

University of New Mexico

UNM Digital Repository

Electrical and Computer Engineering ETDs

Engineering ETDs

Fall 12-15-2023

PUF-Based Digital Money with Propagation-of-Provenance and Offline Transfers Between Two Parties

Benjamin Bean

University of New Mexico

Cyrus Minwalla

Bank of Canada

Eirini Eleni Tsiropoulou

University of New Mexico

Jim Plusquellic

University of New Mexico

Follow this and additional works at: https://digitalrepository.unm.edu/ece_etds



Part of the [Electrical and Computer Engineering Commons](#)

Recommended Citation

Bean, Benjamin; Cyrus Minwalla; Eirini Eleni Tsiropoulou; and Jim Plusquellic. "PUF-Based Digital Money with Propagation-of-Provenance and Offline Transfers Between Two Parties." (2023).

https://digitalrepository.unm.edu/ece_etds/622

This Thesis is brought to you for free and open access by the Engineering ETDs at UNM Digital Repository. It has been accepted for inclusion in Electrical and Computer Engineering ETDs by an authorized administrator of UNM Digital Repository. For more information, please contact disc@unm.edu.

Benjamin Gabriel Bean

Candidate

Electrical and Computer Engineering

Department

This thesis is approved, and it is acceptable in quality and form for publication:

Approved by the Thesis Committee:

Dr. Jim Plusquellic, Chair

Dr. Eirini Eleni Tsiropoulou, Member

Dr. Cyrus Minwalla, Member

PUF-Based Digital Money with Propagation-of-Provenance and Offline Transfers Between Two Parties

by

Benjamin Gabriel Bean

B.S., Computer Science and Engineering, New Mexico Institute for
Mining and Technology, NM, 2014

THESIS

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Science
Computer Engineering

The University of New Mexico

Albuquerque, New Mexico

December, 2023

Dedication

To my wife Carly, who is always there to help me find my way.

Acknowledgments

I would like to thank the many amazing instructors that helped me get here, especially Dr. Plusquellic and Dr. Tsiropoulou, whose guidance and friendship are a beacon to the students in our department.

PUF-Based Digital Money with Propagation-of-Provenance and Offline Transfers Between Two Parties

by

Benjamin Gabriel Bean

B.S., Computer Science and Engineering, New Mexico Institute for Mining and Technology, NM, 2014

M.S., Computer Engineering, University of New Mexico, 2023

Abstract

Building on prior concepts of electronic money (eCash), we introduce a digital currency where a physical unclonable function (PUF) engenders devices with the twin properties of being verifiably enrolled as a member of a legitimate set of eCash devices and possessing of a hardware-based root-of-trust. A hardware-obfuscated secure enclave (HOSE) is proposed as a means of enabling a PUF-based propagation-of-provenance (POP) mechanism, which allows eCash tokens (eCt) to be securely signed and validated by recipients without incurring any third party dependencies at transfer time. The POP scheme establishes a chain of custody starting with token creation, extending through multiple bilateral in-field transactions, and culminating in redemption at the token-issuing authority. A light-weight mutual-zero-trust (MZT) authentication protocol establishes a secure channel between any two fielded devices. The POP and MZT protocols, in combination with the HOSE, enables transitivity and anonymity of eCt transfers between online and offline devices.

Contents

List of Figures	ix
List of Tables	x
Glossary	xi
1 Introduction	1
2 Literature Review	4
2.1 Mutual Authentication and Session Key Exchange	4
2.2 Electronic Money	5
3 Hardware Security Primitives	7
3.1 SiRF PUF Architecture	7
3.2 Challenge and Response Construction	8
3.3 Strong Timing-Based Authentication and Key Generation Protocol Primitives	10

Contents

4	Mutual-Zero-Trust Protocol	12
4.0.1	MZT Enrollment Operations	13
4.0.2	MZT In-Field Operations	15
5	PUF-Cash Protocol	19
5.0.1	PUF-Cash Overview	20
5.0.2	Propagation of Provenance (POP) Qualities	22
5.0.3	PUF-Cash Protocol Details	23
6	Experimental Setup	35
6.0.1	Device Resource Utilization	35
6.0.2	Database Size Overhead	37
7	Experiment Results	39
7.0.1	Run Time Analysis	39
7.0.2	POP LLK Statistical Analysis	41
8	Security Analysis	43
8.0.1	CRP Space Analysis	44
8.0.2	Model-Building Attack Analysis	44
8.0.3	Protocol Security Properties	47
8.0.4	Ecosystem Attacks	48

Contents

9 Conclusion and Future Work 50

References 52

List of Figures

3.1	SiRF PUF overview	8
4.1	MZT enrollment and authentication process	17
4.2	MZT in-field authentication and SKG	18
5.1	PUF-Cash overview	20
5.2	POP databases and vector generation	22
5.3	PUF-Cash bootstrap operation	24
5.4	PUF-Cash withdrawal operation	27
5.5	PUF-Cash transfer operation	30
5.6	PUF-Cash deposit operation	32
6.1	Experiment set for the evaluation of PUF-Cash.	35
7.1	Transaction run times and NIST test results	40
8.1	Adversarial attack interface	44

List of Tables

6.1	PL-side resource utilization on the Zynq 7010.	36
6.2	Database Size Overhead	37
7.1	SiRF PUF Primitive Run Time Analysis	40
7.2	PUF-Cash Primitive Run Time Analysis	41

Glossary

Chlng or C	Challenge: Input values such as v, SF, p, HD which are used with a PUF to (re)produce or (re)generate a bitstring.
CRP	Challenge-Response-Pair: The Chlng values and expected response for a given PUF.
CD	Customer Device: A HPUF that allows for eCt transactions and can be deployed to the field.
DV	Delay Value: The time it takes for a rising or falling edge to propagate in the SiRF-PUF.
DVD	Delay Value Difference: The delta between a pairs of DVs.
eCt	Electronic Cash Token: The smallest discrete unit of currency.
FI	Financial Institution: One of several banks that manages customer accounts and communicates with the TI.
HPUF	Hardware PUF: Implementation of a PUF in hardware such as an FPGA or ASIC.
HOSE	Hardware-Oriented Secure Enclave: The trusted execution environment on a CD.

Glossary

HD	Helper Data: The bitstring used to categorize weak and strong bits for reproducing exact responses.
LLK	Long-Lived Key: A key that can be regenerated by a SiRF-PUF from locally stored information, for use in the HOSE.
MA	Mutual Authentication: Two party authentication, where each party authenticates the other.
MZT	Mutual-Zero-Trust MA and SKG: Non-anonymous MA and SKG between parties.
n	Nonce: A random number.
p	Parameters: Controls the DVD in the SiRF-PUF.
POP	Propagation-of-Provenance: The extension of the TI's providence
SK	Session Key: The key that is used to encrypt messages using AES.
SKG	Session Key Generation: Creation of a shared session key for encrypted data exchange.
SiRF	Shift-Register Reconvergent-Fanout: The type of Physically Unclonable Function (PUF) which we use for PUF-Cash.
SPUF	Software PUF: Emulation of a PUF in software using the HPUF's timings.
SF	SpreadFactors: Information to correct population and individual statistical DV differences.
TB	Timing-Based MA and SKG: Anonymous MA and SKG between a HPUF and SPUF.

Glossary

TI	Token Issuer: The bank that enrolls devices, authenticates devices, and generates tokens.
ZHK	Zero-Trust Secure Hash Long-Lived Key: The authentication token used in MZT.

Chapter 1

Introduction

An electronic money (eCash) ecosystem defines a set of protocols and security properties for enabling the creation and validation of electronic cash tokens (**eCt**), as well as secure transfers between end-users. A typical incarnation consists of a token-issuer (TI), e.g., a central bank, a set of financial institutions (FI), e.g., commercial banks, and a consumer population. The TI defines the root-of-trust for the entire ecosystem, while the FIs provide accounting services to subsets of the consumer population.

Electronic money presents unique challenges not commonly associated with security protocols providing, e.g., a secure communication channel. An inherent property of paper money is anonymity, which hides the identity of the purchaser of goods and services. Although anonymity is a desirable feature in eCash, it also substantially reduces the barrier for adversaries to create counterfeits. By contrast, fiat money is naturally resistant to counterfeiting, owing to the cost of the paper and inks, but copying digital representations of **eCt** is trivial and nearly cost-free. Similarly, dishonest users who engage in the double spending, i.e., use the same **eCt** to purchase goods and services from different vendors, is difficult to prevent, particularly for offline value transfer operations where users are not able to consult with a trusted-third

Chapter 1. Introduction

party (TTP).

Supplemental to these primary goals of an eCash system is a mechanism for identifying malicious actors, who might engage in attacks on other devices or who may collude to launder money. Yet another desirable property is a mechanism to recover lost funds, which would occur if the device is lost or stolen, or is destroyed. The implementation of these supplemental services requires disclosure of the user's identity, and therefore, eCash protocols must resolve these issues through coordination across multiple TTPs, as a means of preventing abuse by the TTP themselves.

In this paper, we propose an eCash ecosystem, called PUF-Cash, that possesses the following characteristics:

- A token-issuer (TI) provides a PUF-based root-of-trust for the entire system, and is able to extend the root-of-trust to trusted intermediaries and in-field devices using a light weight mutual-zero-trust (MZT) protocol.
- End-user (in-field) devices incorporate a strong PUF hardware security primitive, encapsulated in hardware-obfuscated secure enclave (HOSE). The HOSE enables the device to perform security functions related to the management and validation of **eCt**. The HOSE is completely independent of the untrusted microprocessor environment which is typically co-located on the same SoC device.
- A propagation-of-provenance (POP) scheme is proposed which enables transactions between entities receiving **eCt** to authenticate the **eCt** of the issuing parties. Moreover, the receiving parties can perform this authentication without the involvement of a TTP. The entire sequence of such authentications establishes provenance back to the entity which created the **eCt**, i.e., the TI.
- The proposed PUF-Cash system supports unlimited transitivity where Alice can pay Bob, and Bob can pay Charlie without the need for any party to interact with a TTP.

Chapter 1. Introduction

- The proposed PUF-Cash protocol supports forensics and recovery of loss of funds. Here, a set of cooperating TTPs can reveal the identity of a malicious actors and bar them from further interactions, as well as act on behalf of the legitimate customers to recover lost funds.

The remainder of this paper is organized as follows. Chapter 2 describes previous related work, and Chapter 3 summarizes relevant features of the shift-register reconvergent-fanout (SiRF) PUF. Chapter 4 describes the proposed mutual-zero-trust authentication and key exchange protocol, while Chapter 5 presents the details of the PUF-Cash message exchange protocol. The experiment setup is described in Chapter 6, with experimental results given in Chapter 7. A security analysis is presented in Chapter 8, and our conclusions in Chapter 9.

Chapter 2

Literature Review

2.1 Mutual Authentication and Session Key Exchange

Recent efforts around authentication utilize PUFs as a local root of trust to augment or substitute asymmetric key pairs. An extensive literature review on previously proposed light-weight PUF-based authentication protocols is given in [1]. The authors of [2] propose a lightweight PUF-based mutual authentication and secret message exchange protocol. The protocol only succeeds in authenticating the server, and further, assumes that the router has a soft model of the PUF and can generate a response to any randomly generated challenge during the refresh phase. Mahalat et al. [3] propose a scheme for secure WiFi authentication of IoT devices. This approach was later expanded to a PUF-based authentication and key sharing scheme that utilizes Pedersens commitment scheme coupled with Shamirs secret sharing [4]. The mutual authentication and key sharing scheme proposed between user (server) and sink nodes can be implemented more easily using challenge-response-pair (CRP) strong PUF-based schemes without the mathematical complexity of the secret shar-

ing schemes [5]. This becomes possible since in both cases the server stores a CRP database that is constructed in a secure environment during provisioning.

A PUF-based authentication and key management protocol for IoT is proposed in [6], improving upon on the attack resilience and performance overhead of the previous method [7]. Elliptic curve cryptography (ECC) is used to create shared keys among IoT nodes with the assistance of a verifier. The scheme requires a trusted setup and tamper-resistant hardware to protect secret keys. Similarly, a controlled PUF that utilizes ECC is proposed as a lightweight authentication and key generation protocol for IoT nodes in [8], relying on zero knowledge proofs for device authentication, however the authentication is one-way, and the server is not authenticated. A PUF-based El-Gamal algorithm is proposed for message encryption as well as a PUF-based digital signature scheme.

In parallel, Yu et al. [9] propose a scheme that is designed to prevent an adversary from obtaining sufficient CRPs to carry out model-building attacks. However, the number of authentications is constrained by the number of CRPs stored in the database, requiring either reuse of entries or re-enrollment of the PUF, a caveat we avoid in our scheme. In [10], the authors propose a crossover ring oscillator (R)O PUF cloning technique that enables a group of IoT devices to all generate the same (shared) key, thereby eliminating the key distribution problem for devices engaging in multi-party shared key encrypted communication.

2.2 Electronic Money

Electronic money and digital currency are synonymous concepts in literature. Viable solutions satisfy the properties of provenance, protection against double-spend attacks, and privacy in payments. Chaum, Fiat, and Naor (CFN) explored anonymous payments through blind signatures [11], [12]. CFN is unsuitable for offline transfers,

Chapter 2. Literature Review

as funds need to be validated at the time of transaction. More recently, blind signatures are incorporated in the GNU Taler system as a privacy-preserving primitive for electronic money [13].

Cash-like anonymity and privacy are defined in [14] as the ability to exchange money without tracing the exchange and without revealing the identities of the payer or payee. In their 2021 paper, they also note that there exist no token-based eCash systems that have all the same peer-to-peer (P2P), offline, and anonymity capabilities as account-based systems, and that offline double spending cannot yet be prevented.

Up to this point, other offline eCash systems [15] [16] [17] [18] have tried to solve the double-spending problem by tracing eCash in bank and wallet records. The authors of [19] address double-spending by executing transactions in an ARM processor's Trusted Execution Environment (TEE) secured by an SRAM PUF. Similarly, the authors of [20] propose a P2P offline payment scheme that prevents double-spending, preserves privacy and anonymity and is reasonably low-powered. Their scheme utilizes an mobile phone's TEE to execute a one-time-program within a sandbox, that executes on a single input and then self-destructs.

Chapter 3

Hardware Security Primitives

The authentication bitstrings, encryption keys and random nonces utilized within the proposed MZT and PUF-Cash protocols are derived from a strong physical unclonable function (PUF) called the shift-register reconvergent-fanout (SiRF) PUF [21]. The SiRF PUF is the physical root-of-trust in the PUF-Cash ecosystem, across TI, the FI and end-user devices.

3.1 SiRF PUF Architecture

The SiRF PUF architecture and algorithm are shown on the left and right sides of Fig. 3.1, respectively. The SiRF PUF utilizes within-die variations in a set of path delays as a source of entropy. Path delays are measured through an engineered netlist of shift-registers and logic gates constructed with fan-in and fan-out, referred to as reconvergent-fanout, to create an exponentially diverse matrix of signal paths that traverse a rectangular region of the FPGA fabric (22x23 CLBs). The architecture incorporates a mode switch to enable the entropy source and algorithm to be used as a true-random number generator (TRNG), which is able to supply an exponentially

large number of random nonces for cryptographic operations [22].




Figure 3.1: SiRF PUF Architecture and Algorithm.

- | | |
|--|---|
| AT_{DB} : Anonymous timing DB | NAT_{DB} : Non-anon. timing DB |
| DV : Delay values | DVD : DV differences |
| DVD_c : Calibrated DV | DVD_{co} : Entropy optimized DV |
| p : LFSR seed, range, thresh. | SF : SpreadFactors 2048 bytes |
| HD : helper data 2048 bits | v : seed specifying Chng vecs |
| $Chng_a\{v\}$: Set of challenge vectors | $Chng_b\{p, SF\}$: One set of p and SF |
| ID_x : Device ID | SK : Session key |
| LLK : Long-lived key | |

3.2 Challenge and Response Construction

The challenge for the SiRF PUF consists of several components, annotated as $Chng_a$ and $Chng_b$ in Fig. 3.1. The v argument to $Chng_a$ designates a vector of bitstrings, which control the configuration of the paths through the netlist of shift-registers and logic gates [21]. Paths are timed by the *Control module*, which launches a set of rising transitions into the SiRF netlist using the **Launch FFs**. The *Control module* selects a path output using the *MUX* shown along the bottom left in the figure, which routes emerging signal transitions to a time-to-digital converter (TDC). The TDC creates high resolution digitized representations (*delay values* or DVs) of the path delays. The *Control module* carries out a sequence of path timing operations using a sequence of v vectors to produce a set of 2048 rising delay values DV_R and a

Chapter 3. Hardware Security Primitives

set of 2048 falling delay values \mathbf{DV}_F . The \mathbf{DV} are stored in a BRAM located within the programmable logic and are used as inputs to a post-processing algorithm used to generate authentication bitstrings and keys.

The LFSR seed argument, \mathbf{v} , of \mathbf{Chng}_a is used to select a sequence of binary vectors, which are extracted from a $Vecs_{DB}$ database stored locally on the device (not shown here but included in the message exchange diagrams in this paper). As described later, the 32-bit LFSR seed is negotiated between the client and server, and serves a dual purpose of reducing the amount of data exchanged between client and server, and additionally obscures the sequence of challenge vectors extracted and applied during authentication from any adversaries observing the exchange.

The remaining components of the challenges are given as \mathbf{p} , \mathbf{SF} and \mathbf{HD} . The \mathbf{p} component specifies parameters to the SiRF PUF algorithm, which are used as input to a sequence of mathematical operations applied to the 4096 \mathbf{DV} values stored in the BRAM. The first operation, carried out by the *DVDiff module*, creates 2048 differences (\mathbf{DVD}) by subtracting unique pairings of \mathbf{DV}_F from \mathbf{DV}_R . Note that there are 2^{22} or 4 million unique \mathbf{DVD} that can be generated from the sets of rise and fall \mathbf{DV} . The LFSR seed parameter component of \mathbf{p} specifies a single unique set of 2048 \mathbf{DVD} to use in subsequent steps, described in detail in [21].

The \mathbf{DVD} are then calibrated using the *GPEVCal module* to remove variations in delay introduced by global performance differences, as well as non-nominal temperature and supply voltage environmental conditions, with the output being calibrated \mathbf{DVD} values, labeled \mathbf{DVD}_c . A third operation carried out by the *SpreadFactors module* removes delay bias introduced by differences in the length of the tested paths, which is accomplished through the application of SpreadFactors, (\mathbf{SF}), to produce \mathbf{DVD}_{co} . These are then converted into an authentication bitstring or key by the *BitGen module*. The \mathbf{HD} component refers to helper data that is needed by the *BitGen module* for regeneration. The components of the challenge, including the \mathbf{SF}

and the helper data \mathbf{HD} produced during enrollment, are stored in a database to enable the device to regenerate the bitstring or key at any point in the future and potentially under adverse environmental conditions.

3.3 Strong Timing-Based Authentication and Key Generation Protocol Primitives

Mutual authentication (MA) and session key generation (SKG) is accomplished using one of two low-level PUF-based protocols. The timing-based (TB) version requires AT_{DB} and NAT_{DB} timing databases stored on the Token Issuer (TI) (see Fig. 3.1). The timing databases encapsulate and compress a subset of the challenge-response space of each device, and are used exclusively by the Token Issuer (TI) to carry out MA and SKA with other entities in PUF-Cash ecosystem.

The TB functions utilized by the TI are annotated in the diagrams as MA_{NA} or MA_A (for mutual authentication) and SKG_{NA} or SKG_A (for session key generation). The subscript NA refers to non-anonymous mutual authentication, where the authenticating device ID, e.g., ID_x , is derived privately by the TI using the NAT_{DB} . The subscript A refers to an anonymous mutual authentication, where the TI is able to confirm that the device has been provisioned, but is not able to identify the end-user's identity.

The AT_{DB} and NAT_{DB} are built with distinct \mathbf{DV} , and the ordering of the device timing data sets in the two databases is scrambled to generate anonymity. The SiRF PUF has 16 million unique paths that can be timed, making it possible to select large numbers of unique \mathbf{DV} per database. For the current experiments, the number of \mathbf{DV} per device stored in the databases is limited to 10,144 (5072 \mathbf{DV}_R and 5072 \mathbf{DV}_F).

Chapter 3. Hardware Security Primitives

Lighter-weight versions of MA and SKG (referred to as mutual-zero-trust or MZT) are annotated in the diagrams as MA_{MZT} and SKG_{MZT} . The MZT versions are based on authentication tokens, and are used between pairs of devices and between a customer device and the FI.

Additional PUF-based security functions leveraged in the proposed MZT and PUF-Cash protocols include true-random-number-generation (TRNG) and long-lived key (**LLK**) generation. As key generation requires reliable reproduction of bitstrings, the functions leverage three reliability enhancing techniques integrated into the SiRF PUF algorithm, namely *GPEVCal*, thresholding and XMR, details of which are available in [23] [24].

Chapter 4

Mutual-Zero-Trust Protocol

An eCash payment system should support trusted payment transactions in any type of environment, including scenarios in which the payer and payee cannot consult with a trusted third party (TTP) to assist with authentication i.e. offline mode. In such cases, trust within the two-party system must be derived from the devices themselves. We use the term mutual-zero-trust (MZT) to describe this scenario, in contrast to environments in which peers (other customers) can be consulted to build a trusted relationship, or a TTP is available to provide authentication credentials for the two entities.

Supporting transactions which might occur between arbitrary pairs of customer devices in an offline context, where neither device has connectivity to a remote service, means that each device must store MZT data about other customer devices. The data composition must be compact to be practical for an embedded system, while simultaneously providing each device with sufficient confidence that the established channel and the counter-party device can be trusted. In this section, we describe a MZT system that meets these requirements.

4.0.1 MZT Enrollment Operations

The security properties of the MZT protocol are derived from the SiRF PUF primitive, a secure hash function and a symmetric encryption algorithm. The PUF serves as the root-of-trust and is the source of entropy, secure hash provides obfuscation and data integrity while the symmetric encryption algorithm provides confidentiality. The integration of the three primitives provides a highly secure and lightweight mechanism for enabling Alice and Bob to communicate sensitive information.

The protocol requires the SiRF PUF to generate a 256-bit long-lived key (**LLK**) and a 256-bit nonce **n**. These bitstrings are used as input to a hash function to create an authentication token (AT), referred to as **ZHK**, with *Z* referring to zero-trust, *H* for secure hash and *K* for **LLK**. In our implementation, SHA-3 is used as the secure hash function. The **ZHK** authentication tokens are created using the relation given by Eq. 4.1.

$$\mathbf{ZHK} := \text{Hash}(\mathbf{LLK} \oplus \mathbf{n}) \tag{4.1}$$

The MZT protocol requires Alice and Bob to carry out an enrollment operation with the Token Issuer (TI). Although shown together here, enrollment is carried out separately and usually at different times by Alice and Bob.

The following sequence of operations, corresponding to the message exchange diagram of Fig. 4.1, defines MZT enrollment.

1. Alice and Bob authenticate non-anonymously via MA_{NA} and generate session keys \mathbf{SK}_{TA} and \mathbf{SK}_{TB} using SKG_{NA} with the TI. As discussed in Section 3.3, Alice and Bob use the TB versions of these security functions since they are interacting with the TI. Non-anonymous authentication allows the TI to identify Alice and Bob as ID_A and ID_B , respectively. Note that the MA_{NA} protocol is privacy-preserving, i.e., the TI derives the IDs based on the re-

Chapter 4. Mutual-Zero-Trust Protocol

sponses provided by Alice and Bob.

2. The TI encrypts and then transmits unique challenges to Alice and Bob, labeled $\mathbf{Chng}_{ZTa/b}$.
3. Alice and Bob decrypt and apply their respective challenges to their hardware PUFs, HPUF_E , in enrollment mode, to generate a long-lived key, $\mathbf{LLK}_{\text{MZT}}$, and helper data \mathbf{HD} . Alice and Bob store the challenge information in the MZT_LLK_{DB} under challenge number CN_1 , which includes the components discussed earlier in reference to Fig. 3.1, i.e., \mathbf{v}_1 , \mathbf{p}_1 , \mathbf{SF}_1 and \mathbf{HD}_1 , to enable regeneration of $\mathbf{LLK}_{\text{MZT}}$ later in the field. As noted above, \mathbf{v}_1 is a 32-bit LFSR seed used to select binary vectors from the Vec_{DB} stored locally on Alice and Bob's devices.
4. Once the $\mathbf{LLK}_{\text{MZT}}$ is generated, the TI sends a request to Alice and Bob to generate x \mathbf{ZHK} s.
5. Alice and Bob construct a set of tuples $\{\mathbf{ZHK}_i, \mathbf{n}_i\}$ by running their PUFs' TRNGs to generate a sequence of nonces, \mathbf{n}_i , which are XORed with $\mathbf{LLK}_{\text{MZT}}$ and used as the input to the AT creation function given by Eq. 4.1.
6. The tuples are encrypted using $\text{SK}_{TA/B}$ and transmitted to the TI. The TI decrypts and stores them along with the field device's identifier $\text{ID}_{A/B}$ in its MZT_AT_{DB} .

The TI collects \mathbf{ZHK}_i from all field devices to build the MZT_AT_{DB} . Each entry consumes only 72 bytes, consisting of a 4-byte integer for the ID_x and CN_x fields, and 32 bytes for each of the \mathbf{ZHK}_i and \mathbf{n}_i fields. Alice and Bob will retrieve one unique \mathbf{ZHK}_i tuple from the TI for each device during the distribution operations, labeled as steps 7 through 10 along the bottom of Fig. 4.1, which they store in their MZT_AT_{DB} databases. Note that the protocol and database structure support multiple challenges, e.g., Bob stores \mathbf{ZHK}_i for a second challenge number, CN_2 , in his \mathbf{ZHK}_i . The MZT_AT_{DB} database can be stored in a standard off-the-shelf NVM, e.g., SD card, where storage capacities of 4 GB or larger are common. Elements

within the MZT_AT_{DB} could be encrypted for additional security, and processing may be carried out in the hardware-obfuscated secure enclave (HOSE) on the device.

4.0.2 MZT In-Field Operations

The MZT in-field process is carried out when Alice contacts Bob for goods or services in an environment where a direct communications channel is established between the two parties. The message exchange protocol is shown in Fig. 4.2, and is described as follows:

1. The transaction begins with Alice sending Bob a request to authenticate and to generate a shared session key.
2. Alice sends Bob an identifier, ID_A , that allows Bob to locate the corresponding AT in his MZT_AT_{DB} .
3. Bob responds to Alice with an ACK or NAK (labeled *status*) on whether or not he possesses an AT for Alice. He also transmits his own identifier, ID_B .
4. Alice determines if she has an AT for Bob in her MZT_AT_{DB} using Bob's ID_B , and transmits a corresponding ACK or NAK to Bob in the *status* message.
5. Assuming both Alice and Bob have ATs for each other (otherwise the transaction is cancelled), Alice and Bob retrieve the AT for the other party from their MZT_AT_{DB} , which is represented by the tuple $\mathbf{ZHK}_x, \mathbf{n}_x$ with $x := b$ or a , respectively.
6. **Shared Key Generation:** Alice and Bob exchange the nonce components, \mathbf{n}_x , of the ATs. Both parties regenerate their long-lived keys $\mathbf{LLK}_{\text{MZTx}}$ using challenge information stored in their MZT_LLK_{DB} (not shown), and then compute a local version of the \mathbf{ZHK}'_x using $\text{Hash}(\mathbf{LLK}_{\text{MZTx}} \oplus \mathbf{n}_x)$. Alice and Bob create a shared key \mathbf{SK}_{AB} by XOR'ing the local copy of \mathbf{ZHK}'_x with the \mathbf{ZHK}_x that they store for the other party in their MZT_AT_{DB} .

7. **Authentication:** Authentication begins with Alice and Bob encrypting the \mathbf{n}_x they received from the other party with the newly created shared key \mathbf{SK}_{AB} to create \mathbf{en}_x . Alice and Bob exchange the encrypted nonces \mathbf{en}_x . They decrypt the \mathbf{en}_x using the shared key and then compare the decrypted \mathbf{n}_x with the ones they store in their MZT_AT_{DB} . The status of the comparison is shared with the other party with each transmitting an ACK or NAK. At this point, they have authenticated and possess a shared key assuming both have acknowledged that the nonces \mathbf{n}_x match their own local copies.
8. **Refresh ZHKs:** Alice and Bob generate new nonces \mathbf{n}_{x_n} by running their TRNGs, and then compute new AT by hashing $(\mathbf{LLK}_{\text{MZT}_x} \oplus \mathbf{n}_{x_n})$. They encrypt the new AT as \mathbf{C}_x with their shared session keys \mathbf{SK}_{AB} . They exchange the \mathbf{C}_x , and then decrypt the \mathbf{C}_x to recover the new AT. They store the new AT in their MZT_AT_{DB} , replacing the existing AT used in this transaction.

The MZT in-field operations are light-weight using only the PUF, secure hash, and symmetric encryption functions. The refresh operation ensures that future transactions can occur between the two parties without either party needing to return to the TI to obtain additional AT. The new AT are generated by Alice and Bob's PUFs and therefore have the same strong security properties as the original AT that are replaced.

Chapter 5

PUF-Cash Protocol

The operating environment used in the PUF-Cash protocol described here is characterized as offline, requiring Alice and Bob to authenticate using the MZT model. By assuming a bilateral, peer-to-peer trust model, we ensure that two parties can transact regardless of the underlying communications medium.

Additionally, transitivity of **eCt** tokens across multiple parties is supported, where transitivity means that a token can be transferred to a new party without requiring synchronization with the token issuer (TI) or any other back-end system. PUF-Cash introduces a novel security primitive called propagation-of-provenance (POP) and a hardware-obfuscated secure enclave (HOSE) as a means of ensuring transitivity in a secure manner. The HOSE prevents Alice from double spending her own **eCt**. POP and HOSE together enable each party in a chain of **eCt** transfers to authenticate the **eCt** that they receive, and to validate provenance back to the point of origin, namely the TI.

The HOSE is implemented as a set of state machines embedded in a hard-wired portion of the device, or, in the case of FPGAs, in the programmable logic. The HOSE limits interactions and potential attacks from malicious software applications,

including the PUF-Cash software components, through an interface consisting of two 32-bit hardware registers. The HOSE incorporates the SiRF PUF for generating keys on-the-fly as needed. Similarly, all sensitive cryptographic operations, specifically AES encryption/decryption and the SHA-3 hash function, are instantiated in the HOSE.

5.0.1 PUF-Cash Overview




Figure 5.1: High level overview of PUF-Cash

A high-level model of the PUF-Cash protocol is presented in Fig. 5.1 to highlight the basic operations. The sequence of numbered operations are described as follows:

1. The bootstrap operation within the PUF-Cash protocol generates and distributes a set of anonymous POP cryptographic tuples (\mathbf{POP}_x) to customer devices, Alice, Bob, Charlie, etc. The set of \mathbf{POP}_x are generated as customer long-lived keys using the TI's AT_{DB} , and are therefore anonymous to the TI.
2. Alice contacts her FI and requests a withdrawal from her account, sending the tuple $\{Amt, ID_A, AID_A\}$, which includes both her non-anonymous ID for the FI and her anonymous AID for the TI. The FI checks her balance and

Chapter 5. PUF-Cash Protocol

responds with an ACK (not shown) if she has sufficient funds or a NAK if she does not.

3. Assuming Alice has sufficient funds, the FI forwards the Amt and AID_A to the TI to generate the eCash tokens (eCt_A). The TI fetches Alice's anonymous POP_A from the POP_{DB} using her AID_A , and adds a withdrawal record to its eCt_{DB} database.
4. The TI generates the eCt_A using a TRNG, and creates signed versions, $heCt_A$, using Alice's POP_A in a XOR-secure-hash operation similar to Eq. 4.1. The TI then encrypts both of the eCt_A and $heCt_A$ as ECt_A using a shared session key that is created between Alice and the TI anonymously (not shown here but included in message exchange diagrams below), and transmits the ECt_A to the FI.
5. The FI simply forwards the ECt_A to Alice. The FI, acting as an intermediary, keeps Alice and her eCt_A anonymous to the TI. Moreover, her ECt_A are also anonymous to the FI because they are encrypted with the Alice-TI session key.
6. Alice contacts Bob for a payment transaction, and then both run the MZT protocol.
7. Assuming authentication succeeds, Alice pays Bob with (a subset of) her ECt_A by encrypting them with the Alice-Bob MZT session key. Although not shown here, Alice and Bob use their POP_x to authenticate and propagate the provenance of the eCt . The MZT security functions and ECt_x transfer operation can be repeated between other pairs of customers, shown here between Bob and Charlie.
8. At some point in the future, Charlie regains internet connectivity and deposits the ECt_C to his FI. Note that although Alice and Charlie use the same FI as shown, this is not a requirement.
9. The FI contacts the TI and asks the TI to validate the ECt_C , as a precursor

to the FI accepting the \mathbf{ECt}_C as a valid deposit.

10. Assuming validation succeeds, the FI credits Charlie’s account, and although not shown, the \mathbf{ECt}_B are marked as ‘redeemed’ in the TI’s eCt_{DB} .

5.0.2 Propagation of Provenance (POP) Qualities




Figure 5.2: Propagation-of-provenance (POP) database storage and challenge vector generation scheme.

The TI uses a specialized process to create the entire set of \mathbf{POP}_x for each of the customers’ \mathbf{POP}_{DB} (distributed in Step 1 of Fig. 5.1) that significantly reduces the storage requirements for the challenges. The \mathbf{POP}_x transmitted to Alice and Bob’s device is composed of two components, a population-based component (PBC) and a device-specific component (DSC), each stored in two different databases labeled \mathbf{POPA}_{DB} and \mathbf{POPB}_{DB} in Fig. 5.2. The PBC component is used for all devices in Alice’s \mathbf{POPB}_{DB} , and is defined as the tuple $(\mathbf{v}_x, \mathbf{SF}_x)$. The \mathbf{v}_x component is a 32-bit LFSR seed that selects the actual vectors \mathbf{vecs}_x from a separate database called \mathbf{Vec}_{DB} while \mathbf{SF}_x represents a set of SpreadFactors for optimizing POP key generation (from Chapter 3).

The DSC components are stored in the \mathbf{POPB}_{DB} , one element for each device, and are defined as a tuple $(\mathbf{AID}_x, \mathbf{p}_x, \mathbf{ePOP}_x, \mathbf{HD}_x)$. The \mathbf{AID}_x is a 32-bit integer representing a device’s anonymous ID, the \mathbf{p}_x component is a 22-bit nonce used to specify the seeds to the LFSRs in the *DVDiff*s module from Chapter 3, the \mathbf{ePOP}_x is the encrypted PUF’s response to the challenge for device x (256-bits), and \mathbf{HD}_x is the helper data (2048-bits). Therefore, the size of each \mathbf{POPB}_{DB} element is 296 bytes, and is larger than the $\mathbf{MZT_AT}_{DB}$ which requires only 74 bytes per device

entry.

The primary benefit of the larger POP database is related to the strength of the authentication process. The challenge information stored by Alice for, e.g., Bob's device, requires Bob's device to generate the response to Alice's stored challenge, and the challenge that Alice's stores (at least initially) has not been exposed a priori to Bob's device. This is true because the TI provides Alice with Bob's response to this challenge (as the $\mathbf{ePOP_B}$ component) by running a 'soft PUF' version of the SiRF PUF algorithm using the anonymous timing data stored in the AT_{DB} . The soft PUF is able to produce the same response that Bob's device generates for this challenge.

Moreover, Alice is able to refresh Bob's entry in her DB after every engagement with Bob in two different ways. In one scenario, Alice generates a new challenge for Bob's device by changing the $\mathbf{p_B}$ component (LFSR seeds) in POP_{DB} , and then asks Bob to generate a new response. She then replaces the DSC component of Bob's entry in her POP_{DB} with the new information. In the second scenario, Alice authenticates with the TI and asks the TI to perform this operation. Although this requires Alice to be online, it prevents Bob from seeing Alice's next challenge for him.

5.0.3 PUF-Cash Protocol Details

The message exchange diagrams for the four PUF-Cash operations are described in this section. The BootStrap operation is carried out after device manufacture, and at any point in the field under the condition that Alice is online, i.e., has internet connectivity. Once Alice has engaged the TI in a bootstrap operation, she can perform any one of the three primary PUF-Cash operations, Withdrawal, Transfer and Deposit (note, her very first transaction must be a Withdrawal).

All transactions with TI are preceded with mutual authentication, e.g., MA_{NA} , and session key generation, e.g., SKG_{NA} . Session key generation produces a shared key SK_x , where x is replaced with the initials of the authenticating parties, e.g., SK_{TF} refers to the shared session key between the TI and FI. Although the details of TB security functions are presented in previous work [25], the enrollment and regeneration functions for the POP_{DB} are similar, and are outlined in Chapter 8 for completeness.

BootStrap




Figure 5.3: PUF-Cash bootstrap operations between the Token Issuer and Alice.

MA_A : Anonymous mutual authen.	SKG_A : Session key generation
AID_x : Anonymous ID	SK_{TA} : TI-Alice, etc. 256-bit session keys
v_x : 32-bit vector specifier	vec_{sA} : 771-bit Chlng vecs
p_x : 22-bit LFSR seed	SF_x : 2048 bytes of SpreadFactors
DV : Set of 16-bit DV	LLK_{hose} : HOSE 256-bit LLK
POP_x : POP LLK	$ePOP_x$: HOSE encrypted POP LLK
HD_x : 2048-bit Helper data	C_x : Cipher text packet

The message exchange diagram for Bootstrap is shown in Fig. 5.3. The sequence of operations performed by Alice and the TI are numbered within circles in the

figure. The shaded regions identify operations carried out within the HOSE, i.e., the programmable logic region of the FPGA.

1. **Generate $\mathbf{LLK}_{\text{hose}}$** : Alice and the Token Issuer (TI) mutually authenticate and generate a shared session key, \mathbf{SK}_{TA} . Mutual authentication is done anonymously. Upon successful authentication, the TI affirms that Alice's device is a genuine SiRF-instantiated PUF device with anonymous ID, AID_A .
2. The TI generates random values for \mathbf{v}_{hose} and \mathbf{p}_{hose} to enable the TI and Alice to enroll and regenerate, respectively, a long-lived key, $\mathbf{LLK}_{\text{hose}}$. The $\mathbf{LLK}_{\text{hose}}$ will be used by TI and Alice to encrypt and decrypt POP data. The term \mathbf{v}_{hose} is used as the seed to pseudo-randomly select vectors, e.g. \mathbf{vecs}_x from the Vec_{DB} , as shown in Fig. 5.2, and \mathbf{p}_{hose} specifies the LFSR seed for the SiRF PUF *DVDiff* module. The anonymous timing data, \mathbf{DV} , for a random subset of devices corresponding to paths timed by \mathbf{vecs}_x , is extracted from the AT_{DB} and used to generate a set of SpreadFactors, $\mathbf{SF}_{\text{hose}}$. The TI runs a software version of the SiRF PUF in enrollment mode, SPUF_E , which produces the $\mathbf{LLK}_{\text{hose}}$ key and the $\mathbf{HD}_{\text{hose}}$ helper data.
3. The TI encrypts the tuple $(\mathbf{v}_{\text{hose}}, \mathbf{SF}_{\text{hose}}, \mathbf{p}_{\text{hose}}, \mathbf{HD}_{\text{hose}})$ with its session key \mathbf{SK}_{TA} as \mathbf{C}_1 and transmits it to Alice. Alice decrypts \mathbf{C}_1 and inserts the challenge data into her HOSE_{DB} .
4. Alice reads the challenge information from her HOSE_{DB} and runs her hardware PUF in regeneration mode, HPUF_R , to produce $\mathbf{LLK}_{\text{hose}}$.
5. **Generate POP**: The TI generates a set of POP elements for a set of devices from the population by first generating random values for \mathbf{v}_A and \mathbf{p}_A and extracting a set of vectors \mathbf{vecs}_A . It then extracts the timing values, \mathbf{DV} , for a random subset of devices from AT_{DB} corresponding to \mathbf{vecs}_A . *GenSF* is called to generate the SpreadFactors, \mathbf{SF}_A^+ , for all possible DVD, that can be created from the selected 2048 DV_R and 2048 DV_F . From the discussion in Section 3.2, the DV_R and DV_F can be paired in 2^{22} possible ways to create

unique DVD. Therefore, the size of \mathbf{SF}_A^+ is 4 million bytes. This enables Alice to refresh her POP challenge and response information after a successful transfer operation with any device in her POP_{DB} .

6. The TI generates a set of 256-bit \mathbf{POP}_i responses to the challenge for a set of devices n , using the \mathbf{DV}_i corresponding to each of the devices i . Alice's response, \mathbf{POP}_A , is stored in the verifier's AT_{DB} for signing and validating \mathbf{eCt} during withdrawals and deposits, respectively.
7. The challenge components, \mathbf{v}_A and \mathbf{SF}_A^+ , and the tuple sets $\{\text{AID}_n, \mathbf{p}_A, \mathbf{POP}_n, \mathbf{HD}_n\}$ are encrypted with Alice's $\mathbf{LLK}_{\text{hose}}$ by the TI. The \mathbf{POP}_A component in Alice's tuple is set to null.
8. The TI sends the encrypted packet C_2 to Alice, which she decrypts within her HOSE. She stores the \mathbf{v}_A and \mathbf{SF}_A^+ in her POPA_{DB} . These components of the challenge are used for all customers she interacts with.
9. She then creates separate records for each of the customers i in the POP_{DB} database and stores the AID_i and \mathbf{HD}_i components. The \mathbf{p}_A component is replicated in each database record and will be used as the initial value during the first eCash transaction with a customer.
10. The \mathbf{POP}_i are re-encrypted with Alice's $\mathbf{LLK}_{\text{hose}}$ and stored as \mathbf{ePOP}_i . This protects the customer responses in the event an adversary gains access to Alice's device and attempts to extract information from the POP_{DB} . Note that Alice's \mathbf{ePOP}_i does not need to be stored (is null) because she can regenerate it with her PUF.

The size of the subset of \mathbf{DV} extracted from the AT_{DB} and used as input to GenSF in Step 5 needs to be large enough to characterize the entire population. In our implementation, we use sets of \mathbf{DV} corresponding to 30 customers. The subset of the customer population for which the TI creates \mathbf{POP}_i for Alice in Step 6 can be selected by Alice's device from preferences she specifies in advance, or learned over time from online eCash transactions she engages in.

Withdrawal

The following series of message exchanges and cryptographic operations are carried out between Alice, the FI and TI to enable Alice to obtain a set of anonymous eCt for payment transactions with end-users or commercial vendors. The corresponding message exchange diagram is shown in Fig. 5.4.




Figure 5.4: PUF-Cash_A message exchange for withdrawal operation between Alice, the Financial Institution and the Token Issuer.

1. The TI and FI mutually authenticate and generate a session key, SK_{TF} using the non-anonymous TB protocol primitives.
2. The FI and Alice mutually authenticate and generate a session key, SK_{FA} , using the MZT protocol primitives.

3. Alice creates a ciphertext packet, \mathbf{C}_1 , containing her anonymous ID called AID_A , and requested withdrawal amount, Amt , and sends the packet to FI.
4. FI decrypts \mathbf{C}_1 with \mathbf{SK}_{FA} , checks Alice's balance and responds to Alice with a ACK if she has enough funds in her account, or a NAK if she does not.
5. If the response is ACK, the transaction continues (otherwise it is cancelled) with FI encrypting AID_A and Amt as \mathbf{C}_2 and transmitting the packet to TI. TI decrypts \mathbf{C}_2 .
6. TI and Alice generate a session key \mathbf{SK}_{TA} using procedure SKG_A (which utilizes the anonymous timing database). The FI serves as a pass-through intermediary. The SKG_A process is nearly identical to the **Generate LLK_{hose}** (Step 2 of BootStrap) where the TI generates a challenge, runs $\mathbf{SK}_{\text{TA}} := \text{SPUF}_E$ and sends the challenge and helper data encrypted to FI (not shown). The FI forwards the packet to Alice, and Alice decrypts and regenerates \mathbf{SK}_{TA} by running her hardware PUF in regeneration mode. Note that Alice remains anonymous to TI because the TI draws the challenge and DV from AT_{DB} using her AID_A .
7. The TI extracts Alice's \mathbf{POP}_A from the AT_{DB} using her AID_A .
8. The TI generates a set of \mathbf{eCt}_A using its *GenNonce* function, equivalent to one 256-bit nonce for each 1 cent token of Alice's requested Amt , and records the \mathbf{eCt}_A in its eCt_{DB} along with her AID_A and \mathbf{POP}_A . The latter two components can be used for recovering lost funds and tracking malicious actors, but are otherwise not needed.
9. The TI creates signed versions of the \mathbf{eCt}_A , labeled as \mathbf{heCt}_A , by XORing each \mathbf{eCt}_A with \mathbf{POP}_A , and then hashing the result. The \mathbf{eCt}_A and \mathbf{heCt}_A are then encrypted with \mathbf{POP}_A to prevent observation or manipulation attacks from the processor-side. The encrypted versions, \mathbf{ECt}_A , are encrypted again with the TI-Alice session key, \mathbf{SK}_{TA} , in a packet \mathbf{C}_3 . The double encryption adds an additional layer of obscurity to the FI and to network packet

eavesdropping.

10. The TI transfers \mathbf{C}_3 through the FI to Alice, who decrypts \mathbf{C}_3 to recover the \mathbf{ECt}_A and passes them to the HOSE for processing.
11. Alice's HOSE fetches challenge information from the POP DBs as \mathbf{Chng}_{Aa} , and runs Alice's hardware PUF in regeneration mode, HPUF_R , to reproduce \mathbf{POP}_{Aa} . The subscript \mathbf{Aa} should be read as "Challenge for Alice from Alice". The \mathbf{ECt}_A are then decrypted and validated by the HOSE. Validation is accomplished by recreating the signatures, \mathbf{heCt}'_A , using the \mathbf{eCt}_A and the locally generated \mathbf{POP}_{Aa} , and comparing them with the received versions, \mathbf{heCt}_A .
12. If any of the signature comparisons fail, then the process is aborted. Otherwise, Alice updates her challenge to \mathbf{Chng}'_{Aa} by incrementing the \mathbf{p}_A component, and then running HPUF_E to generate a new \mathbf{POP}'_{Aa} . She updates her POPB_{DB} record with \mathbf{p}'_A and \mathbf{HD}'_A , and creates new signatures \mathbf{heCt}^*_A for the \mathbf{eCt}_A . She encrypts the \mathbf{eCt}_A and \mathbf{heCt}^*_A and adds the new \mathbf{ECt}'_A to her \mathbf{eCt}_{DB} . She also re-encrypts any existing elements with the new \mathbf{POP}'_{Aa} . This type of key rolling scheme reduces the likelihood of a successful differential power analysis attack carried out on Alice's device.
13. If the previous step succeeds, Alice encrypts the new \mathbf{POP}'_{Aa} with the original \mathbf{POP}_{Aa} as \mathbf{C}_4 , and then sends an ACK with the \mathbf{C}_4 to TI through the FI. Otherwise, she sends a NAK only.
14. If FI receives an ACK, it updates the Acct_{DB} by deducting the Amt from her account.
15. If TI receives an ACK, it decrypts \mathbf{C}_4 with \mathbf{POP}_A to recover \mathbf{POP}'_{Aa} . The TI updates Alice's POP in its AT_{DB} database with the new version. If a NAK is received, it removes the \mathbf{eCt}_A .

Transfer

The PUF-Cash protocol supports the transfer of \mathbf{eCt} between offline entities, as illustrated by the message exchange diagram for two devices labeled Alice and Bob in Fig. 5.5. The transfer protocol is carried out exclusively between Alice and Bob, and is transitive i.e., Bob can pay Charlie, etc.




Figure 5.5: PUF-Cash transfer operation between Alice and Bob.

1. Alice sends Bob a request to transfer Amt to Bob. Bob accepts the transfer request. Alice and Bob mutually authenticate and generate a session key, \mathbf{SK}_{AB} , using MZT protocol.
2. Bob extracts a challenge from his POP DBs, with subscript \mathbf{Ab} indicating "Challenge for Alice from Bob". Bob encrypts the challenge \mathbf{Chng}_{Ab} with his session key \mathbf{SK}_{AB} as C_1 and transmits it to Alice.
3. Alice decrypts \mathbf{Chng}_{Ab} and passes it into her HOSE. Alice's HOSE fetches \mathbf{Chng}_{Aa} and regenerates both \mathbf{POP}_{Aa} and \mathbf{POP}_{Ab} by running HPUF_R . Al-

ice then fetches a subset of her encrypted \mathbf{ECt}_A from her \mathbf{eCt}_{DB} and decrypts it with her regenerated \mathbf{POP}_{Aa} to recover the tuple $\{\mathbf{eCt}_{As}, \mathbf{heCt}_{As}\}$. Following that, she validates the database stored \mathbf{eCt}_{As} by creating \mathbf{heCt}'_{As} with \mathbf{POP}_{Aa} and compares them with the \mathbf{heCt}_{As} extracted from the database. The process is aborted if a mismatch is detected. If not, the \mathbf{eCt}_{As} subset is signed with \mathbf{POP}_{Ab} as \mathbf{heCt}^*_{As} , and the tuple, \mathbf{ECt}'_{As} , is encrypted with AES using \mathbf{POP}_{Ab} as the secret key. The re-signing of Alice's \mathbf{eCt} with a key that Bob stores and only Alice can generate is a key feature of the scheme.

4. Alice encrypts the \mathbf{ECt}'_{As} with \mathbf{SK}_{AB} as \mathbf{C}_2 and transmits it to Bob, who decrypts and transfers it to his HOSE.
5. Bob regenerates his \mathbf{LLK}_{hose} and \mathbf{POP}_{Bb} using challenges from the \mathbf{HOSE}_{DB} and POP DBs, respectively. He uses his \mathbf{LLK}_{hose} to decrypt Alice's \mathbf{ePOP}_A as \mathbf{POP}_{Ab} , and then uses \mathbf{POP}_{Ab} to decrypt \mathbf{ECt}'_{As} to recover Alice's \mathbf{eCt}_{As} and \mathbf{heCt}^*_{As} . He validates them by creating \mathbf{heCt}'_{As} and comparing with the \mathbf{heCt}^*_{As} received from Alice.
6. Bob records the result of the validation process as $\text{Resp} := \text{ACK}$ or NAK . Bob also updates Alice's challenge by randomly modifying the \mathbf{p}_A parameter in his \mathbf{POP}_{DB} to produce a new challenge, \mathbf{Chng}'_{Ab} . He encrypts the Resp and the new challenge with \mathbf{POP}_{Ab} as \mathbf{C}_3 and transmits it to Alice.
7. Alice transfers \mathbf{C}_3 into her HOSE, decrypts and aborts if the Resp is NAK . Similarly, Bob aborts if Resp is NAK . Otherwise, Alice removes the spent tuples $\{\mathbf{ECt}_{As}, \mathbf{heCt}_{As}\}$ from \mathbf{ECt}_A and stores the updated \mathbf{ECt}^*_A back to her \mathbf{eCt}_{DB} .
8. Alice generates a new response, \mathbf{POP}'_{Ab} , using the new challenge, \mathbf{Chng}'_{Ab} , encrypts it along with the new helper data \mathbf{HD}'_A using the original \mathbf{POP}_{Ab} as \mathbf{C}_4 and transmits it to Bob.
9. Bob hashes Alice's \mathbf{eCt}_{As} as \mathbf{heCt}''_{As} and encrypts them with the \mathbf{eCt}_{As} using his \mathbf{POP}_{Bb} as \mathbf{ECt}_B , which he stores to his \mathbf{eCt}_{DB} . He then decrypts

C_4 to recover Alice's POP'_{Ab} and HD'_A . Bob encrypts POP'_{Ab} with his LLK_{hose} as ePOP'_A , and updates his POP_{DB} with the new challenge (p'_A), encrypted response ePOP'_A and helper data HD'_A .

Deposit

The deposit operation involves the TI, FI and an end-user device, e.g., Charlie, where anonymity is preserved between the TI and Charlie. The message exchange sequence is shown in Fig. 5.6, and is described in the following numbered sequence.




Figure 5.6: PUF-Cash deposit operation between Charlie, the Financial Institution and the Token Issuer.

1. The TI and FI mutually authenticate and generate a session key, SK_{TF} , using the TB protocol.
2. The FI and Alice mutually authenticate and generate a session key, SK_{FA} , using MZT protocol.

3. Charlie fetches challenge information and runs his hardware PUF to regenerate $\mathbf{LLK}_{\text{hose}}$ and \mathbf{POP}_{C_c} . Bob selects a subset of his \mathbf{ECt}_C and then decrypts them using his \mathbf{POP}_{C_c} as \mathbf{eCt}_{C_s} and \mathbf{heCt}_{C_s} . Charlie then validates his \mathbf{eCt} and aborts if validation fails.
4. Charlie constructs a packet consisting of his AID_C , Amt and \mathbf{ECt}_{C_s} . The packet is encrypted with \mathbf{SK}_{FA} as \mathbf{C}_1 , and transmitted to the FI. The FI recovers the elements in \mathbf{C}_1 , re-encrypts them with \mathbf{SK}_{TF} as \mathbf{C}_2 and transmits them to the TI.
5. The TI decrypts \mathbf{C}_2 to recover AID_C , Amt , and the corresponding tokens \mathbf{ECt}_{C_s} . It then reads \mathbf{POP}_C from AT_{DB} using AID_C , and decrypts the \mathbf{ECt}_{C_s} using Charlie's \mathbf{POP}_C .
6. The TI searches for each of the \mathbf{eCt}_{C_s} in the master eCt_{DB} and checks if they are marked *valid*, i.e., still in circulation. If all \mathbf{eCt}_{C_s} are valid, then signatures are created using Charlie's \mathbf{POP}_C as \mathbf{heCt}_{A_s} . For convenience, it is assumed that some of the tokens in Charlie's possession were originally withdrawn by Alice (\mathbf{eCt}_A), to show the complete cycle for any given \mathbf{eCt}_A .
7. The TI assigns $\text{Resp} := \text{ACK}$ if all (Amt) of the \mathbf{eCt} and \mathbf{heCt} validation processes succeed, and NAK otherwise. The TI encrypts Resp as \mathbf{C}_3 with \mathbf{SK}_{TF} and transmits it to the FI.
8. FI decrypts \mathbf{C}_3 , records Resp , re-encrypts Resp with \mathbf{SK}_{FA} as \mathbf{C}_4 , and transmits it to Charlie.
9. Charlie's HOSE decrypts \mathbf{C}_4 and aborts if Resp is a NAK . Otherwise, he deducts the deposited subset of \mathbf{eCt} and \mathbf{heCt} from his database and generates a new challenge and corresponding \mathbf{POP}'_{C_c} . He encrypts and stores the new \mathbf{ECt}'_C element to his eCt_{DB} , and re-encrypts other elements in his database with the new \mathbf{POP}'_{C_c} if any exist. He encrypts \mathbf{POP}'_{C_c} as \mathbf{C}_5 using the original \mathbf{POP}_{C_c} and transmits it to FI.
10. Once FI receives \mathbf{C}_5 , if $\text{Resp} == \text{ACK}$, it credits Charlie's account with Amt .

Chapter 5. PUF-Cash Protocol

11. Once TI receives \mathbf{C}_5 and $\text{Resp} == \text{ACK}$, it marks all of Charlie's \mathbf{eCt} as invalid (redeemed) in the \mathbf{eCt}_{DB} . Otherwise it aborts.
12. If $\text{Resp} \neq \text{NAK}$, the TI decrypts \mathbf{C}_5 and updates Charlie's POP in its AT_{DB} .

Chapter 6

Experimental Setup

We first present the hardware and software overheads associated with the implementation of the PUF-Cash protocol on our test bed and then present a run time analysis.




Figure 6.1: Experiment set for the evaluation of PUF-Cash.

6.0.1 Device Resource Utilization

The test bed consists of one server and three devices as shown in Fig. 6.1. The token issuer (TI) is implemented on a Dell PowerEdge T440 Server with 32 1.8 GHz processors and 128 GB of main memory. The financial institution (FI), Alice

Chapter 6. Experimental Setup

and Bob run on a set of Digilent ZYBO-Z7 boards, which utilize Xilinx Zynq 7010 SoCs [26]. The processor system (PS) incorporates a 667 MHz dual-core Cortex-A9 processor, which has access to 1 GB of DRAM and a 1 Gbit/sec ethernet port. The programmable logic side (PL) has 17,600 LUTs and 240 KB of PL-side block RAM (BRAM).

The TI and FI applications are implemented as multi-threaded C programs, while customer device applications are single-threaded. The FI and device implementations are co-design applications, where network and database functions are implemented in C on the PS, while encryption, secure hash, SiRF PUF authentication bitstring and key generation functions are implemented as state machines in the PL. The TI emulates the SiRF PUF in software. The server and devices are connected to a network switch with 1 Gbit/sec of available bandwidth.

Table 6.1: PL-side resource utilization on the Zynq 7010.

Resource	Engine	SiRF		Crypto			Available	Overall
		Netlist	Total	AES	SHA-3	Total		
LUT	5842	796	37.72%	3128	3809	39.41%	17,600	77.13%
LUTRAM	60	96	2.60%	-	-	-	6000	2.60%
FF	4377	32	12.53%	2992	2244	14.88%	35,200	27.40%
BRAM blocks	5 (20 KB)	-	8.33%	-	-	-	60 (240 KB)	8.33%
DSP	2	-	2.50%	-	-	-	80	2.50%
BUFG	2	-	6.25%	-	-	-	32	6.25%

The PL resources used by the SiRF PUF Engine and SiRF PUF netlist (Entropy source) are given in Table 6.1. The percentage of resources used by all SiRF PUF components is shown in the 4th column. The overheads of other components of the HOSE, i.e., AES and SHA-3 cryptographic functions, are also presented in columns 5 and 6. The LUT row shows that SiRF PUF utilization is 37.72%, while utilization of the cryptographic functions is slightly larger at 39.41%. However, the two 32-bit multipliers and 20 KB BRAM used by the SiRF PUF, which are mapped into hardwired components on the FPGA, will make the SiRF PUF larger than the cryptographic functions in an ASIC implementation.

Table 6.2: Database Size Overhead.

DB name	Used By	Base (MB)	Rec (B)	Per Device (KB)	All Devices (MB)
NAT _{DB}	TI	0	60	594.3	81.2
AT _{DB}	TI	0	60	594.3	81.2
Chlng _{DB}	TI/FI/CD	1	-	-	1
MZT_AT _{DB}	FI/CD	0	80	10.9	0.011
POP _{DB}	FI/CD	4	296	40.5	4.040
HOSE/MZT _{DB}	FI/CD	0	2372	4.6	0.005

6.0.2 Database Size Overhead

The TI, FI and field-device applications utilize the sqlite3 database management software [27]. The database size overheads for each of the databases shown in Fig. 6.1 are given in Table 6.2, with CD used as an abbreviation for the customer device. The TI stores SiRF provisioning data in the NAT_{DB} and AT_{DB} for 140 ZYBO devices, each with 5072 DV_R and 5072 DV_F records (Rec), which supports a challenge-response space of more than 2^{24} bits per device (Note that the entire CRP space is 2^{35} bits per device using the SiRF netlist configuration shown in Fig. 3.1, which is discussed further in Chapter 8). Although each DV can be represented as a 16-bit fixed point integer with 4 digits of precision, indexing and other book-keeping within the database engine increase the size to 60 bytes per DV record. Therefore, with 10,144 DV per device, the total storage is 594.3 KB per device. The right-most column shows the total size of these databases after populating them with provisioning data from 140 FPGAs.

The Challenge_{DB}, utilized by all entities in the PUF-Cash system, is 1 MB in size. Individual record sizes within the MZT_AT_{DB} is given as 80 bytes per entry. From Fig. 4.1, a record consists of a 32-byte **ZHK** and **nonce**, and two 4-byte integers, leaving 8 bytes for database overhead. Assuming each device stores a record for 140 other devices in the population, the overhead is 10.9 KB.

The POP_{DB} is composed of two component tables. The POPA_{DB} stores 4 MB of **SF**, which is used for all customers, and is therefore fixed in size. The POPB_{DB} record size is 296 bytes, leading to a size overhead of 40.5 KB assuming each device

Chapter 6. Experimental Setup

stores data for 140 other devices. Combined with the fixed overhead of the POPA_{DB} , the storage per device is 4.04 MB. The size of the challenges stored in the MZT and HOSE DBs is small at 4.6 KB per device. The total overhead for a CD is the sum of the last four rows, and is slightly larger than 5 MB.

Chapter 7

Experiment Results

Hardware experiments are used to measure the performance of the PUF-Cash protocol operations. In each experiment, the timing information is collected as the protocol executes the withdraw-transfer-deposit sequence using an increasing number of **eCts** of sizes 1, 5, 10, 50, 100, up to 500,000, for a total of 12 sequences. In each sequence, the entire set of tokens are processed through all operations, thereby tokens are created, transacted and deleted, leaving the database in an empty state at the beginning of the next sequence. In addition, the POP **LLKs** generated in each experiment are stored to a file to allow an assessment of their statistical quality. The experiment as described is performed 40 times for statistical significance and to determine the limits on the transfer times.

7.0.1 Run Time Analysis

The processing and transfer times associated with the withdrawal, transfer and deposit operations within the PUF-Cash protocol are plotted in Fig. 7.1a. The processing times, and 3σ limits, are plotted as a function of the number of **eCt** processed. The run times include the time taken to complete all operations specified in the

Chapter 7. Experiment Results

message exchange diagrams, including the network transfer times between devices.

The processing time of all three operations is linear with respect to the number of tokens. The transfer operation, which occurs exclusively between Alice and Bob’s devices, possesses the largest processing time overhead. This is due largely to the limited processing capability of the devices. The processing times associated with the smaller value transfer operations, e.g., from 1 cent to 10,000 (\$100), are upper bounded at approximately 6 seconds. Processing times for amounts larger than 10,000 begin to diverge, but remain less than 20 seconds up to 100,000 **eCt** (\$1000). These processing times are competitive with existing state of the art payment systems where settlement occurs at the end of the transaction.




Figure 7.1: (a) Run times (y-axis) plotted against the number of **eCt** transferred (x-axis) for Withdrawal, Transfer and Deposit operations, as measured on the devices. The number of **eCt** processed varies from 1 (1 cent) to 500,000 (\$5,000). Three- σ variations in the measurements are shown as shaded regions around the mean values. (b) NIST test results using 1,400 POP LLKs collected during execution of the protocol.

Table 7.1: SiRF PUF Primitive Run Time Analysis. Mean and $\pm 3\sigma$ values are given. All times in seconds (s).

DA	VA	SE	LLK Enroll	LLK Regen
0.764 ± 0.064	0.618 ± 0.057	0.715 ± 0.071	1.400 ± 0.025	0.919 ± 0.025

Table 7.1 gives mean run times and 3σ limits for the timing-based (TB) authentication and key generation operations. The mean time for each operation is less than

Chapter 7. Experiment Results

Table 7.2: PUF-Cash Primitive Run Time Analysis. Average transfer times per 100 elements plus fixed overhead time. All times in seconds (s).

BootStrap (For 100 POP LLKs)	MZT Enroll (For 100 MZT ATs)
5.9 per 100 + 3.5	0.025 per 100 + 3.0

1 second for all operations except LLK Enroll. Table 7.2 gives the run time for the BootStrap operation carried out at device startup (see Fig. 5.3). The value given corresponds to the total transfer time for 100 **POP_x** elements, which includes a fixed overhead of 3.5 seconds. The transfer times correspond to approximately 17 **POP_x** per second. MZT enrollment time (also performed during BootStrap) translates to approximately 4000 MZT authentication tokens per second, with a fixed overhead of 3.0 seconds. A typical sequence performed at startup includes DA, VA, SE, two LLK regeneration operations and the BootStrap and MZT Enroll operations. The authentication and key generation operations (DA, VA, SE and LLK Regen) take approximately 4 seconds. For 100 POP **LLKs** and MZT **ATs**, the total startup time is $4 + 9.4 + 3.025 = 16.36$ seconds.

7.0.2 POP LLK Statistical Analysis

In this section, we assess the statistical quality of the **LLKs** generated during runs of the PUF-Cash protocol. The MZT protocol utilizes only one **LLK** in the protocol, and therefore nearly all of the **LLK** analyzed are produced by the POP protocol.

The message exchange diagrams for withdrawal, transfer and deposit shown in Figs. 5.4, 5.5 and 5.6 incorporate **LLK** refresh operations, which are annotated as **POP'_{xy}**. In each withdraw-transfer-deposit sequence, Alice generates a new **LLK** for herself during withdrawals and a second new **LLK** on behalf of Bob during transfers, and Bob generates a new **LLK** during deposits. Therefore, during the execution of the 12 sequences in each experiment, 36 **LLKs** are generated. The 40 repeated runs of the experiment yields a total of 1440 256-bit **LLKs**.

Chapter 7. Experiment Results

The 1440 **LLKs** are used as input to the NIST statistical test suite. Given the limited size of the bitstrings (256 bits), only six of the NIST statistical tests are applicable, as shown in Fig. 7.1b. The minimum pass threshold for a set size of 1440 bitstrings, and with α set to the default value 0.01, is 1414. All six tests passed, with the Runs test representing the worst case pass, i.e., at 1418 passing bitstrings.

Inter-bitstring Hamming distance (Inter-HD) is commonly used to measure uniqueness among the bitstrings. Inter-HD is computed by pairing bitstrings under all combinations and then counting the number of bits that differ in each pairing. The ideal result occurs when half of the bits in any pairing of the bitstrings differ.

$$\text{Inter-HD}_{i,j} = \frac{\sum_{k=1}^{|bs|} bs_{i,k} \oplus bs_{j,k}}{|bs|} \quad (7.1)$$

Eq. 7.1 gives the expression for Inter-HD for a pair of devices (i, j) of length $|bs| = 256$ bits. Inter-HD is computed across all possible pairings of 1440 bitstrings, i.e., $1440 \cdot 1439 / 2 = 1,036,080$ combinations, and averaged. The mean Inter-HD is 49.9992%, which is very close to the ideal value of 50%. No failures of any type were observed in **LLK** regeneration over the 8+ hour run of the experiments.

Chapter 8

Security Analysis

The TB authentication, session key generation and long-lived key generation PUF primitives act as the root-of-trust for the entire system, and as such, represent the primary targets of an attack. The TI extends the root-of-trust to customer devices using two different lighter-weight mechanisms. The security properties of the MZT and POP protocols, characterized as single **LLK** and challenge-response-based multi-**LLK**, respectively, draw from the root-of-trust by virtue of the amount and type of security related information transferred and stored on the devices. From the message-exchange diagrams, the security-sensitive data utilized by the MZT and POP protocols always originates from the TI, and is authenticated and encrypted in transit by the TB primitives. This requires adversaries to focus their attack on the PUF-based authentication and session key generation functions localized to customer devices. We assume a Dolev-Yao adversary model [28] where the attacker has significant computational resources, can eavesdrop on communication channels, and can intercept and manipulate packets as desired.

8.0.1 CRP Space Analysis

The TB security properties are rooted in the SiRF PUF’s physical source of entropy, which are discussed at length in [21]. The size of the CRP space when utilizing DV sets of size 5072 (the number of \mathbf{DV}_R and \mathbf{DV}_F stored in each of the NAT_{DB} and AT_{DB} databases in our experiments) is given by $5072^2 * 2^{23}/5072$, yielding a total CRP space size of more than 2^{35} possibilities in cases where the TI stores all possible sets in its timing databases.




Figure 8.1: Adversarial attack model to read out device responses.

8.0.2 Model-Building Attack Analysis

With the CRP space defined, we next describe an attack scenario where the adversary collects CRPs from a device-under-attack (DUA), as a precursor to building a machine-learning model. The message exchange diagram shown in Fig. 8.1 gives the sequence of operations carried out when the device requests authentication, which occurs as the first action in mutual authentication ($\text{MA}_{NA/A}$) in Figs. 4.1, 5.3, 5.4 and 5.6. Note that LLK generation and transmission is gated and encrypted by the $\text{MA}_{NA/A}$ and $\text{SKG}_{NA/A}$ functions. Therefore, the adversary must first defeat these security functions to extract response-based bitstrings, namely LLKs, from the PUF.

The following sequence describes the attack:

Chapter 8. Security Analysis

1. The device controls the start of the transaction by sending an *Authentication request* message to the adversary on the right. The adversary is assumed to be in possession of the device to initiate this request.
2. The adversary then specifies two parameters, v_x and p'_x . The v_x parameter specifies a 32-bit seed to an LFSR, that is used to pseudo-randomly select a set of vectors from the Vec_{DB} . We assume the adversary is in possession of a copy of the Vec_{DB} and algorithm used by the DUA. The v_x and p'_x parameters are sent to the DUA.
3. The device defines p_x by XOR'ing the adversary-generated p'_x with the output of its TRNG. The p_x parameter is used by both the DUA and adversary to seed another LFSR in the *DVDiffs module* that defines the pairing combinations of \mathbf{DV}_R and \mathbf{DV}_F used to create the **DVD**. We assume the adversary knows the LFSR primitive used by the device to create the **DVD**. The adversary cannot, however, control the final value of p_x . The p_x parameter is transmitted to the adversary.
4. The adversary extracts the vector sequence from the Vec_{DB} using a copy of the algorithm that the device uses. The adversary accesses an instance of the NAT_{DB} to extract timing data for the paths tested by the vectors. Since the adversary does not have access to the TI version of this database, he/she must create a version using simulation experiments, which assumes he/she has layout information related to the SiRF PUF instantiated on the DUA. The adversary constructs a 2048-byte array of SpreadFactors, \mathbf{SF}_x , from the data extracted from the NAT_{DB} . If SiRF PUF layout information is available, the \mathbf{SF}_x may represent good approximations to what the TI provides to the device during a genuine exchange, otherwise the \mathbf{SF}_x are random guesses. The adversary transmits the \mathbf{SF}_x to the device.
5. The device selects \mathbf{vecs}_x and runs the SiRF PUF in enrollment mode, HPUF_E , to produce helper data \mathbf{HD}_x . The \mathbf{HD}_x is a bitstring of length 2048 bits,

which reflects which of the device-computed \mathbf{DVD}_{co} are strong (1), and which are weak (0). The \mathbf{HD}_x is transmitted to the adversary for storage and analysis.

From this description, the adversary controls the challenge sequence parameter, v_x , and optimization parameters, \mathbf{SF}_x , and receives a helper data bitstring, \mathbf{HD}_x , as the response. The v_x can be manipulated by the adversary to force specific paths to be timed while the \mathbf{SF}_x can be manipulated incrementally to force weak bits to become strong, and vice versa. However, the actual response bits are not revealed (or even computed) by the SiRF PUF. Also note that the device can limit the rate of authentication attempts, which would limit the amount of data the adversary can collect over a given period of time.

It is the adversary's goal to produce an \mathbf{HD}_x bitstring that is highly correlated to the \mathbf{HD}_x that would be produced by the device under any arbitrary challenge, as a means of spoofing the identity of the genuine device to the TI. This requires the adversary to send accurate estimates of the \mathbf{SF}_x to the device during the model-building phase, to enable it to learn the bit classification, weak or strong, corresponding to each of the 2048 \mathbf{DVD}_{co} .

The effectiveness of the model-building attack can be evaluated during the verifier authentication stage, where the device compares the \mathbf{HD}_x bitstring input by the adversary with the version it creates internally. The adversary again controls the v_x and \mathbf{SF}_x transmissions to the device but must now provide an \mathbf{HD}_x that is highly correlated to \mathbf{HD}_x produced internally by the device. To date, we have been unsuccessful in attempts to carry out this type of attack.

Session key generation, $\text{SKG}_{NA/A}$, is gated by the success of verifier authentication and is not performed unless the adversary succeeds in convincing the device it is communicating with an authentic server. Although response bitstrings are sent in

the clear over the network for $SKG_{NA/A}$, the gating of this operation by $MA_{NA/A}$, the size of the CRP space and the lack of control over the parameter p_x will make it difficult to gather sufficient information to build an accurate model. If the adversary instead just eavesdrops on genuine device-TI authentications, the adversary will be tasked with identifying data communicated to/from a specific device because the $MA_{NA/A}$ functions are privacy-preserving, and therefore, they do not reveal the identity of the device in the exchanged messages.

8.0.3 Protocol Security Properties

The security properties of MZT enrollment operation (Fig. 4.1), as well as all communication between customer devices and the TI during the bootstrap, withdrawal and deposit transactions utilize $MA_{NA/A}$, and therefore, they inherit the security properties of the TB functions, which were covered above in the context of model-building attacks.

In contrast, the MZT in-field operations shown in Fig. 4.2 are not preceded with $MA_{NA/A}$, and therefore, warrant a separate assessment. In Step 6, a shared key is created by XOR-combining a ZHK_y key stored in the MZT_AT_{DB} for customer y and a SiRF PUF regenerated key (LLK_{MZTx}). Therefore, the shared key is constructed using information that is not stored in a database. Moreover, although not shown, the data stored in the MZT_AT_{AT} databases can be encrypted by the LLK_{hose} , and all actions carried out in the HOSE to provide an additional layer of security for the database-stored component of the shared key. In Step 7, Alice and Bob exchange nonces, which, by definition, are used only once so their exposure to adversaries is incidental to its security properties of the protocol. In Step 8, the key refresh data is AES-encrypted and is therefore protected by the strong security properties of the AES algorithm.

8.0.4 Ecosystem Attacks

The PUF-Cash protocol is designed to be resistant to all classes of attacks. For example, signatures and encryption keys are used only once, which increases resistance to side-channel and replay attacks. The HOSE provides a secure environment in the programmable logic (PL-side) by limiting access to the HOSE to two 32-bit memory-mapped registers, and by limiting interactions with the HOSE to a small set of valid operations. The HOSE avoids a wide range of software-based attacks that are possible within internet-connected processor environments. The self-validation, cross-validation and update protocol related to the **heCt** provides on-the-spot tamper detection capabilities during in-field, offline value transactions. The SiRF PUF challenge characteristics in combination with the proposed POP database scheme gives each fielded device access to large pool of entropy in a highly compressed format, which is key to enabling the aforementioned strengths of PUF-Cash. The HOSE embeds secrets that the untrusted processor-side cannot access, and serves as the primary mechanism to prevent duplication and double-spending of **eCt**. The security properties of the protocol that protect against network-based attacks, e.g., maninthemiddle, replay attacks and Sybil, as well as properties related to user privacy are similar to those described in previous work [29] and [30].

However, several attack scenarios are still possible. The most straightforward attack is a database copy-and-replace operation carried out on Alice’s device, where software running on Alice’s processor makes a copy of the eCt_{DB} before a payment transaction and then overwrites the HOSE updated version with the copy. This malicious operation is equivalent to double-spending because it restores **eCt** that Alice has just transferred to Bob. Although the TI will eventually detect the double-spent **eCt**, Alice’s recipients, and others in a transitive sequence of value transfer operations, will be prevented from depositing the duplicated **eCt**. The root of the problem stems from the inability of the device to 1) maintain state across power

Chapter 8. Security Analysis

cycles, or 2) to restrict updates to the databases stored in its NVM to the HOSE only.

Chapter 9

Conclusion and Future Work

A PUF-Cash protocol is proposed, and assessed in FPGA hardware experiments, that leverages two novel security protocols, called propagation-of-provenance (POP) and mutual-zero-trust (MZT), both of which can take place exclusively between two untrusted parties and both constructed to eliminate the need to interact with a trusted authority. POP is applied in this paper to secure the propagation of eCash tokens (**eCt**) from one customer to another. It leverages the exponential challenge-response space of a strong PUF, called the SiRF PUF, to allow recipients of **eCt** to authenticate their origin back to the token issuing authority. Successive hand-offs between devices utilize the combination of the payer's SiRF-instantiated device and the recipients stored PUF responses to authenticate **eCt** signatures.

A hardware-obfuscated secure enclave (HOSE) is introduced as a means of alleviating software security issues which are difficult to address in microprocessor environments, and as a means of preventing end-users from attempting to double spend the **eCt** which they store. The HOSE is implemented entirely in the programmable logic of an FPGA embedded within a system-on-chip (SoC) device. A PUF-Cash protocol implementing bootstrap, withdrawal, transfer and deposit transactions is

Chapter 9. Conclusion and Future Work

built on top of the HOSE, and POP and MZT protocols. Future work will investigate database-cloning attacks, network attack scenarios, and will further explore the model-building resistance of the SiRF PUF.

References

- [1] J. Delvaux, R. Peeters, D. Gu, and I. Verbauwhede, “A survey on lightweight entity authentication with strong pufs,” *ACM Comput. Surv.*, vol. 48, oct 2015.
- [2] T. Idriss and M. Bayoumi, “Lightweight highly secure puf protocol for mutual authentication and secret message exchange,” in *2017 IEEE Int. Conf. on RFID Technology Application*, pp. 214–219, 2017.
- [3] M. H. Mahalat, S. Saha, A. Mondal, and B. Sen, “A puf based light weight protocol for secure wifi authentication of iot devices,” in *2018 8th Int. Symp. on Embed. Comp. and System Design*, pp. 183–187, 2018.
- [4] M. H. Mahalat, D. Karmakar, A. Mondal, and B. Sen, “Puf based secure and lightweight authentication and key-sharing scheme for wireless sensor network,” *Journal of Emerging Technologies in Computer Systems*, vol. 18, Sep 2021.
- [5] W. Che, M. Martin, G. Pocklassery, V. K. Kajuluri, F. Saqib, and J. Plusquellic, “A privacy-preserving, mutual puf-based authentication protocol,” *Cryptography*, vol. 1, no. 1, 2017.
- [6] U. Chatterjee, V. Govindan, R. Sadhukhan, D. Mukhopadhyay, R. S. Chakraborty, D. Mahata, and M. M. Prabhu, “Building puf based authentication and key exchange protocol for iot without explicit crps in verifier database,” *IEEE Transactions on Dependable and Secure Computing*, vol. 16, no. 3, pp. 424–437, 2019.
- [7] U. Chatterjee, R. S. Chakraborty, and D. Mukhopadhyay, “A puf-based secure communication protocol for iot,” *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 16, no. 3, pp. 1–25, 2017.
- [8] J. R. Wallrabenstein, “Practical and secure iot device authentication using physical unclonable functions,” in *2016 IEEE 4th Int. Conference on Future Internet of Things and Cloud*, pp. 99–106, 2016.

References

- [9] M.-D. Yu, M. Hiller, J. Delvaux, R. Sowell, S. Devadas, and I. Verbauwhede, “A lockdown technique to prevent machine learning on pufs for lightweight authentication,” *IEEE Transactions on Multi-Scale Computing Systems*, vol. 2, no. 3, pp. 146–159, 2016.
- [10] J. Zhang and G. Qu, “Physical unclonable function-based key sharing via machine learning for iot security,” *IEEE Transactions on Industrial Electronics*, vol. 67, no. 8, pp. 7025–7033, 2020.
- [11] D. Chaum, “Blind signatures for untraceable payments,” in *Advances in Cryptology*, (Boston, MA), pp. 199–203, Springer US, 1983.
- [12] D. Chaum, A. Fiat, and M. Naor, “Untraceable electronic cash,” in *Advances in Cryptology* (S. Goldwasser, ed.), pp. 319–327, Springer, 1990.
- [13] O. Kesim, C. Grothoff, F. Dold, and M. Schanzenbach, “Zero-knowledge age restriction for gnu taler,” (Berlin, Heidelberg), p. 110129, Springer-Verlag, 2022.
- [14] I. H. Hanna Armelius, Carl Andreas Claussen, “On the possibility of a cash-like cbdc,” tech. rep., Sveriges Riksbank Payments Department and Research Division, 2021.
- [15] J. Zhang, L. Ma, and Y. Wang, “Fair e-cash system without trustees for multiple banks,” in *2007 Intl.1 Conf. on Computational Intelligence and Security Workshops (CISW 2007)*, pp. 585–587, 2007.
- [16] B. Lian, G. Chen, and J. Li, “Provably secure e-cash system with practical and efficient complete tracing,” *International Journal of Information Security*, no. 13, 2014.
- [17] N. A. Akbar, A. Muneer, N. ElHakim, and S. M. Fati, “Distributed hybrid double-spending attack prevention mechanism for proof-of-work and proof-of-stake blockchain consensus,” *Future Internet*, no. 11, 2021.
- [18] J. Ni, M. H. Au, W. Wu, X. Luo, X. Lin, and X. S. Shen, “Dual-anonymous off-line electronic cash for mobile payment,” *IEEE Transactions on Mobile Computing*, vol. 22, no. 6, pp. 3303–3317, 2023.
- [19] B. Yang, K. Yang, Z. Zhang, Y. Qin, and D. Feng, “Aep-m: Practical anonymous e-payment for mobile devices using arm trustzone and divisible e-cash,” in *Information Security* (A. C. A. N. Matt Bishop, ed.), (Cham), pp. 130–146, Springer International Publishing, 2016.

References

- [20] L. Mainetti, M. Aprile, E. Mele, and R. Vergallo, “A sustainable approach to delivering programmable peer-to-peer offline payments,” *Sensors*, vol. 23, no. 3, 2023.
- [21] J. Plusquellic, “Shift register, reconvergent-fanout (sirf) puf implementation on an fpga,” *Cryptography*, vol. 6, no. 4, 2022.
- [22] N. Irtija, E. Tsiropoulou, C. Minwalla, and J. Plusquellic, “True random number generation with the shift-register reconvergent-fanout (sirf) puf,” in *2022 IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, 2022.
- [23] J. Ju, R. Chakraborty, C. Lamech, and J. Plusquellic, “Stability analysis of a physical unclonable function based on metal resistance variations,” in *2013 IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, pp. 143–150, 2013.
- [24] D. Heeger and J. Plusquellic, “Analysis of iot authentication over lora,” in *2020 16th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pp. 458–465, 2020.
- [25] J. Plusquellic, E. E. Tsiropoulou, and C. Minwalla, “Privacy-preserving authentication protocols for iot devices using the sirf puf,” *IEEE Transactions on Emerging Topics in Computing*, pp. 1–16, 2023.
- [26] Diligent Corporation, Pullman, WA, *ZYBO-Z7 Reference Manual*, May 2023.
- [27] Hipp, Wyrick And Company, Inc, Charlotte, NC, *Sqlite3 webpage*, May 2023.
- [28] D. Dolev and A. Yao, “On the security of public key protocols,” *IEEE Transactions on Information Theory*, vol. 29, no. 2, pp. 198–208, 1983.
- [29] J. Calhoun, C. Minwalla, C. Helmich, F. Saqib, W. Che, and J. Plusquellic, “Physical unclonable function (puf)-based e-cash transaction protocol (puf-cash),” *Cryptography*, vol. 3, no. 3, 2019.
- [30] G. Fragkos, C. Minwalla, E. E. Tsiropoulou, and J. Plusquellic, “Enhancing privacy in puf-cash through multiple trusted third parties and reinforcement learning,” *J. Emerg. Technol. Comput. Syst.*, vol. 18, sep 2021.