Space

In this section, we look at various variants of the attack and evaluate their effectiveness against different modulation schemes and receivers. This provides an overview of the attack space, and guides the design of countermeasures in Section 5.3. We assume that the ToA estimation is performed only during preamble reception.

We proceed in growing attack complexity. We first explore the simplest *cicada attack* that is designed against modulation schemes with no preamble time-hopping, and demonstrate that it is effective against most receiver. Furthermore, we show that this attack is effective against most receivers even if preamble time-hopping is applied. Next, we look at the generalized version, the *coded cicada attack*, and show how this attack can defeat the vanilla rake receiver that is robust to the basic cicada attack. Then, we discuss how an *reactive* adversary can improve the attack effectiveness. Finally, we consider the most sophisticated version of the attack, where ATX is tightly synchronized and knows the preamble code. We show that this attack is highly effective, even against strongest-path ToA estimation algorithms. Note that we omit some of the less interesting (but valid) combinations of adversarial capabilities (e.g., a tightly synchronized adversary that does not know the code).

5.2.1 The "Cicada" Attack

The *cicada attack*[2] is mounted by a blind adversary transmitting an infinite sequence of identical equally spread pulses: $a_i^{\mathrm{A}} = 1$, $t_i^{\mathrm{A}} = 0$. The adversarial frame duration is $T_f^{\mathrm{A}} = \frac{1}{\rho} T_f^{\mathrm{P}}$, where the attack *rate* ρ is an integer. We term this the *constant ρ-attack*. Note that this attack does not require the knowledge of any codes. Recall that there is no synchronization with the honest transmitter, hence the adversarial pulses are located randomly with respect to the honest signal. The attack's target is broadcast.

This attack is tailored to modulation schemes without preamble time-hopping. The principle is illustrated in Figure 5.1. The signal of the honest transmitter T (Figure 5.1(a)) and the adversarial signal (Figure 5.1(b)) interfere at the honest receiver R (Figure 5.1(c)). If the adversarial signal is weaker than T's signal, R should correctly detect T's signal. However, there is a good chance that the fine synchronization algorithm will incorrectly find the "first arriving path" in the adversarial signal (Figure 5.1(d)). The estimated ToA is then significantly lower than the actual ToA, resulting in a distance-decrease.

If either the honest signal, or the adversarial signal, or both use (random) time-hopping, the honest and adversarial pulses are not going to be aligned. Rather, from the perspective of the honest receiver, the adversarial pulses are randomly spread over time in any of these 3 cases. Still, if the adversary transmits with appropriate power, such random interference turns out to be sufficient, for the most part, to introduce false peaks into fine synchronization, and not disrupt the other reception stages. To show this, we only consider the case where the honest modulation scheme uses random preamble time-hopping, as the other two cases are essentially equivalent.

Preamble without Time-Hopping

The main factor determining the attack outcome is $\mathrm{SNR_A}$. This can be seen in Figure 5.2(a), which shows the performance of the constant 1-attack against the vanilla energy-detection receiver at $\mathrm{SNR_H} = 20\mathrm{dB}$. For low $\mathrm{SNR_A}$, the adversarial signal is too weak to influence the receiver operation. From $\mathrm{SNR_A} \approx 0\mathrm{dB}$ distance-decrease begins, and it reaches its maximum of around 36% for $\mathrm{SNR_A} \approx 15\mathrm{dB}$. Beyond the maximum point, denial begins to take over, and for $\mathrm{SNR_A} \approx 25\mathrm{dB}$, it reaches 100% – partially due to coarse synchronization failure ($\mathsf{Denial}_{\mathsf{SYNC}}$), and partially due to failure of subsequent reception stages ($\mathsf{Denial}_{\mathsf{postSYNC}}$). More generally, the distance-decrease begins at $\mathrm{SNR_A} \approx 0\mathrm{dB}$ and ends at $\mathrm{SNR_A} \approx \mathrm{SNR_H}$, as confirmed by Figure 5.2(b).

To increase the probability of distance-decrease, the adversary can increase ρ. This is because with low ρ, the adversarial signal, even spread by the channel (Figure 5.1c), is not always present in the back-search window. The results are illustrated in Figure 5.3. Note that increasing ρ also results in distance-decrease ending at lower $\mathrm{SNR_A}$ – this is because there is more interference that disrupts other (than fine synchronization) stages of the receiver operation, causing denial. In subsequent experiments, unless otherwise stated, we fix $\rho = 8$,

[2] We use the term cicada attack, because the ATX signal is reminiscent of a cicada song.

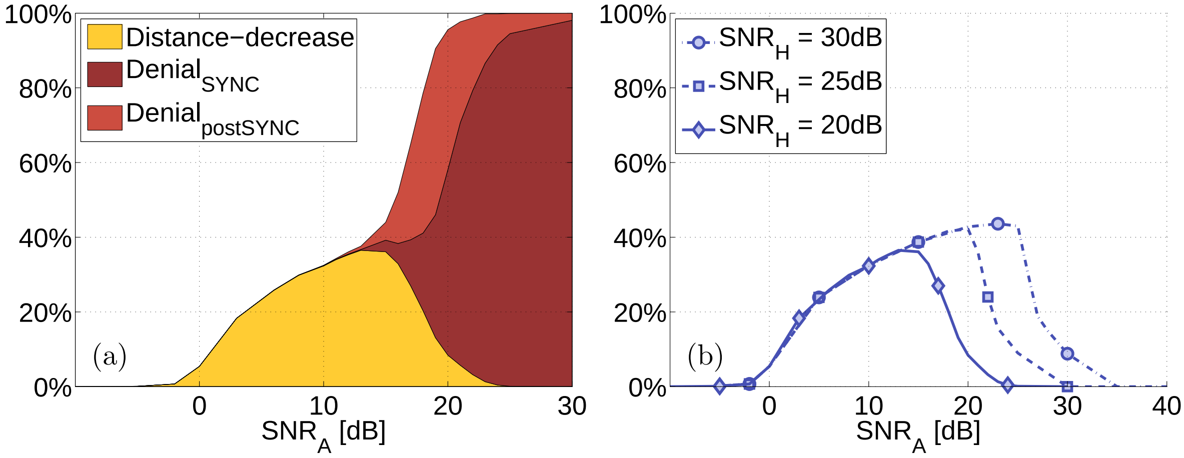

Figure 5.2: Performance of constant 1-attack mounted against a vanilla energy-detection receiver. (a) Ratio of packets for which distance-decrease and denial occurs at $\mathrm{SNR_H} = 20\mathrm{dB}$ (b) Success ratio of distance-decrease at varying $\mathrm{SNR_H}$.

which strikes a balance between achieving a high maximum distance-decrease rate and not interfering too much with other stages of the receiver operation.

We observe a similar attack performance for the MINF and PID energy-detection receivers (Figure 5.4). Both methods were designed with benign interference in mind but, as expected, neither can prevent the attack. In case of the MINF receiver, this is because the min filter cannot remove the adversarial signal present in *every* frame. Distance-decrease is less pronounced than for vanilla, because of a more conservative back-search threshold (inherent to MINF). In the case of the PID receiver, distance-decrease ends at $\mathrm{SNR_A}$ approximately 5dB lower than for vanilla. This is mostly due to coarse synchronization failure. The reason is the PID method: As soon as the adversarial signal peaks rise above the PID noise threshold, they are converted to 1. The honest signal peaks, even if they are stronger, also account for only 1. Hence, essentially, the adversarial signal drowns T's signal sooner than for the vanilla receiver (causing coarse synchronization to fail).

The adversary can overcome this limitation by not transmitting some of the adversarial pulses, thus lowering the PID correlator output for the adversarial signal. In Figure 5.5, we show the performance of such *intermittent* 8-attack, in which the adversary transmits only every 3rd frame.[3] As expected, the intermittent attack ends at considerably higher $\mathrm{SNR_A}$ than the constant attack. Compared to the constant attack, the attack begins at a higher $\mathrm{SNR_A}$, because overall ATX introduces 67% less interference. This is also why the attack ends later than the constant attack against the vanilla receiver.

Overall, we note that the performance of the attack improves slightly if the modulation scheme uses time-hopping for data. This is most pronounced for the MINF receiver, hence we show only the curve for this receiver (Figure 5.4). The reason is that the interference with the data part of the packet is the main factor limiting the distance-decrease performance. With time-hopping, not all data pulses are aligned with the adversarial pulses, which prevents losing some ranging packets to data demodulation failure.

[3] Transmitting only 1/3 of the frames is sufficient to spoof the ToA, and at the same time low enough not to interfere with a median-based channel estimation, improving the attack performance slightly.

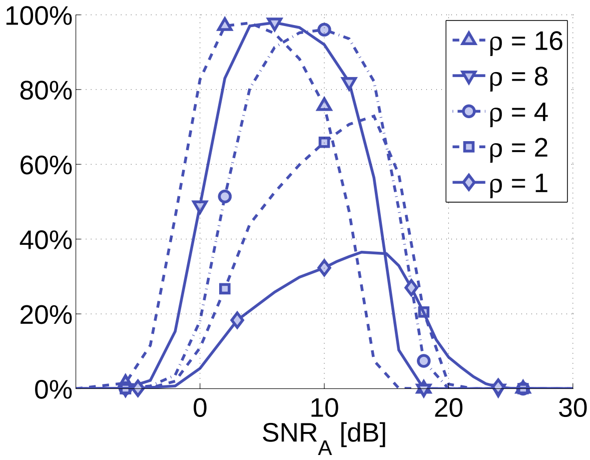

Figure 5.3: Success ratio of distance-decrease achieved by constant ρ-attack mounted against a vanilla energy-detection receiver at $\mathrm{SNR_H} = 20\mathrm{dB}$

For the vanilla rake receiver, the effectiveness of the attack depends on the sum of the amplitudes of the fine synchronization template. On one hand, if the sum is non-zero, distance-decrease will occur. This is confirmed in Figure 5.6 ("no pTHS"). The figure shows the attack performance at $\mathrm{SNR_H} = 10\mathrm{dB}$ when the template amplitudes follow the IEEE 802.15.4a preamble code 5. In this code, for every 10 frames with $a_i = 1$, there are only 6 frames with $a_i = -1$. The attack performance follows the familiar pattern, and the distance-decrease reaches 80% - 100% depending on the data modulation scheme (not shown in the figure). As with the MINF receiver, data demodulation is the main liming factor for attack performance. Hence, applying time-hopping in the data improves attack performance.

On the other hand, if the fine synchronization template sums to zero, the constant cicada code causes the adversarial pulses to cancel each other out, and no distance-decrease occurs.

Preamble with Time-Hopping

The cicada attack is effective against the vanilla and PID energy-detection receivers even with preamble time-hopping, as shown in Figure 5.7. Compared to the case without preamble time-hopping, distance-decrease begins at higher SNR. This is because the adversarial pulses are not aligned, but rather spread from the perspective of the receiver. Hence, it takes more power to raise them above the fine synchronization threshold. For vanilla, distance-decrease ends at the same SNR with and without preamble time-hopping, because the limiting factor is data demodulation. However, for PID, distance-decrease ends significantly later for the case *with* preamble time-hopping. This is because for PID *without* time-hopping, the limiting factor is coarse synchronization (see Section 5.2.1). Preamble time-hopping circumvents this limitation.

For the vanilla rake receiver, the attack performance is shown in Figure 5.6, and it follows the same pattern as for the vanilla energy detector, with data demodulation being the limiting factor. However, it should be noted that with preamble time-hopping the attack works even with a zero-sum template.

For MINF, the min filter is relatively efficient in diminishing the effects of the attack once the adversarial pulses are randomly spread due to the time-hopping. In Figure 5.7 we show

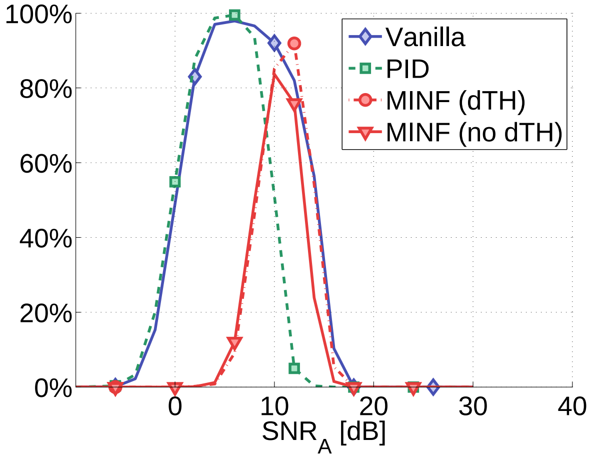

Figure 5.4: Success ratio of distance-decrease achieved by constant 8-attack mounted against a vanilla, PID, and MINF energy-detection receivers $\mathrm{SNR_H} = 20\mathrm{dB}$; for the latter we shown performance against modulation with and without data time-hopping (dTH).

the performance of one instance of the attack, which manages to reach almost 20% probability of distance-decrease. This is significantly lower than for the other receivers, but still far from negligible.

Ranging Error

To illustrate the magnitude of the distance-decrease, we show in Figure 5.8 the median absolute error of the ToA under the 8-attack taken over packets for which distance-decrease occurs. The error increases with $\mathrm{SNR_A}$, because more adversarial peaks rise above the fine-synchronization threshold. We can also observe that even for high $\mathrm{SNR_A}$, the variance of the error is quite high. This implies that with the blind attack, the adversary has relatively low control over the amount of distance-decrease introduced. All receivers experience similar ToA error patterns.

Clock Drift

An unsophisticated or blind adversary will not have the means to adjust its clock speed to exactly match the clock of HTX. We hence explore the effect of clock drift on the distance-decrease performance. We introduce a clock drift of 40ppm (the largest allowed by IEEE 802.15.4a) between the HTX and ATX, while the HRX's clock runs at the speed of the HTX's clock. For the rake receiver, this drift causes the attack to begin at an $\mathrm{SNR_A}$ approximately 5dB larger than for the no-drift case, whereas for energy detectors there is no difference. For both receiver classes, the drift shifts the end of distance-decrease to higher $\mathrm{SNR_A}$. This is because between channel estimation and the subsequent stages that use the channel estimate, the adversarial signal drifts out of sync with the channel estimate – essentially diminishing the weight of the adversarial signal.

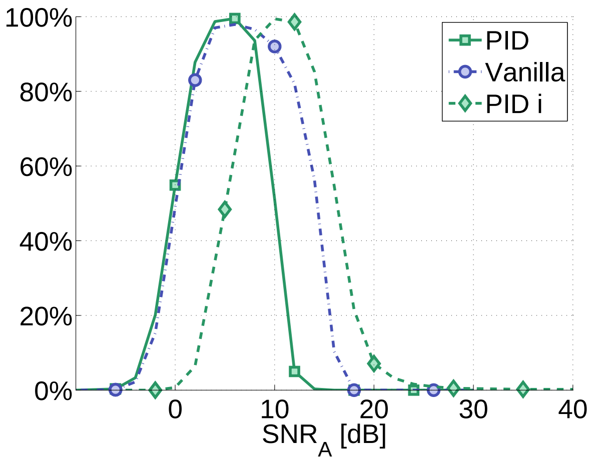

Figure 5.5: Success ratio of distance-decrease achieved by an intermittent ("i") 8-attack against a PID energy-detection receiver at $\text{SNR}_{\text{H}} = 20\text{dB}$. For comparison, we also show the performance of the constant 8-attack against PID and vanilla energy detectors.

5.2.2 Coded "Cicada" Attack

In Section 5.2.1 we showed that the adversary can use time-hopping to mount a successful attack against a vanilla rake receiver with zero-sum fine synchronization template (without time-hopping). Alternative, the adversary can achieve this with a *coded ρ-attack*, in which he modulates the amplitudes of the pulses using a non-constant code. We term this attack the *coded cicada attack*' except for the non-constant code it is identical to the basic cicada attack: It is blind, broadcast, and it does not require the knowledge of any codes.

We evaluate the probability of success of this attack against a vanilla rake receiver analytically. Assume that the number of non-zero elements in the fine synchronization template is N. In fine synchronization, the ToA is estimated as the first time-offset t for which:

$$|\sum_{i=0}^{N-1} a_i^{\text{P}} y_i(t)| > N\Theta \tag{5.5}$$

where $y_i(t)$ are the rake samples taken at time instances corresponding to the preamble code and time-offset t, and Θ is a (noise-based) threshold for a single frame.

Assuming that offset t contains only the adversarial signal, and that noise is negligible, (5.5) boils down to:

$$|\sum_{i=0}^{N-1} a_i^{\text{P}} a_i^{\text{A}} x| > N\theta \tag{5.6}$$

where x is the power with which the adversary transmits, $a_i^{\text{A}} \in \{-1, 0, 1\}$ is the adversarial code, and $\theta = C\Theta$ is a threshold normalized by a factor C.

We chose a pseudo-random preamble code, which is the worst case for the adversary. This implies $\mathbb{P}(a_i^{\text{P}} a_i^{\text{A}} = -1 \mid a_i^{\text{P}} \neq 0) = \mathbb{P}(a_i^{\text{P}} a_i^{\text{A}} = 1 \mid a_i^{\text{P}} \neq 0) = 0.5$. Assuming that there is a single adversarial frame in the back-search window, the probability of spoofing the ToA for

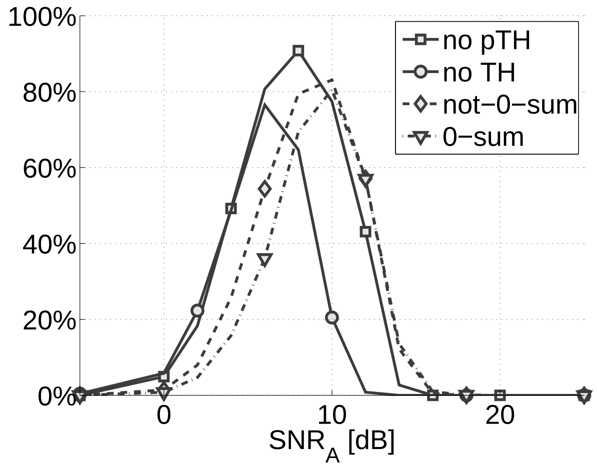

Figure 5.6: Success ratio of distance-decrease achieved by constant 8-attack mounted against the vanilla rake receiver at $\mathrm{SNR_H} = 10\mathrm{dB}$. We show the effectiveness of the attack when no time-hopping is used ("no TH"), when only data time-hopping is used ("no pTH"), and when both data and preamble time-hopping is used; for the latter we distinguish between a non-zero-sum template ("not-0-sum") and a zero-sum template ("0-sum").

time-offset t is:

$$\begin{aligned}\mathbb{P}(A) &= \mathbb{P}(|2\mathcal{B}(n,0.5) - n| > N\theta x^{-1}) \\ &= 2 \cdot \mathbb{P}(2\mathcal{B}(n,0.5) - n > N\theta x^{-1}) \\ &= 2 \cdot F_{\mathrm{BIN}}(0.5n - 0.5N\theta x^{-1}|n, 0.5) \\ &\leq 2\exp(-N^2\theta n^{-1}x^{-1}) \leq 2\exp(-N\theta x^{-1})\end{aligned} \tag{5.7}$$

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

Although the probability of spoofing decreases exponentially fast with N, the transmission power x is under the control of the adversary. By increasing x, the adversary can achieve a reasonable spoofing probability for practical values of N. Indeed, Figure 5.9 shows practical instances of this attack with $N = 120$. To increase the attack success probability, the attack rate is $\rho = 16$, meaning that 4 to 5 adversarial frames fall into the back-search window. Although (5.7) suggests to set $n = N$, this does not take into account the interference created by the adversarial signal during other reception stages. Indeed, Figure 5.9 shows that the adversary achieves better results with $n = 12$, rather than $n = 120$, where $N = 120$ in both cases.

5.2.3 Reactive Attacks

From the adversarial perspective, one of the drawbacks of the cicada attack is the sensitivity of attack effectiveness to $\mathrm{SNR_A}$. Indeed, notably for the MINF receiver (Figure 5.4), the $\mathrm{SNR_A}$ region for which the attack has high probability of success is quite narrow. This is because at high $\mathrm{SNR_A}$ the adversarial signal causes too much interference during coarse

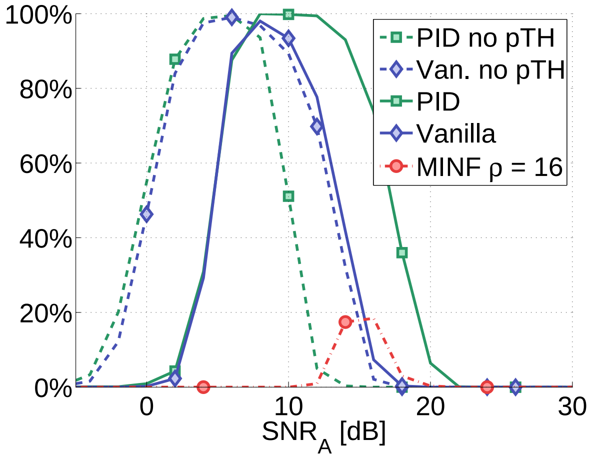

Figure 5.7: Modulation with preamble time-hopping: Success ratio of distance-decrease achieved by constant 8-attack (16-attack for MINF) mounted against energy-detection receivers at $\mathrm{SNR_H} = 20\mathrm{dB}$. Performance without preamble time-hopping ("no pTH") is shown for comparison.

synchronization, channel estimation, SFD detection and/or data demodulation, and prevents the ranging packet from being received correctly.

A reactive adversary can try to overcome this. It is reasonable to assume that the adversary, knowing the algorithms run by the honest receiver, will know (roughly) when the receiver is performing which stage. Then, the reactive adversary can transmit the adversarial signal only when fine synchronization is performed, with arbitrary high power and high rate ρ, which guarantees that the ToA is spoofed. If no other stage is performed on this signal, then the adversary will not interfere with those stages, which guarantees correct packet reception.

However, a receiver can perform two or more stages on the same signal as fine synchronization. On one hand, if that stage is coarse synchronization, $\mathrm{SNR_A}$ cannot exceed $\mathrm{SNR_H}$. Otherwise the honest signal will be overshadowed by the adversarial signal, and coarse synchronization would fail. On the other hand, if that stage is channel estimation, $\mathrm{SNR_A}$ can be increased more substantially, but not indefinitely. Intuitively, the ratio between $\mathrm{SNR_A}$ and $\mathrm{SNR_H}$ in channel estimation determines how much weight is put on the bogus part versus the honest, information-bearing part of the data symbols. The upper limit on $\mathrm{SNR_A}$ depends on what the adversary does in the data part. We elaborate on this in Section 5.2.4.

5.2.4 Known Code Attacks

We now look at a targeted version of the malicious interference attack, in which the adversary knows the preamble code, and is tightly synchronized with HTX. Under these assumptions, the adversary can spoof the ToA easily, by transmitting a copy of the preamble that arrives at the receivers k nanoseconds before the honest preamble. HRX has no means of detecting this attack. In the data part, the adversary knows the time-hopping sequence, but not the amplitudes which are determined by the challenge/response bits. Hence, ATX transmits a pulse k nanoseconds before each honest pulse, with $\mathrm{SNR_A^D}$. Without loss of generality, we assume that the adversarial pulse has a constant amplitude, whereas the honest pulse has amplitude 0/1 for OOK modulation, or ±1 for BPSK modulation.

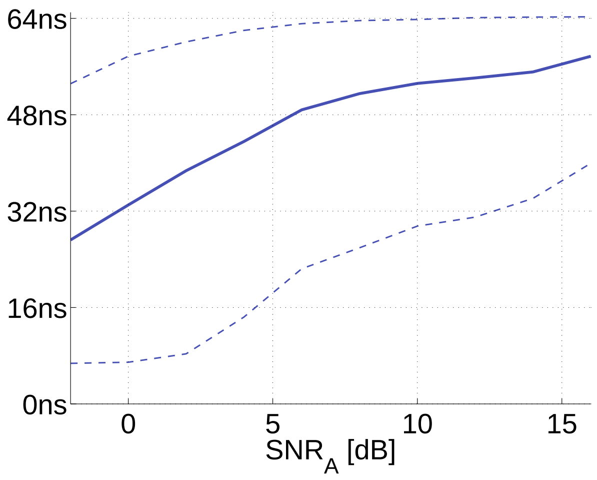

Figure 5.8: ToA median absolute error under constant 8-attack against the vanilla energy-detection receiver at $\mathrm{SNR_H} = 20\mathrm{dB}$ (over packets for which distance-decrease was successful), modulation without time-hopping. 5% and 95% percentiles are also shown.

Note that in case of a distance-decreasing relay attack or a malicious prover attack, the adversary would transmit a signal which is the combination of the honest signal and the adversarial signal described above. We hence focus on the malicious interference attack in what follows.

We can see in Figure 5.10 that this attack achieves virtually 100% distance-decrease across a relatively broad range of $\mathrm{SNR_A}$. (By default, we set $\mathrm{SNR_A^D} = -\infty$.) Furthermore, the distance-decrease is much more deterministic that with the blind attacks. With the offset k set to $k = 20\mathrm{ns}$, the 5% and 95% percentiles of absolute ToA error are 10ns and 22ns at $\mathrm{SNR_A} = 10\mathrm{dB}$ and 17ns and 22ns at $\mathrm{SNR_A} = 15\mathrm{dB}$. The adversary can further improve the ToA spoofing performance by making the ATX to HRX channel as close as possible to a single-tap channel (i.e., no multipath), e.g., by moving very close to the victim receiver and/or using a highly directional antenna. In this case, the absolute ranging error percentiles are between 19ns and 22ns. Furthermore, with this attack, the adversary can succeed in spoofing not only the 1st arriving path but even the strongest path, albeit with a lower success probability. This would circumvent a simple "countermeasure" that estimates ToA based on the strongest path (Section 5.3).

In general, the adversary can choose the values of $\mathrm{SNR_A}$ and $\mathrm{SNR_A^D}$ that optimize the performance on the data. Not surprisingly a good choice is to set $\mathrm{SNR_A}$ large enough to achieve a desired success ratio of distance-decrease, but not larger. With this choice, the adversarial contribution to the channel estimate, and the resulting influence on data demodulation are minimized. For such small $\mathrm{SNR_A}$, the choice of $\mathrm{SNR_A^D}$ plays a negligible role. Small values, e.g., $\mathrm{SNR_A^D} = -\infty\mathrm{dB}$ give good attack performance, meaning that denial rate is low. Indeed, consider as a baseline the $\mathrm{SNR_H}$ necessary to achieve a packet error rate below 1% if no attack takes place, denote it $\mathrm{SNR_H^{ben}}$. Under attack, denote by $\mathrm{SNR_H^{att}}$ the $\mathrm{SNR_H}$ necessary to achieve a packet error rate below 1%. Then, the difference between $\mathrm{SNR_H^{att}}$ and the baseline $\mathrm{SNR_H^{ben}}$ is below 0.2dB for energy-detection receivers and rake receivers alike.

However, the adversary can decide to choose a large $\mathrm{SNR_A}$. This makes sense notably if the adversary is not targeted, but mounts a broadcast attack and wants to extend the

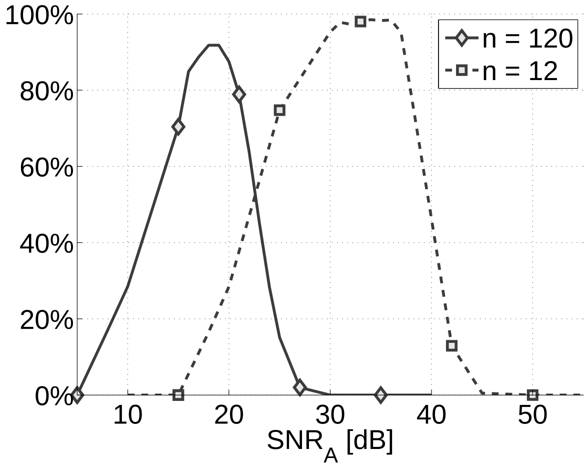

Figure 5.9: Success ratio of distance-decrease achieved by a coded 16-attack against a vanilla rake receiver at $\mathrm{SNR_H} = 20\mathrm{dB}$. The number of frames during which adversary transmits (out of 120) is denoted by n.

geographical area effected. In that case, the optimal adversarial strategy depends on the receiver algorithms and $\mathrm{SNR_H}$. For example, we established via simulations that for the vanilla energy detector at $\mathrm{SNR_H} = 20\mathrm{dB}$ the optimal strategy is to set $\mathrm{SNR_A^D} = \mathrm{SNR_A} - 5\mathrm{dB}$. In Figure 5.10 we can see that for this attack the upper limit on $\mathrm{SNR_A}$ is roughly 7dB greater than the attack with $\mathrm{SNR_A^D} = -\infty\mathrm{dB}$.

In contrast, for the rake receiver it is easy to show that the optimal attack is to set $\mathrm{SNR_A^D} = -\infty\mathrm{dB}$. Under such an attack, at $\mathrm{SNR_H} = 10\mathrm{dB}$ the success ratio of distance decrease drops below 5% at $\mathrm{SNR_A} = 15\mathrm{dB}$ (not shown in a figure) compared to $\mathrm{SNR_A} = 12\mathrm{dB}$, 14dB and 15dB for the cicada attack mounted against various modulation schemes (Figure 5.6).

5.3 Countermeasures

In this section, we use the insight we gained in Section 5.2 to evaluate and propose countermeasures that can thwart the ToA attack vector. We start with a few naive countermeasures. Then we discuss two groups of countermeasures that are secure against the attack. The first group is based on *secure ToA estimation*. We consider two algorithms, the PID method for rake receivers and a novel PIDH for energy-detection receivers. The second group relies on insecure ToA estimation, and verification of the ToA estimate in the data part. The verification is performed with early detection countermeasure from Chapter 4 with very short detection time (*very early detection*, VED). We study the performance of both countermeasures in the benign case, and show that VED has better ranging precision but is more prone to fail. We then propose an extension of VED that we term VEDG (*very early detection with graceful degradation*) and show that it combines the best performance properties of PIDH and VED.

5.3.1 Naive Countermeasures

An obvious way of countering the attack is to disable fine synchronization, and estimate ToA based on the strongest path. This has, however, the significant disadvantage of decreasing

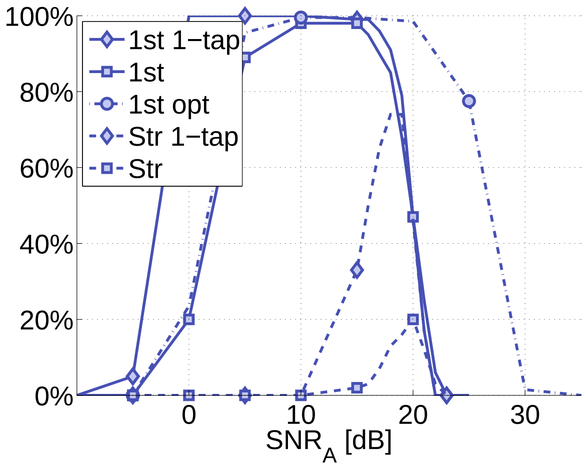

Figure 5.10: Success ratio of distance-decrease achieved by known-code attack mounted against a vanilla energy-detection receiver at $\mathrm{SNR_H} = 20\mathrm{dB}$. We show the performance under two adversarial channel models: the default one and a single-tap channel ("1-tap"). We distinguish between the usual success in spoofing the 1st path ("1st"), and spoofing of the strongest path ("Str"). In the data, we fix $\mathrm{SNR_A^D} = -\infty\mathrm{dB}$, except for "opt" where we use the optimal attack, i.e., $\mathrm{SNR_A^D} = \mathrm{SNR_A} - 5\mathrm{dB}$.

the ranging precision. Furthermore, we have seen in Section 5.2.4 that a more sophisticated adversary that knows the preamble codes and that is precisely synchronized with the honest transmitter can circumvent this method.

Another way to counter the attack is to repeat the ranging over multiple packets, and detect the unusually high variability of the measured distance caused by the attack (Figure 5.8). This again has the disadvantage that it can be circumvented by a precise reactive adversary.

Finally, in Section 5.2.1 we have seen that the MINF energy-detection receiver, if coupled with random preamble time-hopping, can offer some degree of protection against the attack. However, the adversary still has a non-negligible probability of success, making this a relatively weak countermeasure.

5.3.2 Secure ToA Estimation

The key to secure ToA estimation is the PID method. Indeed, we have seen in Section 5.2.2 that a rake receiver that does not use PID is vulnerable to an attack by an adversary that transmits with high power.

PID Rake

We first consider the PID rake reciever. Assume that the amplitudes of non-zero frames are independent and $\mathbb{P}(a_i = -1) = \mathbb{P}(a_i = 1) = 0.5$. Then, for the PID rake, we can apply an almost identical reasoning as in (5.7) to estimate the probability of spoofing the ToA:

$$\begin{aligned} \mathbb{P}(A) &= \mathbb{P}(|2\mathcal{B}(n, 0.5) - n| > N\theta) \\ &= 2 \cdot F_{\mathrm{BIN}}(0.5n - 0.5N\theta | n, 0.5) \leq 2\exp(-N\theta) \end{aligned} \tag{5.8}$$

where N is the number of non-zero frames in the fine synchronization template, $n \leq N$ is the number of non-zero frames in the adversarial code a_i^{A}, and θ is a noise-based threshold. The difference with (5.7) is that the power factor x disappears. This is crucial, because the adversary can no longer increase the transmission power to compensate for an exponentially fast decline of $\mathbb{P}(A)$. This implies that the PID rake is secure against the attack with an arbitrary transmission power, if the values of N and θ are chosen appropriately and if the amplitudes are unknown to the adversary.

PIDH Countermeasure

We now propose a secure ToA estimation algorithm for energy-detection receivers. It also relies to the PID method. The input to algorithm is a sequence of samples $y_{i=1,\dots,N} \in \{0,1\}$ (after applying the PID filter) corresponding to some time-offset t, and the known template $a_{i=1,\dots,N} \in \{0,1\}$. Then, the time-offset t is considered a valid ToA candidate if:

$$d(y_i, a_i) \leq \tau \tag{5.9}$$

where d is the Hamming distance and τ is a threshold. All time-offsets t_i in the back-search window are validated, and the resulting ToA is the first valid ToA. Essentially, in this algorithm we replace correlation with the Hamming distance, hence we denote it *Power Independent Detection with the Hamming distance* (PIDH).

For optimal security, we assume that the template a_i is a pseudo-random binary sequence ($\mathbb{P}(a_i = 0) = \mathbb{P}(a_i = 1) = 0.5$), unknown in advance to the adversary. Under this assumption, for a single time-offset t (without benign signal contribution), the adversary can spoof the ToA with probability:

$$\mathbb{P}(A(t)) = F_{\mathrm{BIN}}(\tau | N, 0.5) \tag{5.10}$$

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

$$\mathbb{P}(A) = \mathbb{P}(\bigcup_j A(t_j)) \leq \sum_j \mathbb{P}(A(t_j) = N_{\mathrm{BS}} \cdot F_{\mathrm{BIN}}(\tau | N, 0.5) \tag{5.11}$$

For a desired security level P_{attack} we can invert (5.11) and obtain the threshold τ that achieves this security level:

$$\tau = F_{\mathrm{BIN}}^{-1}(P_{\mathrm{attack}} \cdot N_{\mathrm{BS}}^{-1} | N, 0.5) \tag{5.12}$$

The security and performance of the PIDH and the PID rake ToA estimation algorithms are independent of time-hopping.

5.3.3 Security through (Very) Early Detection

VED Countermeasure

An alternative to secure ToA estimation is to perform ToA estimation in a non-secure fashion, and rely on data demodulation with a very short detection time ($t_{\mathrm{ED}} = T_{\mathrm{int}}$) to detect attacks. This is a natural extension of the early detection countermeasure advocated in Section 4.4.1, and we term in *(very) early detection* (VED). It proceeds as follows:

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

The values of N_{nonce} and N_{err} can be determined by (3.3).

If is not hard to see the similarities between VED and the PIDH ToA estimation algorithm. Indeed, applying the PID filter is similar to performing OOK demodulation with detection time equal to the integration time ($t_{\text{det}} = T_{\text{int}}$). Furthermore, computing the Hamming distance to the template is identical to counting the number of errors in a demodulated bit sequence (Section 3.1.4), and (3.3) is equivalent to (5.12).

Naturally, there are also differences between the PIDH method and the VED method. These have an effect on benign-case performance. First, in the former method N_{BS} time-offsets are evaluated versus only one time-offset in VED. Hence, to obtain the same probability of attack success P_{attack}, the PIDH threshold τ is lower than the corresponding N_{err} threshold for VED. This means that that PIDH is more likely to reject a valid time-offset. Second, the threshold of the PID filter in PIDH is determined using the estimated noise level only, whereas in the VED method the OOK demodulation threshold can be determined based on the noise level and the signal level. This means that VED has a higher reliability of demodulation. In both cases, the VED method has a slight performance advantage over PIDH, which is confirmed in Section 5.3.4.

However, the VED method is less robust than the PIDH method, in the sense that it is more likely to fail in benign conditions. This is because if the first pulse detected by the insecure ToA estimation algorithm in VED is low, there is a chance that verification will fail, causing packet rejection. In contrast, in such a case PIDH does not reject the packet, but rather reports a ToA obtained at a stronger, but later pulse. In other words, PIDH offers graceful degradation of service, reporting a ToA that is both secure, and as close as possible to the actual ToA (but never lower, i.e., distance-decrease does not occur).

VEDG Countermeasure

To address the above drawback of VED, we propose an extension, *very early detection with graceful degradation* (VEDG). In this method, instead of performing data demodulation at only one time-offset (the first detected pulse), data demodulation is performed independently at multiple time-offsets (the first few detected pulses). The resulting ToA is the fist time-offset for which verification is successful. With multiple time-offsets, the acceptable number of errors (N_{err}) must be decreased to maintain a fixed security level P_{attack}. This decreases the ranging precision. To minimize this negative effect we select time-offsets such that their number is low compared to the PIDH method. The VEDG method proceeds as follows:

1. **ToA estimation.** Perform non-secure fine synchronization that selects multiple time-offsets t_i:
   (a) t_1 is selected as the first offset in the back-search window that is above the noise-based threshold (regular ToA estimation).
   (b) t_{i+1} is selected as the first offset in the back-search window that is above the ith time-offset.
2. **Verification.** Perform data demodulation with detection time $t_{\text{det}} = T_{\text{int}}$ (on the time-

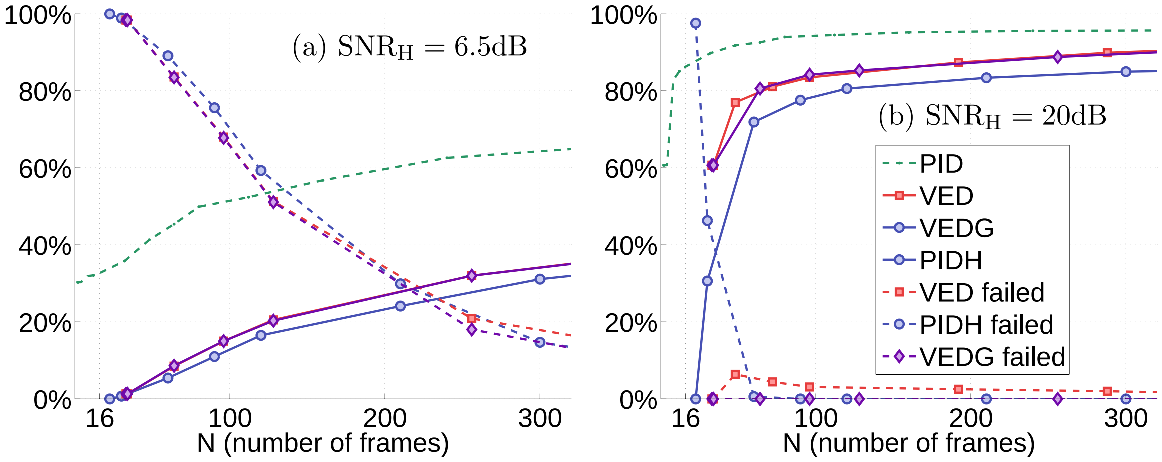

Figure 5.11: Countermeasures benign case performance: We show the precise-ToA rate ("PIDH", "VED", "VEDG"), as a function of the number of frames N at (a) $\mathrm{SNR_H} = 6.5\mathrm{dB}$ (the same as in Figure 4.15) (b) $\mathrm{SNR_H} = 20\mathrm{dB}$. For PID (the non-secure ToA estimation algorithm used in VED), $N = N_{\mathrm{toa}}$ (the template length). For PIDH, N is the template length. For VED/VEDG, $N = N_{\mathrm{toa}} + N_{\mathrm{nonce}}$, the latter being the length of the data. In addition, we show the failure rate of all methods ("failed").

offset determined by fine synchronization) for every time-offset t_i; the ToA estimate is the first t_i for which the number of errors is below N_{err}.

With this method, the number of time-offsets is kept small compared to PIDH.

5.3.4 Performance Evaluation

The PIDH, VED and VEDG methods offer the same level of security, but perform differently with respect to other metrics. We evaluate the performance of the countermeasures in the benign case. Our primary metric pertains to ranging precision: We count the percentage of packets for which the receiver finds the first path, i.e., the ranging error is below 2ns. We term this metric *precise-ToA rate*. We find it to be more informative than the rather coarse-grained mean absolute ranging error, or mean root square error. Furthermore, we count the percentage of packets which are not received properly, and we term this metric *failure rate*. Finally, the *imprecise-ToA rate* is the percentage of packets that are received, but with a ToA estimate greater than the actual ToA (by more than 2ns). We do not report the values for the last metric, as it sums to 100% with the other two metrics. This is because it is almost impossible (probability below P_{attack}) that the ToA estimate is below the actual ToA, i.e., the *distance-decrease rate* is approximately zero. The exact values that we obtain are specific to the channel model (the residential NLOS model from [99]), but we are more interested in relative performance of different algorithms.

We assume $P_{\mathrm{attack}} = 2^{-16}$ and derive the thresholds for PIDH, VED and VEDG according to (5.12) and (3.6). Compared to the non-secure PID ToA estimation algorithm[4] used in VED/VEDG, the PIDH ToA estimation has a relatively low benign-case probability of false detection of P_{FD} (approximately, the per-time-offset $P_{\mathrm{FD}} \approx 10^{-9}$ for PIDH, compared to $P_{\mathrm{FD}} \approx 10^{-4}$ for PID). Hence, the probability of missed detection, P_{MD}, is higher for PIDH. This means that for the same template duration, the PIDH algorithm misses some of the lower

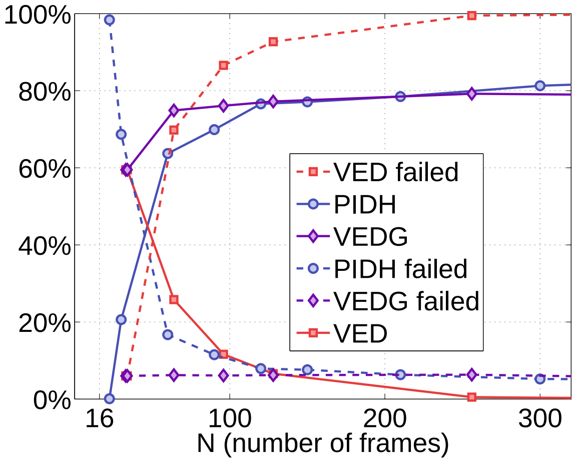

Figure 5.12: Countermeasures performance under constant 8-attack (at $\mathrm{SNR_A} = 8\mathrm{dB}$). We show the precise-ToA rate (“PIDH”, “VED”, “VEDG”), as a function of the number of frames N at $\mathrm{SNR_H} = 20\mathrm{dB}$. For PIDH, N is the template length. For VED/VEDG, $N = N_{\mathrm{toa}} + N_{\mathrm{nonce}}$, where N_{toa} is the template length and N_{data} is length of the data. We also show the failure rate of all methods (“failed”).

peaks, captured by the PID algorithm. Thus, the latter exhibits better ranging precision. This is confirmed in Figure 5.11, which shows the performance of the PID and the PIDH algorithms for the same template length N. PID performs significantly better.

However, we need to consider the full VED/VEDG operation. In the PIDH method, all detected pulse are automatically secure. In contrast, in VED/VEDG a pulse detected by the non-secure ToA estimation algorithm still needs to be verified as secure. Hence, a considerable number of data frames (N_{nonce}) needs to be devoted to provide verification with an acceptable probability of false rejection (as hinted by Figure 4.15). In fact, increasing N_{toa} decreases the power of pulses that can be detected by the PID algorithm. Reliably verifying such low pulses requires an increase of N_{nonce}.

To compare PIDH with VED/VEDG, we assume that for VED/VEDG the number of frames N is equal to $N_{\mathrm{toa}} + N_{\mathrm{nonce}}$. For a fixed N, we set $N_{\mathrm{data}} = kN_{\mathrm{toa}}$, where k ranges from 5 to 7 depending on $\mathrm{SNR_H}$. We established this ratio experimentally. In Figure 5.11 we see that for both low and high SNR ($\mathrm{SNR_H} = 6.5\mathrm{dB}$ and $\mathrm{SNR_H} = 20\mathrm{dB}$) the chosen ratio of N_{toa} and N_{data} allows VED to outperform PIDH in terms of ranging precision. Furthermore, VEDG achieves an almost identical ranging precision as VED. This is because the number of time-offsets for which VEDG demodulates is relatively small: below 6 for 99% of packets. At $\mathrm{SNR_H} = 20\mathrm{dB}$, to achieve precise-ToA rate of 80%, VED/VEDG requires $N \approx 65$, compared to $N \approx 115$ required for PIDH. In contrast, PID requires a fine synchronization template of length only $N \approx 8$. The cost of precision and security is indeed quite high.

The second metric of interest is the failure rate, i.e., the number of ranging packets that are not received correctly either due to synchronization failure or due to rejection by the VED/VEDG verification. For $\mathrm{SNR_H} = 6.5\mathrm{dB}$, all three methods experience similar failure rate. However, for $\mathrm{SNR_H} = 20\mathrm{dB}$, and $N > 70$ the failure rate of VED is in the order of 2 to 3%. In contrast, for PIDH the failure rate of PIDH is negligible for $N > 70$, whereas for

[4] Other ToA estimation algorithms can also be used.

VEDG it is negligible altogether. This confirms our expectation that PIDH and VEDG are more robust than VED in the benign case.

We also evaluate both countermeasure under attack. More specifically, in Figure 5.12 we show the countermeasures' performance at $\mathrm{SNR_H} = 20\mathrm{dB}$ against a constant 8-attack with $\mathrm{SNR_A} = 8\mathrm{dB}$. (We choose $\mathrm{SNR_A}$ according to Figure 5.4.) Compared to benign case operation at $\mathrm{SNR_H} = 20\mathrm{dB}$, the PIDH method experiences a relatively minor degradation of performance (decrease of precision, increase of failure). VEDG performs almost as well as the PIDH for large N, and better for small N. In contrast, the VED method experiences a significant performance degradation, notably as N increases. The reason for this is that as N increases, the template length N_{toa} increases. Hence, the insecure PID ToA estimation algorithm is more likely to fall victim to distance-decrease. If this happens, VED verification must fail. This confirms that the PIDH and VEDG countermeasures are more robust than the VED countermeasure under attack. However, note that by increasing $\mathrm{SNR_A}$, the adversary can cause a denial against both PIDH and VEDG.

5.4 Experiments

To further demonstrate the feasibility of the cicada attack, we test it on an Impulse-Radio test-bed that is being developed in the scope of the MICS project [56]. The primary goal of the test-bed is time-difference-of-arrival (TDOA) localization of a mobile transmitter. However, at the current stage of development, the test-bed has certain limitations, which do not allow us to replicate the experiments simulated in Section 5.2. We hence design a different set of experiments compatible with the test-bed. With these experiments, we are able to show the distance-decreasing effect of the cicada attack.

5.4.1 IR Test-Bed

The test-bed is comprised of a set of static receivers and a set of mobile transmitters. To perform TDOA localization of a mobile transmitter, the receivers should be synchronized. However, at the current stage of development, there is no synchronization between the receivers or, more generally, any pair of receiver/transmitter devices. Hence, we are not able to implement ranging, or even pseudo-ranging. This is the primary reason for a new experiment design.

Transmitter A transmitter is composed of an analog front-end described in [36] and of a simple modulator module. The transmitter can send a train of pulses with pulse repetition frequency (PRF) of 10MHz. The train of pulses can be modulated with a binary sequence, i.e., with on/off keying modulation. Pulses have a bandwidth of 500MHz at 4.25GHz central frequency.

Receiver A receiver front-end is described in [56]. It is an energy-detection receiver, i.e., the received signal is squared. However, in contrast to a traditional energy-detector assumed in rest of this Chapter, the test-bed receiver uses I/Q demodulation. The goal of I/Q demodulation is to prevent a beat that could occur due to a small frequency mismatch between the receiver and transmitter local oscillators. However, from the perspective of the attack performance, such a difference in architecture is of little consequence, as observed in Section 5.4.3. This confirms that the universal nature of the attack.

The receiver back-end is composed of a dual 1.5GS/s ADC08D1500 analog-to-digital converter (ADC). It samples the signal at a rate of 2.842GS/s with 8 bit resolution. These samples are then handled by an Altera Cyclone II EP2C70 Field-Programmable Gate Array (FPGA). The FPGA has relatively limited memory of 140Kb, which influences the experiment design. The FPGA is connected to a PC with a serial-to-USB converter.

5.4.2 Experiment Setup

As explained in Section 5.4.1, the test-bed has two limitations that make it necessary to design a new set of experiments. The first is the lack of a synchronization, i.e., a common clock base, between any two devices. This causes the following difficulty: When the receiver estimates the ToA, there is no *ground truth* ToA that the estimate could be compared with. Hence, we cannot immediately determine if a distance-decrease has occurred.

We can resolve this issue as follows. We make HTX transmit a number of consecutive ranging packets. At the same time, ATX periodically turns the cicada attack on and off, for periods of time longer than the ranging packet length (see Figure 5.13 (b) and (c)). As a result, some of the ranging packets will be affected by the cicada signal, some will be affected partially, and some will not be affected at all. Assuming that we can (roughly) determine when the cicada signal is on/off, HRX can use the ToA computed for unaffected ranging packets as the ground truth.

The second limitation is the modest memory on the FPGA. To overcome this issue, we refrain from using complete ranging packets, with a preamble and payload. Rather, HTX only transmits short preambles, and HRX only performs coarse synchronization (simplified) and fine synchronization.

Configurations The experiments involve the usual set of devices: an honest receiver (HRX), an honest transmitter (HTX) and an adversarial transmitter (ATX). We consider two configurations: In configuration (A), there are no obstacles between either transmitter and the receiver. In configuration (B) we attenuate the ATX signal by putting a tin-foil obstacle between ATX and HRX. In both configurations, the distances between HTX and HRX and ATX and HRX are in the order of one meter. The test-bed is located in the corner of a 10m by 10m electrical engineering lab room filled with working benches and other equipment.

We define SNR as the ratio between the signal power and the noise power.[5] In both configurations, $\mathrm{SNR_H} \approx 22\mathrm{dB}$. In configuration (A), $\mathrm{SNR_A} \approx 21\mathrm{dB}$. In configuration (B), $\mathrm{SNR_A}$ varies between 6dB and 17dB with mean 11dB.

Experiment In a single experiment, we capture $12\mu s$ of samples from the ADC (maximum that the FPGA memory allows), and send them to the PC. There, we perform the remaining receiver operations in MATLAB. This allows us to minimize implementation overhead, compared to implementing the receiver on the FPGA. Note that the results are not affected. It also allows us to process the samples off-line, and run different receiver algorithms on the

[5] Recall that in other sections of the thesis we define SNR as the ratio between the pulse energy and the noise spectral density, E_p/N_0. Converting E_p/N_0 to $P_{\mathrm{signal}}/P_{\mathrm{noise}}$ entails subtracting roughly 24dB. For example, $\mathrm{SNR_H} = 10log_{10}(E_p/N_0) = 20\mathrm{dB}$ in Figure 5.4 translates to $10log_{10}(P_{\mathrm{signal}}/P_{\mathrm{noise}}) \approx -4\mathrm{dB}$. Hence, the signals in the experiments are significantly stronger than the signals used in the simulations. This is because of the short distances between devices in the experiment setup.

same input. For each configuration, we perform 1000 experiments. During an experiment both HTX and ATX are transmitting.

Honest Transmitter HRX transmits an infinite sequence of preamble symbols, almost identical[6] to the preamble symbols used in IEEE 802.15.4a. We consider each preamble symbol to be a separate ranging packet. In more detail, a preamble symbol is composed of 31 frames, modulated according to the IEEE 802.15.4a preamble code no 5. The frame duration is 100ns. Hence, the duration of a preamble symbol is 3.1μs. Figure 5.13(a) shows the received samples for one experiment with only HTX transmitting.

Adversarial Transmitter ATX transmits periodically a sequence of $4 \cdot 31$ frames modulated with a sequence of $2 \cdot 31$ ones followed by $2 \cdot 31$ zeros. The frame duration is 100ns. Examples of the ATX signal are shown in Figure 5.13(b) and (c). Note that this is a simple variant of the cicada attack with rate $\rho = 1$.

Honest Receiver Operation HRX operation is adjusted to the signal transmitted by HTX. Most notably, the receiver only performs synchronization, but no channel estimation, SFD detection or data demodulation. *Coarse synchronization* can be significantly simplified compared to Section 5.1.2, as HRX can a priori assume that the HTX signal is always present. Recall that a a single preamble symbol is the entire "ranging packet". Hence, HRX simply correlates the received signal with a template that corresponds to one preamble symbol. Then, the receiver finds the maximum in the correlator output of duration 3.1μs, which is the "ranging packet" duration (i.e., one preamble symbol). The index at which the maximum is found is deemed the coarse ToA estimate.

Fine synchronization is performed on the same signal as coarse synchronization. It follows exactly the description in Section 5.1.2, searching back from the coarse ToA estimate for a time-offset above a threshold. The only differences are the duration of the back-search window (32ns in place of 64ns, corresponding to the observed channel spread) and the MINF window size (4 in place of 8, to account for honest signal imperfections).

Given the number of samples in a single experiment, and the preamble symbol duration, in one experiment we obtain a ToA estimate 3 times (dashed lines in Figure 5.13(a)). Because there is no synchronization between HTX and HRX, there is no way of knowing what should be the ToA estimate, i.e., the *ground truth*. However, we can compare the 3 ToA estimates within an experiment, and obtain a *relative* ToA estimation error. In other words, we assume that one of the 3 ToA estimates is the ground truth. More precisely, this is the estimate corresponding to the "ranging packet" (preamble symbol) not affected by the cicada signal.

[6] The differences are in the frame duration, and in a signal modulated by a binary, rather than a ternary signal. The latter is enforced by the available transmitter and is irrelevant for an energy detector receiver because of the squaring operation.

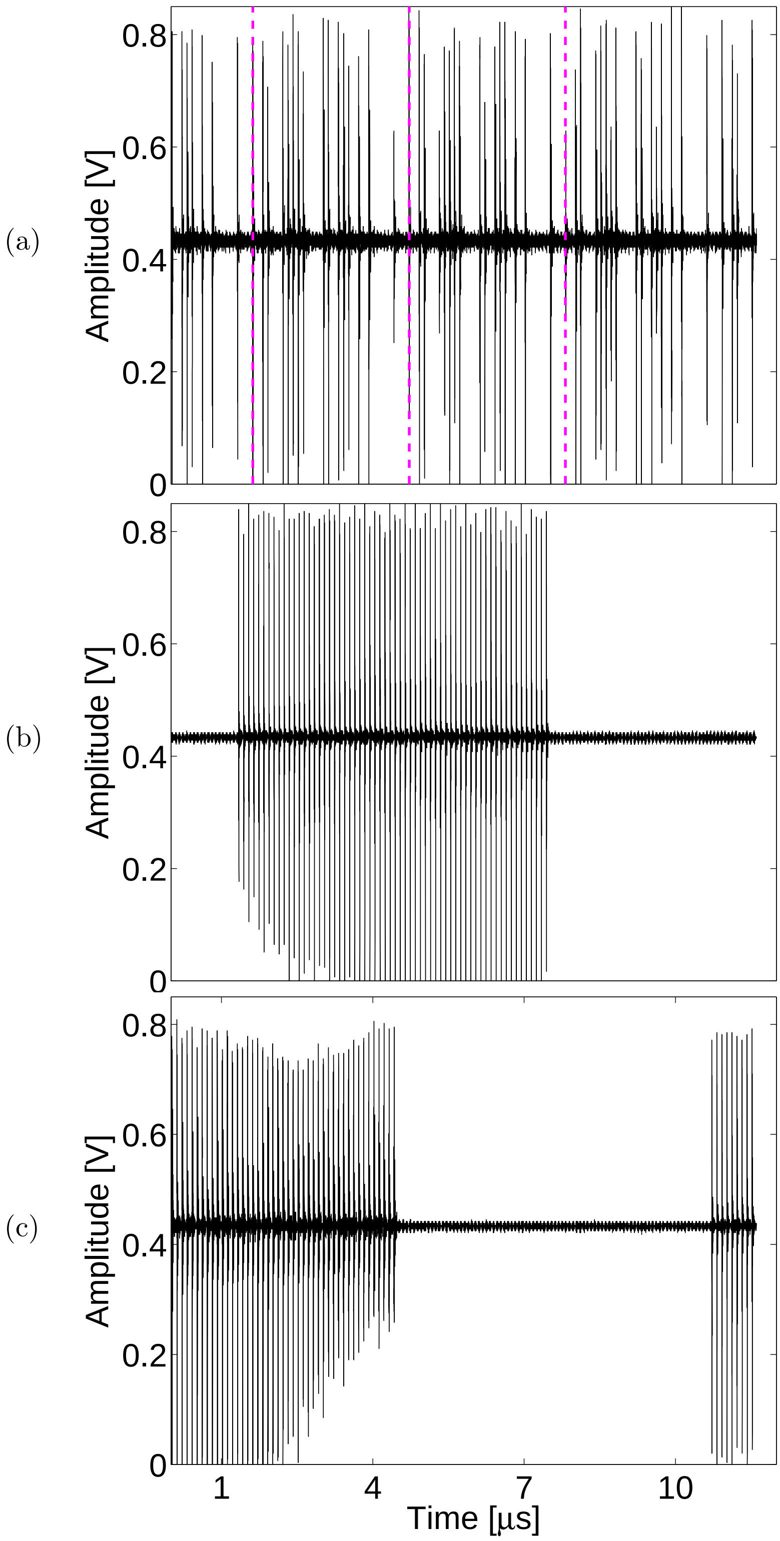

Figure 5.13: Samples from the HRX's ADC for a single experiment. (a) HTX signal only, the dashed lines show the start of preamble symbols. (b), (c) ARX signal only.

Metrics For a preamble symbol received by HRX we define *coverage* as the percentage of the preamble symbol ("ranging packet") that is covered by the ATX signal. For example, in Figure 5.13, if (a) and (b) are received simultaneously, than the coverage of the 1st symbol is 1, the 2nd symbol is around 0.8 and the 3rd symbol is 0. If (a) and (c) are received simultaneously, the coverage is roughly 0.8, 0 and 0.2, for the 1st, 2nd and 3rd preamble symbol, respectively. In practice, we compute the coverage based on the position of the cicada signal, which we detect with the PID method combined with computing the Hamming distance from the cicada template (rather than correlation). We use this detection method due to its robustness. We note that the detection is not extremely accurate (it can be off by a few frames), but this is good enough for our purposes.

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

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

Note that contrary to simulations performed in Section 5.2, denial can occur due to a benign failure. This is because we are working with prototype-grade hardware, which sometime results in imperfections in the received signals. (For example, note the variation across pulse amplitude in Figure 5.13, notably (c)).

5.4.3 Experimental Results

We show the distance-decrease rate and the denial rate as a function of coverage in Figure 5.14. The results follow our expectations. Consider first configuration (A), where the ATX signal is roughly as strong as the HTX signal. For low coverage, the vanilla receiver shows some distance degradation which increases up to roughly 20% as the coverage increases, but for coverage beyond 0.3 denial starts to take over. To explain this behavior, recall that the vanilla receiver *sums* the samples from the frames indicated by the template. Hence, the coverage effectively works as a multiplying factor for the total adversarial signal power; e.g., at coverage 0.5, $\mathrm{SNR_A}$ can be considered approximately 3dB lower than $\mathrm{SNR_A}$ at coverage 1 (where $\mathrm{SNR_A} \approx \mathrm{SNR_H}$). Thus, the distance-decrease rate and denial rate follow the same pattern that we have seen in Section 5.2. A similar coverage interpretation can be applied to MINF receiver, although it should be noted that the min filter provides some non-linear distortion. In particular, at low coverage, the min filter removes the cicada signal completely, and thus no distance-decrease is observed. In general, the performance pattern of the attack is similar as for vanilla, but the distance-decrease rate reaches only roughly 10% due to a conservative threshold value.

In contrast, for the PID receiver, there is no simple parallel between the coverage and $\mathrm{SNR_A}$, but we can find an equally simple interpretation. The output of the PID correlator at a time-offset corresponding to the true ToA (the start of the "ranging packet") is equal to N, the number of non-zero frames in the template (in noiseless conditions). With coverage x, the output of the correlator at a time-offset with adversarial signal contribution is roughly

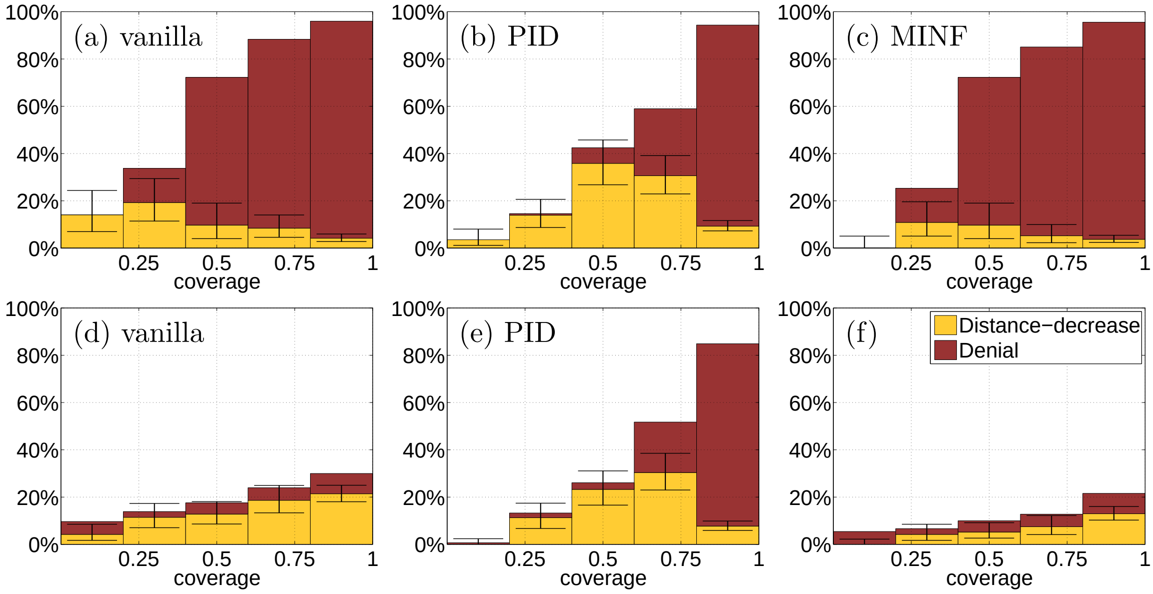

Figure 5.14: Distance-decrease rate and denial rate as a function of coverage. Configuration (A) results are shown in (a) for vanilla, (b) for PID and (c) for MINF. Configuration (B) results are shown in (d) for vanilla, (e) for PID and (f) for MINF. We shown 95% confidence intervals for distance-decrease rate.

$x \cdot N$. This explains why in Figure 5.14(b) denial becomes dominant only at coverage close to 1. This allows the distance-decrease rate to reach roughly 35%.

In configuration (B) the ATX signal is weaker then the HTX signal. Hence, for the vanilla and MINF filters, Figure 5.14 can be considered to show the performance for a low $\mathrm{SNR}_{\mathrm{A}}$ only. This explains why denial occurs marginally, and the distance-decrease rate increases with coverage. In contrast, for the PID receiver the power of the adversarial signal plays a smaller role, as long as it is above the PID filter threshold. Hence, the attack performance in configuration (B) is close to the performance in configuration (A).

Finally, we verify that no distance-decease occurs with the PIDH fine synchronization algorithm.

5.5 Conclusion

We have identified a novel attack vector against IR-UWB ranging, based on disrupting the time-of-arrival (ToA) estimation. This *ToA attack vector* allows an adversary to decrease the measured distance in malicious prover attacks and distance-decreasing relay attacks. It also creates a novel type of PHY attack based on malicious interference. We have demonstrated, with simulations and experiments on an IR test-bed, that even a simple-to-mount variant of the malicious interference attack (the cicada attack) is effective against a number of modulation schemes and receivers that are designed for precise ranging, i.e., attempt to find the first arriving path rather than the strongest path. We have also shown that more sophisticated variants of the malicious interference attack can decrease the distance measured by less precise receivers that only detect the strongest path.

Furthermore, we have investigated countermeasures that mitigate this attack vector and thus achieve distance bounding that is both *precise* and *secure*. "Precise" means that the DB protocol can achieve a ranging precision close to the optimal ranging precision of the receiver, even in weak NLOS conditions. "Secure" means that the DB protocol prevents the adversary from decreasing the measured distance (by more than the ranging precision of the receiver). We have shown that this can be achieved by *secure ToA estimation* algorithms. One example of such an algorithm is the known PID algorithm for rake receivers with appropriately adjusted parameters. Another is the novel PIDH algorithm for energy-detection receivers. We have also revisited the early detection countermeasure from Chapter 4, decreasing substantially the detection time. We term the resulting countermeasure VED (for *very early detection*). We have shown that this countermeasure is as secure as PIDH and exhibits higher ranging precision than PIDH, but has a higher failure rate. Finally, we have proposed a hybrid countermeasure, VEDG, that achieves the ranging precision of VED and has a failure rate close to PIDH.

Compare these countermeasures with the countermeasures considered in Chapter 4. Foremost, the latter countermeasures allow an adversary to achieve a distance-decrease in the order of the channel spread, in contrast to a negligible distance-decrease allowed by Chapter 5 countermeasures. However, in terms of *benign case* ranging precision, the countermeasures in Chapter 4 offer the ranging precision of insecure ranging. In comparison, the countermeasures in Chapter 5 offer lower benign case precision, even with ranging packets orders of magnitude longer; The cost of having precision and security at the same time is quite high. Nevertheless, using the countermeasures in Chapter 5, a reasonable ranging precision can be achieved while keeping the packet duration below 100μs; this should not be prohibitive for most applications. It is an interesting open question whether it is possible to design countermeasures that have lower communication cost.

Acknowledgments

I would like to extend my thanks to the designers and implementers of the MICS IR test-bed: Alexander Feldman, Alexander Bahr, James Colli-Vignarelli, Stephan Robert, Catherine Dehollain, and Alcherio Martinoli. They give me access to the test-bed and provided me with invaluable help in performing the experiments.

Conclusion

In this thesis, we focus on two fundamental elements of secure wireless networking: neighbor discovery and distance bounding. We investigate two aspects crucial for the security of these protocols, which nevertheless have received relatively little attention: *formal analysis* and *security on the physical-communication-layer*. Our motivation stems from a basic security principle: A system is as secure as its weakest link. Indeed, if physical-communication-layer attacks are not properly addressed, any cryptographic protocol that attempts to provide distance bounding or neighbor discovery based on message propagation time can be defeated. Furthermore, informal arguments for protocol correctness have repeatedly proven insufficient as new attacks are discovered against published protocols.

Our first contribution is a formal framework for reasoning about time- and location-based neighbor discovery (ND) protocols. The framework is among the first to incorporate notions such as propagation time, node location, or the neighbor relation into formal reasoning about security protocols. We consider two general classes of protocols: time-based protocols (T-protocols) and time- and location-based protocols (TL-protocols). The framework allows us to analyze and design concrete provably secure communication ND protocols in both classes. It also allows us to formally prove a fundamental limitation: No T-protocol can provide two-part communication ND if adversarial nodes can relay messages faster than a particular threshold. The threshold is determined by the honest nodes' communication range. This result is a useful measure of the security achieved by ND T-protocols that we believe are one of the more practical and universal solutions to ND.

Our second contribution is an investigation of physical-communication-layer (PHY) attacks against ranging and distance bounding, and of countermeasures against such attacks. We focus on a technology particularly well suited for ranging and distance-bounding: Impulse-Radio Ultra-wideband (IR-UWB). We demonstrate that de facto standard for IR-UWB, IEEE 802.15.4a, is vulnerable to distance-decreasing attacks, first introduced in [35]. In particular, if honest devices use energy-detection receivers without appropriate countermeasures, an adversary can decrease the measured distance by hundreds of meters, with a success rate arbitrarily close to 100%. This is in part because of features of the standard that improve benign-case performance. We propose minor patches to IEEE 802.15.4a, which improve its security while retaining the benefits introduced by these features. With appropriate countermeasures implemented by honest receivers, these patches thwart distance-decreasing attacks: At a mild cost in terms of packet length, the maximum distance decrease can be limited to values in the order of the channel spread. Furthermore, we discover a new attack vector that disrupts time-of-arrival estimation algorithms. It allows an adversary to decrease the measured distance by a value in the order of the channel spread; hence, it is a threat for precise ranging. This vector also creates a new type of PHY attack based on malicious interference; it is much simpler to mount than alternative external PHY attacks. Finally, we propose countermeasures against the ToA attack vector; they enable ranging that is both precise and secure, although at a non-negligible cost in terms of packet length. In summary, the countermeasures we propose allow a system designer to trade off security, ranging precision and packet length.

Future Work

The formal framework we propose in Part I is used for reasoning about communication ND and can be easily extended to physical ND and distance bounding. However, the assumptions on which the framework is constructed are reasonable only if both nodes participating in the ND/DB protocol are honest. To capture internal attacks, we need a framework that models messages at the bit level and that is probabilistic, making it possible to express attacks that work with non-negligible probability. (The only work we are aware of that takes a step in this direction is [119].) Furthermore, the framework should also model PHY attacks in some way. A simple solution would be to incorporate those as ToA estimation error, but a more direct modeling could also be explored: This could reveal a yet undiscovered hybrid attack that exploits the interaction between PHY attacks and cryptographic protocols. Ideally, the framework should allow for mechanization or even automation.

In Part II, we evaluate PHY attacks and countermeasures with simulations. We also provide a simple proof-of-concept implementation of the attacks in an IR-UWB test-bed, limited to the synchronization phase. The next logical step is to create a more elaborate implementation that would include all receiver operations. This would allow for a complete validation of the attacks, as well as the VED and VEDG countermeasures. The ultimate goal is a complete implementation of an IR DB protocol secure from PHY attacks. The most challenging step in such an implementation is to build a prover able to perform the rapid response in the rapid-bit-exchange. (An interesting proposal for an analog rapid response, along with a prototype implementation, can be found in [136]).

In Chapter 5 we propose countermeasures that limit achievable distance-decrease to values below the channel spread. However, they introduce a large cost in terms of packet length, compared to countermeasures in Chapter 4. The existence and design of less costly countermeasures is an interesting open question.

In Part II we focus on one particular technology, Impulse-Radio Ultra-wideband. Any other PHY technology envisioned to be used for time-based secure ND or for DB, should be evaluated for feasibility of PHY attacks. This, in particular, includes an investigation of feasibility of the malicious interference attack we discovered.

Bibliography

[1] http://online.wsj.com/article/SB10001424052748703576204576226722412152678.html.

[2] https://foursquare.com/.

[3] https://www.facebook.com/places/.

[4] http://wigle.net/.

[5] http://www.bloomberg.com/news/2011-03-29/microsoft-is-said-to-plan-mobile-payments-in-next-windows-phone.html.

[6] http://www.businessweek.com/technology/content/jan2011/tc20110125_174308.htm.

[7] http://www.google.com/mobile/navigation/.

[8] http://www.skyhookwireless.com/.

[9] http://www.toll-collect.de/.

[10] http://www.tomtom.com/.

[11] Nokia Instant Community project, http://research.nokia.com/news/9391.

[12] G. Ács, L. Buttyán, and I. Vajda. Provably Secure On-Demand Source Routing in Mobile Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 5(11), 2006.

[13] Z. Ahmadian and L. Lampe. Performance Analysis of the IEEE 802.15.4a UWB System. *IEEE Transactions on Communications*, 57(5), 2009.

[14] H. Alzaid, S. Abanmi, S. Kanhere, and C. T. Chou. A Wireless Sensor Networks Testbed for the Wormhole Attack. *International Journal of Digital Content Technology and its Applications*, 3(3), 2009.

[15] C. Arackaparambil, S. Bratus, A. Shubina, and D. Kotz. On the Reliability of Wireless Fingerprinting Using Clock Skews. In *Proceedings of the 3rd ACM Conference on Wireless Network Security (WiSec)*, 2010.

[16] G. Avoine, M. A. Bingöl, S. Kardaş, C. Lauradoux, and B. Martin. A Framework for Analyzing RFID Distance Bounding Protocols. *Journal of Computer Security – Special Issue on RFID System Security*, 19(2), 2010.

[17] G. Avoine, C. Floerkemeier, and B. Martin. RFID Distance Bounding Multistate Enhancement. In *Proceedings of the 10th International Conference on Cryptology in India: Progress in Cryptology*, 2009.

[18] G. Avoine, C. Lauradoux, and B. Martin. How Secret-sharing Can Defeat Terrorist Fraud. In *Proceedings of the 4th ACM Conference on Wireless Network Security (WiSec)*, 2011.

[19] G. Avoine and A. Tchamkerten. An Efficient Distance Bounding RFID Authentication Protocol: Balancing False-Acceptance Rate and Memory Requirement. In *Proceedings of the 12th International Conference on Information Security*, 2009.

[20] D. Basin, S. Čapkun, P. Schaller, and B. Schmidt. Let's Get Physical: Models and Methods for Real-world Security Protocols. In *Proceedings of the 22nd International Conference on Theorem Proving in Higher Order Logics (TPHols)*, 2009.

[21] A. Benfarah, B. Miscopein, J. M. Gorce, C. Lauradoux, and B. Roux. Distance Bounding Protocols on TH-UWB Radios. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOMM)*, 2010.

[22] T. Beth and Y. Desmedt. Identification Tokens — or: Solving The Chess Grandmaster Problem. In *Advances in Cryptology*, volume 537 of *Lecture Notes in Computer Science*, 1991.

[23] S. Brands and D. Chaum. Distance-bounding protocols. In *Advances in Cryptology - Workshop on the Theory and Application of Cryptographic Techniques (EUROCRYPT)*, volume 765 of *Lecture Notes in Computer Science*, 1994.

[24] V. Brik, S. Banerjee, M. Gruteser, and S. Oh. Wireless Device Identification with Radiometric Signatures. In *Proceedings of the 14th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2008.

[25] L. Bussard. *Trust establishment protocols for communicating devices.* PhD thesis, 2004.

[26] L. Bussard and Y. Roudier. Authentication in Ubiquitous Computing. In *Proceedings of the Workshop on Security and Ubiquitous Computing (Ubicomp)*, 2002.

[27] L. Bussard and W. Bagga. Distance-Bounding Proof of Knowledge to Avoid Real-Time Attacks. In *Security and Privacy in the Age of Ubiquitous Computing*, volume 181 of *IFIP Advances in Information and Communication Technology*, 2005.

[28] L. Buttyán, L. Dóra, and I. Vajda. Statistical Wormhole Detection in Sensor Networks. In *Proceedings of the 2nd European Workshop on Security and Privacy in Ad hoc and Sensor Networks (ESAS)*, 2005.

[29] I. Cervesato, C. Meadows, and D. Pavlovic. An Encapsulated Authentication Logic for Reasoning about Key Distribution Protocols. *Proceedings of the 18th IEEE Workshop on Computer Security Foundations*, 2005.

[30] H. Chen, W. Lou, and Z. Wang. Conflicting-Set-Based Wormhole Attack Resistant Localization in Wireless Sensor Networks. In *Proceedings of the 6th International Conference on Ubiquitous Intelligence and Computing*, 2009.

[31] J. T. Chiang, J. J. Haas, and Y.-C. Hu. Secure and Precise Location Verification using Distance Bounding and Simultaneous Multilateration. In *Proceedings of the 2nd ACM Conference on Wireless Network Security (WiSec)*, 2009.

[32] J. Chiang, J. Haas, Y.-C. Hu, P. Kumar, and J. Choi. Fundamental Limits on Secure Clock Synchronization and Man-In-The-Middle Detection in Fixed Wireless Networks. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM)*, 2009.

[33] H. S. Chiu and K.-S. Lui. DelPHI: Wormhole Detection Mechanism for Ad Hoc Wireless Networks. In *Proceedings of 1st International Symposium on Wireless Pervasive Computing*, 2006.

[34] H. Choe, E. Poole, A. Yu, and H. Szu. Novel Identification of Intercepted Signals from Unknown Radio Transmitters. In *Proceedings of Conference on Society of Photo-Optical Instrumentation Engineers (SPIE)*, 1995.

[35] J. Clulow, G. P. Hancke, M. G. Kuhn, and T. Moore. So Near and Yet So Far: Distance-Bounding Attacks in Wireless Networks. In *Proceedings of the 3rd European Workshop on Security and Privacy in Ad hoc and Sensor Networks (ESAS)*, 2006.

[36] J. Colli-Vignarelli and C. Dehollain. A Discrete-Components Impulse-Radio Ultrawide-Band (IR-UWB) Transmitter. *IEEE Transactions on Microwave Theory and Techniques*, 59(4), 2011.

[37] M. Conti, R. D. Pietro, L. V. Mancini, and A. Mei. A Randomized, Efficient, and Distributed Protocol for the Detection of Node Replication Attacks in Wireless Sensor Networks. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2007.

[38] C. Cremers, K. B. Rasmussen, and S. Capkun. Distance Hijacking Attacks on Distance Bounding Protocols. Cryptology ePrint Archive, Report 2011/129, 2011.

[39] A. A. D'Amico, U. Mengali, and E. Arias-De-Reyna. Energy-Detection UWB Receivers with Multiple Energy Measurements. *IEEE Transactions on Wireless Communications*, 6(7), 2007.

[40] B. Danev, T. Heydt-Benjamin, and S. Čapkun. Physical-layer Identification of RFID Devices. In *Proceedings of the 18th USENIX Security Symposium*, 2009.

[41] B. Danev, H. Luecken, S. Čapkun, and K. El Defrawy. Attacks on Physical-layer Identification. In *Proceedings of the 3rd ACM Conference on Wireless Network Security (WiSec)*, 2010.

[42] B. Danev and S. Čapkun. Transient-based Identification of Wireless Sensor Nodes. *Proceedings of the 8th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2009.

[43] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Win. Ranging With Ultrawide Bandwidth Signals in Multipath Environments. *Proceedings of the IEEE*, 97(2), 2009.

[44] D. Dardari, A. Giorgetti, and M. Win. Time-of-Arrival Estimation of UWB Signals in the Presence of Narrowband and Wideband Interference. In *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, 2007.

[45] R. de Graaf, I. Hegazy, J. Horton, and R. Safavi-Naini. Distributed Detection of Wormhole Attacks in Wireless Sensor Networks. In *Ad Hoc Networks Journal*, volume 28 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, 2010.

[46] Y. Desmedt. Major Security Problems with the "Unforgeable"(Feige-)Fiat-Shamir Proofs of Identity and How to Overcome Them. In *Proceedings of SecuriCom*, 1988.

[47] T. Dimitriou and A. Giannetsos. Wormholes No More? Localized Wormhole Detection and Prevention in Wireless Networks. In *Distributed Computing in Sensor Systems*, volume 6131 of *Lecture Notes in Computer Science*, 2010.

[48] D. Dolev and A. Yao. On the Security of Public Key Protocols. *IEEE Transactions on Information Theory*, 29(2), 1983.

[49] D. Dong, M. Li, Y. Liu, X.-Y. Li, and X. Liao. Topological Detection on Wormholes in Wireless Ad Hoc and Sensor Networks. In *Proceedings of the 17th IEEE International Conference on Network Protocols (ICNP)*, 2009.

[50] P. F. Driessen and L. J. Greenstein. Modulation Techniques for High-speed Wireless Indoor Systems Using Narrowbeam Antennas. *IEEE Transactions on Communications*, 43(10), 1995.

[51] S. Drimer and S. Murdoch. Keep Your Enemies Close: Distance Bounding Against Smartcard Relay Attacks. In *Proceedings of the 16th USENIX Security Symposium*, 2007.

[52] A. El Fawal and J.-Y. Le Boudec. A Power Independent Detection Method for UWB Impulse Radio Networks. In *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, 2005.

[53] A. El Fawal and J.-Y. Le Boudec. A Robust Signal Detection Method for Ultra Wide Band (UWB) Networks with Uncontrolled Interference. *IEEE Transactions on Microwave Theory and Techniques*, 54(4), 2006.

[54] J. Eriksson, M. Faloutsos, and S. Krishnamurthy. Routing amid Colluding Attackers. In *Proceedings of the 15th IEEE International Conference on Network Protocols (ICNP)*, 2007.

[55] J. Eriksson, S. V. Krishnamurthy, and M. Faloutsos. TrueLink: A Practical Countermeasure to the Wormhole Attack in Wireless Networks. In *Proceedings of the 2006 14th IEEE International Conference on Network Protocols (ICNP)*, 2006.

[56] A. Feldman, A. Bahr, J. Colli-Vignarelli, S. Robert, C. Dehollain, and A. Martinoli. Toward the Deployment of an Ultra-Wideband Localization Test Bed. In *Procedding of the 74th Vehicular Technology Conference (VTC)*, 2011.

[57] M. Flury, R. Merz, and J.-Y. Le Boudec. An Energy Detection Receiver Robust to Multi-User Interference for IEEE 802.15.4a Networks. In *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, 2008.

[58] M. Flury, R. Merz, and J.-Y. Le Boudec. Robust Non-Coherent Timing Acquisition in IEEE 802.15.4a IR-UWB Networks. In *Proceedings of the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2009.

[59] M. Flury, M. Poturalski, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec. Effectiveness of Distance-Decreasing Attacks Against Impulse Radio Ranging. In *Proceedings of the 3rd ACM Conference on Wireless Network Security (WiSec)*, 2010.

[60] M. Flury. *Interference Robustness and Security of Impulse-Radio Ultra-Wide Band Networks*. PhD thesis, Lausanne, 2010.

[61] A. Francillon, B. Danev, and S. Čapkun. Relay Attacks on Passive Keyless Entry and Start Systems in Modern Cars. In *Proceedings of the 18th Annual Network and Distributed System Security Symposium (NDSS)*, 2011.

[62] L. Francis, G. Hancke, K. Mayes, and K. Markantonakis. Practical NFC Peer-to-peer Relay Attack using Mobile Phones. In *Proceedings of the 6th International Conference on Radio Frequency Identification*, 2010.

[63] W. Galuba, P. Papadimitratos, M. Poturalski, K. Aberer, Z. Despotovic, and W. Kellerer. Castor: Scalable Secure Routing for Ad Hoc Networks. In *Proceedings of the 29th IEEE International Conference on Computer Communications (INFOCOM)*, 2010.

[64] S. Gezici, Z. Tian, G. Giannakis, H. Kobayashi, A. Molisch, H. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *Signal Processing Magazine, IEEE*, 22(4), 2005.

[65] I. Guvenc, Z. Sahinoglu, P. Orlik, and H. Arslan. Searchback Algorithms for TOA Estimation in Non-coherent Low-rate IR-UWB Systems. *Wireless Personal Communications*, 48(4), 2009.

[66] G. Hancke and M. Kuhn. An RFID Distance Bounding Protocol. In *Proceedings of the 1st International Conference on Security and Privacy for Emerging Areas in Communications Networks (SecureComm)*, 2005.

[67] G. P. Hancke. Practical Attacks on Proximity Identification Systems. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2006.

[68] G. P. Hancke and M. G. Kuhn. Attacks on time-of-flight Distance Bounding channels. In *Proceedings of the 1st ACM Conference on Wireless Network Security (WiSec)*, 2008.

[69] G. Hancke. Design of a Secure Distance-Bounding Channel for RFID. *Elsevier Journal of Network and Computer Applications*, 2010.

[70] G. Hancke, K. Mayes, and K. Markantonakis. Confidence in Smart Token Proximity: Relay Attacks Revisited. *Computers & Security*, 28(7), 2009.

[71] Y. Hanna, H. Rajan, and W. Zhang. Slede: A Domain-specific Verification Framework for Sensor Network Security Protocol Implementations. In *Proceedings of the 1st ACM Conference on Wireless Network Security (WiSec)*, 2008.

[72] C. He, M. Sundararajan, A. Datta, A. Derek, and J. C. Mitchell. A Modular Correctness Proof of IEEE 802.11i and TLS. In *Proceedings of the 12th ACM Conference on Computer and Communications Security (CCS)*, 2005.

[73] L. Hu and D. Evans. Using Directional Antennas to Prevent Wormhole Attacks. In *Proceedings of the 11th Annual Symposium on Network and Distributed Systems Security (NDSS)*, 2004.

[74] Y.-C. Hu, A. Perrig, and D. B. Johnson. Packet Leashes: A Defense against Wormhole Attacks in Wireless Networks. In *Proceedings of the 22nd IEEE International Conference on Computer Communications INFOCOM*, 2003.

[75] Y.-C. Hu, D. Johnson, and A. Perrig. SEAD: Secure Efficient Distance Vector Routing for Mobile Wireless Ad Hoc Networks. In *Proceedings 4th IEEE Workshop on Mobile Computing Systems and Applications*, 2002.

[76] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner, Jr. Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer. In *Proceedings of the ION GNSS Conference*, 2008.

[77] IEEE Computer Society, LAN/MAC Standard Committee. IEEE P802.15.4a/D7 (Amendment of IEEE Std 802.15.4), Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks, 2007.

[78] S. Jana and S. K. Kasera. On Fast and Accurate Detection of Unauthorized Wireless Access Points Using Clock Skews. In *Proceedings of the 14th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2008.

[79] G. Kapoor, W. Zhou, and S. Piramuthu. Distance Bounding Protocol for Multiple RFID Tag Authentication. In *Proceedings of the IEEE/IFIP International Conference on Embedded and Ubiquitous Computing (EUC)*, 2008.

[80] O. Kara, S. Kardaş, M. A. Bingöl, and G. Avoine. Optimal Security Limits of RFID Distance Bounding Protocols. In *Proceedings of the 6th Workshop on RFID Security (RFIDSec)*, 2010.

[81] Z. Kfir and A. Wool. Picking Virtual Pockets using Relay Attacks on Contactless Smartcard. In *Proceedings of the 1st International Conference on Security and Privacy for Emerging Areas in Communications Networks (SecureComm)*, 2005.

[82] I. Khalil, S. Bagchi, and N. B. Shroff. LiteWorp: A Lightweight Countermeasure for the Wormhole Attack in Multihop Wireless Networks. *Proceedings of the International Conference on Dependable Systems and Networks (DNS)*, 2005.

[83] I. Khalil, S. Bagchi, and N. B. Shroff. LiteWorp: Detection and Isolation of the Wormhole Attack in Static Multihop Wireless Networks. *Computer Networks*, 51(13), 2007.

[84] I. Khalil, S. Bagchi, and N. B. Shroff. MobiWorp: Mitigation of the Wormhole Attack in Mobile Multihop Wireless Networks. *Ad Hoc Networks*, 6(3), 2008.

[85] C. H. Kim and G. Avoine. RFID Distance Bounding Protocol with Mixed Challenges to Prevent Relay Attacks. In *Proceedings of the 8th International Conference on Cryptology and Network Security (CANS)*, 2009.

[86] C. H. Kim, G. Avoine, F. Koeune, F.-X. Standaert, and O. Pereira. The Swiss-Knife RFID Distance Bounding Protocol. In *Proceedings of the 11th International Conference on Information Security and Cryptology (ICISC)*, 2008.

[87] M. Kuhn, H. Luecken, and N. Tippenhauer. UWB Impulse Radio Based Distance Bounding. In *Proceedings of the Workshop on Positioning, Navigation and Communication (WPNC)*, 2010.

[88] L. Lazos and R. Poovendran. HiRLoc: High-resolution Robust Localization for Wireless Sensor Networks. *IEEE Journal on Selected Areas in Communications*, 24(2), 2006.

[89] L. Lazos and R. Poovendran. SeRLoc: Robust Localization for Wireless Sensor Networks. *ACM Transactions on Sensor Networks*, 1(1), 2005.

[90] L. Lazos, R. Poovendran, and S. Čapkun. ROPE: Robust Position Estimation in Wireless Sensor Networks. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks*, 2005.

[91] R. Liu and W. Trappe. *Securing Wireless Communications at the Physical Layer*. Springer Publishing Company, Incorporated, 1st edition, 2009.

[92] G. Lowe. Breaking and Fixing the Needham-Schroeder Public-Key Protocol Using FDR. In *Proceedings of the Second International Workshop on Tools and Algorithms for Construction and Analysis of Systems*, 1996.

[93] R. Maheshwari, J. Gao, and S. R. Das. Detecting Wormhole Attacks in Wireless Networks using Connectivity Information. In *Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM)*, 2007.

[94] S. Malladi, B. Bezawada, and K. Kothapalli. Automatic Analysis of Distance Bounding Protocols. In *Proceedings of the Workshop on Foundations of Computer Security*, 2009.

[95] C. Meadows, R. Poovendran, D. Pavlovic, L.-W. Chang, and P. Syverson. Distance Bounding Protocols: Authentication Logic Analysis and Collusion Attacks. In *Secure Localization and Time Synchronization for Wireless Sensor and Ad Hoc Networks*, 2007.

[96] C. Meadows, P. Syverson, and L. Chang. Towards More Efficient Distance Bounding Protocols for Use in Sensor Networks. In *Proceedings of the 2nd International Conference on Security and Privacy for Emerging Areas in Communications Networks (SecureComm)*, 2006.

[97] J. Millen and V. Shmatikov. Constraint Solving for Bounded-process Cryptographic Protocol Analysis. In *Proceedings of the 8th ACM Conference on Computer and Communications Security (CCS)*, 2001.

[98] A. Mitrokotsa, C. Dimitrakakis, P. Peris-Lopez, and J. C. Hernandez-Castro. Reid et al.'s Distance Bounding Protocol and Mafia Fraud Attacks over Noisy Channels. *IEEE Commmunication Letters*, 14(2), 2010.

[99] A.-F. Molisch, K. Balakrishnan, C.-C. Chong, S. Emami, A. Fort, J. Karedal, J. Kunisch, H. Schantz, U. Schuster, and K. Siwiak. IEEE 802.15.4a Channel Model - Final Report, Document 04/662r1, 2004.

[100] R. Mulloy. Ultrawide-band RFID Technology. FCC Radio Frequency Identification Workshop, 2004.

[101] J. Munilla and A. Peinado. Security Analysis of Tu and Piramuthu's Protocol. In *Proceedings of the 1st IFIP Conference on New Technologies, Mobility and Security (NTMS)*, 2008.

[102] J. Munilla and A. Peinado. Distance Bounding Protocols for RFID Enhanced by using Void-Challenges and Analysis in Noisy Channels. *Wireless Communications and Mobile Computing*, 8(9), 2008.

[103] J. Munilla and A. Peinado. Attacks on a Distance Bounding Protocol. *Computer Communications Journal*, 33(7), 2010.

[104] J. Munilla and A. Peinado. Enhanced low-cost RFID Protocol to detect relay attacks. *Wireless Communications and Mobile Computing Journal*, 10(3), 2010.

[105] F. Nait-Abdesselam, B. Bensaou, and T. Taleb. Detecting and Avoiding Wormhole Attacks in Wireless Ad Hoc Networks. *IEEE Communications Magazine*, 46(4), 2008.

[106] S. Nanz and C. Hankin. A Framework for Security Analysis of Mobile Wireless Networks. *Theoretical Computer Science*, 367(1), 2006.

[107] R. M. Needham and M. D. Schroeder. Using Encryption for Authentication in Large Networks of Computers. *Communications of the ACM*, 21(12), 1978.

[108] J. Newsome, E. Shi, D. Song, and A. Perrig. The Sybil Attack in Sensor Networks: Analysis & Defenses. In *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks (IPSN)*, 2004.

[109] V. Nikov and M. Vauclair. Yet Another Secure Distance-Bounding Protocol. In *Proceedings of the Conference on Security and Cryptography (SECRYPT)*, 2008.

[110] T. Nipkow, L. C. Paulson, and M. Wenzel. A Proof Assistant For Higher Order Logic, 2008.

[111] P. Papadimitratos, Z. Haas, and J.-P. Hubaux. How to Specify and How to Prove Correctness of Secure Routing Protocols for MANET. In *Proceedings of the 3rd IEEE-CS International Conference on BroadBand Communications, Networks, and Systems*, 2006.

[112] P. Papadimitratos and Z. J. Haas. Secure Routing for Mobile Ad Hoc Networks. In *Proceedings of the SCS Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS)*, 2002.

[113] P. Papadimitratos, M. Poturalski, P. Schaller, P. Lafourcade, D. Basin, S. Čapkun, and J.-P. Hubaux. Secure Neighborhood Discovery: A Fundamental Element for Mobile Ad Hoc Networking. *IEEE Communications Magazine*, 46(2), 2008.

[114] P. Papadimitratos and Z. J. Haas. Secure Message Transmission in Mobile Ad Hoc Networks. *Ad Hoc Networks Journal*, 1(1), 2003.

[115] S. Paquelet, L. M. Aubert, and B. Uguen. An Impulse Radio Asynchronous Transceiver for High Data Rates. In *Proceedings of the IEEE Joint Conference on Ultra Wideband Systems and Technologies & International Workshop on Ultra Wideband Systems*, 2004.

[116] B. Parno, A. Perrig, and V. Gligor. Distributed Detection of Node Replication Attacks in Sensor Networks. *Proceedings of the IEEE Symposium on Security and Privacy*, 2005.

[117] L. C. Paulson. The Inductive Approach to Verifying Cryptographic Protocols. *Journal of Computputer Security*, 6(1-2), 1998.

[118] D. Pavlovic and C. Meadows. Deriving Secrecy Properties in Key Establishment Protocols. In *Proceedings of the 11th European Symposium on Research in Computer Security (ESORICS)*, 2006.

[119] D. Pavlovic and C. Meadows. Quantifying Pervasive Authentication: the case of the Hancke-Kuhn Protocol. arXiv:0910.5745, 2009.

[120] D. Pavlovic and C. Meadows. Bayesian Authentication: Quantifying Security of the Hancke-Kuhn Protocol. In *Proceedings of the 26th Conference on the Mathematical Foundations of Programming Semantics (MFPS)*, 2010.

[121] P. Peris-Lopez, J. Hernandez-Castro, J. Tapiador, E. Palomar, and J. van der Lubbe. Cryptographic Puzzles and Distance-Bounding Protocols: Practical Tools for RFID Security. In *Proceedings of the IEEE International Conference on RFID*, 2010.

[122] P. Peris-Lopez, J. Hernández-Castro, J. M. Estévez-Tapiador, and J. van der Lubbe. Shedding Some Light on RFID Distance Bounding Protocols and Terrorist Attacks. *CoRR*, abs/0906.4618, 2009.

[123] C. Perkins and E. Royer. Ad-hoc On-demand Distance Vector Routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1999.

[124] T. V. Phuong, N. T. Canh, Y.-K. Lee, S. Lee, and H. Lee. Transmission Time-Based Mechanism to Detect Wormhole Attacks. *Proceedings of the IEEE Asia-Pacific Conference on Services Computing*, 2007.

[125] S. Piramuthu. Protocols for RFID Tag/Reader Authentication. *Decision Support Systems Journal*, 43(3), 2007.

[126] R. Poovendran and L. Lazos. A Graph Theoretic Framework for Preventing the Wormhole Attack in Wireless Ad Hoc Networks. *Wireless Networks Journal*, 13(1), 2007.

[127] M. Poturalski, M. Flury, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec. The Cicada Attack: Degradation and Denial of Service in IR Ranging. In *Proceedings of the IEEE International Conference on Ultra-Wideband (ICUWB)*, 2010.

[128] M. Poturalski, M. Flury, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec. Distance Bounding with IEEE 802.15.4a: Attacks and Countermeasures. *IEEE Transactions on Wireless Communications*, 10(4), 2011.

[129] M. Poturalski. Formal Analysis of Secure Neighbor Discovery (Isabelle/HOL Mechanization). http://infoscience.epfl.ch/record/166660.

[130] M. Poturalski, P. Papadimitratos, and J.-P. Hubaux. Secure Neighbor Discovery in Wireless Networks: Formal Investigation of Possibility. In *Proceedings of the 3rd ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2008.

[131] M. Poturalski, P. Papadimitratos, and J.-P. Hubaux. Towards Provable Secure Neighbor Discovery in Wireless Networks. In *Proceedings of the 6th ACM workshop on Formal Methods in Security Engineering*, 2008.

[132] J. G. Proakis and M. Salehi. *Digital Communications*. McGraw–Hill, 5th edition, 2008.

[133] L. Qian, N. Song, and X. Li. Detection of Wormhole Attacks in Multi-path Routed Wireless Ad Hoc Networks: A Statistical Analysis Approach. *Journal of Network and Computer Applications*, 30(1), 2007.

[134] K. B. Rasmussen and S. Čapkun. Implications of Radio Fingerprinting on the Security of Sensor Networks. *Proceedings of the 3rd IEEE Conference on Security and Privacy in Communication Networks (SecureComm)*, 2007.

[135] K. B. Rasmussen and S. Čapkun. Location privacy of Distance Bounding Protocols. In *15th ACM conference on Computer and Communications Security (CCS)*, 2008.

[136] K. B. Rasmussen and S. Čapkun. Realization of RF Distance Bounding. In *Proceedings of the USENIX Security Symposium*, 2010.

[137] J. Reid, J. M. G. Nieto, T. Tang, and B. Senadji. Detecting Relay Attacks with Timing-based Protocols. In *Proceedings of the 2nd ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2007.

[138] R. Robles, J. Haas, J. Chiang, Y.-C. Hu, and P. Kumar. Secure Topology Discovery through Network-wide Clock Synchronization. In *Proceedings of the 3rd International Conference on Signal Processing and Communications (SPCOM)*, 2010.

[139] H. P. Romero, K. A. Remley, D. F. Williams, and C. M. Wang. Electromagnetic Measurements for Counterfeit Detection of Radio Frequency Identification Cards. *IEEE Transactions on Microwave Theory and Techniques*, 57(5), 2009.

[140] Z. Sahinoglu and I. Guvenc. Multiuser Interference Mitigation in Noncoherent UWB Ranging via Nonlinear Filtering. *EURASIP Journal on Wireless Communications and Networking*, 2006.

[141] P. Schaller, B. Schmidt, D. Basin, and S. Čapkun. Modeling and Verifying Physical Properties of Security Protocols for Wireless Networks. In *Proceedings of the 2009 22nd IEEE Computer Security Foundations Symposium (CSF)*, 2009.

[142] R. Shokri, M. Poturalski, G. Ravot, P. Papadimitratos, and J.-P. Hubaux. A Practical Secure Neighbor Verification Protocol for Wireless Sensor Networks. In *Proceedings of the 2nd ACM Conference on Wireless Network Security (WiSec)*, 2009.

[143] D. Singelée and B. Preneel. Distance Bounding in Noisy Environments. In *Proceedings of the 4th European Workshop on Security and Privacy in Ad hoc and Sensor Networks (ESAS)*, 2007.

[144] D. Singelee and B. Preneel. Limitations on the Usage of Noise Resilient Distance Bounding Protocols. Technical report, COSIC, 2008.

[145] F. J. Thayer, J. C. Herzog, and J. D. Guttman. Strand Spaces: Why is a Security Protocol Correct? In *Proceedings of the IEEE Symposium on Security and Privacy*, 1998.

[146] F. J. Thayer, V. Swarup, and J. D. Guttman. Metric Strand Spaces for Locale Authentication Protocols. In *Proceedings of the 4th IFIP International Conference on Trust Management*, 2010.

[147] N. O. Tippenhauer, K. B. Rasmussen, C. Pöpper, and S. Capkun. Attacks on Public WLAN-based Positioning. In *Proceedings of the ACM/Usenix International Conference on Mobile Systems, Applications and Services (MobiSys)*, 2009.

[148] N. O. Tippenhauer, K. B. Rasmussen, and S. Čapkun. Secure Ranging With Message Temporal Integrity. In *Cryptology ePrint Archive: Report 2009/602*, 2009.

[149] N. O. Tippenhauer and S. Čapkun. ID-based Secure Distance Bounding and Localization. In *In Proceedings of 14th European Symposium on Research in Computer Security (ESORICS)*, 2009.

[150] F. Trösch and A. Wittneben. MLSE Post-Detection for ISI Mitigation and Synchronization in UWB Low Complexity Receivers. In *Proceedings of the 65th IEEE Conference on Vehicular Technology*, 2007.

[151] R. Trujillo Rasua, B. Martin, and G. Avoine. The Poulidor Distance-Bounding Protocol. In *Proceedings of the 6th Workshop on RFID Security (RFIDSec)*, 2010.

[152] Y.-J. Tu and S. Piramuthu. RFID Distance Bounding Protocols. In *Proceedings of the 1st International EURASIP Workshop on RFID Technology*, 2007.

[153] O. Ureten and N. Serinken. Detection of Radio Transmitter Turn-on Transients. *Electronics Letters*, 35(23), 1999.

[154] O. Ureten and N. Serinken. Wireless Security through RF Fingerprinting. *Canadian Journal of Electrical and Computer Engineering*, 32(1), 2007.

[155] S. Čapkun and J. Hubaux. Secure Positioning in Wireless Networks. *IEEE Journal on Selected Areas in Communications (JSAC)*, 24(2), 2006.

[156] S. Čapkun, M. Čagalj, and M. Srivastava. Securing Localization With Hidden and Mobile Base Stations. In *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM)*, 2006.

[157] S. Čapkun, L. Buttyán, and J.-P. Hubaux. SECTOR: Secure Tracking of Node Encounters in Multi-hop Wireless Networks. In *Proceedings of the 1st ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, 2003.

[158] W. Wang and B. Bhargava. Visualization of Wormholes in Sensor Networks. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, 2004.

[159] W. Wang, B. Bhargava, Y. Lu, and X. Wu. Defending against Wormhole Attacks in Mobile Ad Hoc Networks. *Wireless Communications and Mobile Computing Journal*, 6(4), 2006.

[160] W. Wang and A. Lu. Interactive Wormhole Detection and Evaluation. *Information Visualization Journal*, 6(1), 2007.

[161] X. Wang and J. Wong. An End-to-end Detection of Wormhole Attack in Wireless Adhoc Networks. In *Proceedings of the 31st Annual International Computer Software and Applications Conference*, 2007.

[162] T. Waraksa, K. Fraley, D. Douglas, and L. Gilbert. Pasive Keyless Entry System. US Patent 4942393, 1990.

[163] J. Wu, H. Chen, W. Lou, Z. Wang, and Z. Wang. Label-Based DV-Hop Localization Against Wormhole Attacks in Wireless Sensor Networks. *Proceedings of the 5th IEEE International Conference on Networking, Architecture, and Storage*, 2010.

[164] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe. A Physical-Layer Technique to Enhance Authentication for Mobile Terminals. In *Proceedings of the IEEE International Conference on Communications (ICC)*, 2008.

[165] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe. Using the Physical Layer for Wireless Authentication in Time-variant Channels. *IEEE Transactions on Wireless Communications*, 7(7), 2008.

[166] L. Xiao, L. J. Greenstein, N. B. Mandayam, and W. Trappe. Fingerprints in the Ether: Using the Physical Layer for Wireless Authentication. In *Proceedings of the IEEE International Conference on Communications (ICC)*, 2007.

[167] Y. Xu, G. Chen, J. Ford, and F. Makedon. Detecting Wormhole Attacks in Wireless Sensor Networks. In *Critical Infrastructure Protection*, volume 253 of *IFIP International Federation for Information Processing*, 2007.

[168] S. Yang and J. S. Baras. Modeling vulnerabilities of Ad Hoc routing protocols. In *Proceedings of the 1st ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, 2003.

[169] D. H. Yum, J. S. Kim, S. J. Hong, and P. J. Lee. Distance Bounding Protocol for Mutual Authentication. *IEEE Transactions on Wireless Communications*, 10(2), 2011.

[170] D. Zanetti, B. Danev, and S. Čapkun. Physical-layer Identification of UHF RFID Tags. *Proceedings of the 16th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2010.

[171] R. Zhang and Y. Zhang. Wormhole-resilient Secure Neighbor Discovery in Underwater Acoustic Networks. In *Proceedings of the 29th IEEE International Conference on Information Communications (INFOCOM)*, 2010.

[172] Y. Zhang, W. Liu, Y. Fang, and D. Wu. Secure Localization and Authentication in Ultra-wideband Sensor Networks. *IEEE Journal on Selected Areas in Communications*, 24(4), 2006.

[173] Y. Zhang, W. Liu, W. Lou, and Y. Fang. Location-based Compromise-tolerant Security Mechanisms for Wireless Sensor Networks. *IEEE Journal on Selected Areas in Communications*, 24(2), 2006.

[174] J. Zhen and S. Srinivas. Preventing Replay Attacks for Secure Routing in Ad Hoc Networks. In *Proceedings of the 2nd Conference on Ad-Hoc, Mobile, and Wireless Networks*, 2003.

[175] W. Znaidi, M. Minier, and J.-P. Babau. Detecting Wormhole Attacks in Wireless Networks using Local Neighborhood Information. In *Proceedings of the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2008.

Index

Marcin Poturalski

EPFL-IC-LCA
Station 14
CH-1015 Lausanne, Switzerland

marcin.poturalski@epfl.ch
+41 21 693 4657

Personal

Born on April 1st, 1981 in Chojnice, Poland. Citizen of Poland.
Languages: Polish (native), English (fluent), French (basic), German (basic)

Education

Ph.D. in Communication Systems, EPFL, *Sep. 2005 – present*
Thesis: "Secure Neighbor Discovery and Ranging in Wireless Networks"
Thesis directors: Prof. Jean-Pierre Hubaux and Prof. Panagiotis Papadimitratos

M.Sc. in Computer Science, Warsaw University, *Oct. 2000 – Jun. 2005*
Thesis: "Observational Equivalence of Timed Processes"
Advisor: Dr Sławomir Lasota

B.Sc. in Mathematics, Warsaw University, *Oct. 2000 – Jan. 2005*
Thesis: "Orientability of Surfaces"
Advisor: Prof. Stanisław Betley

Professional Experience

Research and Teaching Assistant *Sep. 2006 – present*
Laboratory of Computer Communications and Applications, EPFL

Software Developer (Internship) *Jun. 2005 – Aug. 2005*
User Interface Strategy, Microsoft, Redmond, VA

Software Developer (Internship) *Jun. 2004 – Aug. 2004*
Asia Center for Interaction Design, MSRA, Beijing, China

Professional Activities

Reviewer for scientific journals and conferences
IEEE WCM, IEEE TMC, IEEE TIFS, IEEE TDSC, IEEE TPDS, Computer Communications, Wireless Networks, Journal of Systems and Software, ACM CCS, ACM SigComm, ACM WiSec, ACM MobiHoc, WiOpt, IEEE WoWMoM, Eurocrypt

Teaching

Teaching Assistant
Computer Networks, Mobile Networks, Security and Cooperation in Wireless Networks

Supervised Master Theses
Yannick Do "Ultra-Wide Band Ranging Security: Distance-decreasing Attacks & Countermeasures" (co-supervised with Manuel Flury)
Michael Jubin "Hide & Seek: Security of Location Based Services" (co-supervised with Nevena Vratonjic)

Publications

M. Poturalski, M. Flury, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec. **Distance Bounding with IEEE 802.15.4a: Attacks and Countermeasures.** *IEEE Transactions on Wireless Communications*, 2011.

M. Poturalski, M. Flury, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec. **The Cicada Attack: Degradation and Denial of Service IR Ranging.** *IEEE International Conference on Ultra-Wideband (ICUWB)*, 2010.

M. Flury, M. Poturalski, P. Papadimitratos, J.-P. Hubaux, and J.-Y. Le Boudec. **Effectiveness of Distance-Decreasing Attacks Against Impulse Radio Ranging.** *ACM Conference on Wireless Network Security (WiSec)*, 2010.

W. Galuba, P. Papadimitratos, M. Poturalski, K. Aberer, Z. Despotovic, and W. Kellerer. **Castor: Scalable Secure Routing for Ad Hoc Networks.** *IEEE International Conference on Computer Communications (INFOCOM)*, 2010.

R. Shokri, M. Poturalski, G. Ravot, P. Papadimitratos, and J.-P. Hubaux. **A Practical Secure Neighbor Verification Protocol for Wireless Sensor Networks.** *ACM* Conference on Wireless Network Security (WiSec), 2009.

M. Poturalski, P. Papadimitratos, and J.-P. Hubaux. **Towards Provable Secure Neighbor Discovery Wireless Networks.** *ACM Workshop on Formal Methods in Security Engineering*, 2008.

M. Poturalski, P. Papadimitratos, and J.-P. Hubaux. **Secure Neighbor Discovery Wireless Networks: Formal Investigation of Possibility.** *ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2008.

P. Papadimitratos, M. Poturalski, P. Schaller, P. Lafourcade, D. Basin, S. Čapkun, and J.-P. Hubaux. **Secure Neighborhood Discovery: A Fundamental Element for Mobile Ad Hoc Networking.** *IEEE Communications Magazine*, 2008.

P. Haghani, P. Papadimitratos, M. Poturalski, K. Aberer, and J.-P. Hubaux. **Efficient and Robust Secure Aggregation for Sensor Networks.** *Workshop on Secure Network Protocols (NPSec)*, 2007.

S. Wang, M. Poturalski, D. Vronay. **Designing a Generalized 3D Carousel View**. *Conference on Human Factors in Computing Systems (CHI)*, 2005