PURE: Payments with UWB REIay-protection

Conference Paper

Author(s):
Coppola, Daniele; Camurati, Giovanni; Anliker, Claudio; Hofmeier, Xenia; Shaller, Patrick; Basin, David; Capkun, Srdjan

Publication date:
2024

Permanent link:
https://doi.org/10.3929/ethz-b-000662474

PURE: Payments with UWB RElay-protection

Daniele Coppola
ETH Zurich

Giovanni Camurati
ETH Zurich

Claudio Anliker
ETH Zurich

Xenia Hofmeier
ETH Zurich

Patrick Schaller
ETH Zurich

David Basin
ETH Zurich

Srdjan Capkun
ETH Zurich

Abstract

Contactless payments are now widely used and are expected to reach $10 trillion worth of transactions by 2027. Although convenient, contactless payments are vulnerable to relay attacks that enable attackers to execute fraudulent payments. A number of countermeasures have been proposed to address this issue, including Mastercard's relay protection mechanism. These countermeasures, although effective against some Commercial off-the-shelf (COTS) relays, fail to prevent physical-layer relay attacks.

In this work, we leverage the Ultra-Wide Band (UWB) radios incorporated in major smartphones, smartwatches, tags and accessories, and introduce PURE, the first UWB-based relay protection that integrates smoothly into existing contactless payment standards, and prevents even the most sophisticated physical layer attacks. PURE extends EMV payment protocols that are executed between cards and terminals, and does not require any modification to the backend of the issuer, acquirer, or payment network. PURE further tailors UWB ranging to the payment environment (i.e., wireless channels) to achieve both reliability and resistance to all known physical-layer distance reduction attacks against UWB 802.15.4z. We implement PURE within the EMV standard on modern smartphones, and evaluate its performance in a realistic deployment. Our experiments show that PURE provides a sub-meter relay protection with minimal execution overhead (41 ms). We formally verify the security of PURE's integration within Mastercard's EMV protocol using the Tamarin prover.

1 Introduction

Contactless payments have greatly improved and streamlined the digital payment experience. Given their convenience, their number is expected to grow by 221% between 2022 and 2026, reaching $10 trillion worth of transactions by 2027 [30]. The preference for contactless payment and withdrawal methods was further accelerated by the need to minimize contact with payment terminals during the COVID pandemic [45].

With the emergence of contactless payments, the awareness and the threat of relay attacks on card payments have also increased [2]; Drimer et al. [47] demonstrated already in 2007 that such attacks were not only possible, but could be accomplished for as little as $500 worth of hardware and with just moderate engineering skills. Since then, the cost of relay attacks has further declined [9].

In the simplest relay attack, the victim does not intend to use their card (e.g., it is in their wallet) but an attacker establishes a relay channel between the victim's card and a terminal of the attacker's choice, thereby charging the victim's card for the purchase. In some countries and for some banks, the damage of such attacks is limited by the contactless payment limit, although not all US banks impose limits [24].

In more sophisticated relay attacks, similar to the chip-and-PIN fraud [13, 17, 47], an attacker plays the role of a (malicious) merchant who presents a fake terminal to the victim customer. By relaying the communication, and even the victim's PIN, the attacker can use the victim's card to pay for more valuable goods at a legitimate terminal or to withdraw a large sum of money from an Automatic Teller Machine (ATM). Mobile payments are especially vulnerable to this kind of fraud because unlocking the phone is generally accepted as a Cardholder Verification Method and thus allows the attacker to arbitrarily increase the amount, e.g., up to the card's daily or monthly limit. With the advent of on-device (smartphone) terminals, like Apple's tap-to-pay [29], the attacker's cost and challenge of creating a fake terminal have been significantly reduced.

Hence, relay attacks on contactless payments, irrespective of whether they are executed against cards, phones, watches or other wearables, remain a concern, and may even increase in the future both in frequency and severity.

There are two main ways of addressing relay attacks. The first involves reducing the usability of payments by requiring additional checks or confirmations by the user on the phone or watch, which is not possible in case of cards or tokens. The second is to prevent relays by technical means using secure distance measurement between the card and the legitimate

terminal. Since usability is one of the main benefits and features of contactless payments, defenses focus on the technical prevention of relay attacks.

The most established technical standard for smart cards and contactless payments is the EMV standard. This standard describes different implementations, called kernels, by which cards and terminals or ATMs can interact. EMVco, the consortium maintaining the EMV standard, has recognized the severity of relay attacks on contactless payment systems and has proposed a relay protection extension for the Mastercard kernel known as the Relay Resistance Protocol (RRP). This proposal has been further improved in [43]. These relay protections, although effective against adversaries equipped with COTS hardware, fail to prevent physical-layer relay attacks. Namely, these protections tolerate relays of μs to $10ms$, allowing physical layer relays over many kilometers.

The recent introduction of UWB ranging radios in many personal devices opens up the possibility of building a more effective relay protection that is not restricted to Near Field Communication (NFC). In particular, smartphones, smart watches, and tags from Apple, Samsung, and Google are now equipped with UWB, allowing sub-decimeter ranging accuracy between devices.

Given this, it is natural to consider using UWB radios to protect payments against relay attacks. However, ensuring that this results in a secure, usable, and deployable relay protection mechanism requires tailoring UWB ranging to contactless payments as well as careful enhancements to the EMV protocols.

In this work, we address this challenge and introduce PURE, a set of relay protection mechanisms that extend EMV contactless payments in a secure, usable, and easy-to-deploy manner. We make the following contributions:

- We propose PURE, the integration of (UWB) physical-layer mechanisms into the EMV kernel, which are the first to provide protection against all relay attacks, including those on the physical layer, for contactless payments.
- We formally verify the integration of PURE mechanisms within EMV protocols in Tamarin [39].
- We implement a prototype of PURE on modern smartphones and evaluate it in realistic deployments.
- We show that PURE provides sub-meter (ca. 50 cm) relay protection, adds very low (41 ms) latency to the payments ($\approx$ 630 ms), and has no noticeable impact on the usability of payments.
- We demonstrate that PURE provides protection against all known distance reduction attacks on UWB IEEE 802.15.4z High Pulse Rate (HRP) secure ranging systems, which are present in modern devices.
- We design PURE so that it can be easily integrated into the existing EMV ecosystem and requires only changes to the payment terminal / device, but no changes to the payment backends. For terminals already supporting UWB (such as an iPhone with Apple's Tap to Pay [29]), only software changes are required.

2 Overview

Contactless payments use NFC to trigger backend financial transfers from the cardholder accounts to accounts of corresponding terminal owners. Having established an NFC connection, the card and terminal negotiate a protocol to initiate the payment and to securely exchange the necessary parameters to trigger the backend financial transfer, following one of the EMVCo Kernels. Currently, there are seven different Kernels for different payment networks, such as Mastercard or Visa. As a result of a successful protocol execution, the card provides a cryptographic token to the terminal that proves to the card's issuer that the card indeed has been involved in a payment at the given terminal.

The proximity of the cardholder to the payment terminal effectively acts as a transaction confirmation. Relay attacks can compromise these payments by relaying communication without compromising payment cards. These include simple attacks that relay communication from the unsuspecting user's card to a payment terminal, and more sophisticated attacks that involve the use of fake terminals and can steal funds from the user up to their daily card limit.

The goal of relay protection mechanisms is to prevent such attacks and ensure that a payment can be successfully executed only if the two devices, the card, and the legitimate terminal, are physically close. In particular, the maximum accepted relay must be sufficiently small ($\approx 1m$) so that relays are prevented simply by the spatial awareness of merchants and cardholders.

PURE achieves this by using UWB ranging, which provides decimeter ranging accuracy. Therefore PURE specifically targets on-device cards (e.g. smartphones) on which UWB radios are already widely deployed [23]. Leveraging UWB ranging for relay prevention in the context of contactless payments is, however, not straightforward. PURE integrates UWB into EMV payment protocols such that it (i) does not noticeably increase the latency of payments, (ii) does not reduce the reliability or convenience of payments, (iii) integrates securely with EMV protocols, and (iv) is easy to deploy within the broader EMV payment network ecosystem (see Figure 1). Finally, UWB radios deployed in most devices implement 802.15.4z HRP, which was shown to be vulnerable to distance decreasing attacks [36]. PURE also addresses these attacks (v).

In what follows, we first specify the attacker model, and then provide an overview of the changes that PURE makes to the EMV protocols and to UWB ranging to achieve (i-v).

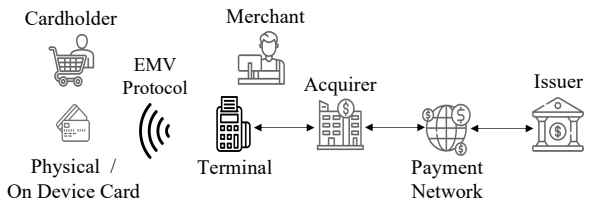

Figure 1: **The Payment Ecosystem.** The stakeholders involved into the payment process.

Attacker Model. The attacker's main goal is to establish an unintended connection between a legitimate card and an attacker-chosen terminal, without triggering any relay protection.

We assume a Dolev-Yao-style network attacker that fully controls the communication between the card and the terminal, both on the NFC and UWB channels. The attacker can eavesdrop, intercept, modify, and inject signals and messages, as well as initiate transactions between the card and the terminal. We further assume that all legitimate parties (i.e., issuer, payment network, acquirer, merchant, registered terminals, card, cardholder) are honest and that all communication channels except the ones between the terminal and the card are secure. We exclude attacker-owned, registered terminals from our threat model. Since the flow of money can be tracked, a malicious owner of a registered terminal would eventually be convicted of fraud charges pressed by cardholders. Further note that attackers who own a registered terminal need not relay communication, but could simply place the terminal close to the victim's card to initiate a payment.

We consider a strong relay attacker that can operate on the physical layer and introduces only negligible processing delays. Francillon et al. demonstrated in [25] that a similar relay can be implemented in practice with simple analog equipment. The only effect introduced by such a relay is a delay proportional to its length divided by the speed of light. The attacker is also capable of all state-of-the-art attacks on round-trip time or distance measurements, including those at the physical layer.

PURE: Protocol Overview. To establish a reliable upper bound on the distance between an honest terminal and a card (and therefore to prevent relay attacks), UWB ranging requires the involved parties (here the payment card and the legitimate terminal) to share a secret key. For this purpose, we have designed and implemented an extension of the existing EMV Kernel that builds on EMV's Public Key Infrastructure (PKI) and allows the card and terminal to establish a shared key. By extending existing (NFC) EMV messages, PURE introduces little overhead to the EMV protocol. Figure 2 illustrates PURE's key exchange extension to the original EMV protocol

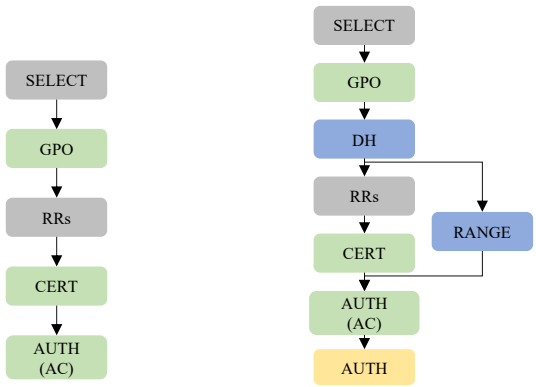

(a) EMV protocol. (b) Extended EMV protocol.

Figure 2: **EMV and PURE.** Gray boxes denote elements of the original standard not affected by PURE. Elements of the original standard used by PURE are colored green. Blue boxes are required and newly added by PURE. The yellow block is required by PURE, but can be smoothly integrated into the existing authentication step of the standard.

and the added UWB ranging phase. Note that PURE makes no changes to the protocols and messages to the payment backend; it only modifies the card-terminal interaction defined by the EMV standard.

We introduce the existing Mastercard kernel together with our extension in more detail in Section 3. The card and terminal negotiate the use of PURE according to their capabilities in an early stage of their NFC-based communication. If either of the two does not support UWB, they fall back and execute the existing protocol. The card's support of PURE is announced as part of an integrity-protected data-element sent by the card. As a consequence, an attacker cannot downgrade PURE-capable devices to omit the relay-protection phase. If the card and terminal are UWB capable and support PURE, they execute a Diffie-Hellman (DH) exchange. At this point, the key shares are not yet authenticated, but the shared secret can already be used for ranging, while the rest of the transaction continues in parallel over NFC. Running the ranging process in parallel with the payment nullifies the additional overhead introduced by the ranging. This ultimately leads to a fast execution of the combined protocol that does not introduce a significant delay for the user to complete a transaction.

As a result of the ranging extension, the terminal derives a reliable upper bound on the distance to the entity it shares the secret with. As part of the authentication/authorization step of the original protocol, the card authenticates to the terminal critical transaction parameters and the DH values. This binds the previously executed ranging to the specific transaction and card. The terminal blocks the transaction if the cryptographic verification fails (possible Machine-in-The-Middle (MITM) attack) or if the measured distance is too large (due to a relay attack).

	Uncertainty Δt	App Relay	Phy Relay
RRP [20]	$\approx$ ms	$\approx km$	$\approx$ 100 km
L1RP [43]	$\approx us$	0	$\approx$ km
This paper	$\approx ns$	0	50 cm

Table 1: **Relay Protections on Contactless Payments.** The larger the measurement uncertainty Δt on the Round-Trip Time (RTT), the larger the error on distance, and the longer the relays that remain undetected. A physical layer relay with negligible delay achieves the longest relay, while an application layer relay has the disadvantage of introducing an additional processing time (e.g., $T_{comp} = 1ms$). Using highly-accurate UWB measurements PURE provides excellent protection from both, and is also secure from distance-reduction attacks.

Comparison to Existing Relay Protections. By specifically targeting on-device cards, PURE provides significantly better guarantees than existing relay protection mechanisms. A comparison can be found in Table 1.

In broad terms, both PURE and previously proposed relay defenses measure the Round-Trip Time (RTT) between the card and the terminal with a nonce exchange. The exchanged nonces are unpredictable and authentic to prevent potential attackers from interfering with the RTT measurement. Based on the signal's Time of Arrival (ToA) measurement and the reply time, the Time of Flight (ToF) and, hence, the mutual distance can be computed. The protocol's reliability and security heavily depend on the accuracy of the ToA measurement. Indeed, measuring an earlier ToA results in a shorter measured distance, which could make a relay undetectable, while a later ToA leads to an overestimated distance measurement, potentially compromising system reliability. Intuitively, even a small uncertainty in the measurement of the RTT (e.g., μs) leads to large errors in the measurement of distance (e.g., km), leaving even long relays undetected (see Appendix B). Existing solutions have low accuracy and tolerate long relays.

The Relay Resistance Protocol (RRP) proposed by Mastercard [20] implements a timed nonce exchange at layer 3 of NFC, achieving milliseconds accuracy. Radu et al. [43] in their Level 1 Relay Protection (L1RP) moved the nonce exchange to the physical layer of NFC, which improved the accuracy to microseconds. Both L1RP and RRP can prevent application-layer relay attacks, which introduce substantial delays due to computation time; however (as recognized by both Mastercard and the authors), it would not stop specialist, expert-built relay/MitM attacks that relay signals on the physical layer [25]. The root cause of vulnerabilities for both solutions is the same: long NFC symbols result in inaccurate ToF measurements and are susceptible to Early-Detect and Late-Commit (ED/LC) attacks [34, 40]. Both solutions also rely on an estimate of the reply time for each card model, which introduces additional errors.

UWB Ranging in PURE. PURE does not simply integrate off-the-shelf UWB ranging into EMV protocols. Widely deployed UWB (IEEE 802.15.4z HRP) chips were found to be vulnerable to physical layer attacks such as Ghost Peak attack [36]. This attack stems from a security-reliability trade-off typical in cyber-physical systems: if the link budget is low, a ranging system will typically sacrifice security in favor of reliability.

With PURE, we show that the short distance and wireless channels between card and terminal during payments are particularly favorable for secure UWB ranging. We devise a simple Leading Edge Detection (LED) algorithm on top of 802.15.4z HRP that maintains the reliability of contactless payments while being secure against known distance reduction attacks. More precisely, we show that PURE limits the success probability of the Ghost Peak attack to 2^{-48}, a suitable security level for payment applications [36].

3 EMV and PURE

We focus on the integration of PURE in the Mastercard kernel, known as the C-2 kernel [20]. Since all kernels rely on the same cryptographic primitives (asymmetric algorithms for offline authentication and symmetric algorithms for online authorization), similar extensions are possible for the other kernels. Moreover, the latest proposal for a new, unifying contactless standard, known as the C-8 kernel [21], includes the central elements needed by PURE.

3.1 The Mastercard Contactless Standard

Conceptually, a contactless payment consists of four stages:

- *Application Selection.* The card and the terminal agree on a kernel that specifies the capabilities of the card and the protocol executed by them.
- *Synchronization.* The card provides a list of data elements required by the issuer to process the transaction, such as the amount and currency. Furthermore, the terminal learns a set of certificates that allows it to validate the card's public key using a hard-coded list of root certificates. Critical data elements, such as the supported *Cardholder Verification Method*s (CVMs), are contained as data elements in the card's certificate and are thus integrity-protected.
- *Cardholder Verification.* Given the CVMs learned in the synchronization step, the terminal verifies the cardholder. For physical cards, this includes PIN or signature verification. For on-device cards, the terminal typically relies on the device's identification methods, e.g., the mobile phone's PIN or facial recognition.

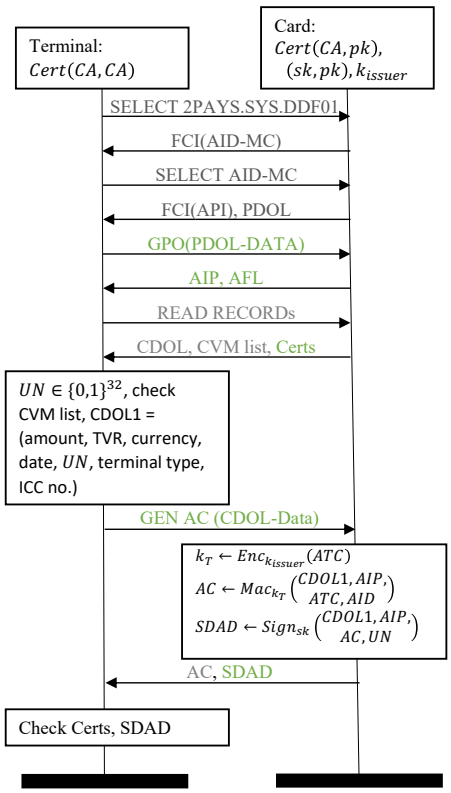

Figure 3: **Mastercard Kernel.** In green are the messages that PURE piggybacks onto. Messages in gray are specific to EMV and are not used by PURE.

- *Authentication/Authorization.* The terminal completes integrity checks of transaction critical data exchanged with the card based on digital signatures, so-called offline data authentication. In addition, the terminal sends transaction-related data (including a Message Authentication Code (MAC) provided by the card) to the issuer for authorization. All involved stakeholders (acquirer, payment network, and issuer) use fraud-detection systems to prevent fraudulent transactions.

The communication between the card and the terminal is depicted in Figure 3. Whereas the application selection and synchronization steps of a contactless payment are integral parts of the protocol executed between the card and the terminal, cardholder verification as well as parts of the authentication/authorization step involve communication between the cardholder and the terminal (or mobile phone) as well as communication between the terminal and back-end services (acquirer, payment-network, or issuer) and are therefore not shown in the figure.

In Figure 3, the terminal starts the communication by requesting the so-called Dictionary Definition File (`2PAYS.SYS.DDF01`), a Payment System Directory that contains one or more *Application Definition Files* (ADFs). The card responds with the *File Control Information* (FCI) containing an *Application Identifier* (AID) that allows the terminal to select the kernel to execute (in the figure depicted as `AID-MC`, where `MC` stands for Mastercard).

The issuer defines transaction-relevant data elements in so-called data object lists (DOLs) on the card. This includes a *Processing Options Data Object List* (PDOL) of terminal resident data elements and the *Card Risk Management Data Object List 1* (CDOL1) that contains transaction-relevant data elements that are later integrity protected using cryptographic algorithms.

The terminal then initiates the transaction by sending a *GET PROCESSING OPTIONS* (GPO) command containing the PDOLs requested by the card in the previous step. The card's response includes the *Application Interchange Profile* (AIP), a byte vector encoding the card's capabilities and an *Application File Locator* (AFL). The AFL enables the terminal to access the CDOL, the CVM list, and the card's certificate using *Read Record* (RR) commands.

In a next step, the terminal creates an Unguessable Number (UN) for the transaction and sends the required CDOL-data as part of the generate *Application Cryptogram* (AC) command, thereby defining the format of the cryptogram required to complete the transaction. The most common and secure form, Combined Data Authentication (CDA), requires successful offline and online authentication/authorization. The AC included in this form is a MAC on transaction-critical data elements that uses a secret key shared between the issuer and the card. Furthermore, it includes the *Signed Dynamic Application Data* (SDAD), a signature on transaction-critical data elements signed with the card's private key that can be verified by the terminal.

If the transaction's transcript passes all cryptographic tests, the cardholder (if necessary) is successfully authenticated, and the transaction does not trigger other security alerts (such as fraud detection system alerts in the payment network or at the issuer), then the payment is approved and triggers a corresponding back-end transaction at the issuer.

3.2 PURE Stand-Alone

We start by describing the stand-alone version of PURE before showing how to smoothly integrate it into the existing Mastercard kernel in a way that ensures backward compatibility with the current version of the standard. PURE comprises two main components: a key exchange and a ranging session. After a successful key exchange, the card and the reader share a secret key. This shared key is then used for ranging, where the reader accurately measures an upper bound on the distance to the card. Note that the term *card* refers to any token able to execute a contactless payment, however, PURE can only be applied to on-device cards because physical cards do not have a UWB radio.

PURE uses a Sign-and-Mac (SIGMA) key exchange [33]

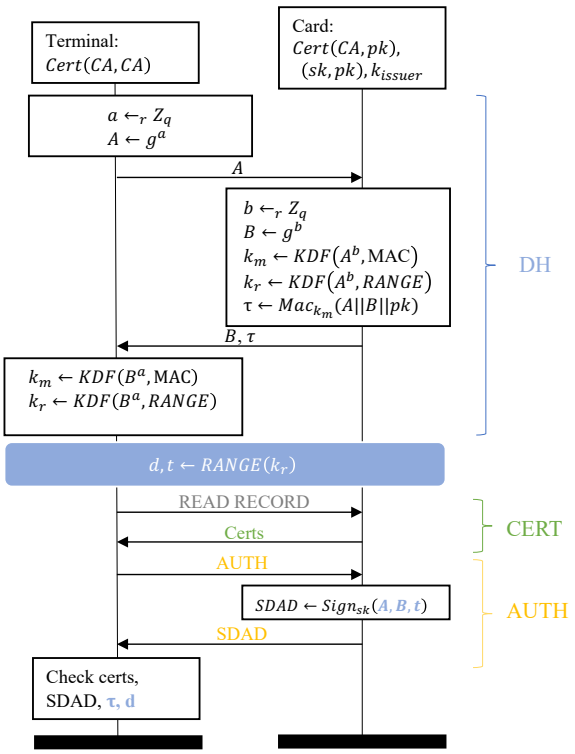

Figure 4: **Stand-Alone PURE.** Messages labeled in green and yellow (CERT and AUTH) can be integrated into data elements of the current Mastercard kernel. The blue elements (DH and RANGE) require specific extensions of the Mastercard kernel.

together with a ranging session to establish a reliable upper bound on the distance between the honest card and the honest terminal. Figure 4 depicts the protocol as a sequence diagram describing messages exchanged between the card and the terminal.

For clarity, we split the protocol into four stages:

- *DH Key Exchange (DH).* The two parties exchange freshly generated DH shares and derive fresh ranging (k_r) and tagging (k_m) keys. Upon reception of the terminal's share, the card computes τ, a MAC over the concatenation of the DH shares and its public key using k_m and appends it to its own DH share.
- *Ranging (RANGE).* Both parties execute UWB ranging using k_r for generating the nonces.
- *Certificate Exchange (CERT).* The card provides a certificate for its public key pk signed by the card's issuer that can be validated by a certificate authority CA known to the terminal.
- *Authentication (AUTH).* The terminal requests a signature over the concatenation of the DH parameters and the timings t. The protocol terminates with the reader checking the signature $SDAD$, the certificate $certs$, and the MAC τ. If any of the cryptographic checks fail, the terminal aborts the protocol.

As a result of a successful protocol execution, the terminal derives a reliable upper bound on the distance to the card with which it is executing the protocol. Distances below a predefined threshold convince the terminal that communication with the card has not been relayed. In contrast, distances above the threshold are considered suspicious and suggest a possible relay attack.

Note that our threat model assumes the terminal completing the transaction with a given acquirer and the card issuer to be honest, thus providing guarantees to the cardholder, since honest terminals would decline transactions that are likely to be relay attacks.

3.3 Integration in the Mastercard Kernel

We present our integration of PURE in the C-2 (Mastercard) kernel. Given that all kernels use a public key infrastructure for offline authentication, and that the AIP is integrity-protected, similar extensions of other kernels are straightforward.

The C-2 kernel integrated with PURE is illustrated in Figure 5. The sequence diagram shows how to integrate the elements of the stand-alone version (compare with Figure 4) into the existing Mastercard standard.

The protocol remains unchanged until the card announces its processing capabilities to the terminal with the AIP bit-vector. Currently, this bit-vector includes three bits reserved for future use (RFU). Our extension uses one of these bits to indicate the card's secure ranging capability. Note that a malicious downgrade is prevented because the AIP's value is integrity protected as a static element by the card's certificate. Adversarial changes of the AIP are thus detected by the terminal resulting in the protocol's abortion. Note also that the terminal initiating the transaction at the issuer is assumed to be honest. Provided that both the card and terminal support UWB relay protection, the terminal executes the DH block in Figure 5. In case UWB relay protection is not supported by either the terminal or the card, the terminal follows the existing C-2 standard.

The next deviation from the original EMV standard is introduced in the AUTH section, where the content of the modified SDAD element is extended with the Diffie-Hellman parameters of the card and the terminal. We execute the DH early in the protocol to ensure that double-sided two-way ranging (DS-TWR) can be run in parallel to NFC communication between card and terminal, thereby nullifying any additional delay possibly caused by the UWB ranging step.

Note that the CERT phase uses the existing certificate exchange of the C-2 kernel and integrates smoothly into the original data elements of the standard.

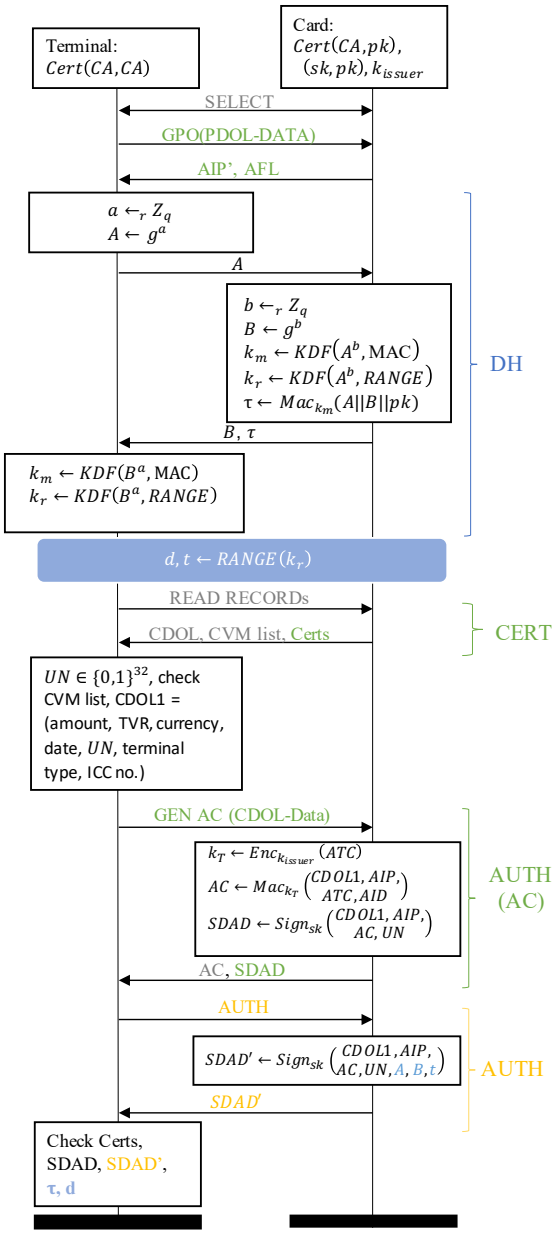

Figure 5: **Mastercard Kernel Extended.** This diagram integrates the stand alone protocol from Figure 4 into the Mastercard kernel in Figure 3. The color coding is consistent with the previous images, see Figure 2.

We emphasize that cryptographic linking of the values from the extension with those identifying the transaction (such as CDOL1, UN, and AIP) is essential. This way, we ensure that the shared secret and the corresponding ranging session are bound to a specific transaction. As a consequence, the UWB and NFC phases of a transaction cannot be decoupled by the attacker. In addition to the cryptographic checks that the terminal makes, depicted in Figure 5, the terminal checks that the measured distance is lower than a given threshold.

Note too that PURE requires offline authentication to prevent a MITM attack within the DH exchange. In theory, protection against MITM attacks, when offline authentication fails, could be achieved in the online authorization phase, where the issuer could validate the correctness of DH parameters in the message authentication code provided by the card and the DH parameters sent by the terminal. This approach would however require changes in the backend logic of the payment process outside of the EMV standard.

Integration in the C-8 Kernel. As explained in the preceding section, the smooth integration of UWB into Mastercard is achieved with three additional messages to establish the DH-based shared secret; all other messages are already in the C-2 kernel. In October 2022, EMVco published a proposal for the new C-8 kernel. This kernel is intended to unify the existing kernels and will become the new standard for future contactless transactions. Without delving into C-8's details, we point out that it includes a key exchange between the card and the terminal to establish an ephemeral Diffie-Hellman key, known as "Blinded Diffie-Hellman." The purpose of the shared secret is to encrypt the messages exchanged between card and terminal to protect sensitive data such as the Primary Account Number (PAN). Similarly to the key exchange we proposed, the reader and card establish a shared secret and the reader authenticates the card.

The established shared secret could be used to derive a ranging key. Hence, extending the C-8 kernel with UWB ranging requires no additional NFC messages. Given that UWB ranging can execute in parallel with the NFC message exchange, the overhead introduced for the C-8 kernel is again negligible.

Formal Model and Security Proofs in Tamarin. To increase the trustworthiness of PURE and formally prove the claimed security properties, we have modeled the extension of the C-2 kernel in the Tamarin prover [39]. Tamarin is a powerful verification tool that allows one to model and reason about security properties of protocols. Tamarin supports, in particular, equational specifications of Diffie-Hellman exponentiation, used by PURE. When using Tamarin, the central goal is either to obtain a formal proof for the security properties stated for the protocol or to find counterexamples, which are attacks that violate the stated properties. Tamarin provides fully automated strategies for constructing proofs or counterexamples but requires manual interaction in complicated cases.

Our model[1] is based on the formal model of the C-2 kernel provided from [7]. We have extended this model with the additional SIGMA key exchange as well as the corresponding

[1] https://github.com/daniCoppola/pure-models

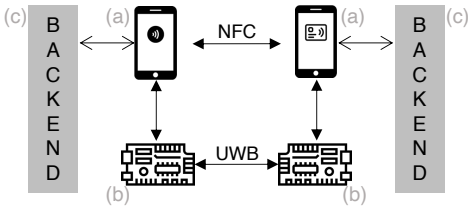

Figure 6: **Proof-of-Concept Setup.** (a) Android phones; (b) UWB chip; (c) backend generating EMV messages.

extension to offline authentication. Although the confidentiality and authenticity of the shared key inherently result from the security of the well-studied SIGMA key exchange [33], our model further establishes the secure integration of SIGMA into the C-2 kernel. Since Tamarin uses a symbolic message model that abstracts away from physical-layer properties fundamental to the security of a ranging scheme [37], our formal model of the protocol extension does not include the ranging phase. We analyze ranging separately in Section 4, given the properties of the derived shared secret proven with Tamarin.

In addition to the security proofs of the original protocol, the following theorems have been proven with Tamarin:

- Authenticity of the DH share: Whenever the terminal concludes that it has received a key share from a card, the card has indeed provided this key share in the current run of the protocol. Actually, we prove a strong version of authenticity, known in the literature as injective agreement [38] that also provides replay protection.
- Secrecy of the shared secret: If the terminal and card have agreed on a shared secret, the attacker cannot derive this secret.
- Authenticity of the AIP: Given that the terminal has successfully executed the protocol with a given card, they both agree on the same value of the card's AIP. Since the card's capabilities are encoded in the AIP, this property ensures downgrade protection of our extension.

The corresponding theorems are proven to hold for the case where CDA has been completed successfully and for the case of offline acceptance, where the card returns a Transaction Certificate (TC).

3.4 Proof-of-Concept Implementation

We show that PURE can be easily integrated into the existing C-2 kernel introducing low overhead and maintaining backward compatibility with a Proof-of-Concept (PoC) implementation.

We have implemented the standard and integrated version of PURE as Android apps on mobile phones. Although UWB is already widely deployed in Android and iOS devices, they lacked the necessary UWB software APIs to implement our extension at the time of writing. We thus have equipped each of the Android phones with an external Qorvo DW3000 board that provides us the necessary low-level access to the chip. Note that software APIs for UWB are constantly improving, for example, as of June 2023, Android allows to set the ranging key[2]. With adequate access to the UWB chip on mobile device, the LED algorithm described later in Section 4.2 could be implemented directly on the UWB chip integrated in phones.

Figure 6 depicts our setup, where the backends are responsible for generating EMV-specific messages. The backends generate EMV transaction data either (i) by replaying a prerecorded real transaction or (ii) by connecting live through NFC to real EMV-executing devices. The backends thus bridge our implementation of PURE to real EMV implementations, thereby proving the correctness of our implementation when connecting to a real terminal and a real card.

The stand-alone implementation follows exactly the protocol described in Section 3.2. Since the card's private key cannot be extracted, the SDAD' signature is created and validated with a self-created public/private key pair. For the real protocol, we can assume the terminal knows and trusts the card's public relying on the successful validation of the corresponding certificate.

Transaction Execution Time. Our PoC implementation shows that PURE adds only a negligible time delay compared to a regular transaction time. As a consequence, the time required for a user to keep his phone close to the terminal is not significantly increased. Table 2 shows the timings of the stand-alone and of the integrated implementation using a prerecorded transaction as backend. We run the apps on a Samsung S21 and a Pixel 4, respectively acting as a card and as a terminal. Ranging requires on average 30 ms; however, it is executed in parallel to the NFC exchange and thus does not introduce any overhead.

Integration with Real Payments. We use the live backends, to prove that PURE can be integrated into existing real-world transactions. The backend of the UWB-enabled terminal consists of a phone that forwards all communication from a real registered terminal to a front-end UWB-enabled phone that in turn executes PURE with a UWB-enabled card (see Appendix A for further details). Similarly the backend of the UWB-enabled card, is implemented with a phone that forwards all NFC communication to a real card. The extended EMV protocol (over NFC) and the UWB ranging are executed between the two front-end phones. If the protocol is

[2] https://developer.android.com/reference/kotlin/androidx/core/uwb/RangingParameters#sessionKeyInfo()

[3] https://gitlab.com/practical_emv/timing-data [43]

	Stand-alone	Integrated A	Integrated B
DH (ms)	46.8 ±9.3	41.0 ±7.5	41.0 ±7.5
CERT (ms)	44.5 ±10.8	-	-
AUTH (ms)	38.6 ±6.9	37.7 ±6.3	-
Overhead	-	10 - 16%	**5-9%**

Table 2: **Transaction Execution Time of PURE.** PURE introduces at most 78 ms delay. The Integrated version B achieves the lowest overhead of 41 ms by authenticating the DH and timings with the existing SDAD. We report the relative increase in time introduce by PURE with respect to the typical execution time of a Mastercard transaction of (630 ±151) ms[3].

executed successfully, a real transaction is executed on behalf of the real card at the real terminal. We could not use this fully integrated implementation for the measurement of the transaction excecution times due to the substantial overhead introduced by the forwarding of messages. However, the PoC with live backends shows in practice that a UWB-enabled phone can successfully execute real transactions with a UWB-enabled terminal. Additionally, it shows PURE's backward compatibility, since normal cards were still able to execute transactions with our UWB-enabled terminal.

4 UWB and PURE

PURE uses UWB to measure the distance between the card and terminal. UWB chips are now largely deployed in smartphones and other devices [23]. In PURE we make this measurement accurate, reliable, and secure against all known attacks [3, 36]. For this, we leverage the favorable wireless channels in the specific case of contactless payments.

4.1 Background on HRP UWB

HRP is the UWB mode of the IEEE 802.15.4z standard [28] deployed in current smartphones. UWB devices measure the distance between themselves based on the ToF of ranging messages. In essence, the first transceiver sends a message and estimates the round-trip time by measuring the ToA of the response. Since the second device provides its local processing time, the first can estimate the ToF with high accuracy.

As a secure ranging scheme, HRP can be used to derive an upper bound on the distance between two trusted ranging devices. An external attacker successfully breaks this scheme if it can convince the ranging devices they are closer than they actually are. Specifically, if the ranging devices measure an upper bound that is smaller than their true distance.

Scrambled Timestamp Sequence. Every HRP message contains a scrambled timestamp sequence (STS) consisting of 4096 pseudo-random pulses. The STS has two purposes: measuring the ToA accurately and, by making the ranging message unpredictable, preventing trivial distance reduction attacks. Following the IEEE 802.15.4z standard [28], given an authentic fresh shared secret, the ranging devices derive the STSs using AES in counter mode (AES-CTR). Note here the connection between the properties proven in Section 3 and the requirements posed by HRP. Despite the messages being unpredictable, the security of the scheme relies on the security of the ToA measurement, which we describe next.

Channel Impulse Response. When an HRP receiver detects a ranging message, it cross-correlates the received signal $RX[t]$ with a local template of the expected STS:

$$CIR[t] = (RX \star STS)[t] = \sum_{\tau} RX[t] \cdot STS[t+\tau]. \quad (1)$$

The result is referred to as the Channel Input Response (CIR), in which correlation peaks indicate timestamps at which STS copies were received. Several such copies can exist due to the multipath effects (e.g., reflections) of the wireless channel. More precisely, the obtained CIR is only an approximation of the real channel: similarities between different parts of the STS (i.e., auto-correlation noise), interference, and other factors may result in non-zero correlation noise. Figure 7 depicts a CIR of a multipath channel.

Leading Edge Detection. To obtain an accurate ToF, it is crucial to find the ToA of the first STS copy. Since the highest peak in the CIR is not necessarily the first, HRP receivers perform LED. Starting from the highest peak, they try to identify a potential first peak (i.e., the *leading edge*) within a pre-defined backsearch window. LED algorithms are closed-source and vendor-specific, but they all have to distinguish between peaks and correlation noise, likely involving some threshold [46]. This is a non-trivial classification problem, especially in severe Non-Line-of-Sight (NLoS) scenarios, where low leading edges are to be expected. Prior work has shown [36, 41] that an attacker signal may result in correlation noise being misclassified as an earlier leading edge, effectively reducing the measured distance.

4.2 UWB Ranging in Contactless Payments

For contactless payments, UWB devices generally benefit from excellent signal reception because they are in close proximity. Nevertheless, transmissions may still be affected by the devices' close surroundings (e.g. reflecting surfaces) and how the user holds the phone.

Consequently, leading edge detection is still indispensable for accurate distance measurements. In PURE, we use these insights to build an LED based on a fixed, absolute CIR threshold, independent of specific channels or CIR noise. While this approach is not applicable to other ranging use cases, we

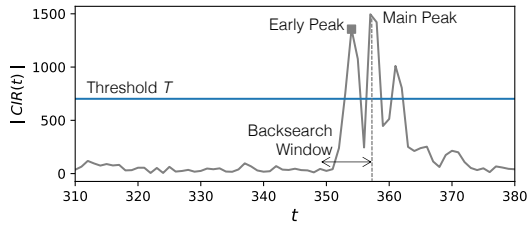

Figure 7: **Leading Edge Detection in PURE.** Each peak in the CIR denotes the arrival time of a copy of the signal. PURE identifies the first copy, corresponding to the distance between transmitter and receiver, using a fixed threshold T. Due to its high threshold independent from the CIR, PURE is robust against any manipulation by an attacker. It is also reliable because channels in contactless payments typically exhibit high CIR peaks.

show how it enables accurate, reliable, and secure ranging for contactless payments. Figure 7 depicts a CIR in which the leading edge is identified using a fixed threshold. We show analytically how such a threshold can be chosen to make the LED robust against distance reduction attacks. We use this CIR threshold within a commercial HRP receiver, and show that only minor adjustments to the threshold are sufficient to build a practically secure and reliable LED.

Intuition. A peak in the CIR is a measure of the similarity between the received signal and the expected STS for a given arrival time. If the contribution of each pulse on the generated peak is limited, then an absolute threshold on the peak height is equivalent to a threshold on the number of correct pulses in the STS. Indeed, each pulse's maximum contribution to a peak is constrained by the receiver's tolerated maximum power. Signals exceeding this limit are clipped by the receiver front-end, as illustrated in Figure 8. In the following, we assume without loss of generality that the output of the ADC is normalized to the range of $[-1,+1]$ and the maximum contribution of a pulse is capped at 1. Under these assumptions, setting an absolute threshold on the amplitude of the peak equates to setting the minimum number of pulses that the attacker has to guess correctly. This is a problem with a well-defined success probability, decreasing in the threshold on the peak height.

LED with an Absolute Threshold. We denote the number of pseudo-random pulses in the STS as n, and the number of received pulses with the correct polarity as n_c. Given that the receiver strictly limits the contribution of a single pulse to the CIR, setting a threshold T on the peak height requires an adversary to guess at least

$$n_c = (n+T)/2 \tag{2}$$

out of n pulses correctly.

The definition of n_c follows from the fact that a signal with $n/2$ correct pulses (i.e., the expected outcome of n random binary guesses) should result in a CIR value of 0, while a fully correct STS can, at most, generate a peak of height n. Therefore, a valid peak can only be generated if the number of correct pulses is sufficiently larger than the expected value $n/2$ of a random guess.

Implications for security. For every pulse the receiver expects, the attacker can add at most 1 to a potential correlation peak. This property forces the attacker to guess at least n_c binary pulses correctly to reach the threshold, even if their transmission power is optimal. Setting a high threshold then limits the attacker's success probability. This result is independent of factors like pulse shapes, signal power, or channel effects. We define the threshold for a desired security level and evaluate it experimentally in Section 4.3.

Implications for reliability. To preserve ranging accuracy, the threshold must be low enough such that the leading edge (i.e, the first peak) of a genuine ranging message exceeds it with high probability. In other words, the link budget must be high enough for the receiver to identify the first STS copy in the overwhelming majority of expected wireless channels. Luckily, this condition is usually met in contactless payment scenarios due to the devices' proximity. We verify this experimentally in Section 4.3, where we implement a LED with an absolute threshold on the Qorvo DWM3000.

Choosing a Threshold. The threshold T is chosen to make the probability of passing the threshold by random guessing negligible. This probability can be written as

$$p(|CIR| > T) = 2 \cdot p(X > n_c) = 2(1 - F_X(n_c)) \tag{3}$$

where $X \sim \mathcal{B}(n, \frac{1}{2})$ is a random variable following the binomial distribution and F_X is its cumulative distribution function.

Note that $n_c > n/2$ n_c correct pulses and n_c wrong pulses are equivalent.

This equation describes the probability of exceeding the threshold at a specific CIR index, but the attacker's signal could cause a peak anywhere in the back-search window. To bound the probability of this event, we can multiply $p(|CIR| > T)$ by our chosen window length of $9 \simeq 2^4$. This is an overestimation in favor of the attacker, since neighboring CIR values are generally not independent. Considering that the attacker could advance the ToA of any of the $3 \simeq 2^2$ ranging messages, we can conservatively bound the probability of a successful attack by

$$p(success) \leq 2^6 \cdot p(|CIR| > T) \tag{4}$$

In Section 4.3 we will estimate the success rate empirically on a real receiver.

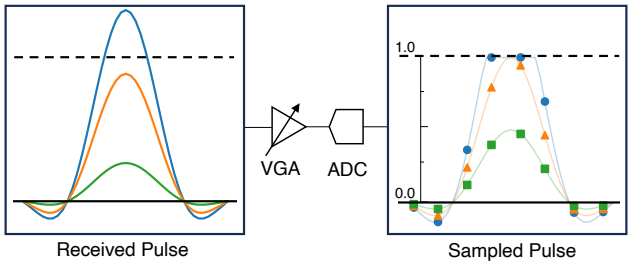

(a) Pulse before and after the receiver frontend.

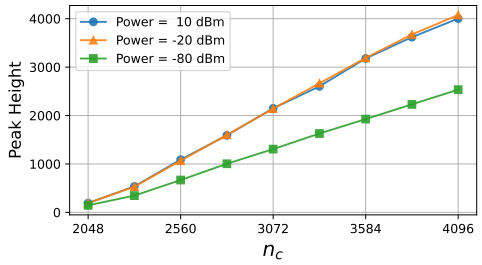

(b) Peak height for different powers and correct pulses n_c.

Figure 8: **Receiver Front-end.** The Variable Gain Amplifier aims at adjusting the received pulse amplitude to match the dynamic range of the Analog-to-Digital Converter. If the pulse is too strong, it is clipped (a). If it is too weak, the VGA fills less than the dynamic range of the ADC. Therefore, each correct pulse contributes less than 1.0 to the CIR peak, which grows linearly with the number of correct pulses (b). Increasing the power from -20dBm to 10dBm does not bring any advantage because of clipping.

4.3 Experimental Analysis

In this section, we prove that the idea of an absolute threshold works in practice. We provide empirical evidence for our security argument and show that, at the same time, reliability and accuracy of the distance measurements do not deteriorate, analyzing its security, reliability and accuracy.

Security Against Distance-Decreasing Attacks. We show that, with a suitable choice of the absolute threshold, PURE is secure against the Ghost Peak distance-reduction attack [35]. We assume an adversary that: (i) transmits a random STS hoping to guess enough pulses to create a fake early peak above the threshold, (ii) has an ideal channel to the victim, and (iii) transmits at the maximum power before clipping occurs at the receiver. Figure 9 shows the results of our empirical evaluation of the success rate for increasing the absolute thresholds on the peak height. Accepting peaks higher than $T = 702$ bounds the success rate to a maximum of 2^{-49}.

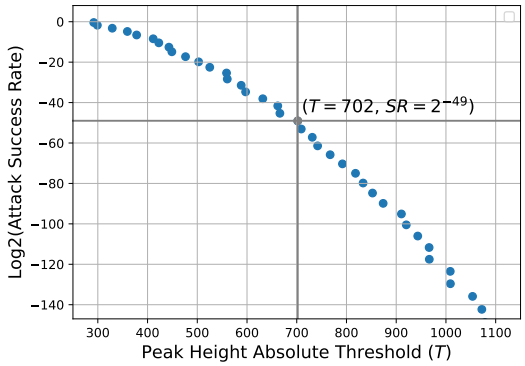

Figure 9: **Success Rate.** Empirical success rate of a Ghost Peak attack for different choices of the absolute threshold. A threshold $T = 702$ limits the success rate to 2^{-49} while still maintaining high reliability and accuracy (see Table 3).

To estimate the success rate of Ghost Peak even when a high threshold makes it extremely low, we take the following strategy, inspired by importance sampling [32] and other similar approaches in security evaluations [11, 16]. First, using a signal generator connected to a Qorvo DWM3000 receiver, we measure the CIR peak amplitude distribution for varying numbers of correct pulses (n_c) in the transmitted signal. The generator uses the maximum power before clipping occurs at the receiver (-21dBm). Then, we identify the minimum n_c necessary to obtain a peak above the threshold (T). More precisely, we make the conservative choice of considering n_c sufficient to pass the threshold as long as the queue of the measured peak distribution for n_c is above T. Finally, following Equation 3 we estimate the probability of having obtained n_c by random guessing. We further multiply it by 2^6 to consider all ranging messages and all indices in the backsearch, as explained in Section 4.2.

Figure 9 depicts the success rate for different thresholds. These results, obtained experimentally, already take into account the effects of measurement noise and other possible nonidealities of the receiver.

PURE is also not vulnerable to Mix-Down [3] because it uses DS-TWR, in which clock drift compensation is done implicitly by measuring the RTT on both sides.

Overall, known distance-reduction attacks are prevented and cannot be used by an attacker to bypass the relay protection provided by PURE.

Reliability and Accuracy in Payment Channels. We show next experimentally that the threshold chosen for security does not hamper the ranging reliability in the contactless payment scenario. To this end, we equipped various readers with a Qorvo DWM3000, and used an iPhone 12 to perform ranging between the two. The current API offered by Apple [4] enables UWB ranging with fixed nonces but does not allow setting the ranging key necessary for secure ranging. However, the API is sufficient to execute ranging with a Qorvo board

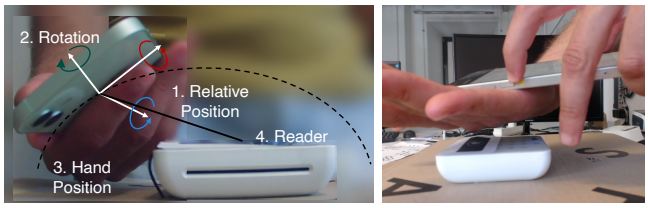

(a) Variability introduced during CIR collection (b) Deliberately obstructed antenna

Figure 10: **CIRs in Payment Scenarios.** We collected CIRs covering a wide range of payment scenarios modifying the parameters shown in Figure 10a. We consider holding the phone with the hand covering the back realistic. For comparison, we generate artificially hard cases by explicitly obstructing the antennas located at the top of the phone (10b).

and extract the measured CIR with a modified version of the Qorvo Nearby Interaction application [42]. This setup was used to collect a total of 300 CIRs by scanning the half-space in front of the reader, modifying the phone angle and position to cover a wide range of payment scenarios (see Figure 10a).

In Figure 11, we contrast the CIRs collected during contactless payments with CIRs collected in the following different scenarios: in Line-of-Sight (LoS) at a 20 cm distance, with deliberately obstructed UWB antennas as shown in Figure 10b and in NLoS at a distance of 3 meters. A qualitative examination already indicates that the peaks in realistic payment channels are visibly easier to detect than the peaks measured in severe NLoS where, in other use cases, UWB is still expected to perform reliably.

We use the extracted CIRs to analyze the performance of our LED with the chosen threshold ($T = 702$) and consider the ToA measured by the Qorvo as the ground truth.

The maximum accepted measured distance d_{th} must be set to ensure reliable payments within their operational range (terminal to card distance $< d_{max} = 5$ cm [19]). Specifically, let Δd^+ represent the maximum overestimation of distance beyond the error introduced by the chip accuracy Δd_{qorvo}. The maximum distance measured from within the payment range is given by $d_{th} = d_{max} + \Delta d^+ + \Delta d_{qorvo}$. Since we never underestimate the distance with respect to the chip, the worst underestimation error is given by the chip's accuracy. The longest undetected relay is located at $d_{relay} = d_{th} + \Delta d_{qorvo}$ distance because, due to an underestimation error, the measured distance falls in the accepted range.

Table 3 relates the accuracy with which we can detect a relay with the False Rejection Rate (FRR) over the measured channels. Our LED achieves 46 cm relay protection maintaining a small FRR (2%) in common payment scenarios. Although we tuned the system not to support unlikely interactions and channels (like the one in Figure 10b), where the hand fully covers the antennas, our system would still work reliably in most of these cases (with a FRR of 9.5% and $d_{relay} = 46$ cm). In general, the few failed measurements can be easily handled in the same way as failed contactless payments: by prompting the user to reposition their phone.

FRR	FRR_{blk}	d_{relay}	Δd^+	d_{max}	Δd_{qorvo}
0.5%	7.2%	95 cm	70 cm	5 cm	10 cm
1%	7.7%	85 cm	55 cm	5 cm	10 cm
2 %	9.5%	**46 cm**	21 cm	5 cm	10 cm

Table 3: **Reliability and Accuracy.** Trade-off between reliability and accuracy in contactless payments with an absolute threshold of 702 on the peak height. The table reports the FRR for realistic payment scenarios and for deliberately blocked antennas (FRR_{blk}).

5 Limitations and Future Work

In general, one limitation of PURE and other relay protections is that they are useful to cardholders only if they are largely deployed in all terminals. Otherwise, cardholders either have to use only a small subset of protected terminals, which is impractical, or cards must support backward compatibility, remaining vulnerable to a relay to an unprotected terminal. In practice, the backward compatibility of PURE can facilitate deployment until old terminals are phased out. The increasing availability of UWB and on-device cards/terminals is also expected to favor a rapid deployment.

A second limitation of PURE is that, while it prevents relays by an external attacker, it does not prevent attacks where one of the parties is dishonest. Malicious terminals could also operate without relays and their attacks are traceable; hence they are generally out of scope in the context of payments. A malicious cardholder could, for example, attempt to extract the ranging key from the card and share it with an accomplice. Merchants/terminals should therefore not use PURE as a proof of the distance of the cardholder. The practical realization of distance bounding in presence of a malicious prover is an open challenge, requiring physical layer capabilities that are currently not available in commercial technologies (e.g., analog processing with negligible delay [44]). At the system level, the use of Trusted Execution Environments (TEEs) could help reduce the trust assumptions on user and device, as it is already the case for on-device cards.

Finally, due to its wide-spread availability, PURE uses HRP UWB, for which a closed-form expression of the security level has not been found yet. We have demonstrated that HRP UWB can be both reliable and secure against known attacks in contactless payments, thanks to favorable channel conditions. However, our results do not automatically apply to other use cases, such as Passive Keyless Entry and Start Systems (PKES). In addition, stronger (e.g., adaptive) attacks on HRP UWB might be identified in the future. Ideally, future efforts in secure ranging will provide a theoretical framework

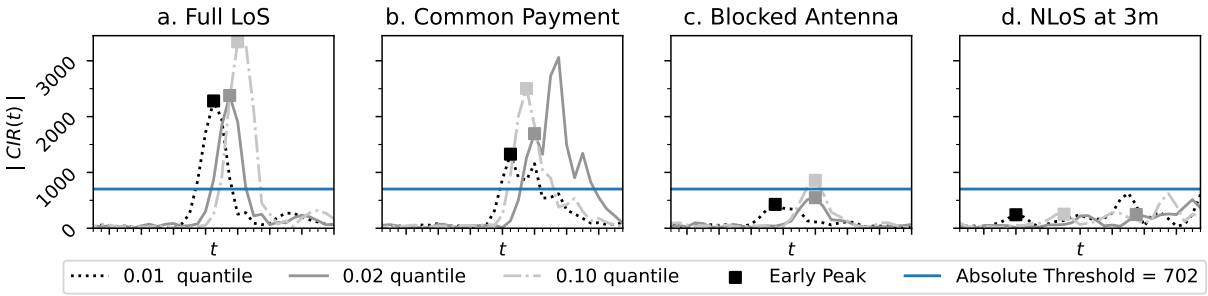

Figure 11: **CIR Comparison.** LoS (a) and Common Payment (b) scenarios show visibly better peaks than in other cases (c, d). In payments, it is thus possible to identify the early peak securely and reliably with an absolute threshold. CIRs were ranked based on the amplitude of the earliest peak. For each setting, we report the 0.01, the 0.02, and the 0.20 quantile CIR.

for the analysis of HRP UWB against any attack.

In the future, PURE could leverage the strengths of upcoming standards. For ranging, the IEEE802.15.4ab standard [27] is being developed with a focus on a higher link budget, for example, thanks to the multi-millisecond packets proposed by Apple [22]. This will facilitate deployment in a wide range of scenarios, further improving usability. It could also help enable novel wireless payment use cases, with seamless user interaction from larger distances. For payments, the C-8 Kernel [21] proposed by EMVCo in 2022 aims at unifying existing kernels and becoming the new standard. To protect sensitive data, it encrypts the card-terminal communication with an ephemeral Diffie-Hellman key. PURE could then derive the ranging key from the existing shared secret.

6 Related Work

Distance Bounding. Distance bounding protocols provide an upper bound on the physical distance between communicating parties under various threat models (e.g., external attacker, honest/dishonest prover) [5, 10, 14]. Narrow-band distance-bounding systems were implemented over Bluetooth [], but their accuracy is significantly lower than in UWB. Moreover, a narrow band system does not fundamentally prevent Early Detect/Late Commit attacks [34, 40]. PURE focuses on integrating HRP UWB ranging in EMV contactless payments, achieving relay protection even against physical-layer attackers.

Relay Protections in EMV. Relay attacks on payment systems were demonstrated in 2005 in [31]. Several works [9, 26, 47] followed, showing their practicality. In the context of payments, distance bounding protocols were proposed to defend against an external attacker [43] or in the setting where readers can be dishonest [12, 18]. We discussed Mastercard RRP [20] and its improved version L1RP [43] in Section 2. Visa's relay protection does include a proper distance bounding protocol and only relies on the assumption that L1 NFC messages are hard to manipulate with COTS devices [48]. The card sends a random byte string to the terminal both at Level 1 and 3 of NFC. The terminal detects a relay if the bytes received at the two layers do not match.

EMV formal modeling. There are various symbolic models of the EMV standards [6, 7, 15]. These models helped to discover and patch protocol flaws that allowed attackers to bypass the CVMs [6, 8, 43]. We prove the security of PURE extending the formal model of the C-2 kernel provided in [7].

Attacks on HRP UWB. An initial security analysis of HRP UWB shows a complex reliability-security trade-off [46]. Attacks that forge early cross-correlation peaks by injecting random pulses [36, 41] questioned the security of cross-correlation for ToA estimation. Single-sided two-way ranging (SS-TWR) is vulnerable to attacks exploiting the non-ideality of clocks in the ranging devices [3]. We address these vulnerabilities in Section 4.

7 Conclusions

Relay attacks have been a constant threat to card payments and are now especially relevant with the increased popularity of contactless payments. This study proposed PURE, the first concrete proposal on how to integrate UWB-ranging techniques into EMV kernels. The integration of PURE within EMV protocols was formally verified in Tamarin. Our prototype showed that PURE can be deployed on real smartphones with negligible overhead on transaction time and without any modification to the backend. We tailored UWB-HRP to be secure and reliable in contactless payment. We provided an analytical and experimental analysis of the introduced absolute threshold demonstrating that PURE is secure against all known distance reduction attacks. In conclusion, PURE provides a practical and effective solution to relay attacks on contactless payments.

References

[1] Aysajan Abidin, Mohieddine El Soussi, Jac Romme, Pepijn Boer, Dave Singelée, and Christian Bachmann. Secure, accurate, and practical narrow-band ranging system. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(2):106–135, 2021.

[2] Ross J. Anderson and Steven J. Murdoch. EMV: why payment systems fail. *Commun. ACM*, 57(6):24–28, 2014.

[3] Claudio Anliker, Giovanni Camurati, and Srdjan Capkun. Time for change: How clocks break UWB secure ranging. In Joseph A. Calandrino and Carmela Troncoso, editors, *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*. USENIX Association, 2023.

[4] Apple. Nearby Interaction. https://developer.apple.com/documentation/nearbyinteraction, September 2023.

[5] Gildas Avoine, Muhammed Ali Bingöl, Ioana Boureanu, Srdjan Capkun, Gerhard P. Hancke, Süleyman Kardas, Chong Hee Kim, Cédric Lauradoux, Benjamin Martin, Jorge Munilla, Alberto Peinado, Kasper Bonne Rasmussen, Dave Singelée, Aslan Tchamkerten, Rolando Trujillo-Rasua, and Serge Vaudenay. Security of distance-bounding: A survey. *ACM Comput. Surv.*, 51(5):94:1–94:33, 2019.

[6] David A. Basin, Ralf Sasse, and Jorge Toro-Pozo. Card brand mixup attack: Bypassing the PIN in non-visa cards by using them for visa transactions. In Michael Bailey and Rachel Greenstadt, editors, *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, pages 179–194. USENIX Association, 2021.

[7] David A. Basin, Ralf Sasse, and Jorge Toro-Pozo. The EMV standard: Break, Fix, Verify. In *2021 IEEE Symposium on Security and Privacy (SP), May 23-27, 2021, San Francisco, CA, US*, pages 1766–1781. IEEE, May 2021.

[8] David A. Basin, Patrick Schaller, and Jorge Toro-Pozo. Inducing authentication failures to bypass credit card pins. In Joseph A. Calandrino and Carmela Troncoso, editors, *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*. USENIX Association, 2023.

[9] Thomas Bocek, Christian Killer, Christos Tsiaras, and Burkhard Stiller. An nfc relay attack with off-the-shelf hardware and software. In Rémi Badonnel, Robert Koch, Aiko Pras, Martin Drašar, and Burkhard Stiller, editors, *Management and Security in the Age of Hyperconnectivity*, pages 71–83, Cham, 2016. Springer International Publishing.

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

[11] Giovanni Camurati, Matteo Dell'Amico, and François-Xavier Standaert. Mcrank: Monte carlo key rank estimation for side-channel security evaluations. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2023(1):277–300, 2023.

[12] Tom Chothia, Ioana Boureanu, and Liqun Chen. Short paper: Making contactless EMV robust against rogue readers colluding with relay attackers. In Ian Goldberg and Tyler Moore, editors, *Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers*, volume 11598 of *Lecture Notes in Computer Science*, pages 222–233. Springer, 2019.

[13] Computerphile. Chip & PIN Fraud Explained - Computerphile. https://www.youtube.com/watch?v=Ks0SOn8hjG8.

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

[15] Joeri de Ruiter and Erik Poll. Formal analysis of the EMV protocol suite. In Sebastian Mödersheim and Catuscia Palamidessi, editors, *Theory of Security and Applications - Joint Workshop, TOSCA 2011, Saarbrücken, Germany, March 31 - April 1, 2011, Revised Selected Papers*, volume 6993 of *Lecture Notes in Computer Science*, pages 113–129. Springer, 2011.

[16] Matteo Dell'Amico and Maurizio Filippone. Monte carlo strength evaluation: Fast and reliable password checking. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, pages 158–169. ACM, 2015.

[17] Saar Drimer and Steven J. Murdoch. Chip & PIN (EMV) relay attacks. https://www.cl.cam.ac.uk/research/security/banking/relay/. Accessed: 2023-09-27.

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

[19] EMVCo. EMV Level 1 Specifications for Payment Systems, EMV Contactless Interface Specification, Version 3.1. https://www.emvco.com/specifications/, December 2020.

[20] EMVCo. EMV Contactless Specifications for Payment Systems, Book C-2, Kernel 2 Specification, Version 2.11. https://www.emvco.com/specifications/, June 2023.

[21] EMVCo. EMV Contactless Specifications for Payment Systems, Book C-8, Kernel 8 Specification, Version 1.1. https://www.emvco.com/specifications/, June 2023.

[22] Ersen Ekrem, Ido Bettesh, and Moche Cohen (Apple Inc.). More on narrowband assisted multi-millisecond UWB. https://mentor.ieee.org/802.15/dcn/21/15-21-0593-00-04ab-more-on-nba-mms.pptx, November 2021. Accessed 2023-10-07.

[23] Unleashing the potential of uwb: Regulatory considerations. https://www.firaconsortium.org/sites/default/files/2022-08/Unleashing-the-Potential-of-UWB-Regulatory-Considerations.pdf.

[24] U.S. Payments Forum. Contactless limits and EMV transaction processing, version 1.0. https://www.uspaymentsforum.org/wp-content/uploads/2020/10/Contactless-Limits-WP-FINAL-October-2020.pdf.

[25] Aurélien Francillon, Boris Danev, and Srdjan Capkun. Relay attacks on passive keyless entry and start systems in modern cars. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2011, San Diego, California, USA, 6th February - 9th February 2011*. The Internet Society, 2011.

[26] Lishoy Francis, Gerhard P. Hancke, Keith Mayes, and Konstantinos Markantonakis. Practical NFC peer-to-peer relay attack using mobile phones. In Siddika Berna Örs Yalçin, editor, *Radio Frequency Identification: Security and Privacy Issues - 6th International Workshop, RFIDSec 2010, Istanbul, Turkey, June 8-9, 2010, Revised Selected Papers*, volume 6370 of *Lecture Notes in Computer Science*, pages 35–49. Springer, 2010.

[27] IEEE. IEEE 802.15 Documents Group TG4AB. https://mentor.ieee.org/802.15/documents?is_dcn=NBA&is_group=04ab. Accessed 2023-10-07.

[28] IEEE Standard for Low-Rate Wireless Networks–Amendment 1: Enhanced Ultra Wideband (UWB) Physical Layers (PHYs) and Associated Ranging Techniques. pages 1–174, 2020.

[29] Apple Inc. Tap to pay on iphone. https://developer.apple.com/tap-to-pay/.

[30] Juniper Research Jordan Rookes. Contactless payments – card vs mobile. https://www.juniperresearch.com/whitepapers/contactless-payments-card-vs-mobile, September 2022.

[31] Ziv Kfir and Avishai Wool. Picking virtual pockets using relay attacks on contactless smartcard. In *First International Conference on Security and Privacy for Emerging Areas in Communications Networks, SecureComm 2005, Athens, Greece, 5-9 September, 2005*, pages 47–58. IEEE, 2005.

[32] Teun Kloek and Herman K. van Dijk. Bayesian estimates of equation system parameters, an application of integration by monte carlo. *Econometrica*, 46:1–19, 1976.

[33] Hugo Krawczyk. Sigma: The 'sign-and-mac' approach to authenticated diffie-hellman and its use in the ike protocols. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, pages 400–425, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

[34] Markus Kuhn. Distance-bounding protocols. https://www.surrey.ac.uk/sites/default/files/2018-05/kuhn.pdf, 2018.

[35] Tian-Fu Lee and Tzonelih Hwang. Three-party authenticated key agreements for optimal communication. *PLoS ONE*, 12, 2017.

[36] Patrick Leu, Giovanni Camurati, Alexander Heinrich, Marc Roeschlin, Claudio Anliker, Matthias Hollick, Srdjan Capkun, and Jiska Classen. Ghost peak: Practical distance reduction attacks against HRP UWB ranging. In Kevin R. B. Butler and Kurt Thomas, editors, *31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10-12, 2022*, pages 1343–1359. USENIX Association, 2022.

[37] Patrick Leu, Mridula Singh, Marc Roeschlin, Kenneth G. Paterson, and Srdjan Capkun. Message Time of Arrival Codes: A Fundamental Primitive for Secure Distance Measurement. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 500–516. IEEE.

[38] Gavin Lowe. A hierarchy of authentication specification. In *10th Computer Security Foundations Workshop (CSFW '97), June 10-12, 1997, Rockport, Massachusetts, USA*, pages 31–44. IEEE Computer Society, 1997.

[39] Simon Meier, Benedikt Schmidt, Cas Cremers, and David A. Basin. The TAMARIN prover for the symbolic analysis of security protocols. In Natasha Sharygina and Helmut Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, pages 696–701. Springer, 2013.

[40] Marcin Poturalski, Manuel Flury, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. Distance bounding with IEEE 802.15.4a: Attacks and countermeasures. *IEEE Trans. Wirel. Commun.*, 10(4):1334–1344, 2011.

[41] Marcin Poturalski, Manuel Flury, Panos Papadimitratos, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. The cicada attack: Degradation and denial of service in IR ranging. In *2010 IEEE International Conference on Ultra-Wideband*, pages 1–4. IEEE.

[42] Qorvo. Qorvo Solutions Interoperable with Apple* U1 Chip for New Ultra-Wideband Enabled Experiences. https://www.qorvo.com/innovation/ultra-wideband/products/uwb-solutions-compatible-with-apple-u1, June 2021.

[43] Andreea-Ina Radu, Tom Chothia, Christopher J. P. Newton, Ioana Boureanu, and Liqun Chen. Practical EMV relay protection. In *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*, pages 1737–1756. IEEE, 2022.

[44] Kasper Bonne Rasmussen and Srdjan Capkun. Realization of RF distance bounding. In *19th USENIX Security Symposium, Washington, DC, USA, August 11-13, 2010, Proceedings*, pages 389–402. USENIX Association, 2010.

[45] RBR Press Release - COVID-19 prompts interest and innovation in cardless ATM withdrawals (6th july 2020). https://www.rbrlondon.com/wp-content/uploads/2020/07/GA25_Press_Release_060720.pdf. Accessed 18-August-2023.

[46] Mridula Singh, Marc Roeschlin, Ezzat Zalzala, Patrick Leu, and Srdjan Capkun. Security analysis of IEEE 802.15.4z/hrp UWB time-of-flight distance measurement. In Christina Pöpper, Mathy Vanhoef, Lejla Batina, and René Mayrhofer, editors, *WiSec '21: 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks, Abu Dhabi, United Arab Emirates, 28 June - 2 July, 2021*, pages 227–237. ACM, 2021.

[47] Chip and PIN Fraud. https://www.youtube.com/watch?v=X7pjUIxKoEc.

[48] CHEN Yuexi, Marc Kekicheff, Mustafa Top, and Hao Ngo. Binding cryptogram with protocol characteristics, September 29 2022. US Patent App. 17/838,561.

A Integrated Implementation Details

This section aims to clarify the setup used to test the functionality of PURE between a real card and a real terminal. As explained in Section 3, we use a backend to generate EMV messages, instead of re-implementing the Mastercard Kernel. Figure 12 shows the setup for the integrated implementation with real, live backends. A UWB Terminal consists of a registered terminal issuing EMV commands to a phone that reports the commands to the front-end phone via WiFi. Similarly, the UWB card consists of a registered card responding to the commands received them from the front-end phone over WiFi. The UWB-protected EMV transaction happens between the UWB card and the terminal identified by the blue boxes. Figure 12 highlights the UWB negotiation showing that the innermost phone of the UWB-Card sets the UWB bit to communicate to the UWB-Reader its ability to execute ranging. The innermost phone in the UWB-Terminal registers that the card can execute ranging and unsets the UWB bit to ensure that the check on the issuer side would not fail. Recall that the AIP is an authenticated value. If the card supports ranging, the innermost phones execute the DH exchange and the ranging. Finally, when the registered terminal issues the `GEN AC` command, the innermost phones exchange the signature over the DH parameters, the transaction parameters, and the timings. The `GEN AC` command is forwarded only if the signature verification succeeds. In conclusion, from the point of view of the registered card and terminal the transaction is a normal Mastercard transaction and the innermost phones are transparent because all the modifications and additions are not forwarded to the backend. This ensures that our integrated system can execute real transactions. We tested this with a Mastercard and a Sumup terminal as registered card and reader. The EMV trace including PURE is highlighted between the two innermost phones. The distance is measured between the two innermost phones. If a relay is detected, the transaction is blocked by sending an error message in response to the `GEN AC` command.

B Analysis of RTT-based Relay Protection

This section provides the derivation of the shortest relay that an RTT-based relay protection can achieve. Let T_{reply} be the processing time of the responder, $(\Delta t^-, \Delta t^+)$ be the maximal

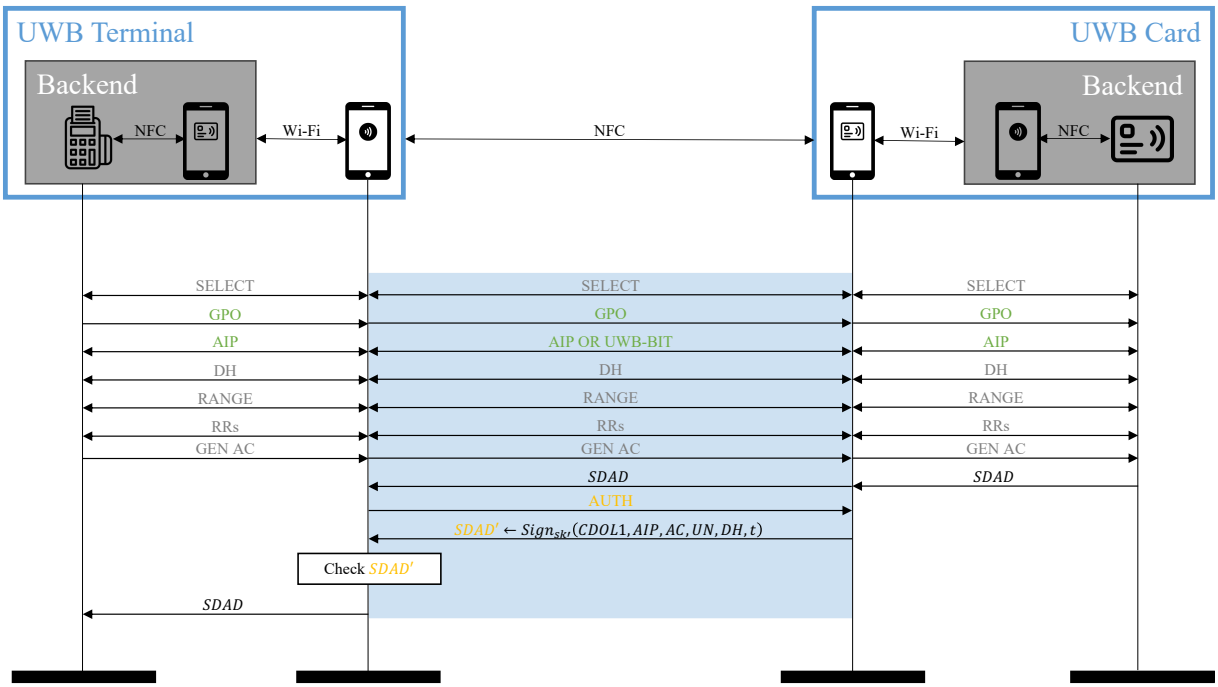

Figure 12: Proof-of-Concept transcript of a transaction using PURE with real card and terminal backend.

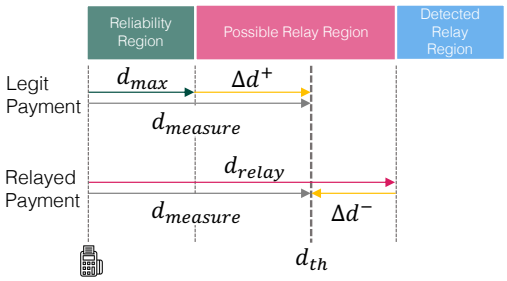

Figure 13: Maximum relay protection achievable by a relay protection system measuring distance with uncertainty. The legitimate payment with maximum distance overestimation Δd^+ determines the maximum accepted measured distance d_{th}. The longest undetected relay happens when the distance measured is affected by the maximal distance underestimation Δd^-.

absolute underestimation and overestimation of the RTT due to the uncertainty of the measurement. In a relay protection system based on RTT, the maximal accepted RTT is

$$T^*_{round} = \frac{2d_{max}}{c} + T_{reply} + \Delta t^+ \tag{5}$$

where d_{max} is the maximum distance accepted by the relay protection system. T^*_{round} must include the uncertainty on the measurement to ensure the reliability of the system. Overall, the relay protection system will have to accept all measured distances d such that

$$d < d_{th} = \frac{c}{2} \cdot (T^*_{round} - T_{reply}). \tag{6}$$

If the messages are relayed over d_{relay} meters by an adversary, who additionally adds T_{proc} seconds of processing time, the minimal round-trip time measured by the victims would be

$$T_{attacker} = 2 \cdot (\frac{d_{relay}}{c} + T_{proc}) + T_{reply} - \Delta t^-. \tag{7}$$

The receiver, burdened with the task of detecting a relay attack on an exchange of messages, can only reliably detect relays such that

$$\begin{aligned} T_{attacker} &> T^*_{round} \\ d_{relay} &> d_{max} + T_{proc} \cdot c + \frac{c}{2}(\Delta t^- + \Delta t^+) \end{aligned} \tag{8}$$

With $\Delta d^+ = \frac{c}{2}\Delta t^+$ and $\Delta d^- = \frac{c}{2}\Delta t^-$ being the maximal distance overestimation and underestimation, Equation 8 can be rewritten as

$$d_{relay} > d_{max} + T_{proc} \cdot c + \Delta d^- + \Delta d^+. \tag{9}$$

Figure 13 provides an intuitive interpretation of the last derived bound.