



The GNU Taler system : practical and provably secure electronic payments

Florian Dold

► To cite this version:

Florian Dold. The GNU Taler system : practical and provably secure electronic payments. Cryptography and Security [cs.CR]. Université de Rennes, 2019. English. NNT : 2019REN1S008 . tel-02138082

HAL Id: tel-02138082

<https://theses.hal.science/tel-02138082>

Submitted on 23 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THESE DE DOCTORAT DE

L'UNIVERSITE DE RENNES 1
COMUE UNIVERSITE BRETAGNE LOIRE

ECOLE DOCTORALE N° 601
*Mathématiques et Sciences et Technologies
de l'Information et de la Communication*
Spécialité : Informatique

Par

« Florian DOLD »

« The GNU Taler System »

« Practical and Provably Secure Electronic Payments »

Thèse présentée et soutenue à « Rennes », le « 25.02.2019 »

Unité de recherche : Inria

Thèse N° : 195897

Rapporteurs avant soutenance :

Philip ROGAWAY
Sarah MEIKLEJOHN

Professeur à l'University of California, Davis
Professeure à l'University College London

Composition du Jury :

Président : Alan SCHMITT

Chercheur à l'Inria Rennes

Examineurs : Philip ROGAWAY
Sarah MEIKLEJOHN
Alex PENTLAND

Professeur à l'University of California, Davis
Professeure à l'University College London
Professeur à Massachusetts Institute of Technology

Dir. de thèse : Christian GROTHOFF

Professeur à Bern University of Applied Sciences

Co-dir. de thèse : Jean-Louis LANET

Directeur de recherche, Inria

Titre : Le système GNU Taler : Paiements électroniques pratiques et sécurisés.

Mots clés : *Monnaie électronique, Cryptographie, Sécurité, Systèmes distribués, Applications pratiques*

Résumé :

Les nouveaux protocoles de réseautage et cryptographiques peuvent considérablement améliorer les systèmes de paiement électroniques en ligne. Le présent mémoire porte sur la conception, la mise en œuvre et l'analyse sécuritaire du GNU Taler, un système de paiement respectueux de la vie privée conçu pour être pratique pour l'utilisation en ligne comme méthode de (micro-)paiement, et en même temps socialement et moralement responsable.

La base technique du GNU Taler peut être dû à l'e-cash de David Chaum. Notre travail va au-delà de l'e-cash de Chaum avec un changement efficace, et la nouvelle notion de transparence des revenus garantissant que les marchands ne peuvent recevoir de manière fiable un paiement d'un payeur non fiable que lorsque leurs revenus du paiement est visible aux autorités fiscales.

La transparence des revenus est obtenue grâce à l'introduction d'un protocole d'actualisation donnant lieu à un changement anonyme pour un jeton partiellement dépensé sans avoir besoin de l'introduction d'une évasion fiscale échappatoire. De plus, nous démontrons la sécurité prouvable de la transparence anonyme de nos revenus e-cash, qui concerne en plus l'anonymat habituel et les propriétés infalsifiables de l'e-cash, ainsi que la conservation formelle des fonds et la transparence des revenus.

Notre mise en œuvre du GNU Taler est utilisable par des utilisateurs non experts et s'intègre à l'architecture du web moderne. Notre plateforme de paiement aborde une série de questions pratiques telles que la prodigue des conseils aux clients, le mode de remboursement, l'intégration avec les banques et les chèques "know-your-customer (KYC)", ainsi que les exigences de sécurité et de fiabilité de la plateforme web. Sur une seule machine, nous réalisons des taux d'opérations qui rivalisent avec ceux des processeurs de cartes de crédit commerciaux globaux.

Pendant que les crypto-monnaies basées sur la preuve de travail à l'instar de Bitcoin doivent encore être mises à l'échelle pour servir de substituant aux systèmes de paiement établis, d'autres systèmes plus efficaces basés sur les blockchains avec des algorithmes de consensus plus classiques pourraient avoir des applications prometteurs dans le secteur financier. Nous faisons dans la conception, la mise en œuvre et l'analyse de la Byzantine Set Union Consensus, un protocole de Byzantine consensus qui s'accorde sur un (Super-)ensemble d'éléments à la fois, au lieu d'accepter en séquence les éléments individuels sur un ensemble. Byzantine Set consensus peut être utilisé comme composante de base pour des chaînes de blocs de permissions, où (à l'instar du style Nakamoto consensus) des blocs entiers d'opérations sont convenus à la fois d'augmenter le taux d'opération.

Title: The GNU Taler System: Practical and Provably Secure Electronic Payments

Keywords: *Electronic Cash, Cryptography, Security, Distributed Systems, Practical Applications*

Abstract:

We describe the design and implementation of GNU Taler, an electronic payment system based on an extension of Chaumian online e-cash with efficient change. In addition to anonymity for customers, it provides the novel notion of income transparency, which guarantees that merchants can reliably receive a payment from an untrusted payer only when their income from the payment is visible to tax authorities.

Income transparency is achieved by the introduction of a refresh protocol, which gives anonymous change for a partially spent coin without introducing a tax evasion loophole. In addition to income transparency, the refresh protocol can be used to implement Camenisch-style atomic swaps, and to preserve anonymity in the presence of protocol aborts and crash faults with data loss by participants.

Furthermore, we show the provable security of our income-transparent anonymous e-cash, which, in addition to the usual anonymity and unforgeability properties of e-cash, also formally models conservation of funds and income transparency.

Our implementation of GNU Taler is usable by non-expert users and integrates with the modern Web architecture. Our payment platform addresses a range of practical issues, such as tipping customers, providing refunds, integrating with banks and know-your-customer (KYC) checks, as well as Web platform security and reliability requirements.

On a single machine, we achieve transaction rates that rival those of global, commercial credit card processors.

We increase the robustness of the exchange—the component that keeps bank money in escrow in exchange for e-cash—by adding an auditor component, which verifies the correct operation of the system and allows to detect a compromise or misbehavior of the exchange early.

Just like bank accounts have reason to exist besides bank notes, e-cash only serves as part of a whole payment system stack. Distributed ledgers have recently gained immense popularity as potential replacement for parts of the traditional financial industry. While cryptocurrencies based on proof-of-work such as Bitcoin have yet to scale to be useful as a replacement for established payment systems, other more efficient systems based on blockchains with more classical consensus algorithms might still have promising applications in the financial industry.

We design, implement and analyze the performance of Byzantine Set Union Consensus (BSC), a Byzantine consensus protocol that agrees on a (super-)set of elements at once, instead of sequentially agreeing on the individual elements of a set. While BSC is interesting in itself, it can also be used as a building block for permissioned blockchains, where—just like in Nakamoto-style consensus—whole blocks of transactions are agreed upon at once, increasing the transaction rate.

Acknowledgements

I would like to thank Moritz Bartl for helping with the funding for this thesis. Bruno Haible provided generous support for the GNU Taler team to visit meetings of the W3C's Web Payment Working Group. I also thank Ashoka, the Tor project and the Donaukurier for their support.

This work benefits from the financial support of the Brittany Region (ARED 9174) and the Renewable Freedom Foundation (RFF).

I want to thank Inria and my team leader Axel Legay for hosting me during the work on my thesis, and Jean-Louis Lanet for agreeing to co-advise my thesis. Special thanks goes to Thomas Given-Wilson, Fabrizio Biondi, Laurent Morin and Nisrine Jafri for their support and company.

I also thank the Bern University of Applied Sciences for providing the hardware that was using during experiments.

Thanks to Marcello Stanisci for his work as an engineer on the GNU Taler project.

Chapter 5 is based on work published in the EURASIP Journal on Wireless Communications and Networking in collaboration with Christian Grothoff. Parts of Chapter 4 have been published in collaboration with Jeff Burdges, Christian Grothoff and Marcello Stanisci at SPACE 2016.

Thanks to Cristina Onete and Jeff Burdges for their collaboration on the provable security of GNU Taler.

I am grateful to the GNU project, in particular Richard Stallman, for their support of this project. I also thank all GNUnet developers and GNU Guix developers, especially Hartmut Goebel, Nils Gillmann, Gabor Toth, Ludovic Courtès and Andreas Enge.

Thanks to the Taler Systems business team, in particular Leon Schumacher and Michael Widmer, for their continuous faith in the project.

I thank my advisor Christian Grothoff for his advice and friendship.

Last but not least I'd like to thank my parents, my oldest friends Tom and Ben and my fiancée Vaish for their relentless support even during the most difficult times.

Contents

1. Introduction	1
1.1. Design Goals for GNU Taler	2
1.2. Features of Value-based Payment Systems	4
1.2.1. Offline vs Online Payments	4
1.2.2. Change and Divisibility	5
1.2.3. Anonymity Control	6
1.2.4. User Suspension	6
1.2.5. Transferability	6
1.2.6. Atomic Swaps	7
1.2.7. Refunds	7
1.3. User Experience and Performance	7
1.4. The Technical Foundation: Anonymous E-Cash	8
1.5. Distributed Ledgers	13
1.5.1. Consensus in Decentralized Blockchains	13
1.5.2. Permissioned Blockchains	14
1.5.3. Blockchains and GNU Taler	14
1.6. Key Contributions	15
1.7. Roadmap	16
2. GNU Taler, an Income-Transparent Anonymous E-Cash System	17
2.1. Design of GNU Taler	17
2.1.1. Entities and Trust Model	17
2.1.2. System Assumptions	18
2.1.3. Reserves	19
2.1.4. Coins and Denominations	20
2.1.5. Partial Spending and Unlinkable Change	21
2.1.6. Refreshing and Taxability	21
2.1.7. Transactions vs. Sharing	22
2.1.8. Aggregation	22
2.1.9. Refunds	23
2.1.10. Fees	23
2.1.11. The Withdraw Loophole and Tipping	23
2.2. Auditing	24
2.2.1. Exchange Compromise Modes	25
2.2.2. Cryptographic Proof	28
2.2.3. Perfect Crime Scenarios	28

2.3.	Related Work	29
2.3.1.	Anonymous E-Cash	29
2.3.2.	Blockchains	33
2.3.3.	Approaches to Micropayments	35
2.3.4.	Walled Garden Payment Systems	37
2.3.5.	Web Integration	37
3.	Security of Income-Transparent Anonymous E-Cash	41
3.1.	Introduction to Provable Security	41
3.1.1.	Algorithms, Oracles and Games	42
3.1.2.	Assumptions, Reductions and Game Hopping	45
3.1.3.	Notation	46
3.2.	Model and Syntax for Taler	46
3.2.1.	Algorithms	48
3.2.2.	Oracles	51
3.3.	Games	54
3.3.1.	Anonymity	54
3.3.2.	Conservation	55
3.3.3.	Unforgeability	56
3.3.4.	Income Transparency	56
3.4.	Security Definitions	57
3.5.	Instantiation	58
3.5.1.	Generic Instantiation	58
3.5.2.	Concrete Instantiation	64
3.6.	Proofs	64
3.6.1.	Anonymity	64
3.6.2.	Conservation	66
3.6.3.	Unforgeability	67
3.6.4.	Income Transparency	67
3.7.	Discussion	69
3.7.1.	Limitations	69
3.7.2.	Other Properties	70
4.	Implementation of GNU Taler	73
4.1.	Overview	75
4.1.1.	Taler APIs	76
4.1.2.	Cryptographic Algorithms	76
4.1.3.	Entities and Public Key Infrastructure	77
4.1.4.	Payments	80
4.1.5.	Resource-based Web Payments	84
4.1.6.	Session-bound Payments and Sharing	86
4.1.7.	Embedded Content	87
4.1.8.	Contract Terms	88
4.1.9.	Refunds	89

4.1.10.	Tipping	89
4.2.	Bank Integration	90
4.2.1.	Wire Method Identifiers	90
4.2.2.	Demo Bank	91
4.2.3.	EBICS and SEPA	92
4.2.4.	Blockchain Integration	92
4.3.	Exchange	92
4.4.	Auditor	94
4.5.	Merchant Backend	95
4.5.1.	Processing payments	96
4.5.2.	Back Office APIs	97
4.5.3.	Example Merchant Frontends	97
4.6.	Wallet	99
4.6.1.	Optimizations	100
4.6.2.	Coin Selection	100
4.6.3.	Wallet Detection	100
4.6.4.	Backup and Synchronization	101
4.6.5.	Wallet Liquidation	101
4.6.6.	Wallet Signaling	102
4.7.	Cryptographic Protocols	103
4.7.1.	Preliminaries	103
4.7.2.	Withdrawing	104
4.7.3.	Payment transactions	106
4.7.4.	Refreshing and Linking	107
4.7.5.	Refunds	114
4.8.	Experimental results	116
4.8.1.	Hardware Setup	117
4.8.2.	Coins Per Transaction	117
4.8.3.	Transaction Rate and Scalability	120
4.8.4.	Latency	121
4.9.	Current Limitations and Future Improvements	122
5.	Byzantine Set-Union Consensus	125
5.1.	Introduction	125
5.2.	Background	126
5.2.1.	The FLP Impossibility Result	127
5.2.2.	Byzantine Consensus in the Partially Synchronous Model	128
5.2.3.	Gradecast	129
5.2.4.	ByzConsensus	130
5.2.5.	Set Reconciliation	130
5.3.	Our Approach	133
5.3.1.	Definition	133
5.3.2.	Byzantine Set-Union Consensus (BSC) Protocol	134

Contents

5.4.	Implementation	139
5.4.1.	The GUNet Framework	139
5.4.2.	Set Reconciliation	139
5.4.3.	Set-Union Consensus	140
5.4.4.	Evaluating Malicious Behavior	141
5.5.	Experimental Results	141
5.5.1.	Bounded Set Reconciliation	142
5.5.2.	Byzantine Set Consensus	146
5.6.	Opportunities for Further Improving BSC	149
5.6.1.	Extension to Partial Synchrony	149
5.6.2.	Persistent Data Structures	150
5.6.3.	Fast Dissemination	150
5.7.	Application to SMC	150
5.7.1.	Bulletin Board for Electronic Voting	151
5.7.2.	Distributed Threshold Key Generation and Cooperative Decryption	152
5.7.3.	Electronic Voting with Homomorphic Encryption	153
5.7.4.	Other Applications of BSC	154
5.8.	Conclusions	154
6.	Future Work	155
7.	Conclusion	157
7.1.	Cryptocurrencies vs. Central-Bank-Issued Currencies	157
7.2.	Electronic Payments	158
	Bibliography	161
A.	Résumé en Français	179
B.	dold-draft-payto	185
C.	Coin Spending Simulation	195

List of Figures

1.1.	The user is prompted to install the wallet.	8
1.2.	The wallet popup shows an empty balance.	8
1.3.	The bank asks for login details.	9
1.4.	Account page of the demo bank.	9
1.5.	Exchange selection dialog in the wallet.	9
1.6.	PIN/TAN dialog of the demo bank.	10
1.7.	After a successful withdrawal, the balance is shown in the wallet. .	10
1.8.	Landing page of a store that sells essays.	11
1.9.	Payment prompt for an essay.	11
1.10.	Essay successfully purchased by the user.	12
2.1.	High-level overview of GNU Taler components.	18
4.1.	Components of GNU Taler in the context of a banking system. . .	75
4.2.	Entities/PKI in Taler	77
4.3.	Example response for /keys	78
4.4.	A denomination's lifetime.	79
4.5.	The contract header that is signed by the merchant.	81
4.6.	The deposit permission signed by the customer's wallet.	81
4.7.	Architecture of the exchange reference implementation	93
4.8.	Data flow for updating the exchange's keys.	94
4.9.	Architecture of the merchant reference implementation	96
4.10.	Code snippet for merchant frontend	98
4.11.	Architecture of the wallet reference implementation	99
4.12.	Withdraw protocol diagram.	105
4.13.	Spend protocol diagram.	108
4.14.	Deposit protocol diagram.	109
4.15.	RefreshDerive algorithm	110
4.16.	Refresh Protocol (Commit Phase)	111
4.17.	Refresh Protocol (Reveal Phase)	112
4.18.	Linking protocol	113
4.19.	Refund protocol	115
4.20.	Coin throughput.	121
4.21.	Comparison of components' CPU usage for the benchmark.	122
4.22.	Effect of artificial network delay on exchange's latency.	123
5.1.	CPU time for the SET service in relation to set size.	142
5.2.	CADET traffic for the SET service in relation to set size.	143

List of Figures

5.3. CPU time for the SET service in relation to set difference.	143
5.4. CADET traffic for the set service in relation to set difference. . . .	144
5.5. CADET traffic for the SET service at boundary for full transmission.	145
5.6. CADET traffic per peer, only correct peers.	146
5.7. CPU usage of BSC, only correct peers.	147
5.8. Runtime of BSC, only correct peers.	147
5.9. CADET traffic for BSC, one malicious peer.	148
5.10. Latency for BSC, one malicious peer.	148
5.11. Total number of extra elements received with one malicious peer. .	149
5.12. Different subsystems related to SMC in GUNet.	152
7.1. Historical market price of Bitcoin.	158

1. Introduction

New networking and cryptographic protocols can substantially improve electronic online payment systems. This thesis is about the design, implementation and security analysis of GNU Taler¹, a privacy-friendly payment system that is designed to be practical for usage as an online (micro-)payment method, and at the same time socially and ethically responsible.

Payment systems can generally be divided into two types: Register-based and value-based [Rik17]. A register-based system associates value with identities (e.g., bank account balances with customers), while a value-based system associates value with typically anonymous, digital or physical tokens (such as cash or prepaid credit cards). In practice, these two types of systems are combined, as different layers have different (and often conflicting) requirements: the payment system used to pay for a cappuccino in a coffee shop is most likely not suitable to buy real estate. Value-based payment systems typically provide more anonymity and convenience but also more risk to consumers (as they are responsible to secure the values they hold), while register-based systems shift risks to the payment service provider who has to authenticate consumers and ensure the integrity of the register.

This thesis covers both categories of payment systems:

- We explain GNU Taler, a design and implementation of a value-based payment system, discussing in-depth how to create a practical, privacy-preserving and secure (micro-)payment protocol that integrates nicely with the modern web. Our value-based payment protocol can in principle operate on top of any existing register-based system.
- For register-based payment systems, we present a new Byzantine consensus protocol. Consensus protocols are a key component of virtually all robust, distributed, register-based systems, as they facilitate agreement on a transaction ledger. Our Byzantine set union consensus (BSC) protocol can be used to achieve consensus in a decentralized and robust manner that tolerates a fraction of actively malicious participants. Our BSC protocol asymptotically speeds up the implementation of such transaction ledgers, compared to classic Byzantine consensus protocols.

GNU Taler is an official package of the GNU project², and the BSC protocol

¹<https://taler.net/>

²<https://gnu.org/>

was implemented in the CONSENSUS subsystem of the GNUnet framework.³ Our free software implementations are freely available from the GNU mirrors.

1.1. Design Goals for GNU Taler

The design of payment systems shapes economies and societies [ZSI13; Dal16]. Payment systems with high transaction costs create an economic burden. Predominantly cash-based societies provide some degree of anonymity for their citizens, but can fail to provide a sound foundation for taxation, facilitate corruption [SB17] and thus risk creating weak governments. On the other hand, systems with too much surveillance eliminate personal freedom.

As the Internet has no standardized payment system, especially not one that is capable of quickly, efficiently and securely settling small transactions (so-called micropayments), the majority of content on the web is financed by advertisements. As a result, advertising (and by implication, collecting data on users) has been a dominant business model on the Internet. This has not only resulted in a loss of independence of publishers—who need to cater to the needs of advertisers—but also in a situation where micro-targeted ads are so wide-spread, that they have been suspected to have influenced multiple major elections [Per17]. Ads are also a vector for malware [Pro+07]. Due to the prevalence of ad blockers, ads are also not guaranteed to be a sustainable business model.

In the world of online payments, credit cards and a sprawling number of smaller, proprietary payment processors are currently dominant, and market shares vary widely between different countries [Ady16; LMS16]. The resulting fragmentation again increases social costs: online shops can either choose to invest in implementing many proprietary protocols, or only implement the most popular ones, thereby reinforcing the dominance of a handful of proprietary payment systems.

Considering these and other social implications of payment systems, we started the development of GNU Taler with a set of high-level design goals that fit our social agenda. They are ranked by the importance we give to them, and when a trade-off must be made, the one that supports the more highly ranked goal is preferred:

1. GNU Taler must be implemented as free software.

Free refers to “free as in free speech”, as opposed to “free as in free beer”. More specifically, the four essential freedoms of free software [Stao2] must be respected, namely users must have the freedom to (1) run the software, (2) study and modify it, (3) redistribute copies, and (4) distribute copies of the modified version.

For merchants this prevents vendor lock-in, as another payment provider can take over, should the current one provide inadequate quality of service. As

³<https://gnunet.org>

the software of the payment provider itself is free, smaller or disadvantaged countries or organizations can run the payment system without being controlled by a foreign company. Customers benefit from this freedom, as the wallet software can be made to run on a variety of platforms, and user-hostile features such as tracking or telemetry could easily be removed from wallet software.

This rules out the mandatory usage of specialized hardware such as smart cards or other hardware security modules, as the software they run cannot be modified by the user. These components can, however, be voluntarily used by merchants, customers or payment processors to increase their operational security.

2. GNU Taler must protect the privacy of buyers.

Privacy should be guaranteed via technical measures, as opposed to mere policies. Especially with micropayments for online content, a disproportionate amount of rather private data about buyers would be revealed, if the payment system does not have privacy protections.

In legislations with data protection regulations (such as the recently introduced GDPR in Europe [VV17]), merchants benefit from this as well, as no data breach of customers can happen if this information is, by design, not collected in the first place. Obviously some private data, such as the shipping address for a physical delivery, must still be collected according to business needs.

The security of the payment systems also benefits from this, as the model shifts from authentication of customers to mere authorization of payments. This approach rules out whole classes of attacks such as phishing [Gar+07] or credit card fraud [SD10].

3. GNU Taler must enable the state to tax income and crack down on illegal business activities.

As a payment system must still be legal to operate and use, it must comply with these requirements. Furthermore, we consider levying of taxes as beneficial to society.

4. GNU Taler must prevent payment fraud.

This imposes requirements on the security of the system, as well as on the general design, as payment fraud can also happen through misleading user interface design or the lack of cryptographic evidence for certain processes.

5. GNU Taler must only disclose the minimal amount of information necessary.

The reason behind this goal is similar to (2). The privacy of buyers is given priority, but other parties such as merchants still benefit from it, for example, by keeping details about the merchant's financials hidden from competitors.

1. Introduction

6. GNU Taler must be usable.

Specifically it must be usable for non-expert customers. Usability also applies to the integration with merchants, and informs choices about the architecture, such as encapsulating procedures that require cryptographic operations into an isolated component with a simple API.

7. GNU Taler must be efficient.

Approaches such as proof-of-work are ruled out by this requirement. Efficiency is necessary for GNU Taler to be used for micropayments.

8. GNU Taler must avoid single points of failure.

While the design we present later is rather centralized, avoiding single points of failure is still a goal. This manifests in architectural choices such as the isolation of certain components, and auditing procedures.

9. GNU Taler must foster competition.

It must be relatively easy for competitors to join the systems. While the barriers for this in traditional financial systems are rather high, the technical burden for new competitors to join must be minimized. Another design choice that supports this is to split the whole system into smaller components that can be operated, developed and improved upon independently, instead of having one completely monolithic system.

1.2. Features of Value-based Payment Systems

There are many different possible features that have been proposed for value-based (sometimes called token-based) payment systems in the past. While we will discuss existing work on e-cash in more detail in Section 2.3.1, we will begin by a brief summary of the possible features that value-based payment systems could provide, and clarify which high-level features we chose to adopt for GNU Taler.

1.2.1. Offline vs Online Payments

Anonymous digital cash schemes since Chaum [Cha83] were frequently designed to allow the merchant to be offline during the transaction, by providing a means to deanonymize customers involved in double-spending, typically by encoding the customer's identity into their coins in a way that makes it only possible to decode the identity with two spending transcripts.

This approach is problematic in practice, as customers that restore a wallet from backup might accidentally double-spend and would then face punishment for it. Enforcing punishment for double-spenders can be rather difficult as well, since the double-spender might have signed up with a false identity or might

already have fled to another country, and a large number of merchants might already have been defrauded with the same coins.

Should the issuer of e-cash be compromised, an attacker could issue coins that fail to identify a culprit or even blame somebody else when they are double-spent. In an offline e-cash system, the detection of such an event is greatly delayed compared to systems with online spending, which can immediately detect when more coins are spent than were issued.

Thus, in GNU Taler, we decided that all coins must be immediately deposited online during a purchase. Only either a merchant or a customer needs to be online, since one of the two can forward messages to the payment service provider for the other.

1.2.2. Change and Divisibility

Customers do not always have the right set of coins available to exactly cover the amount to be paid to a merchant. With physical cash, the store would give the customer change. For e-cash, the situation is more complex, as the customer would have to make sure that the change has not already been spent, does not violate their anonymity and the merchant does not have a digital “copy” of the change tokens that the merchant can spend before the customer. Note that it would be unwise to always withdraw the correct amount of e-cash directly before a purchase, as it creates a temporal correlation between the non-anonymous withdrawal event and the spending event.

Most modern e-cash schemes instead deal with exact spending by providing *divisibility* of coins, where the customer can decide to only spend part of a coin. A significant chunk of the e-cash literature has been concerned with developing schemes that allow the individual, divided parts of a coin to be unlinkable (thus preserving anonymity) and to optimize the storage costs for wallets and the communication cost of withdrawals.

The current state of the art for divisible e-cash [PST17] achieves constant-time withdrawal and wallet storage cost for coins that can be split into an arbitrary but fixed (as a system parameter) number of pieces. A continuous “chunk” of the smallest pieces of a coin can be spent with constant-time communication complexity.

While this sounds attractive in theory, these results are mostly of academic interest, as the storage and/or computational complexity for the party that is checking for double spending of coins remains enormous: each smallest piece of every coin needs to be recorded and checked individually. When paying \$10.00 with a coin that supports division into cent pieces, 1000 individual coin pieces must be checked for double spending and recorded, possibly in compressed form to trade storage costs for more computation.

For GNU Taler, we use a rather simple and practical approach, made possible by requiring participants to be online during spending: coins can be fractionally spent without having divisible, unlinkable parts. The remaining value on a coin

1. Introduction

that was spend (and thus revealed) can be used to withdraw fresh, unlinkable coins. The protocol to obtain change takes additional measures to ensure that it cannot be misused to facilitate untaxed transactions. Giving change for e-cash has been proposed before [BGK95; TW01], but to the best of our knowledge, the idea of income-transparent change is novel.

1.2.3. Anonymity Control

Some proposed e-cash protocols contain mechanisms to allow selective deanonymization of transactions for scenarios involving crime [ST99], specifically blackmailing and tax evasion.

Unfortunately this does not really work as a countermeasure against blackmailing in practice. As noted in the paper that initially described such a mechanism for blind signatures [SPC95], a blackmailer could simply request to be paid directly with plain, blindly signed coins, and thereby completely circumvent the threat of revocable anonymity.

GNU Taler provides *income transparency* as a measure against tax evasion. We furthermore describe a different approach in a blackmailing scenario in Section 2.2.3, which we believe is more practical in dissuading blackmailers in practice.

1.2.4. User Suspension

Anonymous user suspension [ASM11] has been proposed as another mechanism to punish users suspected in illicit activities by preventing them from making further transactions until the suspension is lifted. Anonymous suspension is based on transactions; the user involved in a particular transaction is suspended, but their identity is not revealed.

While the approach is interesting, it is not practical, as it requires a single permanent key pair to be associated with each user. If a user claims they lost their private key and requests a new key pair, their suspension would be effectively lifted. Suspending users from a dominant payment system is also socially problematic, as excluding them from most commercial activities would likely be a disproportionate and cruel punishment.

1.2.5. Transferability

Transferability is a feature of certain e-cash systems that allows transfer of e-cash between two parties without breaking anonymity properties [FPV09]. Contemporary systems that offer this type of disintermediation attract criminal activity [Ric16].

GNU Taler specifically provides roughly the *opposite* of this property, namely *income transparency*, to guarantee that e-cash is not easily abused for tax evasion. Mutually trusting users, however, can share ownership of a coin.

1.2.6. Atomic Swaps

Atomic swaps (often called “fair exchange” in the e-cash literature) are a feature of some e-cash systems that allows e-cash to be exchanged against some service or (digital) product, with a trusted third party ensuring that the payee receives the payment if and only if they correctly provided the merchandise.

GNU Taler supports Camenisch-style atomic swaps [CLM07], as explained in Section 3.7.2.

1.2.7. Refunds

GNU Taler allows merchants to provide refunds to customers during a limited time after the coins for the payment were deposited. The merchant signs a statement that effectively allows the customer to reclaim a previously spent coin. Customers can request anonymous change for the reclaimed amount.

While this is a rather simple extension, we are not aware of any other e-cash system that supports refunds.

1.3. User Experience and Performance

For adoption of a payment system, the user experience is critical. Thus, before diving into *how* GNU Taler is implemented, we begin by showing how GNU Taler *looks* from the perspective of an end user in the context of web payments, in a desktop browser (Chromium).

To use GNU Taler, the user must first install a browser extension (Figure 1.1). Once installed, the user can open a pop-up window by clicking on the Taler logo, to see the initially empty wallet balance (Figure 1.2).

The customer logs into their online bank—a simple demo bank in our case—to withdraw digital cash from their bank account into their wallet (Figures 1.3 and 1.4). Our demo uses KUDOS as an imaginary currency. Before the user is asked to confirm, they are given the option to view details about or change the default exchange provider, the GNU Taler payment service provider (Figure 1.5).

With a real bank, a second factor (such as a mobile TAN) would now be requested from the user. Our demo instead asks the user to solve a simple CAPTCHA (Figure 1.6). The amount withdrawn—minus withdrawal fees—is now available as e-cash in the wallet (Figure 1.7).

The customer can now go to an online shop to spend their digital cash. We’ve implemented a shop that sells single chapters from Richard Stallman’s essay collection “Free Software, Free Society” [Stao2] (Figure 1.8). The user selects an essay, and is then immediately presented with a confirmation page rendered by the wallet (Figure 1.9). After paying, the user can immediately read the article (Figure 1.10).

Our benchmarks, discussed in Chapter 4 show that a single machine can support around 1000 payments per second, and our implementation is easily

1. Introduction




Figure 1.1.: The user is prompted to install the wallet.

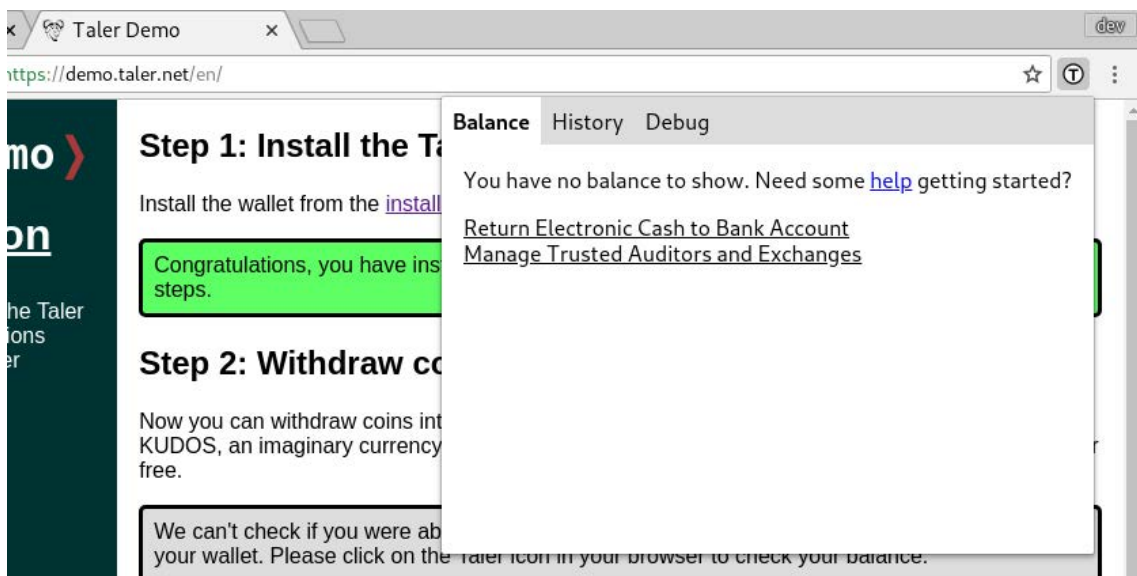


Figure 1.2.: The wallet popup shows an empty balance.

amenable to further scaling.

The extra computation required in the customer's wallet is in the order of a few hundred milliseconds even on typical mobile or tablet devices, and thus barely noticeable.

1.4. The Technical Foundation: Anonymous E-Cash

GNU Taler is based on anonymous e-cash. Anonymous e-cash was invented by David Chaum in the 1980s [Cha83]. The idea behind Chaumian e-cash is both simple and ingenious, and can be best illustrated with the carbon paper⁴ analogy: A long, random serial number is generated, for example, by throwing a die a few dozen times, and written on a piece of paper. A carbon paper is placed on top, with the pigmented side facing down, and both pieces of paper

⁴Carbon paper is a paper coated with pigment (originally carbon) on one side. When put face-down between two sheets of normal paper, the pressure from writing with a pen or typewriter on the first layer causes pigment to be deposited on the paper beneath, allowing a copy to be made.

1.4. The Technical Foundation: Anonymous E-Cash

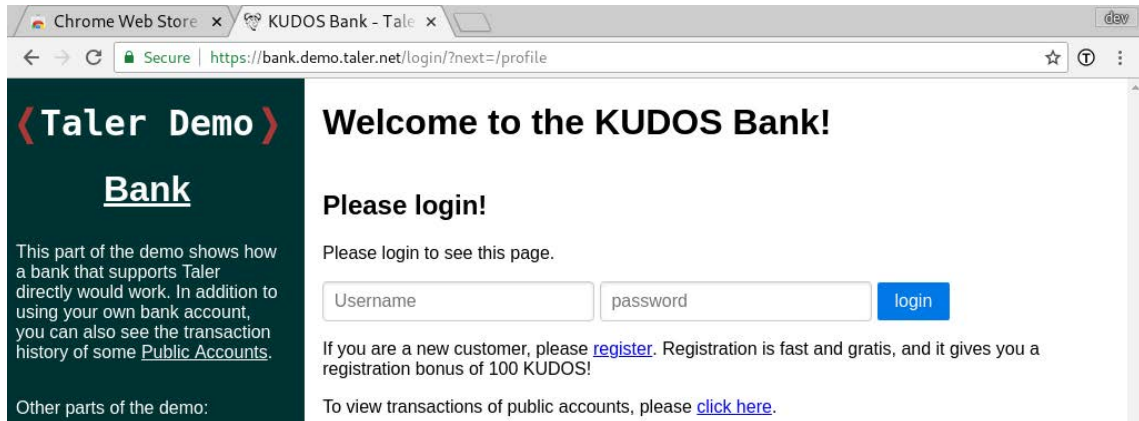


Figure 1.3.: The bank asks for login details.

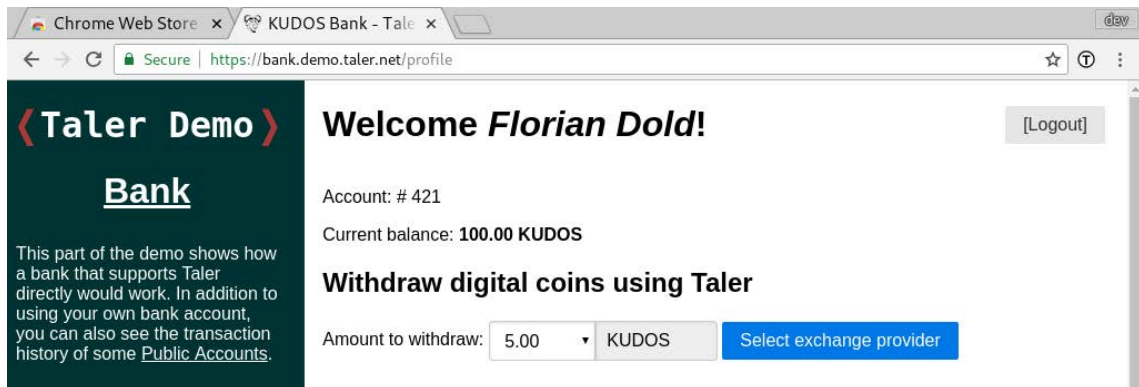


Figure 1.4.: Account page of the demo bank.

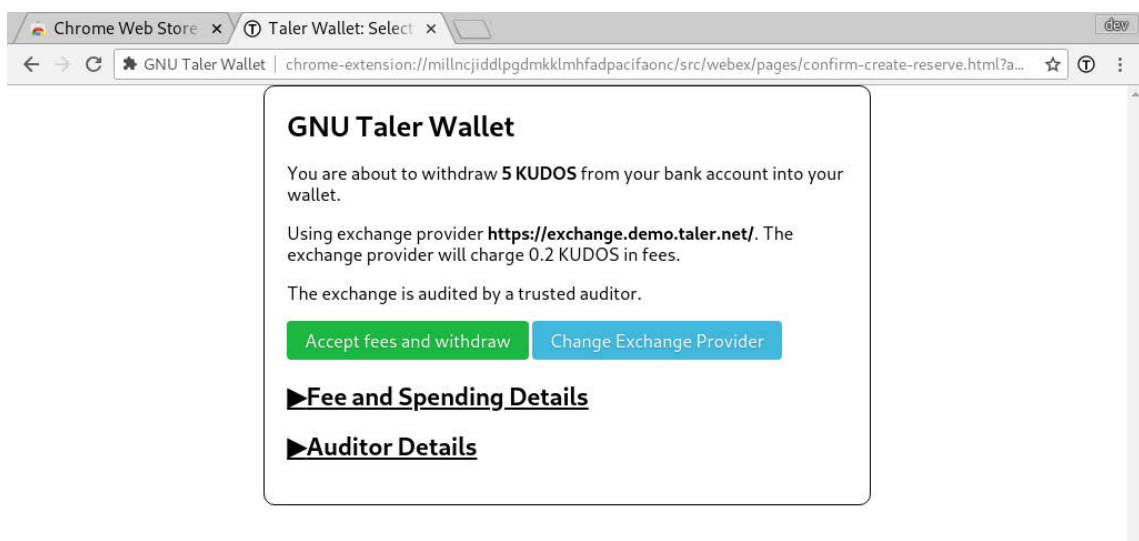


Figure 1.5.: Exchange selection dialog in the wallet.

1. Introduction

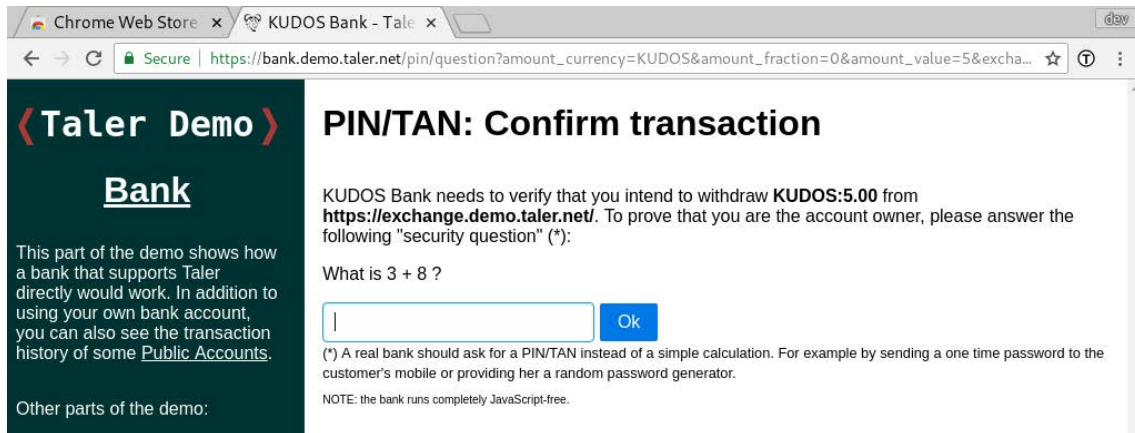


Figure 1.6.: PIN/TAN dialog of the demo bank.




Figure 1.7.: After a successful withdrawal, the balance is shown in the wallet.

1.4. The Technical Foundation: Anonymous E-Cash

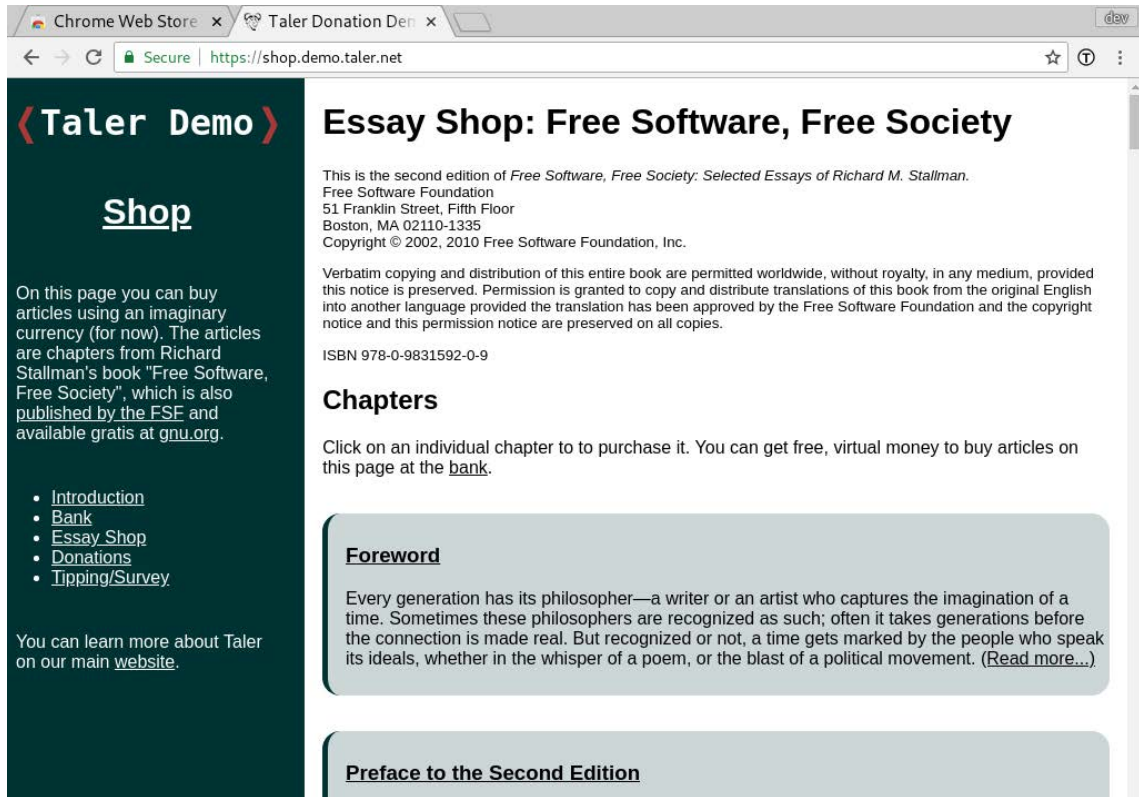


Figure 1.8.: Landing page of a store that sells essays.

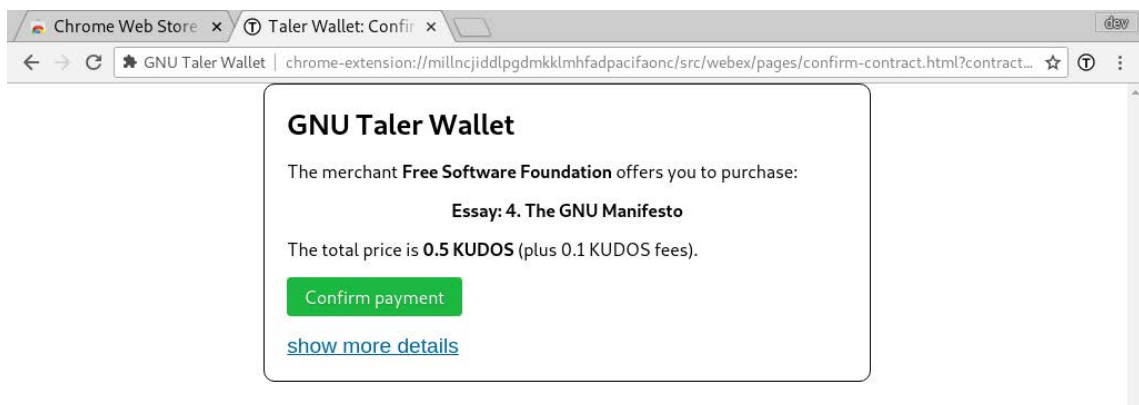


Figure 1.9.: Payment prompt for an essay. Rendered by the wallet.

1. Introduction




Figure 1.10.: Essay successfully purchased by the user.

are put into an opaque envelope. The envelope is now sealed and brought to a bank. The bank draws a signature on the outside of the envelope, which presses through to the piece of paper with the serial number. In exchange for the signed envelope, the bank deducts a fixed amount (say five dollars) from the customer's bank account. Under the (admittedly rather strong) assumption that the bank's signature cannot be forged, the signed piece of paper with the serial number is now an untraceable bank note worth five dollars, as the bank signed it without seeing the serial number inside the envelope! Since the signed paper can be easily copied, merchants that accept it as payment must check the bank's signature, call the bank and transmit the serial number. The bank keeps a register of all serial numbers that have been used as payment before. If the serial number is already in the bank's register, the bank informs the merchant about the attempted double spending, and the merchant then rejects the payment.

The digital analogue of this process is called a *blind signature*, where the signer knows that it gave a digital signature, but does not know the contents of the message that it signed.

In this document, we use *coin* to refer to a token of value in an e-cash system. Note that the analogy of a coin does not always hold up, as certain types of operations possible in some e-cash schemes, such as partial spending, divisibility, etc., do not transfer to physical coins.

We have the following security and correctness properties for GNU Taler (formally defined in Chapter 3):

- *Anonymity* guarantees that transactions cannot be correlated with withdrawals or other transactions made by the same customer.
- *Unforgeability* guarantees that users cannot spend more e-cash than they withdrew.
- *Conservation* guarantees that customers do not lose money due to inter-

rupted protocols or malicious merchants; they can always obtain anonymous change or a proof of successful spending.

- *Income transparency* guarantees that mutually distrusting parties are unable to reliably transfer e-cash between them without the income of participants being visible to tax auditors.

While anonymity and unforgeability are common properties of e-cash, we are not aware of any other treatments of income transparency and conservation.

1.5. Distributed Ledgers

The main purpose of blockchains, including those implementing cryptocurrencies, is to maintain a distributed ledger that holds state, together with rules on how this state can be updated. The name “blockchain” derives from its structure: A list of updates (“transactions”) is bundled into a so-called block, and each block contains a hash of the previous block. Cryptocurrencies use blockchains to remember the amount of currency controlled by a particular account (\equiv private key). Thus, while cryptocurrencies use the term “coin” (creating potentially misleading associations with cash), they actually realize a decentralized register-based payment system with the blockchain storing the register⁵ using private keys to authenticate account owners.

Cryptocurrencies based on blockchains gained immense popularity over the last years on the promise of a universal, global and decentralized payment system that is independent from country boundaries and legislations. In practice, however, current incarnations of these technologies can only handle a handful of transactions, have high transaction fees and are surprisingly centralized [BS15; Böh+15]. Bitcoin, the most popular cryptocurrency, can handle around 3-7 transactions per second, globally. While there are various plans to make blockchains more scalable [GM16], there is no concrete evidence that any of them will work without further sacrificing decentralization.

1.5.1. Consensus in Decentralized Blockchains

In decentralized blockchains, multiple parties must agree on the current state of the ledger by agreeing on a “head” of the chain of blocks. How to advance this head to include new transactions is thus a critical design choice.

⁵Anonymous cryptocurrencies such as ZeroCash [Ben+14] have special accounts (called shielded addresses) that can “hide” their balance, and require the owner to prove in zero-knowledge that their balance is sufficient for a transaction. As such, anonymous transactions in these systems (which are typically only a small subset of all transactions) are closer to value-based systems. However, currently only a small percentage ($\approx 5\%$) of all funds in ZCash, the most widely used anonymous cryptocurrency, belong to shielded addresses (<https://explorer.zcha.in/statistics/value>).

1. Introduction

With proof-of-work blockchains such as Bitcoin, each block contains the solution to a computationally expensive puzzle that is derived from the contents of the block. The block that, together with its ancestors, contains the most expensive accumulated work (and respects the rules of the blockchain with regards to what transactions are valid) is considered the head of the chain. All participants of the network can “mine” a block by collecting transactions and trying to solve the corresponding computational puzzle. Successful miners are rewarded with a mining reward and transaction fees. This type of agreement on a ledger is also called “Nakamoto Consensus”, after the inventor of Bitcoin. The result of the agreement is not final: if a branch originating from an earlier block of the chain accumulates more work, it becomes the canonical head. While this type of consensus has some attractive properties—there is no fixed set of members, and remains secure as long as an adversary has less than $1/4$ of computational power [ES18]—it consumes a huge amount of energy to provide for computation of the proof-of-work puzzles.

After Bitcoin popularized the concept of blockchains, alternative consensus mechanisms were proposed to replace or augment proof-of-work. In proof-of-stake blockchains, a single node is selected as a validator. The validator must provide a safety deposit (the “stake”), and if any misbehavior is detected, the safety deposit is destroyed. If the validator behaves correctly, they earn transaction fees and get back their safety deposit. Currently proof-of-stake protocols are still in development, and often require falling back to other consensus mechanisms in certain situations.

1.5.2. Permissioned Blockchains

Permissioned blockchains have a known, relatively small set of participants, and can rely on more traditional and cheaper consensus algorithms. When resilience against actively malicious members is required, a so-called Byzantine consensus protocol must be used. Byzantine consensus protocols typically agree on a single value at once.

In Chapter 5 we introduce a Byzantine consensus algorithm that can be used to agree directly on a (super-)set of all transaction that honest peers proposed. This allows for implementations of permissioned blockchains where transactions are accumulated into blocks, and the transactions within a block are agreed upon in a way that’s asymptotically faster than agreeing on every transaction sequentially.

This protocol could be used in the future to implement an efficient and robust implementation of the register-based layer of a payment system, with GNU Taler e-cash as the value-based layer above it.

1.5.3. Blockchains and GNU Taler

Blockchains today fail to satisfy most of our design goals for payment systems. While most blockchains are implemented as free software, they often manage to

both fail to adequately protect the privacy of buyers *and* to enable the state to crack down on illegal activities: With most non-permissioned blockchains, the transaction history of all participants is publicly available, creating serious privacy risks [Mei+13; Jaw+18]. At the same time, as accounts are simply private keys, states have a hard time tracking down users [LI16]. Design variations that do offer reasonable privacy generally have even more atrocious performance characteristics and create additional traceability problems for law enforcement [Ben+14]. Additionally, blockchain-based cryptocurrencies suffer from usability and performance problems.

With our BSC protocol, we focus on improving the performance of the consensus protocol for permissioned blockchains. Permissioned blockchains can be given rules that enforce Know-Your-Customer (KYC) and Anti-Money-Laundering (AML) regulations [SWP16]. When deployed in the context of centrally-banked fiat currencies, such a permissioned blockchain can then effectively recreate the semantics of a classical distributed banking system. As mentioned before, GNU Taler’s value-based protocol can be integrated with any kind of register-based banking—including those based on blockchains—improving performance and privacy for value-based transactions.

1.6. Key Contributions

We claim the following key contributions for this thesis:

- We design, implement and analyze an efficient Byzantine consensus protocol on set structures that allows an optimized implementation of distributed transaction ledgers.
- We introduce the notion of income transparency for e-cash, with an instantiation in Chaum-style e-cash and proofs.
- We design the GNU Taler payment system under consideration of practical aspects of e-cash including aborts, network failures, refunds, multi-coin payments, faults from wallet synchronization and their effects on anonymity; showing the necessity of a refresh operation.
- We propose a modification to our protocol that provides protection against certain blackmailing and kidnapping scenarios.
- We design and implement the seamless, native integration of e-cash into the web architecture, and discuss security and privacy aspects of this integration.
- We implemented the GNU Taler payment system and evaluate its performance.

1.7. Roadmap

Chapter 2 describes the high-level design of GNU Taler, and compares it to payment systems found in the academic literature and real-world usage. Chapter 3 first gives a gentle introduction to provable security (which can be skipped by readers with a background in cryptography), and then defines security properties for income-transparent, anonymous e-cash. The cryptographic protocols for GNU Taler are defined in detail, and proofs are given that our protocols satisfy the security properties defined earlier. In Chapter 4, the implementation of GNU Taler is described, and the performance and scalability is evaluated. Chapter 5 is about the design, implementation and evaluation of our Byzantine set union consensus protocol. Chapter 6 discusses future work and missing pieces to deploy GNU Taler in production. Chapter 7 concludes with an outlook on the potential impact and practical relevance of this work.

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

This chapter gives a high-level overview of the design of GNU Taler, based on the requirements discussed in Chapter 1. The cryptographic protocols and security properties are described and analyzed in detail in Chapter 3. A complete implementation with focus on Web payments is discussed in Chapter 4.

2.1. Design of GNU Taler

GNU Taler is based on the idea of Chaumian e-cash [Cha83], with some differences and additions explained in the following sections. Other variants and extensions of anonymous e-cash and blind signatures are discussed in Section 2.3.1.

2.1.1. Entities and Trust Model

GNU Taler consists of the following entities (see 2.1):

- The *exchanges* serve as payment service provider for a financial transaction between a customer and a merchant. They hold bank money in escrow in exchange for anonymous digital *coins*.
- The *customers* keep e-cash in their electronic *wallets*.
- The *merchants* accept digital coins in exchange for digital or physical goods and services. The digital coins can be deposited with the exchange, in exchange for bank money.
- The *banks* receive wire transfer instructions from customers and exchanges. A customer, merchant and exchange involved in one GNU Taler payment do not need to have accounts with the same bank, as long as wire transfers can be made between the respective banks.
- The *auditors*, typically run by trusted financial regulators, monitor the behavior of exchanges to assure customers and merchants that exchanges operate correctly.

In GNU Taler, the exchanges can be separate entities from the banks. This fosters competition between exchanges, and allows Taler to be deployed in an environment with legacy banks that do not support Taler directly.

2. GNU Taler, an Income-Transparent Anonymous E-Cash System




Figure 2.1.: High-level overview of the different components of GNU Taler, banks are omitted.

If a customer wants to pay a merchant, the customer needs to hold coins at an exchange that the merchant trusts. To make the selection of trusted exchanges simpler, merchants and customers can choose to automatically trust all exchanges audited by a certain auditor.

The exchange is trusted to hold funds of its customers in escrow and to make payments to merchants when digital coins are deposited. Customer and merchant can have assurances about the exchange’s liquidity and operation through the auditor, which would typically be run by financial regulators or other trusted third parties.

2.1.2. System Assumptions

We assume that an anonymous, bi-directional communication channel¹ is used for all communication between the customer and the merchant, as well as for obtaining unlinkable change for partially spent coins from the exchange and for retrieving the exchange’s public keys used in verifying and blindly signing coins. The withdrawal protocol, on the other hand, does not require an anonymous channel to preserve the anonymity of electronic coins.

During withdrawal, the exchange knows the identity of the withdrawing customer, as there are laws, or bank policies, that limit the amount of cash that an individual customer can withdraw in a given time period [Bad15; Reu15]. GNU Taler is thus only anonymous with respect to *payments*. While the exchange does know their customer (KYC), it is unable to link the known identity of the

¹An anonymization layer like Tor [DMS04] can provide a practical approximation of such a communication channel, but does not provide perfect anonymity [Joh+13].

customer that withdrew anonymous digital coins to the *purchase* performed later at the merchant.

While customers can make untraceable digital cash payments, the exchange will always learn the merchants' identity, which is necessary to credit their accounts. This information can also be used for taxation, and GNU Taler deliberately exposes these events as anchors for tax audits on merchants' income. Note that while GNU Taler *enables* taxation, it does not *implement* any automatic taxation.

GNU Taler assumes that each participant has full control over their system². We assume the contact information of the exchange is known to both customer and merchant from the start, and the customer can authenticate the merchant, for example, by using X.509 certificates [Yee13]. A GNU Taler merchant is expected to deliver the service or goods to the customer upon receiving payment. The customer can seek legal relief to achieve this, as the customer receives cryptographic evidence of the contract and the associated payment.

2.1.3. Reserves

A *reserve* refers to a customer's non-anonymous funds at an exchange, identified by a reserve public key. Suppose a customer wants to convert money into anonymized digital coins. To do that, the customer first creates a reserve private/public key pair, and then transfers money via their bank to the exchange. The wire transfer instruction to the bank must include the reserve public key. To withdraw coins from a reserve, the customer authenticates themselves using the corresponding reserve private key.

Typically, each wire transfer is made with a fresh reserve public key and thus creates a new reserve, but making another wire transfer with the same reserve public key simply adds funds to the existing reserve. Even after all funds have been withdrawn from a reserve, customers should keep the reserve key pair until all coins from the corresponding reserve have been spent, as in the event of a denomination key revocation (see Section 2.2.1) the customer needs this key to recover coins of revoked denominations.

The exchange automatically transfers back to the customer's bank account any funds that have been left in a reserve for an extended amount of time, allowing customers that lost their reserve private key to eventually recover their funds. If a wire transfer to the exchange does not include a valid reserve public key, the exchange transfers the money back to the sender.

Instead of requiring the customer to manually generate reserve key pairs and copy them onto a wire transfer form, banks can offer tight integration with the GNU Taler wallet software. In this scenario, the bank's website or banking app provides a "withdraw to GNU Taler wallet" action. After selecting this action,

²Full control goes both ways: it gives the customer the freedom to run their own software, but also means that the behavior of fraudulent customers cannot be restricted by simpler technical means such as keeping balances on tamper-proof smart cards, and thus can lead to an overall more complex system.

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

the user is asked to choose the amount to withdraw from their bank account into the wallet. The bank then instructs the GNU Taler wallet software to create record of the corresponding reserve; this request contains the anticipated amount, the reserve key pair and the URL of the exchange to be used. When invoked by the bank, the wallet asks the customer to select an exchange and to confirm the reserve creation. The exchange chosen by the customer must support the wire transfer method used by the bank, which will be automatically checked by the wallet. Typically, an exchange is already selected by default, as banks can suggest a default exchange provider to the wallet, and additionally wallets have a pre-defined list of trusted exchange providers. Subsequently, the wallet hands the reserve public key and the bank account information of the selected exchange back to the bank. The bank—typically after asking for a second authentication factor from the customer—will then trigger a wire transfer to the exchange with the information obtained from the wallet.

When the customer's bank does not offer tight integration with GNU Taler, the customer can still manually instruct their wallet to create a reserve. The public key must then be included in a bank transaction to the exchange. When the customer's banking app supports pre-filling wire transfer details from a URL or a QR code, the wallet can generate such a URL or QR code that includes the pre-filled bank account details of the exchange as well as the reserve public key. The customer clicks on this link or scans the QR code to invoke their banking app with pre-filled transaction details. Since there currently is no standardized format for pre-filled wire transfer details, we are proposing the `payto://` URI format explained in Section 4.2.1, currently under review for acceptance as an IETF Internet Standard.

2.1.4. Coins and Denominations

Unlike plain Chaumian e-cash, where a coin just contains a serial number, a *coin* in Taler is a public/private key pair where the private key is only known to the owner of the coin.

A coin derives its financial value from a blind signature on the coin's public key. The exchange has multiple *denomination key* pairs available for blind-signing coins of different financial values. Other approaches for representing different denominations are discussed in Section 2.3.1.

Denomination keys have an expiration date, before which any coins signed with it must be spent or exchanged into newer coins using the refresh protocol explained in Section 2.1.6. This allows the exchange to eventually discard records of old transactions, thus limiting the records that the exchange must retain and search to detect double-spending attempts. If a denomination's private key were to be compromised, the exchange can detect this once more coins are redeemed than the total that was signed into existence using that denomination key. Should such an incident occur, the exchange can allow authentic customers to redeem their unspent coins that were signed with the compromised private

key, while refusing further deposits involving coins signed by the compromised denomination key (see Section 2.2.1). As a result, the financial damage of losing a private signing key is limited to at most the amount originally signed with that key, and denomination key rotation can be used to bound that risk.

To prevent the exchange from deanonymizing users by signing each coin with a fresh denomination key, exchanges publicly announce their denomination keys in advance with validity periods that imply sufficiently strong anonymity sets. These announcements are expected to be signed with an offline long-term private *master signing key* of the exchange and the auditor. Customers should obtain these announcements using an anonymous communication channel.

After a coin is issued, the customer is the only entity that knows the private key of the coin, making them the *owner* of the coin. Due to the use of blind signatures, the exchange does not learn the public key during the withdrawal process. If the private key is shared with others, they become co-owners of the coin. Knowledge of the private key of the coin and the signature over the coin's public key by an exchange's denomination key enables spending the coin.

2.1.5. Partial Spending and Unlinkable Change

Customers are not required to have exact change ready when making a payment. In fact, it should be encouraged to withdraw a larger amount of e-cash beforehand, as this blurs the correlation between the non-anonymous withdrawal event and the anonymous spending event, increasing the anonymity set.

A customer spends a coin at a merchant by cryptographically signing a *deposit permission* with the coin's private key. A deposit permission contains the hash of the *contract terms*, i.e., the details of the purchase agreement between the customer and merchant. Coins can be *partially* spent, and a deposit permission specifies the fraction of the coin's value to be paid to the merchant. As digital coins are trivial to copy, the merchant must immediately deposit them with the exchange, in order to get a deposit confirmation or an error that indicates double spending.

When a coin is used in a completed or attempted/aborted payment, the coin's public key is revealed to the merchant/exchange, and further payments with the remaining amount would be linkable to the first spending event. To obtain unlinkable change for a partially spent (or otherwise revealed coin), GNU Taler introduces a *refresh protocol*. The refresh protocol allows the customer to obtain new coins for the remaining amount on a coin. The old coin is marked as spent after it has been refreshed into new coins. Using blind signatures to withdraw the refreshed coins makes them unlinkable from the old coin.

2.1.6. Refreshing and Taxability

One goal of GNU Taler is to make merchants' income transparent to state auditors, so that income can be taxed appropriately. Naively implemented, however, a simple refresh protocol could be used to evade taxes: the payee of an untaxed

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

transaction would generate the private keys for the coins that result from refreshing a (partially spent) old coin, and send the corresponding public keys to the payer. The payer would execute the refresh protocol, provide the payee's coin public keys for blind signing, and provide the signatures to the payee, who would now have exclusive control over the coins.

To remedy this, the refresh protocol introduces a *link threat*: coins are refreshed in such a way that the owner of the old coin can always obtain the private key and exchange's signature on the new coins resulting from refreshes, using a separate *linking protocol*. This introduces a threat to merchants that try to obtain untaxed income. Until the coins are finally deposited at the exchange, the customer can always re-gain ownership of them and could deposit them before the merchant gets a chance to do so. This disincentivizes the circulation of unreported income among untrusted parties in the system.

In our implementation of the refresh and linking protocols, there is a non-negligible success chance ($\frac{1}{\kappa}$, depending on system parameter κ , typically ≥ 3) for attempts to cheat during the refresh protocol, resulting in refreshed coins that cannot be recovered from the old coin via the linking protocol. Cheating during refresh, however, is still not *profitable*, as an unsuccessful attempt results in completely losing the amount that was intended to be refreshed.

For purposes of anti-money-laundering and taxation, a more detailed audit of the merchant's transactions can be desirable. A government tax authority can request the merchant to reveal the business agreement details that match the contract terms hash recorded with the exchange. If a merchant is not able to provide these values, they can be subjected to financial penalties by the state in relation to the amount transferred by the traditional currency transfer.

2.1.7. Transactions vs. Sharing

Sharing—in contrast to a transaction—happens when mutually trusted parties simultaneously have access to the private keys and signatures on coins. Sharing is not considered a transaction, as subsequently both parties have equal control over the funds. A useful application for sharing are peer-to-peer payments between mutually trusting parties, such as families and friends.

2.1.8. Aggregation

For each payment, the merchant can specify a deadline before which the exchange must issue a wire transfer to the merchant's bank account. Before this deadline occurs, multiple payments from deposited coins to the same merchant can be *aggregated* into one bigger payment. This reduces transaction costs from underlying banking systems, which often charge a fixed fee per transaction. To incentivize merchants to choose a longer wire transfer deadline, the exchange can charge the merchant a fee per aggregated wire transfer.

2.1.9. Refunds

The aggregation period also opens the opportunity for cheap *refunds*. If a customer is not happy with their product, the merchant can instruct the exchange to give the customer a refund before the wire transfer deadline has occurred. This effectively “undoes” the deposit of the coin, and restores the available amount left on it. The refresh protocol is then used by the customer on the coins involved in a refund, so that payments remain unlinkable.

2.1.10. Fees

In order to subsidize the operation of the exchange and enable a sustainable business model, the exchange can charge fees for most operations. For withdrawal, refreshing, deposit and refunds, the fee is dependent on the denomination, as different denominations might have different key sizes, security and storage requirements.

Most payment systems hide fees from the customer by putting them to the merchant. This is also possible with Taler. As different exchanges (and denominations) can charge different fees, the merchant can specify a maximum amount of fees it is willing to cover. Fees exceeding this amount must be explicitly paid by the customer.

Another consideration for fees is the prevention of denial-of-service attacks. To make “useless” operations, such as repeated refreshing on coins (causing the exchange to use relatively expensive storage), unattractive to an adversary, these operations must charge a fee. Again, for every refresh following a payment, the merchant can cover the costs up to a limit set by the merchant, effectively hiding the fees from the customer.

Yet another type of fee are the *wire transfer fees*, which are charged by the exchange for every wire transfer to a merchant in order to compensate for the cost of making a transaction in the underlying bank system. The wire transfer fees encourage merchants to choose longer aggregation periods, as the fee is charged per transaction and independent of the amount.

Merchants can also specify the maximum wire fee they are willing to cover for customers, along with an *amortization rate* for the wire fees. In case the wire fees for a payment exceed the merchant’s chosen maximum, the customer must additionally pay the excess fee divided by the amortization rate. The merchant should set amortization rate to the expected number of transactions per wire transfer aggregation window. This allows the merchant to adjust the total expected amount that it needs to pay for wire fees.

2.1.11. The Withdraw Loophole and Tipping

The withdraw protocol can be (ab)used to illicitly transfer money, when the receiver generates the coin’s secret key, and gives the public key to the party executing the withdraw protocol. We call this the “withdraw loophole”. This

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

is only possible for one “hop”, as money can still not circulate among mutually distrusted parties, due to the properties of the refresh protocol.

A “benevolent” use of the withdraw loophole is *tipping*, where merchants give small rewards to customers (for example, for filling out a survey or installing an application), without any contractual obligations or digitally signed agreement.

Fixing the Withdraw Loophole

In order to discourage the usage of the withdraw loophole for untaxed payments, the following approach would be possible: Normal withdraw operations and unregistered reserves are disabled, except for special tip reserves that are registered by the merchant as part of a tipping campaign. Customers are required to pre-register at the exchange and obtain a special withdraw key pair against a small safety deposit. Customer obtain new coins via a refresh operation from the withdraw key to a new coin. If customers want to abuse Taler for untaxed payments, they either need to risk losing money by lying during the execution of the refresh protocol, or share their reserve private key with the payee. In order to discourage the latter, the exchanges gives the safety deposit to the first participant who reports the corresponding private key as being used in an illicit transaction, and requires a new safety deposit before the customer is allowed to withdraw again.

However since the withdraw loophole allows only one additional “payment” (without any cryptographic evidence that can be used in disputes) before the coin must be deposited, these additional mitigations might not even be justified considering their additional cost.

2.2. Auditing

The auditor is a component of GNU Taler which would typically be deployed by a financial regulator, fulfilling the following functionality:

- It regularly examines the exchange’s database and bank transaction history to detect discrepancies.
- It accepts samples of certain protocol responses that merchants received from an audited exchange, to verify that what the exchange signed corresponds to what it stored in its database.
- It certifies exchanges that fulfill the operational and financial requirements demanded by regulators.
- It regularly runs anonymous checks to ensure that the required protocol endpoints of the exchange are available.
- In some deployment scenarios, merchants need to pre-register with exchanges to fulfill know-your-customer (KYC) requirements. The auditor

provides a list of certified exchanges to merchants, to which the merchant then can automatically KYC-register.

- It provides customers with an interface to submit cryptographic proof that an exchange misbehaved. If a customer claims that the exchange denies service, it can execute a request on behalf of the customer.

2.2.1. Exchange Compromise Modes

The exchange is an attractive target for hackers and insider threats. We now discuss different ways that the exchange can be compromised, how to reduce the likelihood of such a compromise, and how to detect and react to such an event if it happens.

Compromise of Denomination Keys and Revocation

When a denomination key pair is compromised, an attacker can “print money” by using it to sign coins of that denomination. An exchange (or its auditor) can detect this when the number of deposits for a certain denomination exceed the number of withdrawals for that same denomination.

We allow the exchange to revoke denomination keys, and wallets periodically check for such revocations. We call a coin of a revoked denomination a revoked coin. If a denomination key has been revoked, the wallets use the *payback* protocol to recover funds from coins of revoked denominations. Once a denomination is revoked, new coins of this denomination can’t be withdrawn or used as the target denomination for a refresh operation. A revoked coin cannot be spent, and can only be refreshed if its public key was recorded in the exchange’s database (as spending/refresh operations) before it was revoked.

The following cases are possible for payback:

1. The revoked coin has never been seen by the exchange before, but the customer can prove via a withdraw protocol transcript and blinding factor that the coin resulted from a legitimate withdrawal from a reserve. In this case, the exchange credits the reserve that was used to withdraw the coin with the value of the revoked coin.
2. The coin has been partially spent. In this case, the exchange allows the remaining amount on the coin to be refreshed into fresh coins of non-revoked denominations.
3. The revoked coin C_R has never been seen by the exchange before, was obtained via the refresh protocol, and the exchange has an existing record of either a deposit or refresh for the ancestor coin C_A that was refreshed into the revoked coin C_R . If the customer can prove this by showing a corresponding refresh protocol transcript and blinding factors, the exchange credits the remaining value of C_R on C_A . It is explicitly permitted for C_A

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

to be revoked as well. The customer can then obtain back their funds by refreshing C_A .

These rules limit the maximum financial damage that the exchange can incur from a compromised denomination key D to $2nv$, with n being the maximum number of D -coins simultaneously in circulation and v the financial value of a single D -coin. Say denomination D was withdrawn by legitimate users n times. As soon as the exchange sees more than n pairwise different D -coins, it must immediately revoke D . An attacker can thus at most gain nv by either refreshing into other non-revoked denominations or spending the forged D -coins. The legitimate users can then request a payback for their coins, resulting in a total financial damage of at most $2nv$.

With one rare exception, the payback protocol does not negatively impact the anonymity of customers. We show this by looking at the three different cases for payback on a revoked coin. Specifically, in case (1), the coin obtained from the credited reserve is blindly signed, in case (2) the refresh protocol guarantees unlinkability of the non-revoked change, and in case (3) the revoked coin C_R is assumed to be fresh. If C_R from case (3) has been seen by a merchant before in an aborted/unfinished transaction, this transaction would be linkable to transactions on C_A . Thus, anonymity is not preserved when an aborted transaction coincides with revoked denomination, which should be rare in practice.

Unlike most other operations, the payback protocol does not incur any transaction fees. The primary use of the protocol is to limit the financial loss in cases where an audit reveals that the exchange's private keys were compromised, and to automatically pay back balances held in a customers' wallet if an exchange ever goes out of business.

To limit the damage of a compromise, the exchange can employ a hardware security module that contains the denomination secret keys, and is pre-programmed with a limit on the number of signatures it can produce. This might be mandated by certain auditors, who will also audit the operational security of an exchange as part of the certification process.

Compromise of Signing Keys

When a signing key is compromised, the attacker can pretend to be a merchant and forge deposit confirmations. To forge a deposit confirmation, the attacker also needs to get a customer to sign a contract from the adversary (which should include the adversary's banking details) with a valid coin. The attack here is that the customer is allowed to have spent the coin already. Thus, a deposit of the resulting deposit permission would result in a rejection from the exchange due to double spending. By forging the deposit confirmation using the compromised signing key, the attacker can thus claim in court that they properly deposited the coin first and demand payment from the exchange.

We note that indeed an evil exchange could simply fail to record deposit permissions in its database and then fail to execute them. Thus, given a merchant

presenting a deposit confirmation, we need a way to establish whether this is a case of an evil exchange that should be compelled to pay, or a case of a compromised signing key and where payouts (and thus financial damage to the exchange) can legitimately be limited.

To limit the financial damage of a compromised signing key, merchants must be required to work with auditors to perform a *probabilistic deposit auditing* of the exchange. Here, the goal is to help detect the compromise of a signing key by making sure that the exchange does indeed properly record deposit confirmations. However, double-checking with the auditor if every deposit confirmation is recorded in the exchange's database would be too expensive and time-consuming. Fortunately, a probabilistic method where merchants only send a small fraction of their deposit confirmations to the auditor suffices. Then, if the auditor sees a deposit confirmation that is not recorded in the exchange's database (possibly after performing the next synchronization with the exchange's database), it signals the exchange that the signing key has been compromised.

At this point, the signing key must be revoked and the exchange will be required to investigate the security of its systems and address the issue before resuming normal operations.

Still, at this point various actors (including the attacker) could still step forward with deposit confirmations signed by the revoked key and claim that the exchange owes them for their deposits. Simply revoking a signing key cannot lift the exchange's payment obligations, and the attacker could have signed an unlimited number of such deposit confirmations with the compromised key. However, in contrast to honest merchants, the attacker will not have participated *proportionally* in the auditor's probabilistic deposit auditing scheme for those deposit confirmations: in that case, the key compromise would have been detected and the key revoked.

The exchange must still pay all deposit permissions it signed for coins that were not double-spent. However, for all coins where multiple merchants claim that they have a deposit confirmation, the exchange will pay the merchants proportionate to the fraction of the coins that they reported to the auditor as part of probabilistic deposit auditing. For example, if 1% of deposits must be reported to the auditor according to the protocol, a merchant might be paid at most say $100+X$ times the number of reported deposits where $X > 0$ serves to ensure proper payout despite the probabilistic nature of the reporting. As a result, honest merchants have an *incentive* to correctly report the deposit confirmations to the auditor.

Given this scheme, the attacker can only report a small number of deposit confirmations to the auditor before triggering the signing key compromise detection. Suppose again that 1% of deposit confirmations are reported by honest merchants, then the attacker can only expect to submit 100 deposit permissions created by the compromised signing key before being detected. The attacker's expected financial benefit from the key compromise would then be the value of $(100 + X) \cdot 100$ deposit permissions.

Thus, the financial benefit to the attacker can be limited by probabilistic deposit

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

auditing, and honest merchants have proper incentives to participate in the process.

Compromise of the Database

If an adversary would be able to modify the exchange, this would be detected rather quickly by the auditor, provided that the database has appropriate integrity mechanisms. An attacker could also prevent database updates to block the recording of spend operations, and then double spend. This is effectively equivalent to the compromise of signing keys, and can be detected with the same strategies.

Compromise of the Master Key

If the master key was compromised, an attacker could de-anonymize customers by announcing different sets of denomination keys to each of them. If the exchange was audited, this would be detected quickly, as these denominations will not be signed by auditors.

2.2.2. Cryptographic Proof

We use the term “proof” in many places as the protocol provides cryptographic proofs of which parties behave correctly or incorrectly. However, as [MA14] point out, in practice financial systems need to provide evidence that holds up in courts. Taler’s implementation is designed to export evidence and upholds the core principles described in [MA14]. In particular, in providing the cryptographic proofs as evidence none of the participants have to disclose their core secrets.

2.2.3. Perfect Crime Scenarios

GNU Taler can be slightly modified to thwart blackmailing or kidnapping attempts by criminals who intend to use the anonymity properties of the system and demand to be paid ransom in anonymous e-cash.

Our modification incurs a slight penalty on the latency for customers during normal use and requires slightly more data to be stored in the exchange’s database, and thus should only be used in deployments where resistance against perfect crime scenarios is necessary. A payment system for a school cafeteria likely does not need these extra measures.

The following modifications are made:

1. Coins can now only be used in either a transaction or in a refresh operations, not a mix of both. Effectively, the customer’s wallet then needs to use the refresh protocol to prepare exact change before a transaction is made, and that transaction is made with exact change.

This change is necessary to preserve anonymity in face of the second modification, but increases storage requirements and latency.

2. The payback protocol is changed so that a coin obtained via refreshing must be recovered differently when revoked: to recover a revoked coin obtained via refreshing, the customer needs to show the transcripts for the chain of all refresh operations and the initial withdrawal operation (including the blinding factor). Refreshes on revoked coins are not allowed anymore.

After an attacker has been paid ransom, the exchange simply revokes all currently offered denominations and registers a new set of denomination with the auditor. Reserves used to pay the attacker are marked as blocked in the exchange's database. Normal users can use the payback protocol to obtain back the money they've previously had in revoked denominations. The attacker can try to recover funds via the (now modified) payback protocol, but this attempt will not be successful, as the initial reserve is blocked. The criminal could also try to spend the e-cash anonymously before it is revoked, but this is likely difficult for large amounts, and furthermore due to income transparency all transactions made between the payment of the ransom and the revocation can be traced back to merchants that might be complicit in laundering the ransom payment.

Honest customers can always use the payback protocol to transfer the funds to the initial reserve. Due to modification (1), unlinkability of transactions is not affected, as only coins that were purely used for refreshing can now be correlated.

We believe that our approach is more practical than the approaches based on tracing, since in a scheme with tracing, the attacker can always ask for a plain blind signature. With our approach, the attacker will always lose funds that they cannot immediately spend. Unfortunately our approach is limited to a kidnapping scenario, and not applicable in those blackmail scenarios where the attacker can do damage after they find out that their funds have been erased.

2.3. Related Work

2.3.1. Anonymous E-Cash

Chaum's seminal paper [Cha83] introduced blind signatures and demonstrated how to use them for online e-cash. Later work [Cha+89; CFN90] introduced offline spending, where additional information is encoded into coins in such a way that double spending reveals the culprit's identity.

Okamoto [Oka95] introduced the first efficient offline e-cash scheme with divisibility, a feature that allows a single coin to be spent in parts. With Okamoto's protocol, different spending operations that used parts of the same coin were linkable. An unlinkable version of divisible e-cash was first presented by Canard [CG07].

Camenisch's compact e-cash [CHL05] allows wallets with 2^ℓ coins to be stored and withdrawn with storage, computation and computational costs in $\mathcal{O}(\ell)$. Each coin in the wallet, however, still needs to be spent separately.

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

The protocol that can currently be considered the state-of-the-art for efficient offline e-cash was introduced by Pointcheval et al. [PST17]. It allows constant-time withdrawal of a divisible coin, and constant-time spending of a continuous “chunk” of a coin. While the pre-determined number of divisions of a coin is independent from the storage, bandwidth and computational complexity of the wallet, the exchange needs to check for double-spending at the finest granularity. Thus, highly divisible coins incur large storage and computational costs for the exchange.

An e-cash system with multiple denominations with different financial values was proposed by Canard and Gouget [CGHo6] in the context of a divisible coupon system.

One of the earliest mentions of an explicit change protocol can be found in [BGK95]. Ian Goldberg’s HINDE system is another design that allows the merchant to provide change, but the mechanism could be abused to hide income from taxation.³ Another online e-cash protocol with change was proposed by Tracz [TWo1]. The use of an anonymous change protocol (called a “refund” in their context) for fractional payments has also been suggested for a public transit fees payment system [Rup+13]. Change protocols for offline e-cash were recently proposed [BY18]. To the best of our knowledge, no change protocol with protections against tax evasion has been proposed so far, and all change protocols suggested so far can be (ab)used to make a payment into another wallet.

Transferable e-cash allows the transfer of funds between customers without using the exchange as in intermediary [FPVo9].

Chaum also proposed wallets with observers [CP92] as a mechanism against double spending. The observer is a tamper-proof hardware security module that prevents double-spending, while at the same time being unable to de-anonymize the user.

Various works propose mechanisms to selectively de-anonymize customers or transactions that are suspected of criminal activities [SPC95; Dav+97]. Another approach suspends customers that were involved in a particular transaction, while keeping the customer anonymous [ASM11].

One of the first formal treatments of the provable security of e-cash was given in [Dam07]. The first complete security definition for blind signatures was given by Pointcheval [PS96] and applied to RSA signatures later [PS00]. While the security proof of RSA signatures requires the random oracle model, many blind signature schemes are provably secure in the standard model [IL13; PST17]. While most literature provides only “human-verified” security arguments, the security of a simple e-cash scheme has been successfully modeled in ProVerif [DKL15], albeit only in the symbolic model.

³Description based on personal communication. HINDE was never published, but supposedly publicly discussed at Financial Crypto ’98.

Implementations

DigiCash was the first commercial implementation of Chaum's e-cash. It ultimately failed to be widely adopted, and the company filed for bankruptcy in 1998. Some details of the implementation are available [Sch98]. In addition to Chaum's infamously paranoid management style [Ano99], reasons for DigiCash's failure could have been the following:

- DigiCash did not allow account-less operations. To use DigiCash, customers had to sign up with a bank that natively supports DigiCash.
- DigiCash did not support change or partial spending, negating a lot of the convenience and security of e-cash by requiring frequent withdrawals from the customer's bank account.
- The technology used by DigiCash was protected by patents, which stifled innovation from competitors.
- Chaum's published design does not clearly limit the financial damage an exchange might suffer from the disclosure of its private online signing key.

To our knowledge, the only publicly available effort to implement anonymous e-cash is Opencoin [DPWo8]. However, Opencoin is neither actively developed nor used, and it is not clear to what degree the implementation is even complete. Only a partial description of the Opencoin protocol is available to date.

Representing Denominations

For GNU Taler, we chose to represent denominations of different values by a different public key for every denomination, together with a mapping from public key to financial value and auxiliary information about fees and expiration dates. This approach has the advantage that coins of higher denominations can be signed by denominations with a larger key size.

Schoenmakers [Sch98] proposes an optimized implementation of multiple denomination that specifically works with RSA keys, which encodes the denomination in the public exponent e of the RSA public key, while the modulus N stays the same for all denominations. An advantage of this scheme is the reduced size of the public keys for a set of denominations. As this encoding is specific to RSA, it would be difficult for future versions of this protocol to switch to different blind signature primitives. More importantly, factoring N would lead to a compromise of all denominations instead of just one.

Partially blind signatures can be used to represent multiple denominations by blindly signing the coin's serial number and including the financial value of the coin in the common information seen by both the signer and signee [AO00].

The compact e-cash scheme of Märtens [Mär15] allows constant-time withdrawal of wallets with an arbitrary number of coins, as long as the number of coins is smaller than some system parameter. This approach effectively dispenses with the need to have different denominations.

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

Comparison

	Year	Implementation	Offline spending	Safe aborts/backups	Key expiration	Income transparency	No trusted setup	Storage for wallet	Storage for exchange	Change/Divisibility	Receipts & Refunds
Chaum [Cha83]	1983	P	✗	✗	?	?	✓	$\log n$	$\log n$	✗	✗
DigiCash [Sch98]	1990	P	✗	✓	✓	✗	✓	$\log n$	$\log n$	✗	✗
Offline Chaum [CFN90]	1990	?	✓	✗	?	?	✓	$\log n$	$\log n$	✗	✗
Tracz [TW01]	2001	E	✗	✓	?	✗	✓	$\log n$	$\log n$	Onl.	✗
Compact [CHLo5]	2005	✗	✓	✗	?	?	✓	$\log n$	n	Off.	✗
Divisible [PST17]	2017	✗	✓	✗	?	?	✗	1	n	Off.	✗
GNU Taler	2017	FS	✗	✓	✓	✓	✓	$\log n$	$\log n$	Onl.	✓

- **Implementation.** Is there an implementation? Is it proprietary (P), experimental (E), or free software (FS).
- **Offline Spending** Can spending happen offline with delayed detection of double spenders, or is double spending detected immediately during spending?
- **Safe abort/backup.** Is anonymity preserved in the presence of interrupted operations or restoration from backups? Inherently conflicts with offline double spending detection in all approaches that we are aware of. We specify “✓” also for schemes that do not explicitly treat aborts/backup, but would allow a safe implementation when aborts/backups happen.
- **Key expiration.** We specify “?” for schemes that do not explicitly discuss key expiration, but do not fundamentally conflict with the concept.
- **Income transparency.** We specify “✓” if income transparency is supported, “✗” if some feature of the scheme conflicts with income transparency and “?” if it might be possible to add income transparency.
- **No trusted setup.** In a trusted setup, some parameters and cryptographic keys are generated by a trusted third party. A compromise of the trusted setup phase can mean loss of anonymity.
- **Storage for wallet/exchange.** The expected storage for coins adding up to arbitrary value n is specified, with some reasonable upper bound for n .
- **Change/Divisibility.** Can customers pay without possessing exact change? If so, is it handled by giving change online (Onl.) or by divisible coins that support offline operation (Off.)?
- **Receipts & Refunds.** The customer either can prove that they paid for a contract, or they can get their (unlinkable) money back. Also merchants can issue refunds for completed transactions. These operations must not introduce linkability or otherwise compromise the customer’s anonymity.

2.3.2. Blockchains

The term “blockchain” refers to a wide variety of protocols and systems concerned with maintaining a ledger—typically involving financial transactions—in a distributed and decentralized manner.⁴

The first and most prominent system that would be categorized as a “blockchain” today⁵ is Bitcoin [Nako8], published by an individual or group under the alias “Satoshi Nakamoto”. The document timestamping service described in [HS90] could be seen as an even earlier blockchain that predates Bitcoin by about 13 years and is still in use today.

As the name implies, blockchains are made up of a chain of blocks, each block containing updates to the ledger and the hash code of its predecessor block. The chain terminates in a “genesis block” that determines the initial state of the ledger.

Some of the most important decisions for the design of blockchains are the following:

- The *consensus mechanism*, which determines how the participants agree on the current state of the ledger.

In the simplest possible blockchain, a trusted authority would validate transactions and publish new blocks as the head of the chain. In order to increase fault tolerance, multiple trusted authorities can use Byzantine consensus to agree on transactions. With classical Byzantine consensus protocols, this makes the system robust with a malicious minority of up to 1/3 of nodes. While fast and appropriate for some applications, classical Byzantine consensus only works with a known set of participants and does not scale well to many nodes.

Bitcoin instead uses Proof-of-Work (PoW) consensus, where the head of the chain that determines the current ledger state is chosen as the block that provably took the most “work” to construct, including the accumulated work of ancestor blocks. The work consists of finding a hash preimage $n||c$, where c are the contents of the block and n is a nonce, such that the hash $H(n||c)$ ends with a certain number of zeroes (as determined by the difficulty derived from previous blocks). Under the random oracle, the only way to find such a nonce is by trial-and-error. This nonce proves to a verifier that the creator of the block spent computational resources to construct it, and the correctness is easily verified by computing $H(n||c)$. The creator of a block is rewarded with a mining reward and transaction fees for transactions within the block.

PoW consensus is not final: an adversary with enough computational power can create an alternate chain branching off an earlier block. Once this alternative, longer chain is published, the state represented by the earlier branch is discarded. This creates a potential for financial fraud, where

⁴Even though there is a centralization tendency from various sources in practice [Wal19].

⁵The paper that introduces Bitcoin does not mention the term “blockchain”

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

an earlier transaction is reversed by publishing an alternate history that does not contain it. While it was originally believed that PoW consensus process is resistant against attackers that have less than a 51% majority of computational power, closer analysis has shown that a 21% majority suffices [ES18].

A major advantage of PoW consensus is that the participants need not be known beforehand, and that Sybil attacks are impossible since consensus decisions are only dependent on the available computational power, and not on the number of participants.

In practice, PoW consensus is rather slow: Bitcoin can currently support 3-7 transactions per second on a global scale. Some efforts have been made to improve Bitcoin's efficiency [Eya+16; Vuk15], but overall PoW consensus needs to balance speed against security.

Proof-of-Stake (PoS) is a different type of consensus protocol for blockchains, which intends to securely reach consensus without depleting scarce resources such as energy for computation [BGM16; Kwo14]. Blocks are created by randomly selected validators, which obtain a reward for serving as a validator. To avoid Sybil attacks and create economic incentives for good behavior, the probability to get selected as a validator is proportional to one's wealth on the respective blockchain. Realizing PoS has some practical challenges with respect to economic incentives: As blocks do not take work to create, validators can potentially benefit from creating forks, instead of validating on just one chain.

Algorand [Gil+17] avoids some of the problems with PoW consensus by combining some of the ideas of PoW with classical Byzantine consensus protocols. Their proposed system does not have any incentives for validators.

Avalanche [Tea18] has been proposed as a scalable Byzantine Consensus algorithm for use with blockchains. It is based on a gossip protocol and is only shown to work in the synchronous model.

- Membership and visibility. Blockchains such as Bitcoin or Ethereum with public membership and public visibility are called *permissionless blockchains*. Opposed to that, *permissioned* blockchains have been proposed for usage in banking, health and asset tracking applications [And+18].
- Monetary policy and wealth accumulation. Blockchains that are used as cryptocurrencies come with their own monetary policy. In the case of Bitcoin, the currency supply is limited, and due to difficulty increase in mining the currency is deflationary. Other cryptocurrencies such as duniters⁶ have been proposed with built-in rules for inflation, and a basic income mechanism for participants.

⁶See <https://duniters.org/>.

- Expressivity of transactions. Transactions in Bitcoin are small programs in a stack-based programming language that are guaranteed to terminate. Ethereum [Woo14] takes this idea further and allows smart contracts with Turing-complete computation and access to external oracles.
- Governance. Blockchain governance [ROH16; Lev17] is a topic that received relatively little attention so far. As blockchains interact with existing legal and social systems across national borders, different sources of “truth” must be reconciled.

Furthermore, consensus is not just internal to the operation of blockchains, but also external in the development of the technology. Currently small groups of developers create the rules for the operation of blockchains, and likewise have the power to change them. There is currently very little research on social and technological processes to find a “meta-consensus” on the rules that govern such systems, and how these rules can be adapted and changed in a consensus process.

- Anonymity and Zero-Knowledge Proofs. Bitcoin transactions are only pseudonymous, the full transaction history is publicly available and leads to reduced anonymity in practice [RH13]. Tumblers [Bon+14; Hei+17] are an approach to increase the anonymity in Bitcoin-style cryptocurrencies by creating additional transactions to cover up the real owner and sources of funds. While newer tumblers such as TumbleBit [Hei+17] provide rather strong security guarantees, mixing incurs transaction costs.

Some cryptocurrencies have direct support for anonymous transactions [Sun+17]. ZeroCash [Ben+14] uses zero-knowledge proofs to hide the sender, receiver and amount of a transaction. While ZeroCash currently relies on a trusted setup for unforgeability of its currency, more recent proposals dispense with that requirement [Ben+18; Wah+18]. As the anonymity provided by ZeroCash facilitates tax evasion and use in other crimes, an additional, optional layer for privacy-preserving policy for taxation, spending limits and identity escrow has been proposed [GGM16].

Practical guidance on what kind of blockchain is appropriate for an application, and if a blockchain is required in the first place, can be found in [WG17].

2.3.3. Approaches to Micropayments

Micropayments refer to payments of very small value. Microtransactions would not be feasible in traditional payment systems due to high transaction costs, which might even exceed that value that is to be transferred.

Peppercoin

Peppercoin [Rivo4] is a microdonation protocol. The main idea of the protocol is to reduce transaction costs by minimizing the number of transactions that

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

are processed directly by the exchange. Instead of always paying, the customer “gambles” with the merchant for each microdonation. Only if the merchant wins, the microdonation is upgraded to a macropayment to be deposited at the exchange. Peppercoin does not provide customer-anonymity. The proposed statistical method by which exchanges detect fraudulent cooperation between customers and merchants at the expense of the exchange not only creates legal risks for the exchange, but would also require that the exchange learns about microdonations where the merchant did not get upgraded to a macropayment. It is therefore unclear how Peppercoin would actually reduce the computational burden on the exchange.

Tick Payments

Tick payments were proposed by Pedersen [Ped96] as a general technique to amortize the cost for small, recurring payments to the same payee. The payer first makes an up-front deposit as one larger payment that involves the payment processor. To make a micropayment, the payer sends a message to the payee that authorizes the payee to claim a fraction of this deposit. Each further micropayment simply increases the fraction of the deposit that can be claimed, and only requires communication between payer and payee. The payee only needs to show the last message received from the payer to the payment processor in order to receive the accumulated amounts received through tick payments.

Payment Channels and Lightning Network

The Lightning Network [PD16] is a proposed payment system that is meant to run on top of Bitcoin and enable faster, cheaper (micro-)transactions. It is based on establishing *payment channels* between Bitcoin nodes. A payment channel is essentially a tick payment where the deposit and settlement happens on a blockchain. The goal of the Lightning network is to route a payment between two arbitrary nodes by finding a path that connects the two routes through payment channels. The protocol is designed in such a way that a node on the path between the initial sender and final receiver can only receive a payment on a payment channel if it correctly forwards it to the next node.

Experimental deployments of the Lightning network recently suffered heavily from denial-of-service attacks.

BOLT [GM16] is an anonymous payment channel for ZeroCash, and is intended to be used as a building block for a second-layer payment protocol like the Lightning Network.

Side-chains

Side-chains are an alternative approach to improve the scalability of blockchains, intended to be useful in conjunction with arbitrary smart contracts. The approach currently developed by the Ethereum project is described in the Plasma white

paper [PB17]. Side-chains are separate blockchains, possibly with different rules and even consensus protocols than the main chain. Side-chains operate in parallel to the main Ethereum chain, and regularly publish “pointers” to the current head of the sidechain on the main chain. Funds can be moved from the main chain to the side-chain, and subsequently be moved off the side-chain by performing an “exit”, during which the main chain verifies claims to funds on the side-chain according to the side-chain’s rules.

At the time of writing, Plasma is not yet implemented. Potential problems with Plasma include the high costs of exits, lack of access to data needed to verify exit claims, and associated potential for denial-of-service attacks.

2.3.4. Walled Garden Payment Systems

Walled garden payment systems offer ease of use by processing payments using a trusted payment service provider. Here, the customer authenticates to the trusted service, and instructs the payment provider to execute a transaction on their behalf. In these payment systems, the provider basically acts like a bank with accounts carrying balances for the various users. In contrast to traditional banking systems, both customers and merchants are forced to have an account with the same provider. Each user must take the effort to establish his identity with a service provider to create an account. Merchants and customers obtain the best interoperability in return for their account creation efforts if they start with the biggest providers. As a result, there are a few dominating walled garden providers, with AliPay, ApplePay, GooglePay, SamsungPay and PayPal being the current oligopoly.

As with card payment systems, these oligopolies are politically dangerous [Run11], and the lack of competition can result in excessive profit taking that may require political solutions [Jon15] to the resulting market failure. The use of non-standard proprietary interfaces to the payment processing service of these providers serves to reinforce the customer lock-in.

2.3.5. Web Integration

Finally, we will discuss software solutions to web payments. We consider other types of payments, including general payments and in particular hardware solutions as out of scope for this thesis.

Web Payments API

The Web Payments API⁷ is a JavaScript API offered by browsers, and currently still under development. It allows merchant to offer a uniform checkout experience across different payment systems. Unlike GNU Taler, the Web Payments API is

⁷See <https://www.w3.org/TR/payment-request/>

2. GNU Taler, an Income-Transparent Anonymous E-Cash System

only concerned with aspects of the checkout process, such as display of a payment request, selection of a shipping address and selection of a payment method.

Currently only basic-card is supported across popular browsers.

The Payment Handler API⁸ supports the registration of user-defined payment method handlers. Unfortunately the only way to add payment method handlers is via an HTTPS URL. This leaks all information to the payment service provider and precludes the implementation of privacy-preserving payment system handlers.

In order to integrate Taler as a payment method, browsers would need to either offer Taler as a native, built-in payment method or allow an extension to register web payment handlers.

The Web Payments Working Group discontinued work on a HTTP-based API for machine-to-machine payments.⁹

Payment Pointers

Payment pointers are a proposed standard syntax for accounts that are able to receive payments. Unlike `payto://` URIs (discussed in Section 4.2.1), payment pointers do not follow the generic URI syntax and only specify a *pointer* to the receiver's bank account in form of a HTTPS URI. Payment pointers do not specify any mechanism for the payment, but instead direct the user's browser to a website to carry out the payment.

3-D Secure

3-D Secure is a complex and widely deployed protocol that is intended to add an additional security layer on top of credit and debit card transactions.

The 3-D Secure protocol requires the use of inline frames on the HTML page of the merchant for extended verification/authentication of the user. This makes it hard or sometimes – such as when using a mobile browser – even impossible to tell whether the inline frame is legitimate or an attempt to steal information from the user.

Traditionally, merchants bear most of the financial risk, and a key “feature” of the 3DS process compared to traditional card payments is to shift dispute *liability* to the issuer of the card—who may then try to shift it to the customer [MA10, §2.4]. Even in cases where the issuer or the merchant remain legally first in line for liabilities, there are still risks customers incur from the card dispute procedures, such as neither them nor the payment processor noticing fraudulent transactions, or them noticing fraudulent transactions past the *deadline* until which their bank would reimburse them. The customer also typically only has a merchant-generated comment and the amount paid in their credit card statement as a proof for the transaction. Thus, the use of credit cards online does not generate any cryptographically *verifiable* electronic receipts for the customer,

⁸See <https://www.w3.org/TR/payment-handler/>

⁹See <https://www.w3.org/TR/webpayments-http-api/>.

which theoretically enables malicious merchants to later change the terms of the contract.

Beyond these primary issues, customers face secondary risks of identity theft from the personal details exposed by the authentication procedures. In this case, even if the financial damages are ultimately covered by the bank, the customer always has to deal with the procedure of *notifying* the bank in the first place. As a result, customers must remain wary about using their cards, which limits their online shopping [ibi14, p. 50].

Other Proprietary Payment APIs

The Electronic Payment Standard URI scheme `epspayment` : is a proprietary / un-registered URI scheme used by predominantly Austrian banks and merchants to trigger payments from within websites on mobile devices. Merchants can register an invoice with a central server. The user's banking app is associated with the `epspayment` URI scheme and will open to settle the invoice. It lies conceptually between `payto://` and `taler:pay` (see Section 4.1.5). A technical problem of `epspayment` is that when a user has multiple bank accounts at different banks that support `epspayment`, some platforms decide non-deterministically and without asking the user which application to launch. Thus, a user with two banking applications on their phone can often not chose which bank account is used for the payment. If `payto` were widely supported, the problem of registering/choosing bank accounts for payment methods could be centrally addressed by the browser / operating system.

PayPal is a very popular, completely proprietary payment system provider. Its offer-based API is similar in the level of abstraction to Taler's reference merchant backend API.

LaterPay is a proprietary payment system for online content as well as donations. It offers similar functionality to session-bound payments in Taler. LaterPay does not provide any anonymity.

3. Security of Income-Transparent Anonymous E-Cash

We so far discussed Taler’s protocols and security properties only informally. In this chapter, we model a slightly simplified version of the system that we have implemented (see Chapter 4), make our desired security properties more precise, and prove that our protocol instantiation satisfies those properties.

3.1. Introduction to Provable Security

Provable security [GM82; Poi05; Shoo4; Cor00] is a common approach for constructing formal arguments that support the security of a cryptographic protocol with respect to specific security properties and underlying assumptions on cryptographic primitives.

The adversary we consider is computationally bounded, i.e., the run time is typically restricted to be polynomial in the security parameters (such as key length) of the protocol.

Contrary to what the name might suggest, a protocol that is “provably secure” is not necessarily secure in practice [KM07; Dam07]. Instead, provable security results are typically based on reductions of the form “if there is an effective adversary \mathcal{A} against my protocol P , then I can use \mathcal{A} to construct an effective adversary \mathcal{A}' against Q ” where Q is a protocol or primitive that is assumed to be secure or a computational problem that is assumed to be hard. The practical value of a security proof depends on various factors:

- How well-studied is Q ? Some branches of cryptography, for example, some pairing-based constructions, rely on rather complex and exotic underlying problems that are assumed to be hard (but might not be) [KM10].
- How tight is the reduction of Q to P ? A security proof may only show that if P can be solved in time T , the underlying problem Q can be solved (using the hypothetical \mathcal{A}) in time, e.g., $f(T) = T^2$. In practice, this might mean that for P to be secure, it needs to be deployed with a much larger key size or security parameter than Q to be secure.
- What other assumptions are used in the reduction? A common and useful but somewhat controversial assumption is the Random Oracle Model (ROM) [BR93], where the usage of hash functions in a protocol is replaced with

3. Security of Income-Transparent Anonymous E-Cash

queries to a black box (called the Random Oracle), which is effectively a trusted third party that returns a truly random value for each input. Subsequent queries to the Random Oracle with the same value return the same result. While many consider ROM a practical assumption [KM15; BR93], it has been shown that there exist carefully constructed protocols that are secure under the ROM, but are insecure with any concrete hash function [CGH04]. It is an open question whether this result carries over to practical protocols, or just certain classes of artificially constructed protocols of theoretical interest.

Furthermore, a provably secure protocol does not always lend itself easily to a secure implementation, since side channels and fault injection attacks [HTI97; Lom+11] are usually not modeled. Finally, the security properties stated might not be sufficient or complete for the application.

For our purposes, we focus on game-based provable security [BR06; Poi05; Sh004; GSM18] as opposed to simulation-based provable security [GMR89; Lin17], which is another approach to provable security typically used for zero-knowledge proofs and secure multiparty computation protocols.

3.1.1. Algorithms, Oracles and Games

In order to analyze the security of a protocol, the protocol and its desired security properties against an adversary with specific capabilities must first be modeled formally. This part is independent of a concrete instantiation of the protocol; the protocol is only described on a syntactic level.

The possible operations of a protocol (i.e., the protocol syntax) are abstractly defined as the signatures of *algorithms*. Later, the protocol will be instantiated by providing a concrete implementation (formally a program for a probabilistic Turing machine) of each algorithm. A typical public key signature scheme, for example, might consist of the following algorithms:

- $\text{KeyGen}(1^\lambda) \mapsto (\text{sk}, \text{pk})$, a probabilistic algorithm which on input 1^λ generates a fresh key pair consisting of secret key sk of length λ and the corresponding public key pk . Note that 1^λ is the unary representation of λ .¹
- $\text{Sign}(\text{sk}, m) \mapsto \sigma$, an algorithm that signs the bit string m with secret key sk to output the signature σ .
- $\text{Verify}(\text{pk}, \sigma, m) \mapsto b$, an algorithm that determines whether σ is a valid signature on m made with the secret key corresponding to the public key pk . It outputs the flag $b \in \{0, 1\}$ to indicate whether the signature was valid (return value 1) or invalid (return value 0).

¹This formality ensures that the size of the input of the Turing machine program implementing the algorithm will be as least as big as the security parameter. Otherwise the run-time complexity cannot be directly expressed in relation to the size of the input tape.

The abstract syntax could be instantiated with various concrete signature protocols.

In addition to the computational power given to the adversary, the capabilities of the adversary are defined via oracles. The oracles can be thought of as the API² that is given to the adversary and allows the adversary to interact with the environment it is running in. Unlike the algorithms, which the adversary has free access to, the access to oracles is often restricted, and oracles can keep state that is not accessible directly to the adversary. Oracles typically allow the adversary to access information that it normally would not have direct access to, or to trigger operations in the environment running the protocol.

Formally, oracles are an extension to the Turing machine that runs the adversary, which allow the adversary to submit queries to interact with the environment that is running the protocol.

For a signature scheme, the adversary could be given access to an \mathcal{OSign} oracle, which the adversary uses to make the system produce signatures, with secret keys that the adversary does not have direct access to. Different definitions of \mathcal{OSign} lead to different capabilities of the adversary and thus to different security properties later on:

- If the signing oracle $\mathcal{OSign}(m)$ is defined to take a message m and return a signature σ on that message, the adversary gains the power to do chosen message attacks.
- If $\mathcal{OSign}(\cdot)$ was defined to return a pair (σ, m) of a signature σ on a random message m , the power of the adversary would be reduced to a known message attack.

While oracles are used to describe the possible interactions with a system, it is more convenient to describe complex, multi-round interactions involving multiple parties as a special form of an algorithm, called an *interactive protocol*, that takes the identifiers of communicating parties and their (private) inputs as a parameter, and orchestrates the interaction between them. The adversary will then have an oracle to start an instance of that particular interactive protocol and (if desired by the security property being modeled) the ability to drop, modify or inject messages in the interaction. The typically more cumbersome alternative would be to introduce one algorithm and oracle for every individual interaction step.

Security properties are defined via *games*, which are experiments that challenge the adversary to act in a way that would break the desired security property. Games usually consist multiple phases, starting with the setup phase where the challenger generates the parameters (such as encryption keys) for the game. In the subsequent query/response phase, the adversary is given some of the parameters (typically including public keys but excluding secrets) from the setup

²In the modern sense of application programming interface (API), where some system exposes a service with well-defined semantics.

3. Security of Income-Transparent Anonymous E-Cash

phase, and runs with access to oracles. The challenger³ answers oracle queries during that phase. After the adversary's program terminates, the challenger invokes the adversary again with a challenge. The adversary must now compute a final response to the challenger, sometimes with access to oracles. Depending on the answer, the challenger decides if the adversary wins the game or not, i.e., the game returns 0 if the adversary loses and 1 if the adversary wins.

A game for the existential unforgeability of signatures could be formulated like this:

$Exp_A^{EUF}(1^\lambda)$:

1. $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$
2. $(\sigma, m) \leftarrow \mathcal{A}^{\mathcal{OSign}(\cdot)}(pk)$
(Run the adversary with input pk and access to the \mathcal{OSign} oracle.)
3. If the adversary has called $\mathcal{OSign}(\cdot)$ with m as argument, return 0.
4. Return $\text{Verify}(pk, \sigma, m)$.

Here the adversary is run once, with access to the signing oracle. Depending on which definition of \mathcal{OSign} is chosen, the game models existential unforgeability under chosen message attack (EUF-CMA) or existential unforgeability under known message attack (EUF-KMA)

The following modification to the game would model selective unforgeability (SUF-CMA / SUF-KMA):

$Exp_A^{SUF}(1^\lambda)$:

1. $m \leftarrow \mathcal{A}()$
2. $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$
3. $\sigma \leftarrow \mathcal{A}^{\mathcal{OSign}(\cdot)}(pk, m)$
4. If the adversary has called $\mathcal{OSign}(\cdot)$ with m as argument, return 0.
5. Return $\text{Verify}(pk, \sigma, m)$.

Here the adversary has to choose a message to forge a signature for before knowing the message verification key.

After having defined the game, we can now define a security property based on the probability of the adversary winning the game: we say that a signature scheme is secure against existential unforgeability attacks if for every adversary \mathcal{A} (i.e., a polynomial-time probabilistic Turing machine program), the success probability

$$\Pr \left[Exp_A^{EUF}(1^\lambda) = 1 \right]$$

of \mathcal{A} in the EUF game is *negligible* (i.e., grows less fast with λ than the inverse of any polynomial in λ).

³The challenger is conceptually the party or environment that runs the game/experiment.

Note that the EUF and SUF games are related in the following way: an adversary against the SUF game can be easily transformed into an adversary against the EUF game, while the converse does not necessarily hold.

Often security properties are defined in terms of the *advantage* of the adversary. The advantage is a measure of how likely the adversary is to win against the real cryptographic protocol, compared to a perfectly secure version of the protocol. For example, let $\text{Exp}_{\mathcal{A}}^{\text{BIT}}()$ be a game where the adversary has to guess the next bit in the output of a pseudo-random number generator (PRNG). The idealized functionality would be a real random number generator, where the adversary's chance of a correct guess is $1/2$. Thus, the adversary's advantage is

$$\left| \Pr \left[\text{Exp}_{\mathcal{A}}^{\text{BIT}}() \right] - 1/2 \right|.$$

Note that the definition of advantage depends on the game. The above definition, for example, would not work if the adversary had a way to “voluntarily” lose the game by querying an oracle in a forbidden way

3.1.2. Assumptions, Reductions and Game Hopping

The goal of a security proof is to transform an attacker against the protocol under consideration into an attacker against the security of an underlying assumption. Typical examples for common assumptions might be:

- the difficulty of the decisional/computational Diffie–Hellman problem (nicely described by [Bon98])
- existential unforgeability under chosen message attack (EUF-CMA) of a signature scheme [GMR88]
- indistinguishability against chosen-plaintext attacks (IND-CPA) of a symmetric encryption algorithm [Bel+98]

To construct a reduction from an adversary \mathcal{A} against P to an adversary against Q , it is necessary to specify a program R that both interacts as an adversary with the challenger for Q , but at the same time acts as a challenger for the adversary against P . Most importantly, R can choose how to respond to oracle queries from the adversary, as long as R faithfully simulates a challenger for P . The reduction must be efficient, i.e., R must still be a polynomial-time algorithm.

A well-known example for a non-trivial reduction proof is the security proof of FDH-RSA signatures [Cor00].

In practice, reduction proofs are often complex and hard to verify. Game hopping has become a popular technique to manage the complexity of security proofs. The idea behind game hopping proofs is to make a sequence of small modifications starting from initial game, until you arrive at a game where the success probability for the adversary becomes obvious, for example, because the winning state for the adversary becomes unreachable in the code that defines the

3. Security of Income-Transparent Anonymous E-Cash

final game, or because all values the adversary can observe to make a decision are drawn from a truly random and uniform distribution. Each hop modifies the game in a way such that the success probability of game G_n and game G_{n+1} is negligibly close.

Useful techniques for hops are, for example:

- Bridging hops, where the game is syntactically changed but remains semantically equivalent, i.e., $\Pr[G_n = 1] = \Pr[G_{n+1} = 1]$.
- Indistinguishability hops, where some distribution is changed in a way that an adversary that could distinguish between two adjacent games could be turned into an adversary that distinguishes the two distributions. If the success probability to distinguish between those two distributions is ϵ , then $|\Pr[G_n = 1] - \Pr[G_{n+1} = 1]| \leq \epsilon$
- Hops based on small failure events. Here adjacent games proceed identically, until in one of the games a detectable failure event F (such as an adversary visibly forging a signature) occurs. Both games most proceed the same if F does not occur. Then it is easy to show [Sho04] that $|\Pr[G_n = 1] - \Pr[G_{n+1} = 1]| \leq \Pr[F]$

A tutorial introduction to game hopping is given by Shoup [Sho04], while a more formal treatment with a focus on “games as code” can be found in [BR06]. A version of the FDH-RSA security proof based on game hopping was generated with an automated theorem prover by Blanchet and Pointcheval [BP06].

3.1.3. Notation

We prefix public and secret keys with pk and sk , and write $x \xleftarrow{\$} S$ to randomly select an element x from the set S with uniform probability.

3.2. Model and Syntax for Taler

We consider a payment system with a single, static exchange and multiple, dynamically created customers and merchants. The subset of the full Taler protocol that we model includes withdrawing digital coins, spending them with merchants and subsequently depositing them at the exchange, as well as obtaining unlinkable change for partially spent coins with an online “refresh” protocol.

The exchange offers digital coins in multiple denominations, and every denomination has an associated financial value; this mapping is not chosen by the adversary but is a system parameter. We mostly ignore the denomination values here, including their impact on anonymity, in keeping with existing literature [CLM07; PST17]. For anonymity, we believe this amounts to assuming that all

customers have similar financial behavior. We note logarithmic storage, computation and bandwidth demands denominations distributed by powers of a fixed constant, like two.

We do not model fees taken by the exchange. Reserves⁴ are also omitted. Instead of maintaining a reserve balance, withdrawals of different denominations are tracked, effectively assuming every customer has unlimited funds.

Coins can be partially spent by specifying a fraction $0 < f \leq 1$ of the total value associated with the coin's denomination. Unlinkable change below the smallest denomination cannot be given. In practice the unspendable, residual value should be seen as an additional fee charged by the exchange.

Spending multiple coins is modeled non-atomically: to spend (fractions of) multiple coins, they must be spent one-by-one. The individual spend/deposit operations are correlated by a unique identifier for the transaction. In practice, this identifier is the hash $\text{transactionId} = H(\text{contractTerms})$ of the contract terms⁵. Contract terms include a nonce to make them unique, that merchant and customer agreed upon. Note that this transaction identifier and the correlation between multiple spend operations for one payment need not be disclosed to the exchange (it might, however, be necessary to reveal during a detailed tax audit of the merchant): When spending the i -th coin for the transaction with the identifier transactionId , messages to the exchange would only contain $H(i \parallel \text{transactionId})$. This is preferable for merchants that might not want to disclose to the exchange the individual prices of products they sell to customers, but only the total transaction volume over time. For simplicity, we do not include this extra feature in our model.

Our system model tracks the total amount (\equiv financial value) of coins withdrawn by each customer. Customers are identified by their public key pkCustomer . Every customer's wallet keeps track of the following data:

- $\text{wallet}[\text{pkCustomer}]$ contains sets of the customer's coin records, which individually consist of the coin key pair, denomination and exchange's signature.
- $\text{acceptedContracts}[\text{pkCustomer}]$ contains the sets of transaction identifiers accepted by the customer during spending operations, together with coins spent for it and their contributions $0 < f \leq 1$.
- $\text{withdrawIds}[\text{pkCustomer}]$ contains the withdraw identifiers of all withdraw operations that were created for this customer.
- $\text{refreshIds}[\text{pkCustomer}]$ contains the refresh identifiers of all refresh operations that were created for this customer.

The exchange in our model keeps track of the following data:

⁴"Reserve" is Taler's terminology for funds submitted to the exchange that can be converted to digital coins.

⁵The contract terms are a digital representation of an individual offer for a certain product or service the merchant sells for a certain price.

3. Security of Income-Transparent Anonymous E-Cash

- $\text{withdrawn}[\text{pkCustomer}]$ contains the total amount withdrawn by each customer, i.e., the sum of the financial value of the denominations for all coins that were withdrawn by pkCustomer .
- The overspending database of the exchange is modeled by $\text{deposited}[\text{pkCoin}]$ and $\text{refreshed}[\text{pkCoin}]$, which record deposit and refresh operations respectively on each coin. Note that since partial deposits and multiple refreshes to smaller denominations are possible, one deposit and multiple refresh operations can be recorded for a single coin.

We say that a coin is *fresh* if it appears in neither the deposited or refreshed lists nor in acceptedContracts . We say that a coin is being overspent if recording an operation in deposited or refreshed would cause the total spent value from both lists to exceed the value of the coin's denomination. Note that the adversary does not have direct read or write access to these values; instead the adversary needs to use the oracles (defined later) to interact with the system.

We parameterize our system with two security parameters: The general security parameter λ , and the refresh security parameter κ . While λ determines the length of keys and thus the security level, using a larger κ will only decrease the success chance of malicious merchants conspiring with customers to obtain unreported (and thus untaxable) income.

3.2.1. Algorithms

The Taler e-cash scheme is modeled by the following probabilistic⁶ polynomial-time algorithms and interactive protocols. The notation $P(X_1, \dots, X_n)$ stands for a party $P \in \{\mathcal{E}, \mathcal{C}, \mathcal{M}\}$ (Exchange, Customer and Merchant respectively) in an interactive protocol, with X_1, \dots, X_n being the (possibly private) inputs contributed by the party to the protocol. Interactive protocols can access the state maintained by party P .

While the adversary can freely execute the interactive protocols by creating their own parties, the adversary is not given direct access to the private data of parties maintained by the challenger in the security games we define later.

- $\text{ExchangeKeygen}(1^\lambda, 1^\kappa, \mathcal{D}) \mapsto (\text{skE}, \text{pkE})$: Algorithm executed to generate keys for the exchange, with general security parameter λ and refresh security parameter κ , both given as unary numbers. The denomination specification $\mathcal{D} = d_1, \dots, d_n$ is a finite sequence of positive rational numbers that defines the financial value of each generated denomination key pair. We henceforth use \mathcal{D} to refer to some appropriate denomination specification, but our analysis is independent of a particular choice of \mathcal{D} .

The algorithm generates the exchange's master signing key pair $(\text{skESig}, \text{pkESig})$ and denomination secret and public keys $(\text{skD}_1, \dots, \text{skD}_n), (\text{pkD}_1, \dots, \text{pkD}_n)$.

⁶Our Taler instantiations are not necessarily probabilistic (except, e.g., key generation), but we do not want to prohibit this for other instantiations

We write $D(\text{pkD}_i)$, where $D : \{\text{pkD}_i\} \rightarrow \mathcal{D}$ to look up the financial value of denomination pkD_i .

We collectively refer to the exchange's secrets by sksE and to the exchange's public keys together with \mathcal{D} by pksE .

- $\text{CustomerKeygen}(1^\lambda, 1^\kappa) \mapsto (\text{skCustomer}, \text{pkCustomer})$: Key generation algorithm for customers with security parameters λ and κ .
- $\text{MerchantKeygen}(1^\lambda, 1^\kappa) \mapsto (\text{skMerchant}, \text{pkMerchant})$: Key generation algorithm for merchants. Typically the same as CustomerKeygen .
- $\text{WithdrawRequest}(\mathcal{E}(\text{sksE}, \text{pkCustomer}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{pkD})) \mapsto (\mathcal{T}_{WR}, \text{wid})$: Interactive protocol between the exchange and a customer that initiates withdrawing a single coin of a particular denomination.

The customer obtains a withdraw identifier wid from the protocol execution and stores it in $\text{withdrawIds}[\text{pkCustomer}]$.

The WithdrawRequest protocol only initiates a withdrawal. The coin is only obtained and stored in the customer's wallet by executing the WithdrawPickup protocol on the withdraw identifier wid .

The customer and exchange persistently store additional state (if required by the instantiation) such that the customer can use WithdrawPickup to complete withdrawal or to complete a previously interrupted or unfinished withdrawal.

Returns a protocol transcript \mathcal{T}_{WR} of all messages exchanged between the exchange and customer, as well as the withdraw identifier wid .

- $\text{WithdrawPickup}(\mathcal{E}(\text{sksE}, \text{pkCustomer}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{wid})) \mapsto (\mathcal{T}_{WP}, \text{coin})$: Interactive protocol between the exchange and a customer to obtain the coin from a withdraw operation started with WithdrawRequest , identified by the withdraw identifier wid .

The first time WithdrawPickup is run with a particular withdraw identifier wid , the exchange increments $\text{withdrawn}[\text{pkCustomer}]$ by $D(\text{pkD})$, where pkD is the denomination requested in the corresponding WithdrawRequest execution. How exactly pkD is restored depends on the particular instantiation.

The resulting coin

$$\text{coin} = (\text{skCoin}, \text{pkCoin}, \text{pkD}, \text{coinCert}),$$

consisting of secret key skCoin , public key pkCoin , denomination public key pkD and certificate coinCert from the exchange, is stored in the customers wallet $\text{wallet}[\text{pkCustomer}]$.

Executing the WithdrawPickup protocol multiple times with the same customer and the same withdraw identifier does not result in any change

3. Security of Income-Transparent Anonymous E-Cash

of the customer's withdraw balance $\text{withdrawn}[\text{pkCustomer}]$, and results in (re-)adding the same coin to the customer's wallet.

Returns a protocol transcript \mathcal{T}_{WP} of all messages exchanged between the exchange and customer.

- $\text{Spend}(\text{transactionId}, f, \text{coin}, \text{pkMerchant}) \mapsto \text{depositPermission}$: Algorithm to produce and sign a deposit permission depositPermission for a coin under a particular transaction identifier. The fraction $0 < f \leq 1$ determines the fraction of the coin's initial value that will be spent.

The contents of the deposit permission depend on the instantiation, but it must be possible to derive the public coin identifier pkCoin from them.

- $\text{Deposit}(\mathcal{E}(\text{sksE}, \text{pkMerchant}), \mathcal{M}(\text{skMerchant}, \text{pksE}, \text{depositPermission})) \mapsto \mathcal{T}_D$: Interactive protocol between the exchange and a merchant.

From the deposit permission we obtain the pkCoin of the coin to be deposited. If pkCoin is being overspent, the protocol is aborted with an error message to the merchant.

On success, we add depositPermission to $\text{deposited}[\text{pkCoin}]$.

Returns a protocol transcript \mathcal{T}_D of all messages exchanged between the exchange and merchant.

- $\text{RefreshRequest}(\mathcal{E}(\text{sksE}), \mathcal{C}(\text{pkCustomer}, \text{pksE}, \text{coin}_0, \text{pkD}_u)) \rightarrow (\mathcal{T}_{RR}, \text{rid})$ Interactive protocol between exchange and customer that initiates a refresh of coin_0 . Together with RefreshPickup , it allows the customer to convert $D(\text{pkD}_u)$ of the remaining value on coin

$$\text{coin}_0 = (\text{skCoin}_0, \text{pkCoin}_0, \text{pkD}_0, \text{coinCert}_0)$$

into a new, unlinkable coin coin_u of denomination pkD_u .

Multiple refreshes on the same coin are allowed, but each run subtracts the respective financial value of coin_u from the remaining value of coin_0 .

The customer only records the refresh operation identifier rid in $\text{refreshIds}[\text{pkCustomer}]$, but does not yet obtain the new coin. To obtain the new coin, RefreshPickup must be used.

Returns the protocol transcript \mathcal{T}_{RR} and a refresh identifier rid .

- $\text{RefreshPickup}(\mathcal{E}(\text{sksE}, \text{pkCustomer}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{rid})) \rightarrow (\mathcal{T}_{RP}, \text{coin}_u)$: Interactive protocol between exchange and customer to obtain the new coin for a refresh operation previously started with RefreshRequest , identified by the refresh identifier rid .

The exchange learns the target denomination pkD_u and signed source coin $(\text{pkCoin}_0, \text{pkD}_0, \text{coinCert}_0)$. If the source coin is invalid, the exchange aborts the protocol.

The first time RefreshPickup is run for a particular refresh identifier, the exchange records a refresh operation of value $D(\text{pkD}_u)$ in $\text{refreshed}[\text{pkCoin}_0]$. If pkCoin_0 is being overspent, the refresh operation is not recorded in $\text{refreshed}[\text{pkCoin}_0]$, the exchange sends the customer the protocol transcript of the previous deposits and refreshes and aborts the protocol.

If the customer \mathcal{C} plays honestly in RefreshRequest and RefreshPickup, the unlinkable coin coin_u they obtain as change will be stored in their wallet $\text{wallet}[\text{pkCustomer}]$. If \mathcal{C} is caught playing dishonestly, the RefreshPickup protocol aborts.

An honest customer must be able to repeat a RefreshPickup with the same rid multiple times and (re-)obtain the same coin, even if previous RefreshPickup executions were aborted.

Returns a protocol transcript \mathcal{T}_{RP} .

- $\text{Link}(\mathcal{E}(\text{sksE}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{coin}_0)) \rightarrow (\mathcal{T}, (\text{coin}_1, \dots, \text{coin}_n))$: Interactive protocol between exchange and customer. If coin_0 is a coin that was refreshed, the customer can recompute all the coins obtained from previous refreshes on coin_0 , with data obtained from the exchange during the protocol. These coins are added to the customer's wallet $\text{wallet}[\text{pkCustomer}]$ and returned together with the protocol transcript.

3.2.2. Oracles

We now specify how the adversary can interact with the system by defining oracles. Oracles are queried by the adversary, and upon a query the challenger will act according to the oracle's specification. Note that the adversary for the different security games is run with specific oracles, and does not necessarily have access to all oracles simultaneously.

We refer to customers in the parameters to an oracle query simply by their public key. The adversary needs the ability to refer to coins to trigger operations such as spending and refresh, but to model anonymity we cannot give the adversary access to the coins' public keys directly. Therefore we allow the adversary to use the (successful) transcripts of the withdraw, refresh and link protocols to indirectly refer to coins. We refer to this as a coin handle \mathcal{H} . Since the execution of a link protocol results in a transcript \mathcal{T} that can contain multiple coins, the adversary needs to select a particular coin from the transcript via the index i as $\mathcal{H} = (\mathcal{T}, i)$. The respective oracle tries to find the coin that resulted from the transcript given by the adversary. If the transcript has not been seen before in the execution of a link, refresh or withdraw protocol; or the index for a link transcript is invalid, the oracle returns an error to the adversary.

In oracles that trigger the execution of one of the interactive protocols defined in Section 3.2.1, we give the adversary the ability to actively control the communication channels between the exchange, customers and merchants; i.e., the adversary can effectively record, drop, modify and inject messages during the

3. Security of Income-Transparent Anonymous E-Cash

execution of the interactive protocol. Note that this allows the adversary to leave the execution of an interactive protocol in an unfinished state, where one or more parties are still waiting for messages. We use \mathcal{I} to refer to a handle to interactive protocols where the adversary can send and receive messages.

- $\mathcal{O}\text{AddCustomer}() \mapsto \text{pkCustomer}$: Generates a key pair $(\text{skCustomer}, \text{pkCustomer})$ using the `CustomerKeygen` algorithm, and sets

$$\begin{aligned} \text{withdrawn}[\text{pkCustomer}] &:= 0 \\ \text{acceptedContracts}[\text{pkCustomer}] &:= \{\} \\ \text{wallet}[\text{pkCustomer}] &:= \{\} \\ \text{withdrawIds}[\text{pkCustomer}] &:= \{\} \\ \text{refreshIds}[\text{pkCustomer}] &:= \{\}. \end{aligned}$$

Returns the public key of the newly created customer.

- $\mathcal{O}\text{AddMerchant}() \mapsto \text{pkMerchant}$: Generate a key pair $(\text{skMerchant}, \text{pkMerchant})$ using the `MerchantKeygen` algorithm.

Returns the public key of the newly created merchant.

- $\mathcal{O}\text{SendMessage}(\mathcal{I}, P_1, P_2, m) \mapsto ()$: Send message m on the channel from party P_1 to party P_2 in the execution of interactive protocol \mathcal{I} . The oracle does not have a return value.
- $\mathcal{O}\text{ReceiveMessage}(\mathcal{I}, P_1, P_2) \mapsto m$: Read message m in the channel from party P_1 to party P_2 in the execution of interactive protocol \mathcal{I} . If no message is queued in the channel, return $m = \perp$.
- $\mathcal{O}\text{WithdrawRequest}(\text{pkCustomer}, \text{pkD}) \mapsto \mathcal{I}$: Triggers the execution of the `WithdrawRequest` protocol. the adversary full control of the communication channels between customer and exchange.
- $\mathcal{O}\text{WithdrawPickup}(\text{pkCustomer}, \text{pkD}, \mathcal{T}) \mapsto \mathcal{I}$: Triggers the execution of the `WithdrawPickup` protocol, additionally giving the adversary full control of the communication channels between customer and exchange.
The customer and withdraw identifier wid are obtained from the `WithdrawRequest` transcript \mathcal{T} .
- $\mathcal{O}\text{RefreshRequest}(\mathcal{H}, \text{pkD}) \mapsto \mathcal{I}$: Triggers the execution of the `RefreshRequest` protocol with the coin identified by coin handle \mathcal{H} , additionally giving the adversary full control over the communication channels between customer and exchange.
- $\mathcal{O}\text{RefreshPickup}(\mathcal{T}) \mapsto \mathcal{I}$: Triggers the execution of the `RefreshPickup` protocol, where the customer and refresh identifier rid are obtained from the `RefreshRequest` protocol transcript \mathcal{T} .

Additionally gives the adversary full control over the communication channels between customer and exchange.

- $\mathcal{OLink}(\mathcal{H}) \mapsto \mathcal{I}$: Triggers the execution of the Link protocol for the coin referenced by handle \mathcal{H} , additionally giving the adversary full control over the communication channels between customer and exchange.
- $\mathcal{OSpend}(\text{transactionId}, \text{pkCustomer}, \mathcal{H}, \text{pkMerchant}) \mapsto \text{depositPermission}$: Makes a customer sign a deposit permission over a coin identified by handle \mathcal{H} . Returns the deposit permission on success, or \perp if \mathcal{H} is not a coin handle that identifies a coin.

Note that \mathcal{OSpend} can be used to generate deposit permissions that, when deposited, would result in an error due to overspending

Adds $(\text{transactionId}, \text{depositPermission})$ to $\text{acceptedContracts}[\text{pkCustomer}]$.

- $\mathcal{OShare}(\mathcal{H}, \text{pkCustomer}) \mapsto ()$: Shares a coin (identified by handle \mathcal{H}) with the customer identified by pkCustomer , i.e., puts the coin identified by \mathcal{H} into $\text{wallet}[\text{pkCustomer}]$. Intended to be used by the adversary in attempts to violate income transparency. Does not have a return value.

Note that this trivially violates anonymity (by sharing with a corrupted customer), thus the usage must be restricted in some games.

- $\mathcal{OCorruptCustomer}(\text{pkCustomer}) \mapsto (\text{skCustomer}, \text{wallet}[\text{pkCustomer}], \text{acceptedContracts}[\text{pkCustomer}], \text{refreshIds}[\text{pkCustomer}], \text{withdrawIds}[\text{pkCustomer}])$:

Used by the adversary to corrupt a customer, giving the adversary access to the customer's secret key, wallet, withdraw/refresh identifiers and accepted contracts.

Permanently marks the customer as corrupted. There is nothing "special" about corrupted customers, other than that the adversary has used $\mathcal{OCorruptCustomer}$ on them in the past. The adversary cannot modify corrupted customer's wallets directly, and must use the oracle again to obtain an updated view on the corrupted customer's private data.

- $\mathcal{ODeposit}(\text{depositPermission}) \mapsto \mathcal{I}$: Triggers the execution of the Deposit protocol, additionally giving the adversary full control over the communication channels between merchant and exchange.

Returns an error if the deposit permission is addressed to a merchant that was not registered with $\mathcal{OAddMerchant}$.

This oracle does not give the adversary new information, but is used to model the situation where there might be multiple conflicting deposit permissions generated via \mathcal{Spend} , but only a limited number can be deposited.

3. Security of Income-Transparent Anonymous E-Cash

We write $\mathcal{O}_{\text{TALER}}$ for the set of all the oracles we just defined, and $\mathcal{O}_{\text{NoShare}} := \mathcal{O}_{\text{TALER}} - \mathcal{O}_{\text{Share}}$ for all oracles except the share oracle.

The exchange does not need to be corrupted with an oracle. A corrupted exchange is modeled by giving the adversary the appropriate oracles and the exchange secret key from the exchange key generation.

If the adversary determines the exchange's secret key during the setup, invoking $\mathcal{O}_{\text{WithdrawRequest}}$, $\mathcal{O}_{\text{WithdrawPickup}}$, $\mathcal{O}_{\text{RefreshRequest}}$, $\mathcal{O}_{\text{RefreshPickup}}$ or $\mathcal{O}_{\text{Link}}$ can be seen as the adversary playing the exchange. Since the adversary is an active man-in-the-middle in these oracles, it can drop messages to the simulated exchange and make up its own response. If the adversary calls these oracles with a corrupted customer, the adversary plays as the customer.

3.3. Games

We now define four security games (anonymity, conservation, unforgeability and income transparency) that are later used to define the security properties for Taler. Similar to [BR06] we assume that the game and adversary terminate in finite time, and thus random choices made by the challenger and adversary can be taken from a finite sample space.

All games except income transparency return 1 to indicate that the adversary has won and 0 to indicate that the adversary has lost. The income transparency game returns 0 if the adversary has lost, and a positive “laundering ratio” if the adversary won.

3.3.1. Anonymity

Intuitively, an adversary \mathcal{A} (controlling the exchange and merchants) wins the anonymity game if they have a non-negligible advantage in correlating spending operations with the withdrawal or refresh operations that created a coin used in the spending operation.

Let b be the bit that will determine the mapping between customers and spend operations, which the adversary must guess.

We define a helper procedure

$$\text{Refresh}(\mathcal{E}(\text{sksE}), \mathcal{C}(\text{pkCustomer}, \text{pksE}, \text{coin}_0)) \mapsto \mathfrak{R}$$

that refreshes the whole remaining amount on coin_0 with repeated application of RefreshRequest and RefreshPickup using the smallest possible set of target denominations, and returns all protocol transcripts in \mathfrak{R} .

$\text{Exp}_{\mathcal{A}}^{\text{anon}}(1^\lambda, 1^\kappa, b):$

1. $(\text{sksE}, \text{pksE}, \text{skM}, \text{pkM}) \leftarrow \mathcal{A}()$
2. $(\text{pkCustomer}_0, \text{pkCustomer}_1, \text{transactionId}_0, \text{transactionId}_1, f) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{NoShare}}}()$

3. Select distinct fresh coins

$$\text{coin}_0 \in \text{wallet}[\text{pkCustomer}_0]$$

$$\text{coin}_1 \in \text{wallet}[\text{pkCustomer}_1]$$

Return 0 if either pkCustomer_0 or pkCustomer_1 are not registered customers with sufficient fresh coins.

4. For $i \in \{0, 1\}$ run

$$\text{dp}_i \leftarrow \text{Spend}(\text{transactionId}_i, f, \text{coin}_{i-b}, \text{pkM})$$

$$\text{Deposit}(\mathcal{A}(), \mathcal{M}(\text{skM}, \text{pksE}, \text{dp}_i))$$

$$\mathfrak{R}_i \leftarrow \text{Refresh}(\mathcal{A}(), \mathcal{C}(\text{pkCustomer}_i, \text{pksE}, \text{coin}_{i-b}))$$

5. $b' \leftarrow \mathcal{A}^{\mathcal{O}\text{NoShare}}(\mathfrak{R}_0, \mathfrak{R}_1)$

6. Return 0 if $\mathcal{O}\text{Spend}$ was used by the adversary on the coin handles for coin_0 or coin_1 or $\mathcal{O}\text{CorruptCustomer}$ was used on pkCustomer_0 or pkCustomer_1 .

7. If $b = b'$ return 1, otherwise return 0.

Note that unlike some other anonymity games defined in the literature (such as [PST17]), our anonymity game always lets both customers spend in order to avoid having to hide the missing coin in one customer's wallet from the adversary.

3.3.2. Conservation

The adversary wins the conservation game if it can bring an honest customer in a situation where the spendable financial value left in the user's wallet plus the value spent for transactions known to the customer is less than the value withdrawn by the same customer through by the exchange.

In practice, this property is necessary to guarantee that aborted or partially completed withdrawals, payments or refreshes, as well as other (transient) misbehavior from the exchange or merchant do not result in the customer losing money.

$\text{Exp}_{\mathcal{A}}^{\text{conserv}}(1^\lambda, 1^\kappa)$:

1. $(\text{sksE}, \text{pksE}) \leftarrow \text{ExchangeKeygen}(1^\lambda, 1^\kappa, M)$
2. $\text{pkCustomer} \leftarrow \mathcal{A}^{\mathcal{O}\text{NoShare}}(\text{pksE})$
3. Return 0 if pkCustomer is a corrupted user.
4. Run WithdrawPickup for each withdraw identifier wid and RefreshPickup for each refresh identifier rid that the user has recorded in withdrawIds and refreshIds . Run Deposit for all deposit permissions in acceptedContracts .
5. Let v_C be the total financial value left on valid coins in $\text{wallet}[\text{pkCustomer}]$, i.e., the denominated values minus the spend/refresh operations recorded

3. Security of Income-Transparent Anonymous E-Cash

in the exchange's database. Let v_S be the total financial value of contracts in $\text{acceptedContracts}[\text{pkCustomer}]$.

6. Return 1 if $\text{withdrawn}[\text{pkCustomer}] > v_C + v_S$.

Hence we ensure that:

- if a coin was spent, it was spent for a contract that the customer knows about, i.e., in practice the customer could prove that they “own” what they paid for,
- if a coin was refreshed, the customer “owns” the resulting coins, even if the operation was aborted, and
- if the customer withdraws, they can always obtain a coin whenever the exchange accounted for a withdrawal, even when protocol executions are intermittently aborted.

Note that we do not give the adversary access to the $\mathcal{O}\text{Share}$ oracle, since that would trivially allow the adversary to win the conservation game. In practice, conservation only holds for customers that do not share coins with parties that they do not fully trust.

3.3.3. Unforgeability

Intuitively, adversarial customers win if they can obtain more valid coins than they legitimately withdraw.

$\text{Exp}_A^{\text{forge}}(1^\lambda, 1^\kappa)$:

1. $(\text{skE}, \text{pkE}) \leftarrow \text{ExchangeKeygen}()$
2. $(C_0, \dots, C_\ell) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{ALL}}}(\text{pkExchange})$
3. Return 0 if any C_i is not of the form $(\text{skCoin}, \text{pkCoin}, \text{pkD}, \text{coinCert})$ or any coinCert is not a valid signature by pkD on the respective pkCoin .
4. Return 1 if the sum of the unspent value of valid coins in $C_0 \dots, C_\ell$ exceeds the amount withdrawn by corrupted customers, return 0 otherwise.

3.3.4. Income Transparency

Intuitively, the adversary wins if coins are in exclusive control of corrupted customers, but the exchange has no record of withdrawal or spending for them. This presumes that the adversary cannot delete from non-corrupted customer's wallets, even though it can use oracles to force protocol interactions of non-corrupted customers.

For practical e-cash systems, income transparency disincentivizes the emergence of “black markets” among mutually distrusting customers, where currency circulates without the transactions being visible. This is in contrast to some other

proposed e-cash systems and cryptocurrencies, where disintermediation is an explicit goal. The Link protocol introduces the threat of losing exclusive control of coins (despite having the option to refresh them) that were received without being visible as income to the exchange.

$Exp_{\mathcal{A}}^{income}(1^\lambda, 1^\kappa)$:

1. $(skE, pkE) \leftarrow \text{ExchangeKeygen}()$
2. $(\text{coin}_1, \dots, \text{coin}_\ell) \leftarrow \mathcal{A}^{O_{\text{ALL}}}(pkExchange)$
(The coin_i must be coins, including secret key and signature by the denomination, for the adversary to win. However these coins need not be present in any honest or corrupted customer's wallet.)
3. Augment the wallets of all non-corrupted customers with their transitive closure using the Link protocol. Mark all remaining value on coins in wallets of non-corrupted customers as spent in the exchange's database.
4. Let L denote the sum of unspent value on valid coins in $(\text{coin}_1, \dots, \text{coin}_\ell)$, after accounting for the previous update of the exchange's database. Also let w' be the sum of coins withdrawn by corrupted customers. Then $p := L - w'$ gives the adversary's untaxed income.
5. Let w be the sum of coins withdrawn by non-corrupted customers, and s be the value marked as spent by non-corrupted customers, so that $b := w - s$ gives the coins lost during refresh, that is the losses incurred attempting to hide income.
6. If $b + p \neq 0$, return $\frac{p}{b+p}$, i.e., the laundering ratio for attempting to obtain untaxed income. Otherwise return 0.

3.4. Security Definitions

We now give security definitions based upon the games defined in the previous section. Recall that λ is the general security parameter, and κ is the security parameter for income transparency. A polynomial-time adversary is implied to be polynomial in $\lambda + \kappa$.

Definition 3.4.1 (Anonymity). We say that an e-cash scheme satisfies *anonymity* if the success probability $\Pr \left[b \xrightarrow{\$} \{0, 1\} : Exp_{\mathcal{A}}^{anon}(1^\lambda, 1^\kappa, b) = 1 \right]$ of the anonymity game is negligibly close to $1/2$ for any polynomial-time adversary \mathcal{A} .

Definition 3.4.2 (Conservation). We say that an e-cash scheme satisfies *conservation* if the success probability $\Pr [Exp_{\mathcal{A}}^{conserv}(1^\lambda, 1^\kappa) = 1]$ of the conservation game is negligible for any polynomial-time adversary \mathcal{A} .

Definition 3.4.3 (Unforgeability). We say that an e-cash scheme satisfies *unforgeability* if the success probability $\Pr [Exp_{\mathcal{A}}^{forge}(1^\lambda, 1^\kappa) = 1]$ of the unforgeability game is negligible for any polynomial-time adversary \mathcal{A} .

3. Security of Income-Transparent Anonymous E-Cash

Definition 3.4.4 (Strong Income Transparency). We say that an e-cash scheme satisfies *strong income transparency* if the success probability $\Pr \left[\text{Exp}_{\mathcal{A}}^{\text{income}}(1^\lambda, 1^\kappa) \neq 0 \right]$ for the income transparency game is negligible for any polynomial-time adversary \mathcal{A} .

The adversary is said to win one execution of the strong income transparency game if the game's return value is non-zero, i.e., there was at least one successful attempt to obtain untaxed income.

Definition 3.4.5 (Weak Income Transparency). We say that an e-cash scheme satisfies *weak income transparency* if, for any polynomial-time adversary \mathcal{A} , the return value of the income transparency game satisfies

$$E \left[\text{Exp}_{\mathcal{A}}^{\text{income}}(1^\lambda, 1^\kappa) \right] \leq \frac{1}{\kappa}. \quad (3.1)$$

In (3.1), the expectation runs over any probability space used by the adversary and challenger.

For some instantiations, e.g., ones based on zero-knowledge proofs, κ might be a security parameter in the traditional sense. However for an e-cash scheme to be useful in practice, the adversary does not need to have only negligible success probability to win the income transparency game. It suffices that the financial losses of the adversary in the game are a deterrent, after all our purpose of the game is to characterize tax evasion.

Taler does not fulfill strong income transparency, since for Taler κ must be a small cut-and-choose parameter, as the complexity of our cut-and-choose protocol grows linearly with κ . Instead we show that Taler satisfies weak income transparency, which is a statement about the adversary's financial loss when winning the game instead of the winning probability. The return-on-investment (represented by the game's return value) is bounded by $1/\kappa$.

We still characterize strong income transparency, since it might be useful for other instantiations that provide more absolute guarantees.

3.5. Instantiation

We give an instantiation of our protocol syntax that is generic over a blind signature scheme, a signature scheme, a combined signature scheme / key exchange, a collision-resistant hash function and a pseudo-random function family (PRF).

3.5.1. Generic Instantiation

Let `BLINDSIGN` be a blind signature scheme with the following syntax, where the party \mathcal{S} is the signer and \mathcal{R} is the signature requester:

- $\text{KeyGen}_{BS}(1^\lambda) \mapsto (\text{sk}, \text{pk})$ is the key generation algorithm for the signer of the blind signature protocol.
- $\text{Blind}_{BS}(\mathcal{S}(\text{sk}), \mathcal{R}(\text{pk}, m)) \mapsto (\bar{m}, r)$ is a possibly interactive protocol to blind a message m that is to be signed later. The result is a blinded message \bar{m} and a residual r that allows to unblind a blinded signature on m made by sk .
- $\text{Sign}_{BS}(\mathcal{S}(\text{sk}), \mathcal{R}(\bar{m})) \mapsto \bar{\sigma}$ is an algorithm to sign a blinded message \bar{m} . The result $\bar{\sigma}$ is a blinded signature that must be unblinded using the r returned from the corresponding blinding operation before verification.
- $\text{UnblindSig}_{BS}(r, m, \bar{\sigma}) \mapsto \sigma$ is an algorithm to unblind a blinded signature.
- $\text{Verify}_{BS}(\text{pk}, m, \sigma) \mapsto b$ is an algorithm to check the validity of an unblinded blind signature. Returns 1 if the signature σ was valid for m and 0 otherwise.

Note that this syntax excludes some blind signature protocols, such as those with interactive/probabilistic verification or those without a “blinding factor”, where the Blind_{BS} and Sign_{BS} and UnblindSig_{BS} would be merged into one interactive signing protocol. Such blind signature protocols have already been used to construct e-cash [CHLo5].

We require the following two security properties for BLINDSIGN:

- *blindness*: It should be computationally infeasible for a malicious signer to decide which of two messages has been signed first in two executions with an honest user. The corresponding game can be defined as in Abe and Okamoto [AO00], with the additional enhancement that the adversary generates the signing key [SU17].
- *unforgeability*: An adversary that requests k signatures with Sign_{BS} is unable to produce $k + 1$ valid signatures with non-negligible probability.

For more generalized notions of the security of blind signatures see, e.g., [FS09; SU17].

Let COINSIGNKx be combination of a signature scheme and key exchange protocol:

- $\text{KeyGenSec}_{CSK}(1^\lambda) \mapsto \text{sk}$ is a secret key generation algorithm.
- $\text{KeyGenPub}_{CSK}(\text{sk}) \mapsto \text{pk}$ produces the corresponding public key.
- $\text{Sign}_{CSK}(\text{sk}, m) \mapsto \sigma$ produces a signature σ over message m .
- $\text{Verify}_{CSK}(\text{pk}, m, \sigma) \mapsto b$ is a signature verification algorithm. Returns 1 if the signature σ is a valid signature on m by pk , and 0 otherwise.
- $\text{Kx}_{CSK}(\text{sk}_1, \text{pk}_2) \mapsto x$ is a key exchange algorithm that computes the shared secret x from secret key sk_1 and public key pk_2 .

3. Security of Income-Transparent Anonymous E-Cash

We occasionally need these key generation algorithms separately, but we usually combine them into $\text{KeyGen}_{\text{CSK}}(1^\lambda) \mapsto (\text{sk}, \text{pk})$.

We require the following security properties to hold for COINSIGNKX :

- *unforgeability*: The signature scheme $(\text{KeyGen}_{\text{CSK}}, \text{Sign}_{\text{CSK}}, \text{Verify}_{\text{CSK}})$ must satisfy existential unforgeability under chosen message attacks (EUF-CMA).
- *key exchange completeness*: Any probabilistic polynomial-time adversary has only negligible chance to find a degenerate key pair $(\text{sk}_A, \text{pk}_A)$ such that for some honestly generated key pair $(\text{sk}_B, \text{pk}_B) \leftarrow \text{KeyGen}_{\text{CSK}}(1^\lambda)$ the key exchange fails, that is $\text{Kex}_{\text{CSK}}(\text{sk}_A, \text{pk}_B) \neq \text{Kex}_{\text{CSK}}(\text{sk}_B, \text{pk}_A)$, while the adversary can still produce a pair (m, σ) such that $\text{Verify}_{\text{BS}}(\text{pk}_A, m, \sigma) = 1$.
- *key exchange security*: The output of Kx_{CSK} must be computationally indistinguishable from a random shared secret of the same length, for inputs that have been generated with $\text{KeyGen}_{\text{CSK}}$.

Let $\text{SIGN} = (\text{KeyGen}_S, \text{Sign}_S, \text{Verify}_S)$ be a signature scheme that satisfies selective unforgeability under chosen message attacks (SUF-CMA).

Let PRF be a pseudo-random function family and $H : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ a collision-resistant hash function.

Using these primitives, we now instantiate the syntax of our income-transparent e-cash scheme:

- $\text{ExchangeKeygen}(1^\lambda, 1^\kappa, \mathcal{D})$:
Generate the exchange's signing key pair $\text{skESig} \leftarrow \text{KeyGen}_S(1^\lambda)$.
For each element in the sequence $\mathcal{D} = d_1, \dots, d_n$, generate denomination key pair $(\text{skD}_i, \text{pkD}_i) \leftarrow \text{KeyGen}_{\text{BS}}(1^\lambda)$.
- $\text{CustomerKeygen}(1^\lambda, 1^\kappa)$: Return key pair $\text{KeyGen}_S(1^\lambda)$.
- $\text{MerchantKeygen}(1^\lambda, 1^\kappa)$: Return key pair $\text{KeyGen}_S(1^\lambda)$.
- $\text{WithdrawRequest}(\mathcal{E}(\text{sksE}, \text{pkCustomer}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{pkD}))$:
Let skD be the exchange's denomination secret key corresponding to pkD .
1. \mathcal{C} generates coin key pair $(\text{skCoin}, \text{pkCoin}) \leftarrow \text{KeyGen}_{\text{CSK}}(1^\lambda)$
2. \mathcal{C} runs $(\bar{m}, r) \leftarrow \text{Blind}_{\text{CSK}}(\mathcal{E}(\text{skCoin}), \mathcal{C}(m))$ with the exchange

The withdraw identifier is then

$$\text{wid} := (\text{skCoin}, \text{pkCoin}, \bar{m}, r)$$

- $\text{WithdrawPickup}(\mathcal{E}(\text{sksE}, \text{pkCustomer}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{wid}))$:
The customer looks up skCoin , pkCoin , pkD \bar{m} and r via the withdraw identifier wid .

1. \mathcal{C} runs $\bar{\sigma} \leftarrow \text{Sign}_{BS}(\mathcal{E}(\text{skD}), \mathcal{C}(\bar{m}))$ with the exchange
 2. \mathcal{C} unblinds the signature $\sigma \leftarrow \text{UnblindSig}_{BS}(\bar{\sigma}, r, \bar{m})$ and stores the coin $(\text{skCoin}, \text{pkCoin}, \text{pkD}, \sigma)$ in their wallet.
- $\text{Spend}(\text{transactionId}, f, \text{coin}, \text{pkMerchant})$: Let $(\text{skCoin}, \text{pkCoin}, \text{pkD}, \sigma_C) := \text{coin}$. The deposit permission is computed as

$$\text{depositPermission} := (\text{pkCoin}, \sigma_D, m),$$

where

$$\begin{aligned} m &:= (\text{pkCoin}, \text{pkD}, \sigma_C, \text{transactionId}, f, \text{pkMerchant}) \\ \sigma_D &\leftarrow \text{Sign}_{CSK}(\text{skCoin}, m). \end{aligned}$$

- $\text{Deposit}(\mathcal{E}(\text{skSE}, \text{pkMerchant}), \mathcal{M}(\text{skMerchant}, \text{pkSE}, \text{depositPermission}))$: The merchant sends depositPermission to the exchange.

The exchange checks that the deposit permission is well-formed and sets

$$(\text{pkCoin}, \text{pkD}, \sigma_C, \sigma_D, \text{transactionId}, f, \text{pkMerchant}) := \text{depositPermission}$$

The exchange checks the signature on the deposit permission and the validity of the coin with

$$\begin{aligned} b_1 &:= \text{Verify}_{CSK}(\text{pkCoin}, \sigma_D, m) \\ b_2 &:= \text{Verify}_{BS}(\text{pkD}, \sigma_C, \text{pkCoin}) \end{aligned}$$

and aborts if $b_1 = 0$ or $b_2 = 0$.

The exchange aborts if spending f would result in overspending pkCoin based on existing deposit/refresh records, and otherwise marks pkCoin as spent for $D(\text{pkD})$.

- $\text{RefreshRequest}(\mathcal{E}(\text{skSE}, \text{pkCustomer}), \mathcal{C}(\text{skCustomer}, \text{pkSE}, \text{coin}_0, \text{pkD}_u))$:

Let skD_u be the secret key corresponding to pkD_u .

We write

$$\text{Blind}_{BS}^*(\mathcal{S}(\text{sk}, \text{skESig}), \mathcal{R}(R, \text{skR}, \text{pk}, m)) \mapsto (\bar{m}, r, \mathcal{T}_{B*})$$

for a modified version of Blind_{BS} where the signature requester \mathcal{R} takes all randomness from the sequence $(\text{PRF}(R, \text{"blind"} \| n))_{n>0}$, the messages from the exchange are recorded in transcript \mathcal{T}_{B*} , all messages sent by \mathcal{R} are signed with skR and all messages sent by \mathcal{S} are signed with skESig .

Furthermore, we write

$$\text{KeyGen}_{CSK}^*(R, 1^\lambda) \mapsto (\text{sk}, \text{pk})$$

for a modified version of the key generation algorithm that takes randomness from the sequence $(\text{PRF}(R, \text{"key"} \| n))_{n>0}$.

For each $i \in \{1, \dots, \kappa\}$, the customer

3. Security of Income-Transparent Anonymous E-Cash

1. generates seed $s_i \xleftarrow{\$} \{1, \dots, 1^\lambda\}$
2. generates transfer key pair $(t_i, T_i) \leftarrow \text{KeyGen}_{CSK}^*(s_i, 1^\lambda)$
3. computes transfer secret $x_i \leftarrow \text{Kx}(t_i, \text{pkCoin}_0)$
4. computes coin key pair $(\text{skCoin}_i, \text{pkCoin}_i) \leftarrow \text{KeyGen}_{CSK}^*(x_i, 1^\lambda)$
5. and executes the modified blinding protocol

$$(\bar{m}_i, r_i, \mathcal{T}_{(B^*, i)}) \leftarrow \text{Blind}_{BS}^*(\mathcal{E}(\text{skD}_u), \mathcal{C}(x_i, \text{skCoin}_0, \text{pkD}_u, \text{pkCoin}_i))$$

with the exchange.

The customer stores the refresh identifier

$$\text{rid} := (\text{coin}_0, \text{pkD}_u, \{s_i\}, \{\bar{m}_i\}, \{r_i\}, \{\mathcal{T}_{(B^*, i)}\}). \quad (3.2)$$

- $\text{RefreshPickup}(\mathcal{E}(\text{sksE}, \text{pkCustomer}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{rid})) \rightarrow \mathcal{T}$: The customer looks up the refresh identifier rid and recomputes the transfer key pairs, transfer secrets and new coin key pairs.

Then customer sends the commitment $\pi_1 = (\text{pkCoin}_0, \text{pkD}_u, h_C)$ together with signature $\text{sig}_1 \leftarrow \text{Sign}_{CSK}(\text{skCoin}_0, \pi_1)$ to the exchange, where

$$\begin{aligned} h_T &:= H(T_1, \dots, T_\kappa) \\ h_{\bar{m}} &:= H(\bar{m}_1, \dots, \bar{m}_\kappa) \\ h_C &:= H(h_T \| h_{\bar{m}}) \end{aligned}$$

The exchange checks the signature sig_1 , and aborts if invalid. Otherwise, depending on the commitment:

1. If the exchange did not see π_1 before, it marks pkCoin_0 as spent for $D(\text{pkD}_u)$, chooses a uniform random $0 \leq \gamma < \kappa$, stores it, and sends this choice in a signed message (γ, sig_2) to the customer, where $\text{sig}_2 \leftarrow \text{Sign}_S(\text{skESig}, \gamma)$.
2. Otherwise, the exchange sends back the same π_2 as it sent for the last equivalent π_1 .

The customer checks if π_2 differs from a previously received π_2' for the same request π_1 , and aborts if such a conflicting response was found. Otherwise, the customer in response to π_2 sends the reveal message

$$\pi_3 = T_\gamma, \bar{m}_\gamma, (s_1, \dots, s_{\gamma-1}, s_{\gamma+1}, \dots, s_\kappa)$$

and signature

$$\text{sig}_{3'} \leftarrow \text{Sign}_{CSK}(\text{skCoin}_0, (\text{pkCoin}_0, \text{pkD}_u, \mathcal{T}_{(B^*, \gamma)}, T_\gamma, \bar{m}_\gamma))$$

to the exchange. Note that $\text{sig}_{3'}$ is not a signature over the full reveal message, but is primarily used in the linking protocol for checks by the customer.

The exchange checks the signature $\text{sig}_{3'}$ and then computes for $i \neq \gamma$:

$$\begin{aligned} (t'_i, T'_i) &\leftarrow \text{KeyGen}_{\text{CSK}}^*(s_i, 1^\lambda) \\ x'_i &\leftarrow \text{Kx}(t_i, \text{pkCoin}_0) \\ (\text{skCoin}'_i, \text{pkCoin}'_i) &\leftarrow \text{KeyGen}_{\text{CSK}}^*(x'_i, 1^\lambda) \\ h'_T &:= H(T'_1, \dots, T'_{\gamma-1}, T_\gamma, T'_{\gamma+1}, \dots, T'_\kappa) \end{aligned}$$

and simulates the blinding protocol with recorded transcripts (without signing each message, as indicated by the dot (\cdot) instead of a signing secret key), obtaining

$$(\bar{m}'_i, r'_i, T_i) \leftarrow \text{Blind}_{\text{BS}}^*(\mathcal{S}(\text{skD}_u), \mathcal{R}(x'_i, \cdot, \text{pkD}_u, \text{skCoin}'_i))$$

and finally

$$\begin{aligned} h'_m &:= H(\bar{m}'_1, \dots, \bar{m}'_{\gamma-1}, \bar{m}_\gamma, \bar{m}'_{\gamma+1}, \dots, \bar{m}'_\kappa) \\ h'_C &:= H(h'_T \| h'_m). \end{aligned}$$

Now the exchange checks if $h_C = h'_C$, and aborts the protocol if the check fails. Otherwise, the exchange sends a message back to \mathcal{C} that the commitment verification succeeded and includes the signature

$$\bar{\sigma}_\gamma := \text{Sign}_{\text{BS}}(\mathcal{E}(\text{skD}_u), \mathcal{C}(\bar{m}_\gamma)).$$

As a last step, the customer obtains the signature σ_γ on the new coin's public key pkCoin_u with

$$\sigma_\gamma := \text{UnblindSig}(r_\gamma, \text{pkCoin}_\gamma, \bar{\sigma}_\gamma).$$

Thus, the new, unlinkable coin is $\text{coin}_u := (\text{skCoin}_\gamma, \text{pkCoin}_\gamma, \text{pkD}_u, \sigma_\gamma)$.

- $\text{Link}(\mathcal{E}(\text{sksE}), \mathcal{C}(\text{skCustomer}, \text{pksE}, \text{coin}_0))$: The customer sends the public key pkCoin_0 of coin_0 to the exchange.

For each completed refresh on pkCoin_0 recorded in the exchange's database, the exchange sends the following data back to the customer: the signed commit message (sig_1, π_1) , the transfer public key T_γ , the signature $\text{sig}_{3'}$, the blinded signature $\bar{\sigma}_\gamma$, and the transcript $\mathcal{T}_{(B^*, \gamma)}$ of the customer's and exchange's messages during the $\text{Blind}_{\text{BS}}^*$ protocol execution.

The following logic is repeated by the customer for each response:

1. Verify the signatures (both from pkESig and pkCoin_0) on the transcript $\mathcal{T}_{(B^*, \gamma)}$, abort otherwise.
2. Verify that sig_1 is a valid signature on π_1 by pkCoin_0 , abort otherwise.

3. Security of Income-Transparent Anonymous E-Cash

3. Re-compute the transfer secret and the new coin's key pair as

$$\begin{aligned} x_\gamma &\leftarrow \text{Kx}_{\text{CSK}}(\text{skCoin}_0, T_\gamma) \\ (\text{skCoin}_\gamma, \text{pkCoin}_\gamma) &\leftarrow \text{KeyGen}_{\text{CSK}}^*(x_\gamma, 1^\lambda). \end{aligned}$$

4. Simulate the blinding protocol with the message transcript received from the exchange to obtain $(\bar{m}_\gamma, r_\gamma)$.
5. Check that $\text{Verify}_{\text{CSK}}(\text{pkCoin}_0, \text{pkD}_u, \text{skCoin}_0, (\mathcal{T}_{(B^*, \gamma)}, \bar{m}_\gamma), \text{sig}_{3'})$ indicates a valid signature, abort otherwise.
6. Unblind the signature to obtain $\sigma_\gamma \leftarrow \text{UnblindSig}(r_\gamma, \text{pkCoin}_\gamma, \bar{\sigma}_\gamma)$
7. (Re-)add the coin $(\text{skCoin}_\gamma, \text{pkCoin}_\gamma, \text{pkD}_u, \sigma_\gamma)$ to the customer's wallet.

3.5.2. Concrete Instantiation

We now give a concrete instantiation that is used in the implementation of GNU Taler for the schemes `BLINDSIGN`, `COINSIGNKx` and `SIGN`.

For `BLINDSIGN`, we use RSA-FDH blind signatures [Cha83; BR96]. From the information-theoretic security of blinding, the computational blindness property follows directly. For the unforgeability property, we additionally rely on the RSA-KTI assumption as discussed in [Bel+03]. Note that since the blinding step in RSA blind signatures is non-interactive, storage and verification of the transcript is omitted in refresh and link.

We instantiate `COINSIGNKx` with signatures and key exchange operations on elliptic curves in Edwards form, where the same key is used for signatures and the Diffie–Hellman key exchange operations. In practice, we use Ed25519 [Ber+12] / Curve25519 [Bero6] for $\lambda = 256$. We caution that some other elliptic curve key exchange implementation might not satisfy the completeness property that we require, due to the lack of complete addition laws.

For `SIGN`, we use elliptic-curve signatures, concretely Ed25519. For the collision-resistant hash function H we use SHA-512 [H306] and HKDF [KE10] as a PRF.

3.6. Proofs

We now give proofs for the security properties defined in Section 3.4 with the generic instantiation of Taler.

3.6.1. Anonymity

Theorem 1. *Assuming*

- the blindness of `BLINDSIGN`,
- the unforgeability and key exchange security of `COINSIGNKx`, and

- the collision resistance of H ,

our instantiation satisfies anonymity.

Proof. We give a proof via a sequence of games $\mathbb{G}_0(b), \mathbb{G}_1(b), \mathbb{G}_2(b)$, where $\mathbb{G}_0(b)$ is the original anonymity game $\text{Exp}_{\mathcal{A}}^{\text{anon}}(1^\lambda, 1^\kappa, b)$. We show that the adversary can distinguish between subsequent games with only negligible probability. Let ϵ_{HC} and ϵ_{KX} be the advantage of an adversary for finding hash collisions and for breaking the security of the key exchange, respectively.

We define \mathbb{G}_1 by replacing the link oracle \mathcal{OLink} with a modified version that behaves the same as \mathcal{OLink} , unless the adversary responds with link data that did not occur in the transcript of a successful refresh operation, but despite of that still passes the customer's verification. In that case, the game is aborted instead.

Observe that in case this failure event happens, the adversary must have forged a signature on sig_3 on values not signed by the customer, yielding an existential forgery. Thus, $|\Pr[\mathbb{G}_0 = 1] - \Pr[\mathbb{G}_1 = 1]|$ is negligible.

In \mathbb{G}_2 , the refresh oracle is modified so that the customer responds with value drawn from a uniform random distribution D_1 for the γ -th commitment instead of using the key exchange function. We must handle the fact that γ is chosen by the adversary after seeing the commitments, so the challenger first makes a guess γ^* and replaces only the γ^* -th commitment with a uniform random value. If the γ chosen by the adversary does not match γ^* , then the challenger rewinds \mathcal{A} to the point where the refresh oracle was called. Note that we only replace the one commitment that will not be opened to the adversary later.

Since $\kappa \ll \lambda$ and the security property of Kx guarantees that the adversary cannot distinguish the result of a key exchange from randomness, the runtime complexity of the challenger still stays polynomial in λ . An adversary that could with high probability choose a γ that would cause a rewind, could also distinguish randomness from the output of Kx .

We now show that $|\Pr[\mathbb{G}_1 = 1] - \Pr[\mathbb{G}_2 = 1]| \leq \epsilon_{KX}$ by defining a distinguishing game $\mathbb{G}_{1 \sim 2}$ for the key exchange as follows:

$\mathbb{G}_{1 \sim 2}(b)$:

1. If $b = 0$, set

$$D_0 := \{(A, B, \text{Kex}(a, B)) \mid (a, A) \leftarrow \text{KeyGen}(1^\lambda), (b, B) \leftarrow \text{KeyGen}(1^\lambda)\}.$$

Otherwise, set

$$D_1 := \{(A, B, C) \mid (a, A) \leftarrow \text{KeyGen}(1^\lambda), (b, B) \leftarrow \text{KeyGen}(1^\lambda), C \xleftarrow{\$} \{1, \dots, 2^\lambda\}\}.$$

2. Return $\text{Exp}_{\mathcal{A}}^{\text{anon}}(b, D_b)$

(Modified anonymity game where the γ -th commitment in the refresh oracle is drawn uniformly from D_b (using rewinding). Technically, we need to draw from D_b on withdraw for the coin secret key, write it to a table, look it up on refresh and use the matching tuple, so that with $b = 0$ we perfectly simulate \mathbb{G}_1 .)

3. Security of Income-Transparent Anonymous E-Cash

Depending on the coin flip b , we either simulate G_1 or G_2 perfectly for our adversary \mathcal{A} against G_1 . At the same time b determines whether \mathcal{A} receives the result of the key exchange or real randomness. Thus, $|\Pr[G_1 = 1] - \Pr[G_2 = 1]| = \epsilon_{KX}$ is exactly the advantage of $G_{1 \sim 2}$.

We observe in G_2 that as x_γ is uniform random and not learned by the adversary, the generation of $(\text{skCoin}_\gamma, \text{pkCoin}_\gamma)$ and the execution of the blinding protocol is equivalent (under the PRF assumption) to using the randomized algorithms KeyGen_{CSK} and Blind_{BS} .

By the blindness of the BLINDSIGN scheme, the adversary is not able to distinguish blinded values from randomness. Thus, the adversary is unable to correlate a Sign_{BS} operation in refresh or withdraw with the unblinded value observed during Deposit.

We conclude the success probability for G_2 must be $1/2$ and hence the success probability for $\text{Exp}_{\mathcal{A}}^{\text{anon}}(1^\lambda, \kappa, b)$ is at most $1/2 + \epsilon(\lambda)$, where ϵ is a negligible function. \square

3.6.2. Conservation

Theorem 2. *Assuming existential unforgeability (EUF-CMA) of COINSIGNKx , our instantiation satisfies conservation.*

Proof. In honest executions, we have $\text{withdrawn}[\text{pkCustomer}] = v_C + v_S$, i.e., the coins withdrawn add up to the coins still available and the coins spent for known transactions.

In order to win the conservation game, the adversary must increase $\text{withdrawn}[\text{pkCustomer}]$ or decrease v_C or v_S . An adversary can abort withdraw operations, thus causing $\text{withdrawn}[\text{pkCustomer}]$ to increase, while the customer does not obtain any coins. However, in step 4, the customer obtains coins from interrupted withdraw operations. Similarly, for the refresh protocol, aborted RefreshRequest / RefreshPickup operations that result in a coin's remaining value being reduced are completed in step 4.

Thus, the only remaining option for the adversary is to decrease v_C or v_S with the $\mathcal{O}\text{RefreshPickup}$ and $\mathcal{O}\text{Deposit}$ oracles, respectively.

Since the exchange verifies signatures made by the secret key of the coin that is being spent/refreshed, the adversary must forge this signature or have access to the coin's secret key. As we do not give the adversary access to the sharing oracle, it does not have direct access to any of the honest customer's coin secret keys.

Thus, the adversary must either compute the coin's secret key from observing the coin's public key (e.g., during a partial deposit operation), or forge signatures directly. Both possibilities allow us to carry out a reduction against the unforgeability property of the COINSIGNKx scheme, by injecting the challenger's public key into one of the coins. \square

3.6.3. Unforgeability

Theorem 3. *Assuming the unforgeability of `BLINDSIGN`, our instantiation satisfies unforgeability.*

Proof. The adversary must have produced at least one coin that was not blindly signed by the exchange. In order to carry out a reduction from this adversary to a blind signature forgery, we inject the challenger's public key into one randomly chosen denomination. Since we do not have access to the corresponding secret key of the challenger, signing operations for this denomination are replaced with calls to the challenger's signing oracle in `OWithdrawPickup` and `ORefreshPickup`. For n denominations, an adversary against the unforgeability game would produce a blind signature forgery with probability $1/n$. \square

3.6.4. Income Transparency

Theorem 4. *Assuming*

- *the unforgeability of `BLINDSIGN`,*
- *the key exchange completeness of `COINSIGNKX`,*
- *the pseudo-random function property of `PRF`, and*
- *the collision resistance of H ,*

our instantiation satisfies weak income transparency.

Proof. We consider the directed forest on coins induced by the refresh protocol. It follows from unforgeability that any coin must originate from some customer's withdraw in this graph. We may assume that all $\text{coin}_1, \dots, \text{coin}_l$ originate from non-corrupted users, for some $l \leq \ell$.

For any $i \leq l$, there is a final refresh operation R_i in which a non-corrupted user could obtain the coin C' consumed in the refresh via the linking protocol, but no non-corrupted user could obtain the coin provided by the refresh, as otherwise coin_i gets marked as spent in step 3. Set $F := \{R_i \mid i \leq l\}$.

During each $R_i \in F$, our adversary must have submitted a blinded coin and transfer public key for which the linking protocol fails to produce the resulting coin correctly, otherwise the coin would have been spent in step 3. In this case, we consider several non-exclusive cases

1. the execution of the refresh protocol is incomplete,
2. the commitment for the γ -th blinded coin and transfer public key is dishonest,
3. a commitment for a blinded coin and transfer public key other than the γ -th is dishonest,

3. Security of Income-Transparent Anonymous E-Cash

We show these to be exhaustive by assuming their converses all hold: As the commitment is signed by skCoin_0 , our key exchange completeness assumption of COIN_SIGN_KX applies to the coin public key. Any revealed values must match our honestly computed commitments, as otherwise a collision in H would have been found. We assumed the revealed γ -th transfer public key is honest. Hence our key exchange completeness assumption of COIN_SIGN_KX yields $\text{Kex}_{\text{CSK}}(t, C') = \text{Kex}_{\text{CSK}}(c', T)$ where $T = \text{KeyGenPub}_{\text{CSK}}(t)$ is the transfer key, thus the customer obtains the correct transfer secret. We assumed the refresh concluded and all submissions besides the γ -th were honest, so the exchange correctly reveals the signed blinded coin. We assumed the γ -th blinded coin is correct too, so customer now re-compute the new coin correctly, violating $R_i \in F$.

We shall prove

$$E \left[\frac{p}{b+p} \middle| F \neq \emptyset \right] = \frac{1}{\kappa} \quad (3.3)$$

where the expectation runs over any probability space used by the adversary and challenger.

We shall now consider executions of the income transparency game with an optimal adversary with respect to maximizing $\frac{p}{b+p}$. Note that this is permissible since we are not carrying out a reduction, but are interested in the expectation of the game's return value.

As a reminder, if a refresh operation is initiated using a false commitment that is detected by the exchange, then the new coin cannot be obtained, and contributes to the lost coins $b := w - s$ instead of the winnings $p := L - w'$. We also note $b + p$ gives the value of refreshes attempted with false commitments. As these are non-negative, $\frac{p}{b+p}$ is undefined if and only if $p = 0$ and $b = 0$, which happens if and only if the adversary does not use false commitments, i.e., $F = \emptyset$.

We may now assume for optimality that \mathcal{A} submits a false commitment for at most one choice of γ in any $R_i \in F$, as otherwise the refresh always fails. Furthermore, for an optimal adversary we can exclude refreshes in F that are incomplete, but that would be possible to complete successfully, as completing such a refresh would only increase the adversaries winnings.

We emphasize that an adversary that loses an R_i loses the coin that would have resulted from it completely, while an optimal adversary who wins an R_i should not gamble again. Indeed, an adversary has no reason to touch its winnings from an R_i .

For any R_i , there are κ game runs identical up through the commitment phase of R_i and exhibiting different outcomes based on the challenger's random choice of γ . If v_i is the financial value of the coin resulting from refresh operation R_i then one of the possible runs adds v_i to p , while the remaining $\kappa - 1$ runs add v_i to b .

We define p_i and b_i to be these contributions summed over the κ possible runs,

i.e.,

$$\begin{aligned} p_i &:= v_i \\ b_i &= (\kappa - 1)v_i \end{aligned}$$

The adversary will succeed in $1/\kappa$ runs ($p_i = v$) and loses in $(\kappa - 1)/\kappa$ runs ($p_i = 0$). Hence:

$$\begin{aligned} E \left[\frac{p}{b+p} \middle| F \neq \emptyset \right] &= \frac{1}{|F|} \sum_{R_i \in F} \frac{p_i}{b_i + p_i} \\ &= \frac{1}{\kappa|F|} \sum_{R_i \in F} \frac{v_i}{0 + v_i} + \frac{\kappa - 1}{\kappa|F|} \sum_{R_i \in F} \frac{0}{v_i + 0} \\ &= \frac{1}{\kappa}, \end{aligned}$$

which yields the equality (3.3).

As for $F = \emptyset$, the return value of the game must be 0, we conclude

$$E \left[\text{Exp}_{\mathcal{A}}^{\text{income}}(1^\lambda, 1^\kappa) \right] \leq \frac{1}{\kappa}.$$

□

3.7. Discussion

3.7.1. Limitations

Not all features of our implementation are part of the security model and proofs. In particular, the following features are left out of the formal discussion:

- **Reserves.** In our formal model, we effectively assume that every customer has access to exactly one unlimited reserve.
- **Offline and online keys.** In our implementation, the exchange has one offline master signing key, and online signing keys with a shorter live span.
- **Refunds** allow merchants to effectively “undo” a deposit operation before the exchange settles the transaction with the merchant. This simple extension preserves unlinkability of payments through refresh.
- **Timeouts.** In practice, a merchant gives the customer a deadline until which the payment for a contract must have been completed, potentially by using multiple coins.

If a customer is unable to complete a payment (e.g., because they notice that their coins are already spent after a restore from backup), a refund for this partial payment can be requested from the merchant.

3. *Security of Income-Transparent Anonymous E-Cash*

Should the merchant become unavailable after a partially completed payment, there are two possibilities: Either the customer can deposit the coins on behalf of the merchant to obtain proof of their on-time payment, which can be used in a later arbitration if necessary. Alternatively, the customer can ask the exchange to undo the partial payments, though this requires the exchange to know (or learn from the customer) the exact amount to be paid for the contract.

- The fees incurred for operations, the protocols for backup and synchronization as well as other possible extensions like tick payments are not formally modeled.

We note that customer tipping (see 2.1.11) basically amounts to an execution of the Withdraw protocol where the party that generates the coin keys and blinding factors (in that case the merchant's customer) is different from the party that signs the withdraw request (the merchant with a "customer" key pair tied to the merchant's bank account). While this is desirable in some cases, we discussed in 2.1.11 how this "loophole" for a one-hop untaxed payment could be avoided.

3.7.2. Other Properties

Exculpability

Exculpability is a property of offline e-cash which guarantees that honest users cannot be falsely blamed for misbehavior such as double spending. For online e-cash it is not necessary, since coins are spent online with the exchange. In practice, even offline e-cash systems that provide exculpability are often undesirable, since hardware failures can result in unintentional overspending by honest users. If a device crashes after an offline coin has been sent to the merchant but before the write operation has been permanently recorded on the user's device (e.g., because it was not yet flushed from the cache to a hard drive), the next payment will cause a double spend, resulting in anonymity loss and a penalty for the customer.

Fair Exchange

The Endorsed E-Cash system by Camenisch et al. [CLM07] allows for fair exchange—sometimes called atomic swap in the context of cryptocurrencies—of online or offline e-cash against digital goods. The online version of Camenisch's protocol does not protect the customer against loss of anonymity from linkability of aborted fair exchanges.

Taler's refresh protocol can be used for fair exchange of online e-cash against digital goods, without any loss of anonymity due to linkability of aborted transactions, with the following small extension: The customer asks the exchange to *lock coins to a merchant* until a timeout. Until the timeout occurs, the exchange

provides the merchant with a guarantee that these coins can only be spent with this specific merchant, or not at all. The fair exchange exchanges the merchant's digital goods against the customer's deposit permissions for the locked coins. On aborted fair exchanges, the customer refreshes to obtain unlinkable coins.

4. Implementation of GNU Taler

This chapter describes the implementation of GNU Taler in detail. Concrete design decisions, protocol details and our reference implementation are discussed.

We implemented the GNU Taler protocol in the context of a payment system for the web, as shown in Figure 2.1. The system was designed for real-world usage with current web technologies and within existing financial systems.

The following technical goals and constraints influenced the design of the concrete protocol and implementation:

- The implementation should allow payments in browsers with hardened security settings. In particular, it must be possible to make a payment without executing JavaScript on a merchant's website and without having to store (session-)cookies or requiring a login.
- Cryptographic evidence should be available to all parties in case of a dispute.
- In addition to the guarantees provided by the GNU Taler protocol, the implementation must take care to not introduce additional threats to security and privacy. Features that trade privacy for convenience should be clearly communicated to the user, and the user must have the choice to deactivate them. Integration with the web should minimize the potential for additional user tracking.
- The integration for merchants must be simple. In particular, merchants should not have to write code involving cryptographic operations or have to manage Taler-specific secrets in their own application processes.
- The web integration must not be specific to a single browser platform, but instead must be able to use the lowest common denominator of what is currently available. User experience enhancements supported for only specific platforms are possible, but fallbacks must be provided for other platforms.
- URLs should be clean, user-friendly and must have the expected semantics when sharing them with others or revisiting them after a session expired.
- Multiple currencies must be supported. Conversion between different currencies is out of scope.
- The implementation should offer flexibility with regards to what context or applications it can be used for. In particular, the implementation must make

4. Implementation of GNU Taler

it possible to provide plugins for different underlying banking systems and provide hooks to deal with different regulatory requirements.

- The implementation must be robust against network failures and crash faults, and recover as gracefully as possible from data loss. Operations must be idempotent if possible, e.g., accidentally clicking a payment button twice should only result in one payment, and refreshing a page should not lead to failures in the payment process.
- Authorization should be preferred to authentication. In particular, there should be no situations in which the user must enter confidential information on a page that cannot be clearly identified as secure.
- No flickering or unnecessary redirects. To complete a payment, the number of request, especially in the user's navigation context, should be minimized.
- While the implementation should integrate well with browsers, it must be possible to request and make payments without a browser. This makes at least part of the implementation completely independent of the extremely complex browser standards, and makes Taler usable for machine-to-machine payments.

We now recapitulate how a GNU Taler payment works, with some more details specific to the implementation.

By instructing their bank to send money to an exchange, the customer creates a (non-anonymous) balance, called a *reserve*, at the exchange. Once the exchange has received and processed the bank transfer, the customer's *wallet* automatically *drains* the reserve by withdrawing coins from it until the reserve is empty. Withdrawing immediately before a purchase should be avoided, as it decreases the customer's anonymity set by creating a correlation between the non-anonymous withdrawal and the spending.

To withdraw coins from the exchange, the customer's wallet authenticates itself using an Ed25519 private key for the customer's reserve. The customer must include the corresponding reserve public key in the payment instruction from the customer's bank to the exchange's bank that funded their reserve. With a bank that directly supports Taler on their online banking website, this process is streamlined for the user, since the wallet automatically creates the key pair for the reserve and adds the public key to the payment instruction.

While browsing a merchant's website, the website can signal the wallet to request a payment from a user. The user is then asked to confirm or reject this proposal. If the user accepts, the wallet spends coins with the merchant. The merchant deposits coins received from the customer's wallet at the exchange. Since bank transfers are usually costly, the exchange delays and aggregates multiple deposits into a bigger wire transfer. This allows GNU Taler to be used even for microtransactions of amounts smaller than usually handled by the underlying banking system.




Figure 4.1.: The different components of the Taler system in the context of a banking system providing money creation, wire transfers and authentication. (Auditor omitted.)

As shown in Figure 4.1, the merchant is internally split into multiple components. The implementation of the Taler protocol and cryptographic operations is isolated into a separate component, called the *merchant backend*, which the merchant accesses through an API or software development kit (SDK) in the programming language of their choice.

Our implementations of the exchange (70,000 LOC) and merchant backend (20,000 LOC) are written in C using PostgreSQL as the database and libgcrypt for cryptographic operations. The *wallet* (10,000 LOC) is implemented in TypeScript as a cross-browser extension using the WebExtensions API, which is available for a majority of widely used browsers. It also uses libgcrypt (compiled to JavaScript) for cryptographic operations as the required primitives are not yet natively supported by web browsers. Sample merchant websites (1,000 LOC) and an example bank (2,000 LOC) with tight Taler integration are provided in Python.

The code is available at <https://git.taler.net/> and a demo is publicly available at <https://demo.taler.net/>.

4.1. Overview

We provide a high-level overview over the implementation, before discussing the respective components in detail.

4.1.1. Taler APIs

The components of Taler communicate over an HTTP-based, RESTful¹ [FT00] API. All request payloads and responses are JSON [Bra17] documents.

Binary data (such as key material, signatures and hashes) is encoded as a base32-crockford [Cro] string. Base32-crockford is a simple, case-insensitive encoding of binary data into a subset of the ASCII alphabet that encodes 5 bits per character. While this is not the most space-efficient encoding, it is relatively resilient against human transcription errors.

Financial amounts are treated as fixed-point decimal numbers. The implementation internally uses a pair of integers (v, f) with value part $0 \leq v \leq 2^{52}$ and fractional part $0 \leq f < 10^8$ to represent the amount $a = v + f \cdot 10^{-8}$. This representation was chosen as the smallest representable amount is equal to one Satoshi (the smallest representable amount in Bitcoin), and the largest possible value part (besides being large enough for typical financial applications) is still accurately representable in 64-bit IEEE 754 floating point numbers. These constraints are useful as some languages such as JavaScript² provide IEEE 753 floating point numbers as the only numeric type. More importantly, fixed-point decimal numbers allow exact representation of decimal values (say 0.10 €), which is not possible with floating point numbers but essential in financial applications.

Signatures are made over custom binary representations of the respective values, prefixed with a 64-bit tag consisting of the size of the message (32 bits) and an integer tag (32 bits) uniquely identifying the purpose of the message. To sign a free-form JSON object, a canonical representation as a string is created by removing all white space and sorting objects' fields.

In the future, more space-efficient representations (such as BSON³ or CBOR [BH13]) could be used. The representation can be negotiated between client and server in a backwards-compatible way with the HTTP "Accept" header.

4.1.2. Cryptographic Algorithms

The following cryptographic primitives are used by Taler:

- SHA512 [H306] as a cryptographic hash function
- Ed25519 [Bero6] for non-blind signing operations
- Curve25519 [Bero6] for the refreshing operation
- HKDF [KE10] as a key derivation function for the refreshing operation
- FDH-RSA blind signatures [Bel+03]

¹Some REST purists might disagree, because the Taler APIs do not follow all REST principles religiously. In particular, the HATEOAS principle is not followed.

²Big integers are currently in the process of being added to the JavaScript language standard.

³<http://bsonspec.org/>

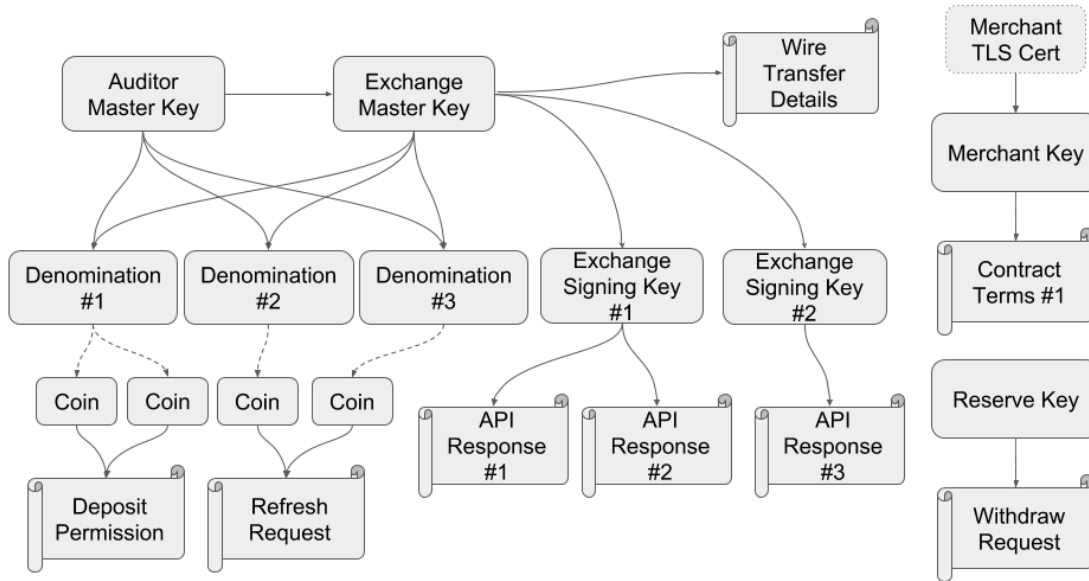


Figure 4.2.: Entities/PKI in Taler. Solid arrows denote signatures, dotted arrows denote blind signatures.

We chose these primitives as they are simple, cheap enough and relatively well studied. Note that other signature schemes that have the syntax and properties described in Section 3.5.1, such as [Bolo3], could be used instead of FDH-RSA.

4.1.3. Entities and Public Key Infrastructure

The public key infrastructure (PKI) used by Taler is orthogonal to the PKI used by TLS [RDo8]. While TLS is used as the transport layer for Taler API messages, we do not rely on TLS for authenticity or integrity of API queries and responses. We do rely on TLS for the confidentiality of digital business contracts and the authenticity, integrity and confidentiality of digital product delivery. For the anonymity properties to hold, the customer must access the merchant and exchange through an anonymity layer (approximated by practical implementations like Tor [DMS04]).

In the case of merchants, we cannot use a trusted auditor or exchange as a trust anchor, since merchants are not required to register within our PKI to accept Taler payments. Here we rely on TLS instead: The merchant is required to include their Taler-specific merchant public key in their TLS certificate. If a merchant fails to do this, the wallet will show a warning when asking the user to confirm a payment.

Auditor

Auditors serve as trust anchors for Taler, and are identified by a single Ed25519 public key. Wallet implementations come with a pre-defined list of trusted

4. Implementation of GNU Taler

```
1  {
2    "version": "2:0:0",
3    "master_public_key": "CQQZ...",
4    "reserve_closing_delay": "/Delay(2419200)/",
5    "signkeys": [
6      {
7        "stamp_start": "/Date(1522223035)/",
8        "stamp_expire": "/Date(1533109435)/",
9        "stamp_end": "/Date(1585295035)/",
10       "master_sig": "842D...",
11       "key": "05XW..."
12     },
13   ],
14   "payback": [],
15   "denoms": [
16     {
17       "master_sig": "BHG5...",
18       "stamp_start": "/Date(1500450235)/",
19       "stamp_expire_withdraw": "/Date(1595058235)/",
20       "stamp_expire_deposit": "/Date(1658130235)/",
21       "stamp_expire_legal": "/Date(1815810235)/",
22       "denom_pub": "51RD...",
23       "value": "TESTKUDOS:10",
24       "fee_withdraw": "TESTKUDOS:0.01",
25       "fee_deposit": "TESTKUDOS:0.01",
26       "fee_refresh": "TESTKUDOS:0.01",
27       "fee_refund": "TESTKUDOS:0.01"
28     },
29     {
30       "master_sig": "QT0T...",
31       "stamp_start": "/Date(1500450235)/",
32       "stamp_expire_withdraw": "/Date(1595058235)/",
33       "stamp_expire_deposit": "/Date(1658130235)/",
34       "stamp_expire_legal": "/Date(1815810235)/",
35       "denom_pub": "51R7",
36       "value": "TESTKUDOS:0.1",
37       "fee_withdraw": "TESTKUDOS:0.01",
38       "fee_deposit": "TESTKUDOS:0.01",
39       "fee_refresh": "TESTKUDOS:0.01",
40       "fee_refund": "TESTKUDOS:0.01"
41     },
42   ],
43   "auditors": [
44     {
45       "denomination_keys": [
46         {
47           "denom_pub_h": "RNTQ...",
48           "auditor_sig": "6SC2..."
49         },
50         {
51           "denom_pub_h": "CP6B...",
52           "auditor_sig": "0GSE..."
53         }
54       ],
55       "auditor_url": "https://auditor.test.taler.net/",
56       "auditor_pub": "BW9DC..."
57     }
58   ],
59   "list_issue_date": "/Date(1530196508)/",
60   "eddsa_pub": "05XW...",
61   "eddsa_sig": "RXCD..."
62 }
```

Figure 4.3.: Example response for /keys

auditors, similar to the certificate store of browsers or operating systems.

Exchange

An exchange is identified by a long term Ed25519 master key and the exchange's base URL. The master key is used as an offline signing key, typically stored on an air-gapped machine. API requests to the exchange are made by appending the name of the endpoint to the base URL.

The exchange uses the master key to sign the following data offline:

- The exchange's online Ed25519 signing keys. The online signing keys are used to sign API responses from the exchange. Each signing key has a validity period.
- The denominations offered by the exchange (explained further in Section 4.1.3).
- The bank accounts supported by the exchange (for withdrawals and deposits) and associated fees.

The `<base-url>/keys` HTTP endpoint of the exchange is used by wallets and merchants to obtain the exchange's signing keys, currently offered denominations and other details. In order to reduce traffic, clients can also request only signing keys and denominations that were created after a specific time. The response to `/keys` is signed by a currently active signing key, so that customers would have proof in case the exchange gave different sets of denomination keys to different customers in an attempt to deanonymize them.




Figure 4.4.: A denomination's lifetime.

Coins and Denominations

Denominations are the RSA public keys used to blindly sign coins of a fixed amount, together with information about their validity and associated fees. The following information is signed by the exchanges master key for every denomination:

- The RSA public key.
- The start date, after which coins of this denomination can be withdrawn and deposited.
- The withdraw expiration date, after which coins cannot be withdrawn anymore, must be after the start date.
- The deposit expiration date, after which coins cannot be deposited anymore, must be after the withdraw expiration date.
- The legal expiration date, after which the exchange can delete all records about operations with coins of this denominations, must be (typically quite a long time!) after the deposit expiration date.
- The fees for a withdraw, deposit, refresh and refund operation with this coin, respectively.

An exchange can be audited by zero, one or multiple auditors. An auditor must monitor all denominations currently offered by the exchange, and an audit of a subset of denominations is not intended in the current design. To allow customers of an exchange to confirm that it is audited properly, the auditor signs an auditing request from the exchange, containing basic information about the exchange as well as all keys offered during the auditing period. In addition to the full auditing request, the auditor also signs an individual certificate for each denomination individually, allowing clients of the exchange to incrementally verify newly offered denominations.

Merchant

The merchant has one Ed25519 key pair that is used to sign responses to the customer and authenticate some requests to the exchange. Depending on the legislation that applies to a particular GNU Taler deployment, merchants might

4. Implementation of GNU Taler

not need to establish an a priori relationship with the exchange, but instead send their bank account information during or after the first deposit of a payment from a customer.

In some jurisdictions, exchanges are required to follow know-your-customer (KYC) regulations and to verify the identity of merchants [Arn+18] using that particular exchange for deposits. Typically, the identity of a merchant only has to be verified if a merchant exceeds a certain threshold of transactions in a given time span. As the KYC registration process can be costly to the exchange, this requirement is somewhat at odds with merchants accepting payments from all exchanges audited by a trusted auditor, since KYC registration needs to be done at every exchange separately. It is, however, unavoidable to run a legally compliant payment system.

A merchant is typically configured with a set of trusted auditors and exchanges, and consequently accepts payments with coins of denominations from a trusted exchange and denominations audited by a trusted auditor.

In order to make the deployment of Taler easier and more secure, the parts that deal with the merchant's private key and cryptographic operations are isolated into a separate service (the merchant backend) with a well-defined RESTful HTTP API. This concept is similar to payment gateways used commonly for credit card payments. The merchant backend can be deployed on-premise by the online shop, or run by a third party provider that is fully trusted by the merchant.

Bank

Since the banks are third parties that are not directly part of Taler, they do not participate directly in Taler's PKI.

Customer

Customers are not registered with an exchange, instead they use the private keys of reserves that they own to authenticate with the exchange. The exchange knows the reserve's public key from the subject/instruction data of the wire transfer. Wire transfers that do not contain a valid public key are automatically reversed.

4.1.4. Payments

Payments in Taler are based on *contract terms*, a JSON object that describes the subject and modalities of a business transaction. The cryptographic hash of such a contract terms object can be used as a globally unique identifier for the business transaction. Merchants must sign the contract terms before sending them to the customer, allowing a customer to prove in case of a dispute the obligations of the merchant resulting from the payment.

Unless a third party needs to get involved in a dispute, it is sufficient (and desirable for data minimization) that only the merchant and the customer know the full content of the contract terms. The exchange, however, must still know the

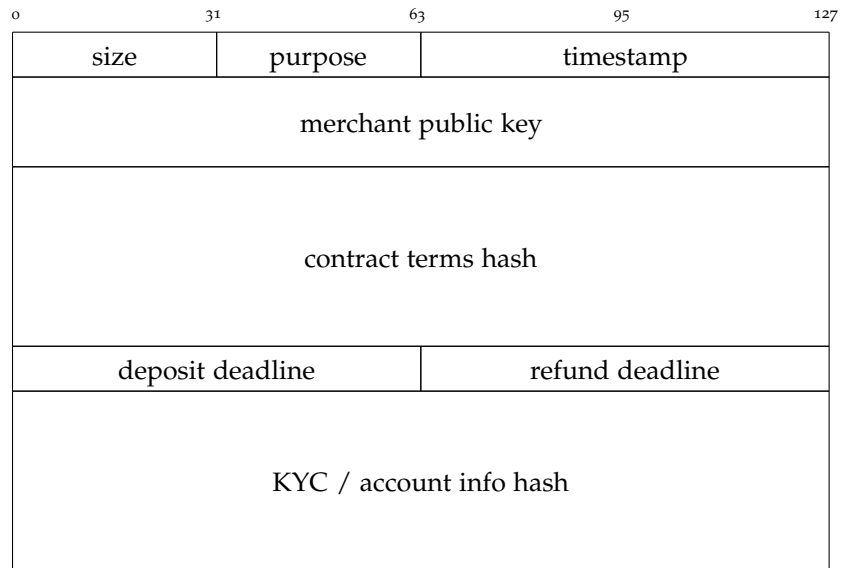


Figure 4.5.: The contract header that is signed by the merchant.

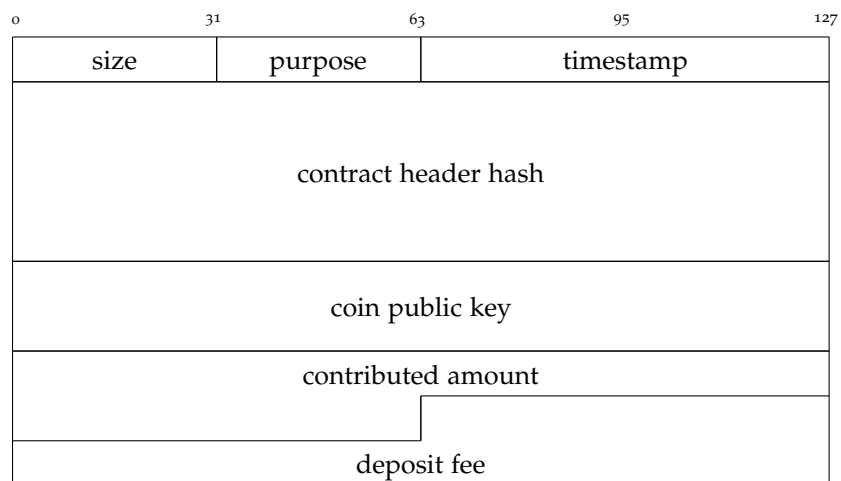


Figure 4.6.: The deposit permission signed by the customer's wallet.

4. Implementation of GNU Taler

parts of the contract terms that specify payment modalities, such as the refund policy, micropayment aggregation deadline and the merchant's KYC registration data (typically a hash to prove the KYC enrollment of the merchant).

Thus, the merchant's signature is made over the *contract header*, which contains the contract terms hash, as well as the payment modalities.

In addition to the data provided by the merchant, the contract terms contain a *claim_pub* field whose value is provided by the customer. This field is an Ed25519 public key, and the customer can use the corresponding private key to prove that they have indeed obtained the individual contract terms from the merchant, and did not copy contract terms that the merchant gave to another customer. Note that this key is not a permanent identity of the customer, but should be freshly generated for each payment.

The signed contract header is created by the merchant's backend from an *order*, which is the "blueprint" for the contract terms. The order is generated by the merchant's frontend and contains a subset of the data contained in the contract terms. Missing data (in particular the merchant's bank account information, public key and accepted auditors/exchanges) and the claim public key obtained from the customer is automatically added by the merchant backend. This allows applications to process payments without having to specify Taler-internal details. In fact, the smallest possible order only needs to contain two fields: the amount to be paid and a human-readable summary of the payment's subject.

An order contains an *order ID*, which is an identifier that is unique within a given merchant and can be a human-friendly identifier such as a booking number. If the order ID is not manually provided, it is automatically filled in by the merchant backend. It can be used to refer to the payment associated with the order without knowing the contract terms hash, which is only available once the customer has provided their claim public key.

To initiate a payment, the merchant sends the customer an *unclaimed* contract terms URL. The customer can download and thereby claim ownership of the contract by appending their claim public key *p* as a query parameter to the unclaimed contract terms URL and making an HTTP GET request to the resulting URL. The customer must then verify that the resulting contract terms are signed correctly by the merchant and that the contract terms contain their claim public key *p*. A malicious customer could try to claim other customers' contracts by guessing contract term URLs and appending their own claim public key. For products that have limited availability, the unclaimed contract URL must have enough entropy so that malicious customers are not able to guess them and claim them before the honest customer.⁴

To give an example, an online shop for concert tickets might allow users to put themselves on a waiting list, and will send them an email once a ticket becomes available. The contract terms URL that allows the customer to purchase the ticket (once they have visited a link in this email), should contain an unguessable nonce,

⁴Note that this URL cannot be protected by a session cookie, as it might be requested from a different session context than the user's browser, namely in the wallet.

as otherwise an attacker might be able to predict the URL and claim the contract for the concert ticket before the customer's wallet can.

In order to settle the payment, the customer must sign a *deposit permission* for each coin that comprises the payment. The deposit permission is a message signed by the coin's private key, containing

- the amount contributed by this coin to the payment,
- the merchant's public key
- the contract header together with the merchant's signature on it,
- the time at which the deposit permission was signed.

After constructing the deposit permissions for a contract, the customer sends them to the merchant by doing an HTTP POST request to the `pay_url` indicated by the merchant in the contract terms. The merchant individually *deposits* each deposit permission with the exchange.

The merchant responds with a payment confirmation to the customer after it has successfully deposited the customer's coins with the exchange. The payment confirmation can be used by the customer to prove that they completed the payment before the payment deadline indicated in the contract terms.

Note that the depositing multiple coins with the exchange deliberately does not have transactional semantics. Instead, each coin is deposited in an individual transaction. This allows the exchange to be horizontally scaled (as discussed in Section 4.9) more easily, as deposit transaction might otherwise have to span multiple database shards.

The lack of transactional semantics, however, means that it must be possible to recover from partially completed payments. There are several cases: If one of the coins that the customer submitted as payment to the merchant is invalid (e.g., because the wallet's state was restored from a backup), the customer can re-try the partially completed payment and provide a different coin instead. If that is not possible or desired by the customer, the merchant may voluntarily give a refund on the coins that have been previously deposited. The reference implementation of the merchant backend offers refunds for partially completed payments automatically.

If refunds were disabled for the payment, the merchant does not cooperate in giving refunds for a partially completed payment, or becomes unresponsive after partially depositing the customer's coin, the customer has two options: They can either complete the deposits on the merchant's behalf, and then use the deposit permissions to prove (either to the merchant or to a court) that they completed the payment.

Another possibility would be to allow customers to request refunds for partially completed payments themselves, directly from the exchange. This requires that the merchant additionally includes the amount to be paid for the contract in the contract header, as the exchange needs to know that amount to decide if a

4. Implementation of GNU Taler

payment with multiple coins is complete. We do not implement this approach, since it implies that the exchange always learns the exact prices of products that the merchant sells, as opposed to just the merchant's total revenue.

The customer could also reveal the contract terms to the exchange to prove that a payment is incomplete, but this is undesirable for privacy reasons, as the exchange should not learn about the full details of the business agreement between customer and merchant.

4.1.5. Resource-based Web Payments

In order to integrate natively with the concepts and architecture of the web, Taler supports paying for a web resource in the form of a URL. In fact all Taler contract terms contain a *fulfillment URL*, which identifies the resource that is being paid for. If the customer is not paying for a digital product (such as an movie, song or article), the fulfillment URL can point to a confirmation page that shows further information, such as a receipt for a donation or shipment tracking information for a physical purchase. A fulfillment URL does not necessarily refer to a single item, but could also represent a collection such as a shopping basket.

The following steps illustrate a typical payment with the online shop `alice-shop.example.com`.

1. The user opens the shop's page and navigates to a paid resource, such as `https://alice-shop.example.com/essay-24.pdf`.
2. The shop sends a response with HTTP status "402 Payment Required" with the headers (↪ marks a continued line)

```
Taler-Contract-Url: _https://alice-shop.example.com/
↪ contract?product=essay-24.pdf
Taler-Resource-Url: _https://alice-shop.example.com/
↪ essay-24.pdf
```

3. Since the user's wallet does not yet contain contract terms with the fulfillment URL `https://alice-shop.example.com/essay-24.pdf` that matches the resources URL, it claims the contract by generating a claim key pair (s, p) and requesting the contract URL with the claim public key p as additional parameter: `https://alice-shop.example.com/contract?product=essay-24.pdf&claim_pub=p`.
4. The wallet displays the contract terms to the customer and asks them to accept or decline. If the customer accepted the contract, the wallet sends a payment to the merchant. After the merchant received a valid payment, it marks the corresponding order as paid.
5. The wallet constructs the extended fulfillment URL by adding the order id from the contract as an additional parameter and navigates the browser

to the resulting URL `https://alice-shop.example.com/esasy-24.pdf?order_id=...`

6. The shop receives the request to the extended fulfillment URL and checks if the payment corresponding to the order ID was completed. In case the payment was successful, it serves the purchased content.

To avoid checking the status of the payment every time, the merchant can instead set a session cookie (signed/encrypted by the merchant) in the user's browser which indicates that `essay-24.pdf` has been purchased.

The resource-based payment mechanism must also handle the situation where a customer navigates again to a resource that they already paid for, without directly navigating to the extended fulfillment URL. In case no session cookie was set for the purchase or the cookie was deleted / has expired, the customer would be prompted for a payment again. To avoid this, the wallet tries to find an existing contract whose plain fulfillment URL matches the resource URL specified in the merchant's HTTP 402 response. If such an existing payment was found, the wallet instead redirects the user to the extended fulfillment URL for this contract, instead of downloading the new contract terms and prompting for payment.

In the example given above, the URL that triggers the payment is the same as the fulfillment URL. This may not always be the case in practice. When the merchant backend is hosted by a third party, say `https://bob.example.com/`, the page that triggers the payment even has a different origin, i.e., the scheme, host or port may differ [Bar11].

This cross-origin operation presents a potential privacy risk if not implemented carefully. To check whether a user has already paid for a particular resource with URL u , an arbitrary website could send an HTTP 402 response with the "Taler-Resource-Url" header set to u and the "Taler-Contract-Url" set to a URL pointing to the attacker's server. If the user paid for u , the wallet will navigate to the extended fulfillment URL corresponding to u . Otherwise, the wallet will try to download a contract from the URL given by the attacker. In order to prevent this attack on privacy, the wallet must only redirect to u if the origin of the page responding with HTTP 402 is the same origin as either the u or the pay URL.⁵

Loose Browser Integration

The payment process we just described does not directly work in browsers that do not have native Taler integration, as the browser (or at least a browser extension) would have to handle the HTTP status code 402 and handle the Taler-specific headers correctly. We now define a fallback, which is transparently implemented in the reference merchant backend.

In addition to indicating that a payment is required for a resource in the HTTP status code and header, the merchant includes a fallback URL in the body of the

⁵This type of countermeasure is well known in browsers as the same origin policy, as also outlined in [Bar11].

4. Implementation of GNU Taler

“402 Payment Required” response. This URL must have the custom URL scheme `taler`, and contains the contract terms URL (and other Taler-specific settings normally specified in headers) as parameters. The above payment would include a link (labeled, e.g., “Pay with GNU Taler”) to the following URL, encoding the same information as the headers:

```
taler:pay?
  ↪ contract_url=
  ↪ https%3A%2F%2Falice-shop.example.com%2Fcontract%3
  ↪ Fproduct%3Dessay-24.pdf
  ↪ &resource_url=
  ↪ https%3A%2F%2Falice-shop.example.com%2Fessay-24.pdf
```

This fallback can be disabled for requests from user agents that are known to natively support GNU Taler.

GNU Taler wallet applications register themselves as a handler for the `taler` URI scheme, and thus following a `taler:pay` link opens the dedicated wallet, even if GNU Taler is not supported by the browser or a browser extension. Registration a custom protocol handler for a URI scheme is possible on all modern platforms with web browsers that we are aware of.

Note that wallets communicating with the merchant do so from a different browsing context, and thus the merchant backend cannot rely on cookies that were set in the customer’s browser when using the shop page.

We chose HTTP headers as the primary means of signaling to the wallet (instead of relying on, e.g., a new content media type), as it allows the fallback content to be an HTML page that can be rendered by all browsers. Furthermore, current browser extension mechanism allow intercepting headers synchronously before the rendering of the page is started, avoiding visible flickering caused by intermediate page loads.

4.1.6. Session-bound Payments and Sharing

As we described the payment protocol so far, an extended fulfillment URL is not bound to a browser session. When sharing an extended fulfillment URL, another user would get access to the same content. This might be appropriate for some types of fulfillment pages (such as a donation receipt), but is generally not appropriate when digital content is sold. Even though it is trivial to share digital content unless digital restrictions management (DRM) is employed, the ability to share links might set the bar for sharing too low.

While the validity of a fulfillment URL could be limited to a certain time, browser session or IP address, this would be too restrictive for scenarios where the user wants to purchase permanent access to the content.

As a compromise, Taler provides *session-bound* payments. For session-bound payments, the seller’s website assigns the user a random session ID, for example, via a session cookie. The extended fulfillment URL for session-bound payments is

constructed by additionally specifying the URL parameter `session_sig`, which contains proof that the user completed (or re-played) the payment under their current session ID.

To initiate a session-bound payment, the HTTP 402 response must additionally contain the “Taler-Session-Id” header, which will cause the wallet to additionally obtain a signature on the session ID from the merchant’s pay URL, by additionally sending the session ID when executing (or re-playing) the payment. As an optimization, instead of re-playing the full payment, the wallet can also send the session ID together with the payment receipt it obtained from the completed payment with different session ID.

Before serving paid content to the user, the merchant simply checks if the session signature matches the assigned session and contract terms. To simplify the implementation of the frontend, this signature check can be implemented as a request to the GNU Taler backend. Using session signatures instead of storing all completed session-bound payments in the merchant’s database saves storage.

While the coins used for the payment or the payment receipt could be shared with other wallets, it is a higher barrier than just sharing a URL. Furthermore, the merchant could restrict the rate at which new sessions can be created for the same contract terms and restrict a session to one IP address, limiting sharing.

For the situation where a user accesses a session-bound paid resource and neither has a corresponding contract in their wallet nor does the merchant provide a contract URL to buy access to the resource, the merchant can specify an *offer URL* in the “Taler-Offer-Url” header. If the wallet is not able to take any other steps to complete the payment, it will redirect the user to the offer URL. As the name suggests, the offer URL can point to a page with alternative offers for the resource, or some other explanation as to why the resource is not available anymore.

4.1.7. Embedded Content

So far we only considered paying for a single, top-level resource, namely the fulfillment URL. In practice, however, websites are composed of many subresources such as embedded images and videos.

We describe two techniques to “paywall” subresources behind a GNU Taler payment. Many other approaches and variations are possible.

1. Visiting the fulfillment URL can set a session cookie. When a subresource is requested, the server will check that the customer has the correct session cookie set.
2. When serving the fulfillment page, the merchant can add an additional authentication token to the URLs of subresources. When the subresource is requested, the validity of the authentication token is checked. If the merchant itself (instead of a Content Delivery Network that supports token

4. *Implementation of GNU Taler*

authentication) is serving the paid subresource, the order ID and session signature can also be used as the authentication token.

It would technically be possible to allow contract terms to refer to multiple resources that are being purchased by including a list or pattern that defines a set of URLs. The browser would then automatically include information to identify the completed payment in the request for the subresource. We deliberately do not implement this approach, as it would require deeper integration in the browser than possible on many platforms. If not restricted carefully, this feature could also be used as an additional method to track the user across the merchant's website.

4.1.8. Contract Terms

The contract terms, only seen by the customer and the merchant (except when a tax audit of the merchant is requested) contain the following information:

- The total amount to be paid,
- the `pay_url`, an HTTP endpoint that receives the payment,
- the deadline until the merchant accepts the payment (repeated in the signed contract header),
- the deadline for refunds (repeated in the signed contract header),
- the claim public key provided by the customer, used to prove they have claimed the contract terms,
- the order ID, which is a short, human-friendly identifier for the contract terms within the merchant,
- the `fulfillment_url`, which identifies the resources that is being paid for,
- a human-readable summary and product list,
- the fees covered by the merchant (if the fees for the payment exceed this value, the customer must explicitly pay the additional fees),
- depending on the underlying payment system, KYC registration information or other payment-related data that needs to be passed on to the exchange (repeated in the signed contract header),
- the list of exchanges and auditors that the merchants accepts for the payment,
- information about the merchant, including the merchant public key and contact information.

4.1.9. Refunds

By signing a *refund permission*, the merchant can “undo” a deposit on a coin, either fully or partially. The customer can then spend (or refresh) the refunded value of the coin again. A refund is only possible before the refund deadline (specified in the contract header). After the refund deadline has passed (and before the deposit deadline) the exchange makes a bank transfer the merchant with the aggregated value from deposits, a refund after this point would require a bank transfer back from the merchant to the exchange.

Each individual refund on each coin incurs fees; the refund fee is subtracted from the amount given back to the customer and kept by the exchange.

Typically a refund serves either one of the following purposes:

- An automatic refund is given to the customer when a payment only partially succeeded. This can happen when a customer’s wallet accidentally double-spends, which is possible even with non-malicious customers and caused by data loss or delayed/failed synchronization between the same user’s wallet on multiple devices. In these cases, the user can choose to re-try the payment with different, unspent coins (if available) or to ask for a refund from the merchant.
- A voluntary refund can be given at the discretion of the merchant, for example, when the customer is not happy with their purchase.

Refunds require a signature by the merchant, but no consent from the customer.

A customer is notified of a refund with the HTTP 402 Payment Required status code and the “Taler-Refund” header. The value of the refund header is a URL. An HTTP GET request on that URL will return a list of refund confirmations that the merchant received from the exchange.

4.1.10. Tipping

Tipping in Taler uses the “withdraw loophole” (see 2.1.11) to allow the merchant⁶ to donate small amounts (without any associated contract terms or legal obligations) into the user’s wallet.

To be able to give tips, the merchant must create a reserve with an exchange. The reserve private key is used to sign blinded coins generated by the user that is being given the tip.

The merchant triggers the wallet by returning an HTTP 402 Payment Required response that includes the “Taler-Tip” header. The value of the tip header (called the tip token) contains

- the amount of the tip,

⁶We still use the term “merchant”, since donations use the same software component as the merchant, but “donor” would be more accurate.

4. Implementation of GNU Taler

- the exchange to use,
- a URL to redirect after processing the tip,
- a deadline for picking up the tip,
- a merchant-internal unique ID for the tip, and
- the *pickup URL* for the tip.

Upon receiving the tip token, the wallet creates coin planchets that sum up to at most the amount specified in the tip token, with denominations offered by the exchange specified in the tip token.

The list of planchets is then sent to the merchant via an HTTP `POST` request to the tip-pickup URL. The merchant creates a withdrawal confirmation signature for each planchet, using the private key of the tipping reserve, and responds to the HTTP `POST` request with the resulting list of signatures. The user then uses these signatures in the normal withdrawal protocol with the exchange to obtain coins “paid for” by the merchant, but anonymized and only spendable by the customer.

4.2. Bank Integration

In order to use Taler for real-world payments, it must be integrated with the existing banking system. Banks can choose to tightly integrate with Taler and offer the ability to withdraw coins on their website. Even existing banks can be used to withdraw coins via a manual bank transfer to the exchange, with the only requirement that the 52 character alphanumeric, case-insensitive encoding of the reserve public key can be included in the transaction without modification other than case folding and white space normalization.⁷

4.2.1. Wire Method Identifiers

We introduce a new URI scheme `payto`, which is used in Taler to identify target accounts across a wide variety of payment systems with uniform syntax.

In its simplest form, a `payto` URI identifies one account of a particular payment system:

```
'payto:/' TYPE '/' ACCOUNT
```

When opening a `payto` URI, the default action is to open an application that can handle payments with the given type of payment system, with the target account pre-filled. In its extended form, a `payto` URL can also specify additional information for a payment in the query parameters of the URI.

⁷Some banking systems specify that the subject of the can be changed, and provide an additional machine-readable “instruction” field.

In the generic syntax for URIs, the payment system type corresponds to the authority, the account corresponds to the path, and additional parameters for the payment correspond to the query parameters. Conforming to the generic URI syntax makes parsing of `payto` URIs trivial with existing parsers.

Formally, a `payto` URI is an encoding of a partially filled out pro forma invoice. The full specification (in the form of an IETF draft currently at the time of writing) of the `payto` URI is in Appendix B.

In the implementation of Taler, `payto` URIs are used in various places:

1. The exchange lists the different ways it can accept money as `payto` URIs. If the exchange uses payment methods that do not have tight Taler integration.
2. In order to withdraw money from an exchange that uses a bank account type that does not typically have tight Taler integration, the wallet can generate a link and a QR code that already contains the reserve public key. When scanning the QR code with a mobile device that has an appropriate banking app installed, a bank transfer form can be pre-filled and the user only has to confirm the transfer to the exchange.
3. The merchant specifies the account it wishes to be paid on as a `payto` URI, both in the configuration of the merchant backend as well as in communication with the exchange.

A major advantage of encoding payment targets as URIs is that URI schemes can be registered with an application on most platforms, and will be “clickable” in most applications and open the right application when scanned as a QR code. This is especially useful for the first use case listed above; the other use cases could be covered by defining a media type instead [FKH13].

As an example, the following QR code would open a banking application that supports SEPA payments, pre-filled with a 15 € donation to the bank account of GNUnet:



4.2.2. Demo Bank

For demonstration purposes and integration testing, we use our toy bank implementation⁸, which might be used in the future for regional currencies or accounting systems (e.g., for a company cafeteria). The payment type identifier is `taler-bank`. The authority part encodes the base URL of the bank, and the path must be the decimal representation of a single integer between 1 and 2^{52} , denoting the internal demo bank account number.

⁸<https://git.taler.net/bank.git>

4.2.3. EBICS and SEPA

The Electronic Banking Internet Communication Standard⁹ (EBICS) is a standard for communicating with banks, and is widely used in Germany, France and Switzerland, which are part of the Single European Payment Area (SEPA). EBICS itself is just a container format. A commonly supported payload for EBICS is ISO 2022, which defines messages for banking-related business processes.

Integration of GNU Taler with EBICS is currently under development, and would allow Taler to be easily deployed in many European countries, provided that the exchange provider can obtain the required banking license.

4.2.4. Blockchain Integration

Blockchains such as Bitcoin could also be used as the underlying financial system for GNU Taler, provided that merchants and customers trust the exchange to be honest.

With blockchains that allow more complex smart contracts, the auditing functionality could be implemented by the blockchain itself. In particular, the exchange can be incentivized to operate correctly by requiring an initial safety deposit to the auditing smart contract, which is distributed to defrauded participants if misbehavior of the exchange is detected.

4.3. Exchange

The exchange consists of three independent processes:

- The `taler-exchange-httpd` process handles HTTP requests from clients, mainly merchants and wallets.
- The `taler-exchange-wirewatch` process watches for wire transfers to the exchange's bank account and updates reserves based on that.
- The `taler-exchange-aggregator` process aggregates outgoing transactions to merchants.

All three processes exchange data via the same database. Only `taler-exchange-httpd` needs access to the exchanges online signing keys and denomination keys.

The database is accessed via a Taler-specific database abstraction layer. Different databases can be supported via plugins; at the time of writing this, only a PostgreSQL plugin has been implemented.

Wire plugins are used as an abstraction to access the account layer that Taler runs on. Specifically, the `wirewatch` process uses the plugin to monitor incoming transfers, and the aggregator process uses the wire plugin to make wire transfers to merchants.

The following APIs are offered by the exchange:

⁹<http://www.ebics.org>

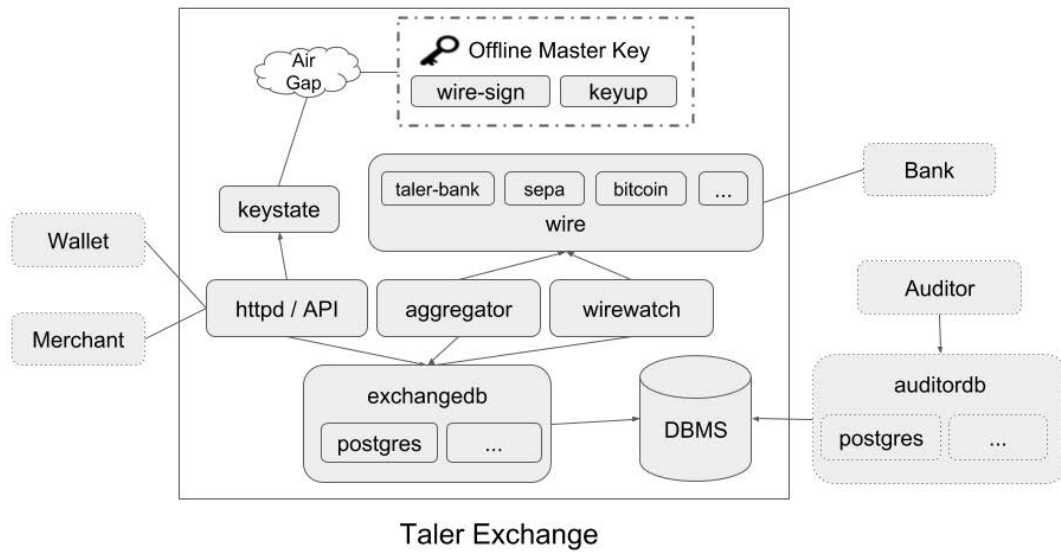


Figure 4.7.: Architecture of the exchange reference implementation

Announcing keys, bank accounts and other public information The exchange offers the list of denomination keys, signing keys, auditors, supported bank accounts, revoked keys and other general information needed to use the exchange's services via the `/keys` and `/wire` APIs.

Reserve status and withdrawal After having wired money to the exchange, the status of the reserve can be checked via the `/reserve/status` API. Since the wire transfer usually takes some time to arrive at the exchange, wallets should periodically poll this API, and initiate a withdrawal with `/reserve/withdraw` once the exchange received the funds.

Deposits and tracking Merchants transmit deposit permissions they have received from customers to the exchange via the `/deposit` API. Since multiple deposits are aggregated into one wire transfer, the merchant additionally can use the exchange's `/track/transfer` API that returns the list of deposits for an identifier included in the wire transfer to the merchant, as well as the `/track/transaction` API to look up which wire transfer included the payment for a given deposit.

Refunds The refund API (`/refund`) can "undo" a deposit if the merchant gave their signature, and the aggregation deadline for the payment has not occurred yet.

Emergency payback The emergency payback API (`/payback`) allows customers to be compensated for coins whose denomination key has been revoked. Customers must send either a full withdrawal transcript that includes their private blinding factor, or a refresh transcript (of a refresh that had the

4. Implementation of GNU Taler

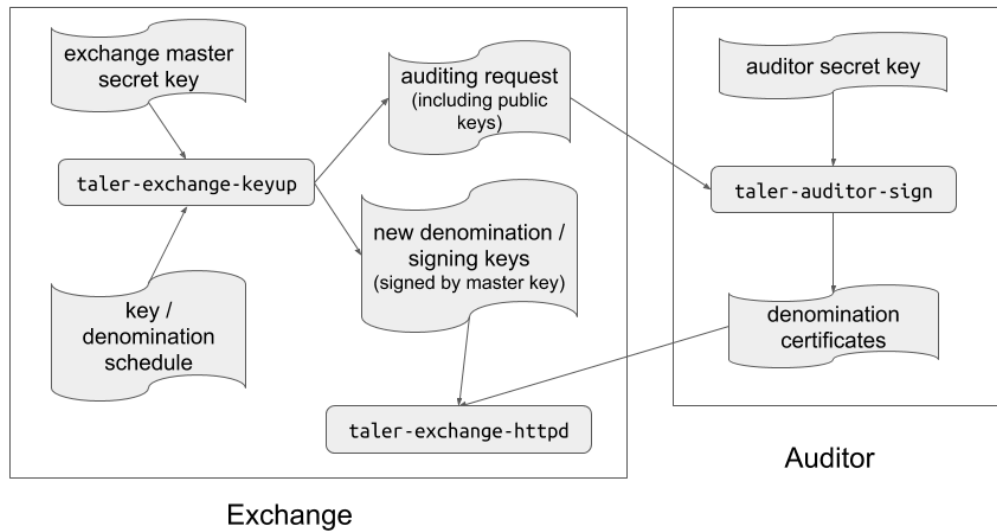


Figure 4.8.: Data flow for updating the exchange's keys.

revoked denominations as one of the targets) that includes blinding factors. In the former case, the reserve is credited, in the latter case, the source coin of the refresh is refunded and can be refreshed again.

New denomination and signing keys are generated and signed with the exchange's master secret key using the `taler-exchange-keyup` utility, according to a key schedule defined in the exchange's configuration. This process should be done on an air-gapped offline machine that has access to the exchange's master signing key.

Generating new keys with `taler-exchange-keyup` also generates an auditing request file, which the exchange should send its auditors. The auditors then certify these keys with the `taler-auditor-sign` tool.

This process is illustrated in Figure 4.8.

4.4. Auditor

The auditor consists of two processes that are regularly run and generate auditing reports. Both processes access the exchange's database, either directly (for exchange-internal auditing as part of its operational security) or over a replica (in the case of external auditors).

The `taler-wire-auditor` process checks that the incoming and outgoing transfers recorded in the exchange's database match wire transfers of the underlying bank account. To access the transaction history (typically recorded by the bank), the wire auditor uses a wire plugin, with the same interface and implementation as the exchange's wire plugins.

The `taler-auditor` process generates a report with the following information:

- Do the operations stored in a reserve's history match the reserve's balance?
- Did the exchange record outgoing transactions to the right merchant for deposits after the deadline for the payment was reached?
- Do operations recorded on coins (deposit, refresh, refund) match the remaining value on the coin?
- Do operations respect the expiration of denominations?
- For a denomination, is the number of pairwise different coin public keys recorded in deposit/refresh operations smaller or equal to the number of blind signatures recorded in withdraw/refresh operations? If this invariant is violated, the corresponding denomination must be revoked.
- What is the income if the exchange from different fees?

The operation of both auditor processes is incremental. There is a separate database to checkpoint the auditing progress and to store intermediate results for the incremental computation. Most database tables used by the exchange are append-only: rows are only added but never removed or changed. Tables that are destructively modified by the exchange only store cached computations based on the append-only tables. Each append-only table has a monotonically increasing row ID. Thus, the auditor's checkpoint simply consists of the set of row IDs that were last seen.

The auditor exposes a web server with the `taler-auditor-httpd` process. Currently, it only shows a website that allows the customer to add the auditor to the list of trusted auditors in their wallet. In future versions, the auditor will also have HTTP endpoints that allow merchants to submit samples of deposit confirmations, which will be checked against the deposit permissions in the exchange's database to detect compromised signing keys or missing writes. Furthermore, in deployments that require the merchant to register with the exchange beforehand, the auditor also offers a list of exchanges it audits, so that the merchant backend can automatically register with all exchanges it transitively trusts.

4.5. Merchant Backend

The Taler merchant backend is a component that abstracts away the details of processing Taler payments and provides a simple HTTP API. The merchant backend handles cryptographic operations (signature verification, signing), secret management and communication with the exchange.

The backend API¹⁰ is divided into two types of HTTP endpoints:

¹⁰See <https://docs.taler.net/api/> for the full documentation

4. Implementation of GNU Taler




Figure 4.9.: Architecture of the merchant reference implementation

1. Functionality that is accessed internally by the merchant. These APIs typically require authentication and/or are only accessible from within the private network of the merchant.
2. Functionality that is exposed publicly on the Internet and accessed by the customer's wallet and browser.

A typical merchant has a *storefront* component that customers visit with their browser, as well as a *back office* component that allows the merchant to view information about payments that customers made and that integrates with other components such as order processing and shipping.

4.5.1. Processing payments

To process a payment, the storefront first instructs the backend to create an *order*. The order contains information relevant to the purchase, and is in fact a subset of the information contained in the contract terms. The backend automatically adds missing information to the order details provided by the storefront. The full contract terms can only be signed once the customer provides the claim public key for the contract.

Each order is uniquely identified by an order ID, which can be chosen by the storefront or automatically generated by the backend.

The order ID can be used to query the status of the payment. If the customer did not pay for an order ID yet, the response from the backend includes a payment redirect URL. The storefront can redirect the customer to this payment redirect

URL; visiting the URL will trigger the customer's browser/wallet to prompt for a payment.

To simplify the implementation of the storefront, the merchant backend can serve a page to the customer's browser that triggers the payment via the HTTP 402 status code and the corresponding headers, and provides a fallback (in the form of a `taler:pay` link) for loosely integrated browsers. When checking the status of a payment that is not settled yet, the response from the merchant backend will contain a payment redirect URL. The storefront redirects the browser to this URL, which is served by the merchant backend and triggers the payment.

The code snippet shown in Figure 4.10 implements the core functionality of a merchant frontend that prompts the customer for a donation (upon visiting `/donate` with the right query parameters) and shows a donation receipt on the fulfillment page with URL `/receipt`. The code snippet is written in Python and uses the Flask library¹¹ to process HTTP requests. The helper functions `backend_post` and `backend_get` make an HTTP POST/GET request to the merchant backend, respectively, with the given request body / query parameters.

4.5.2. Back Office APIs

The back office API allows the merchant to query information about the history and status of payments, as well as correlate wire transfers to the merchant's bank account with the respective GNU Taler payment. This API is necessary to allow integration with other parts of the merchant's e-commerce infrastructure.

4.5.3. Example Merchant Frontends

We implemented the following applications using the merchant backend API.

Blog Merchant The blog merchant's landing page has a list of article titles with a teaser. When following the link to the article, the customer is asked to pay to view the article.

Donations The donations frontend allows the customer to select a project to donate to. The fulfillment page shows a donation receipt.

Codeless Payments The codeless payment frontend is a prototype for a user interface that allows merchants to sell products on their website without having to write code to integrate with the merchant backend. Instead, the merchant uses a web interface to manage products and their available stock. The codeless payment frontend then creates an HTML snippet with a payment button that the merchant can copy-and-paste integrate into their storefront.

Survey The survey frontend showcases the tipping functionality of GNU Taler. The user fills out a survey and receives a tip for completing it.

¹¹<http://flask.pocoo.org/>

4. Implementation of GNU Taler

```
@app.route("/donate")
def donate():
    donation_amount = expect_parameter("donation_amount")
    donation_donor = expect_parameter("donation_donor")
    fulfillment_url = flask.url_for("fulfillment", _external=True)
    order = dict(
        amount=donation_amount,
        extra=dict(donor=donation_donor, amount=donation_amount),
        fulfillment_url=fulfillment_url,
        summary="Donation_to_the_GNU_Taler_project",
    )
    # ask backend to create new order
    order_resp = backend_post("order", dict(order=order))
    order_id = order_resp["order_id"]
    return flask.redirect(flask.url_for("fulfillment", order_id=order_id))

@app.route("/receipt")
def fulfillment():
    order_id = expect_parameter("order_id")
    pay_params = dict(order_id=order_id)

    # ask backend for status of payment
    pay_status = backend_get("check-payment", pay_params)

    if pay_status.get("payment_redirect_url"):
        return flask.redirect(pay_status["payment_redirect_url"])

    if pay_status.get("paid"):
        # The "extra" field in the contract terms can be used
        # by the merchant for free-form data, interpreted
        # by the merchant (avoids additional database access)
        extra = pay_status["contract_terms"]["extra"]
        return flask.render_template(
            "templates/fulfillment.html",
            donation_amount=extra["amount"],
            donation_donor=extra["donor"],
            order_id=order_id,
            currency=CURRENCY)

    # no pay_redirect but article not paid, this should never happen!
    err_abort(500, message="Internal_error,_invariant_failed", json=pay_status)
```

Figure 4.10.: Code snippet with core functionality of a merchant frontend to accept donations.

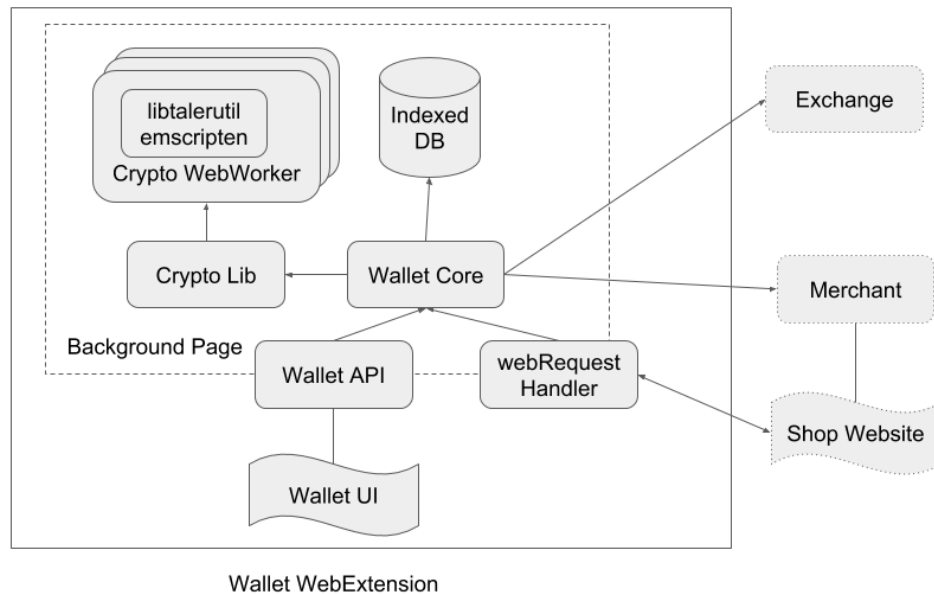


Figure 4.11.: Architecture of the wallet reference implementation

Back office The example back-office application shows the history and status of payments processed by the merchant.

The code for these examples is available at <https://git.taler.net/> in the repositories `blog`, `donations`, `codeless`, `survey` and `backoffice` respectively.

4.6. Wallet

The wallet manages the customer’s reserves and coins, lets the customer view and pay for contracts from merchants. It can be seen in operation in Section 1.3.

The reference implementation of the GNU Taler wallet is written in the TypeScript language against the WebExtension API¹², a cross-browser mechanism for browser extensions. The reference wallet is a “tightly integrated” wallet, as it directly hooks into the browser to process responses with the HTTP status code “402 Payment Required”.

Many cryptographic operations needed to implement the wallet are not commonly available in a browser environment. We cross-compile the GNU Taler utility library written in C as well as its dependencies (such as `libgcrypt`) to `asm.js` (and WebAssembly on supported platforms) using the LLVM-based `emscripten` toolchain [Zak11].

¹²<https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions>

4. Implementation of GNU Taler

Cryptographic operations run in an isolated process implemented as a WebWorker.¹³ This design allows the relatively slow cryptographic operations to run concurrently in the background in multiple threads. Since the crypto WebWorkers are started on-demand, the wallet only uses minimal resources when not actively used.

4.6.1. Optimizations

To improve the perceived performance of cryptographic operations, the wallet optimistically creates signatures in the background while the user is looking at the “confirm payment” dialog. If the user does not accept the contract, these signatures are thrown away instead of being sent to the merchant. This effectively hides the latency of the most expensive cryptographic operations, as they are done while the user consciously needs to make a decision on whether to proceed with a payment.

4.6.2. Coin Selection

The wallet hides the implementation details of fractionally spending different denomination from the user, and automatically selects which denominations to use for withdrawing a given amount, as well as which existing coins to (partially) spend for a payment.

Denominations for withdrawal are greedily selected, starting with the largest denomination that fits into the remaining amount to withdraw. Coin selection for spending proceeds similarly, but first checks if there is a single coin that can be partially spent to cover the whole amount. After each payment, the wallet automatically refreshes coins with a remaining amount large enough to be refreshed. We discuss a simple simulation of the current coin selection algorithm in Section 4.8.2.

A more advanced coin selection would also consider the fee structure of the exchange, minimizing the number of coins as well as the fees incurred by the various operations. The wallet could additionally learn typical amounts that the user spends, and adjust withdrawn denominations accordingly to further minimize costs. An opposing concern to the financial cost is the anonymity of customers, which is improved when the spending behavior of wallets is as similar as possible.

4.6.3. Wallet Detection

When websites such as merchants or banks try to signal the Taler wallet—for example, to request a payment or trigger reserve creation—it is possible that the customer simply has no Taler wallet installed. To accommodate for this situation in a user-friendly way, the HTTP response containing signaling to wallet should

¹³<https://html.spec.whatwg.org/>

contain as response body an HTML page with (1) a `taler:` link to manually open loosely integrated wallets and (2) instructions on how to install a Taler wallet if the user does not already have one.

It might seem useful to dynamically update page content depending on whether the Taler wallet is installed, for example, to hide or show a “Pay with Taler” or “Withdraw to Taler wallet” option. This functionality cannot be provided in general, as only the definitive presence of a wallet can be detected, but not its absence when the wallet is only loosely integrated in the user’s browser via a handler for the `taler:` URI scheme.

We nevertheless consider the ability to know whether a customer has definitely installed a Taler wallet useful (if only for the user to confirm that the installation was successful), and expose two APIs to query this. The first one is JavaScript-based and allows to register a callback for the when presence/absence of the wallet is detected. The second method works without any JavaScript on the merchant’s page, and uses CSS [CSS11] to dynamically show/hide element on the page marked with the special `taler-installed-show` and `taler-installed-hide` CSS classes, whose visibility is changed when a wallet browser extension is loaded.

Browser fingerprinting [Mul+13] is a concern with any additional APIs made available to websites, either by the browser itself or by browser extensions. Since a website can simply try to trigger a payment to determine whether a tightly integrated Taler wallet is installed, one bit of additional fingerprinting information is already available through the usage of Taler. The dynamic detection methods do not, however, expose any information that is not already available to websites by signaling the wallet through HTTP headers.

4.6.4. Backup and Synchronization

While users can manually import and export the state of the wallet, at the time of writing this, automatic backup and synchronization between wallets is not implemented yet. We discuss the challenges with implementing backup and synchronization in a privacy-preserving manner in Chapter 6.

4.6.5. Wallet Liquidation

If a customer wishes to stop using GNU Taler, they can deposit the remaining coins in their wallet back to their own bank account. We call this process *liquidation*.

In deployments with relatively lenient KYC regulation, the normal deposit functionality used by merchants is used for wallet liquidation. The wallet simply acts as a merchant for one transaction, and asks the exchange to deposit the coins into the customer’s bank account.

However in deployments with strict KYC regulations, the customer would first have to register and complete a KYC registration procedure with the exchange. To

4. Implementation of GNU Taler

avoid this, liquidation can be implemented as a modified deposit, which restricts the payment to the bank account that was used to create a reserve of the customer.

The exchange cannot verify that a coin that is being liquidated actually originated the reserve that the customer claims it originated from, unless the user reveals the protocol transcripts for withdrawal and refresh operations on that coin, violating their privacy. Instead, each reserve tracks the amount that was liquidated into it, and the exchange rejects a liquidation request if the liquidated amount exceeds the amount that was put into the reserve. Note that liquidation does not refill the funds of a reserve, but instead triggers a bank transfer of the liquidated amount to the bank account that created the reserve.

4.6.6. Wallet Signaling

We now define more precisely the algorithm that the wallet executes when a website signals to that wallet that an operation related to payments should be triggered, either by opening a `taler:pay` URL or by responding with HTTP 402 and at least one Taler-specific header.

The steps to process a payment trigger are as follows. The algorithm takes the following parameters: `current_url` (the URL of the page that raises the 402 status or null if triggered by a `taler:pay` URL), `contract_url`, `resource_url`, `session_id`, `offer_url`, `refund_url`, `tip_token` (from the “Taler-...” headers or `taler:pay` URL parameters respectively)

1. If `resource_url` is a non-empty string, set `target_url` to `resource_url`, otherwise set `target_url` to `current_url`.
2. If `target_url` is empty, stop.
3. If there is an existing payment p whose fulfillment URL equals `target_url` and either `current_url` is null or `current_url` has the same origin as either the fulfillment URL or payment URL in the contract terms, then:
 - 3.1. If `session_id` is non-empty and the last session ID for payment p was recorded in the wallet with session signature sig , construct a fulfillment instance URL from sig and the order ID of p .
 - 3.2. Otherwise, construct an extended fulfillment URL from the order ID of p .
 - 3.3. Navigate to the extended fulfillment URL constructed in the previous step and stop.
4. If `contract_url` is a non-empty URL, execute the steps for processing a contract URL (with `session_id`) and stop.
5. If `offer_url` is a non-empty URL, navigate to it and stop.
6. If `refund_url` is a non-empty URL, process the refund and stop.

7. If `tip_url` is a non-empty URL, process the tip and stop.

For interactive web applications that request payments, such as games or single page apps (SPAs), the payments initiated by navigating to a page with HTTP status code 402 are not appropriate, since the state of the web application is destroyed by the navigation. Instead the wallet can offer a JavaScript-based API, exposed as a single function with a subset of the parameters of the 402-based payment: `contract_url`, `resource_url`, `session_id` `refund_url`, `offer_url`, `tip_token`. Instead of navigating away, the result of the operation is returned as a JavaScript promise (either a payment receipt, refund confirmation, tip success status or error). If user input is required (e.g., to ask for a confirmation for a payment), the page's status must not be destroyed. Instead, an overlay or separate tab/window displays the prompt to the user.

4.7. Cryptographic Protocols

In this section, we describe the main cryptographic protocols for Taler in more detail. The more abstract, high-level protocols from Section 3.5.1 are instantiated and embedded in concrete protocol diagrams that can hopefully serve as a reference for implementors.

For ease of presentation, we do not provide a bit-level description of the cryptographic protocol. Some details from the implementation are left out, such as fees, additional timestamps in messages and checks for the expiration of denominations. Furthermore, we do not specify the exact responses in the error cases, which in the actual implementation should include signatures that could be used during a legal dispute. Similarly, the actual implementation contains some additional signatures on messages sent that allow to prove to a third party that a participant did not follow the protocol.

As we are dealing with financial transactions, we explicitly describe whenever entities need to safely write data to persistent storage. As long as the data persists, the protocol can be safely resumed at any step. Persisting data is cumulative, that is an additional persist operation does not erase the previously stored information.

The implementation also records additional entries in the exchange's database that are needed for auditing.

4.7.1. Preliminaries

In our protocol definitions, we write **check** COND to abort the protocol with an error if the condition COND is false.

We use the following algorithms:

- `Ed25519.Keygen()` $\mapsto \langle \text{sk}, \text{pk} \rangle$ to generate an Ed25519 key pair.
- `Ed25519.GetPub(sk)` $\mapsto \text{pk}$ to derive the public key from an Ed25519 public key.

4. Implementation of GNU Taler

- $\text{Ed25519.Sign}(\text{sk}, m) \mapsto \sigma$ to create a signature σ on message m using secret key sk .
- $\text{Ed25519.Verify}(\text{pk}, \sigma, m) \mapsto b$ to check if σ is a valid signature from pk on message m .
- $\text{HKDF}(n, k, s) \mapsto m$ is the HMAC-based key derivation function [KE10], producing an output m of n bits from the input key material k and salt s .

We write \mathbb{Z}_N^* for the multiplicative group of integers modulo N . Given an $r \in \mathbb{Z}_N^*$, we write r^{-1} for the multiplicative inverse modulo N of r .

We write $H(m)$ for the SHA-512 hash of a bit string, and $\text{FDH}(N, m)$ for the full domain hash that maps the bit string m to an element of \mathbb{Z}_N^* .

The expression $x \xleftarrow{\$} X$ denotes uniform random selection of an element x from set X . We use $\text{SelectSeeded}(s, X) \mapsto x$ for pseudo-random uniform selection of an element x from set X and seed s . Here, the result is deterministic for fixed inputs s and X .

The exchange's denomination signing key pairs $\{\langle \text{skD}_i, \text{pkD}_i \rangle\}$ are RSA keys pairs, and thus $\text{pkD}_i = \langle e_i, N_i \rangle$, $\text{skD}_i = d_i$. We write $D(\text{pkD}_i)$ for the financial value of the denomination pkD_i .

4.7.2. Withdrawing

The withdrawal protocol is defined in Figure 4.12. The following additional algorithms are used, which we only define informally here:

- $\text{CreateBalance}(W_p, v) \mapsto \perp$ is used by the exchange, and has the side-effect of creating a reserve record with balance v and reserve public key (effectively the identifier of the reserve) W_p .
- $\text{GetWithdrawR}(\rho) \mapsto \{\perp, \bar{\sigma}_C\}$ is used by the exchange, and checks if there is an existing withdraw request ρ . If the existing request exists, the existing blind signature $\bar{\sigma}_C$ over coin C is returned. On a fresh request, \perp is returned.
- $\text{BalanceSufficient}(W_s, \text{pkD}_t) \mapsto b$ is used by the exchange, and returns true if the balance in the reserve identified by W_s is sufficient to withdraw at least one coin if denomination pkD_t .
- $\text{DecreaseBalance}(W_s, \text{pkD}_t) \mapsto \perp$ is used by the exchange, and decreases the amount left in the reserve identified by W_s by the amount $D(\text{pkD}_t)$ that the denomination pkD_t represents.

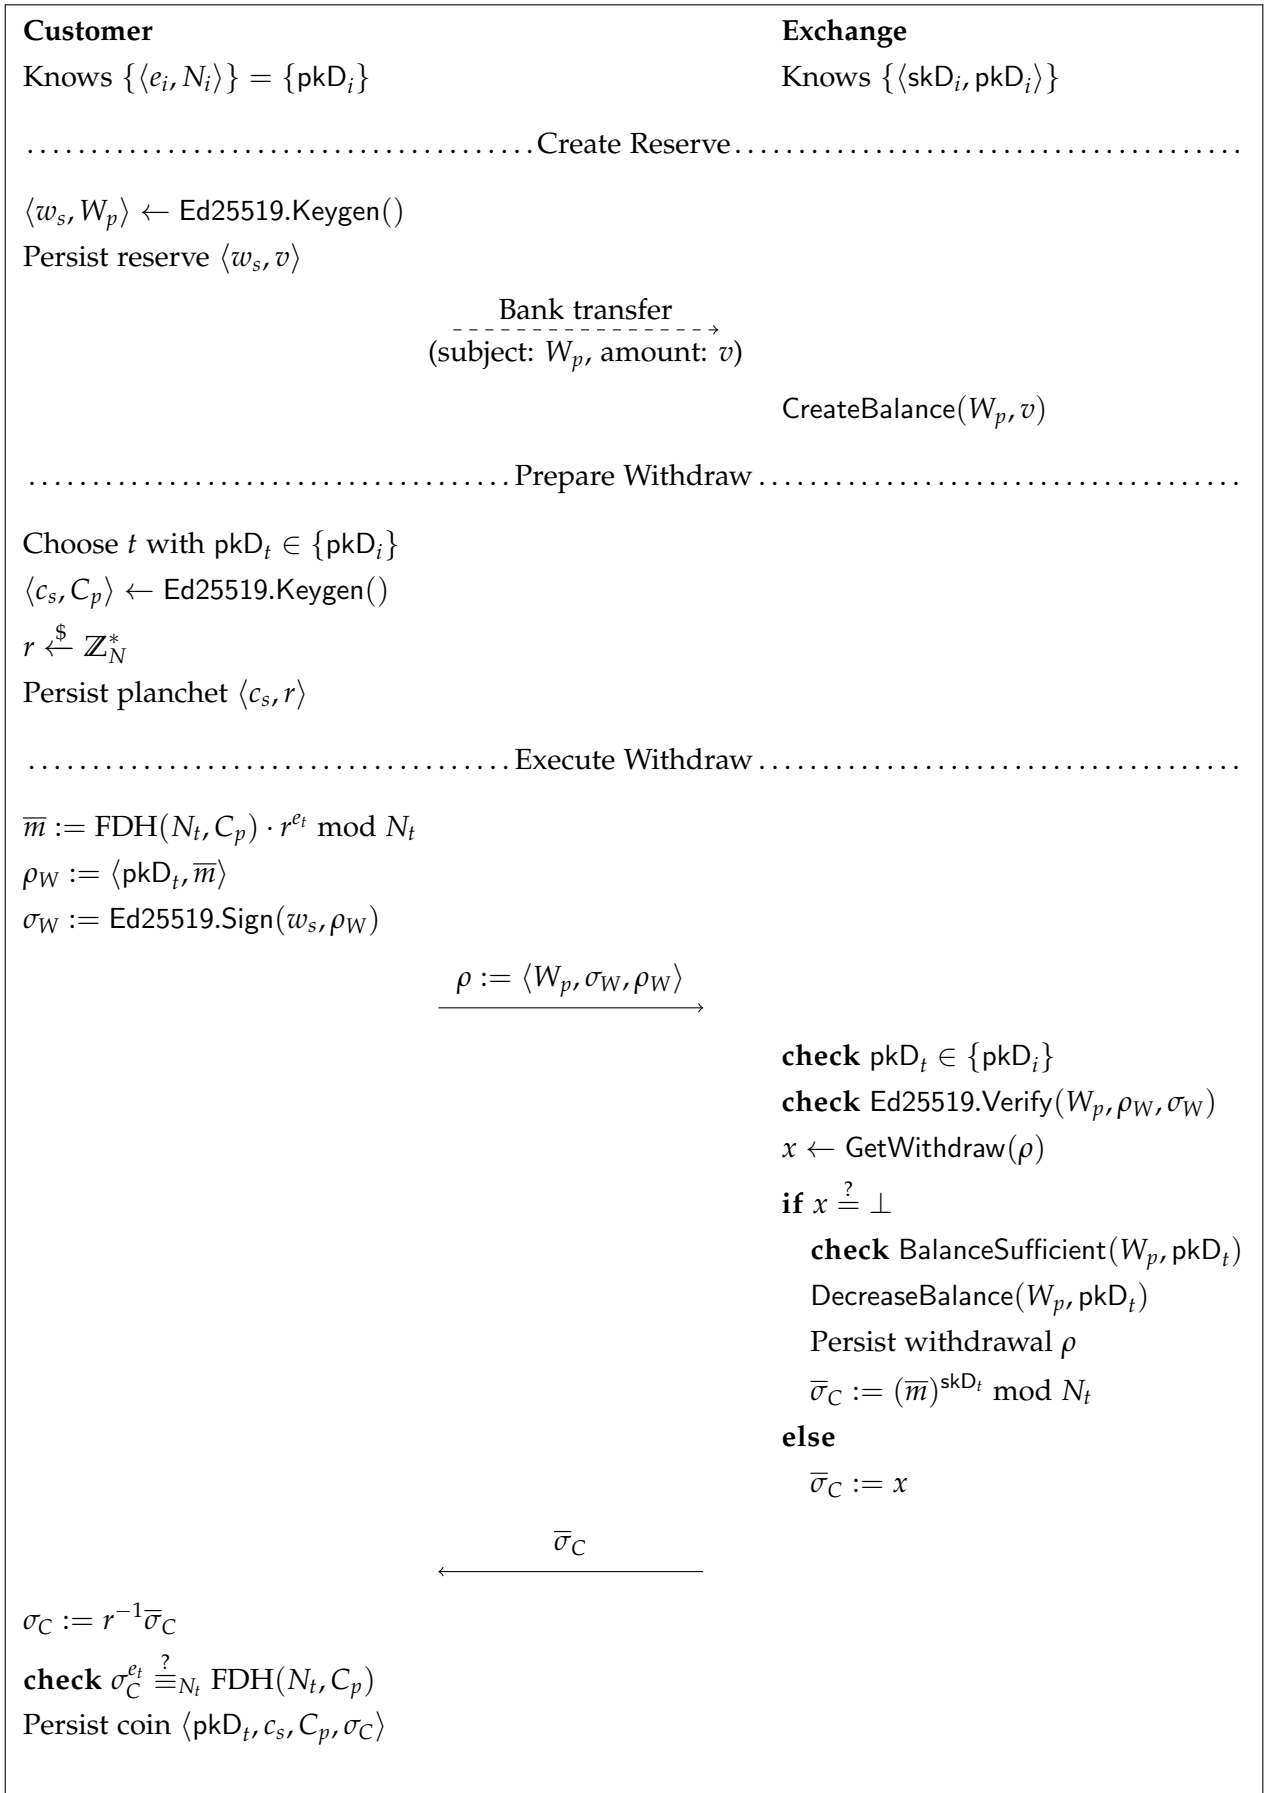


Figure 4.12.: Withdrawal protocol diagram.

4.7.3. Payment transactions

The payment protocol is defined in two parts. First, the spend protocol in Figure 4.13 defines the interaction between a merchant and a customer. The customer obtains the contract terms (as ρ_P) from the merchant, and sends the merchant deposit permissions as a payment. The deposit protocol in Figure 4.14 defines how subsequently the merchant sends the deposit permissions to the exchange to detect double-spending and ultimately to receive a bank transfer from the exchange.

Note that in practice the customer can also execute the deposit protocol on behalf of the merchant. This is useful in situations where the customer has network connectivity but the merchant does not. It also allows the customer to complete a payment before the payment deadline if a merchant unexpectedly becomes unresponsive, allowing the customer to later prove that they paid on time.

We limit the description to one exchange here, but in practice, the merchant communicates to the customer the exchanges that it supports, in addition to the account information A_M that might differ between exchanges.

We use the following algorithms, defined informally here:

- $\text{SelectPayCoins}(v, E_M) \mapsto \{\langle \text{coin}_i, f_i \rangle\}$ selects fresh coins (signed with denomination keys from exchange E_M) to pay for the amount v . The result is a set of coins together with the fraction of each coin that must be spent such that the amounts contributed by all coins sum up to v .
- $\text{MarkDirty}(\text{coin}, f) \mapsto \perp$ subtracts the fraction f from the available amount left on a coin, and marks the coin as dirty (to trigger refreshing in case f is below the denomination value). Thus, assuming the coin has any residual value, the customer's wallet will do a refresh on coin and not use it for further payments. This provides unlinkability of transactions made with change arising from paying with fractions of a coin's denomination.
- $\text{Deposit}(E_M, D_i) \mapsto b$ executes the second part of the payment protocol (i.e., the deposit) with exchange E_M , using deposit permission D_i .
- $\text{GetDeposit}(C_p, h) \mapsto \{\perp, \rho_{(D,i)}\}$ checks in the exchange's database for an existing processed deposit permission on coin C_p for the contract identified by h . The algorithm returns the existing deposit permission $\rho_{(D,i)}$, or \perp if a matching deposit permission is not recorded.
- $\text{IsOverspending}(C_p, \text{pkD}, f) \mapsto b$ checks in the exchange's database if there is at least the fraction f of the coin C_p of denomination pkD is still available for use, based on existing spend/withdraw records of the exchange.
- $\text{MarkFractionalSpend}(C_p, f) \mapsto \perp$ adds a spending record to the exchanges database, indicating that fraction f of coin C_p has been spent (in addition to existing spending/refreshing records).

- $\text{ScheduleBankTransfer}(A_M, f, \text{pkD}, h_c) \mapsto \perp$ schedules a bank transfer from the exchange to the account identified by A_M , for subject h_c and for the amount $f \cdot D(\text{pkD})$.

4.7.4. Refreshing and Linking

The refresh protocol is defined in Figures 4.16 and 4.17. The refresh protocol allows the customer to obtain change for the remaining fraction of the value of a coin. The change is generated as a fresh coin that is unlinkable to the dirty coin to anyone except for the owner of the dirty coin.

A naïve implementation of a refresh protocol that just gives the customer a new coin could be used for peer-to-peer transactions that hides income from tax authorities. Thus, (with probability $(1 - 1/\kappa)$) the refresh protocol records information that allows the owner of the original coin to obtain the refreshed coin from the original coin via the linking protocol (illustrated in Figure 4.18).

We use the following algorithms, defined informally here:

- RefreshDerive is defined in Figure 4.15.
- $\text{GetOldRefresh}(\rho_{RC}) \mapsto \{\perp, \gamma\}$ returns the past choice of γ if ρ_{RC} is a refresh commit message that has been seen before, and \perp otherwise.
- $\text{IsConsistentChallenge}(\rho_{RC}, \gamma) \mapsto \{\perp, \top\}$ returns \top if no refresh-challenge has been persisted for the refresh operation by commitment ρ_{RC} or γ is consistent with the persisted (and thus previously received) challenge; returns \perp otherwise.
- $\text{LookupLink}(C_p) \mapsto \{\langle \rho_L^{(i)}, \sigma_L^{(i)}, \bar{\sigma}_C^{(i)} \rangle\}$ looks up refresh records on coin with public key C_p in the exchange's database and returns the linking message $\rho_L^{(i)}$, linking signature $\sigma_L^{(i)}$ and blinded signature $\bar{\sigma}_C^{(i)}$ for each refresh record i .

4. Implementation of GNU Taler

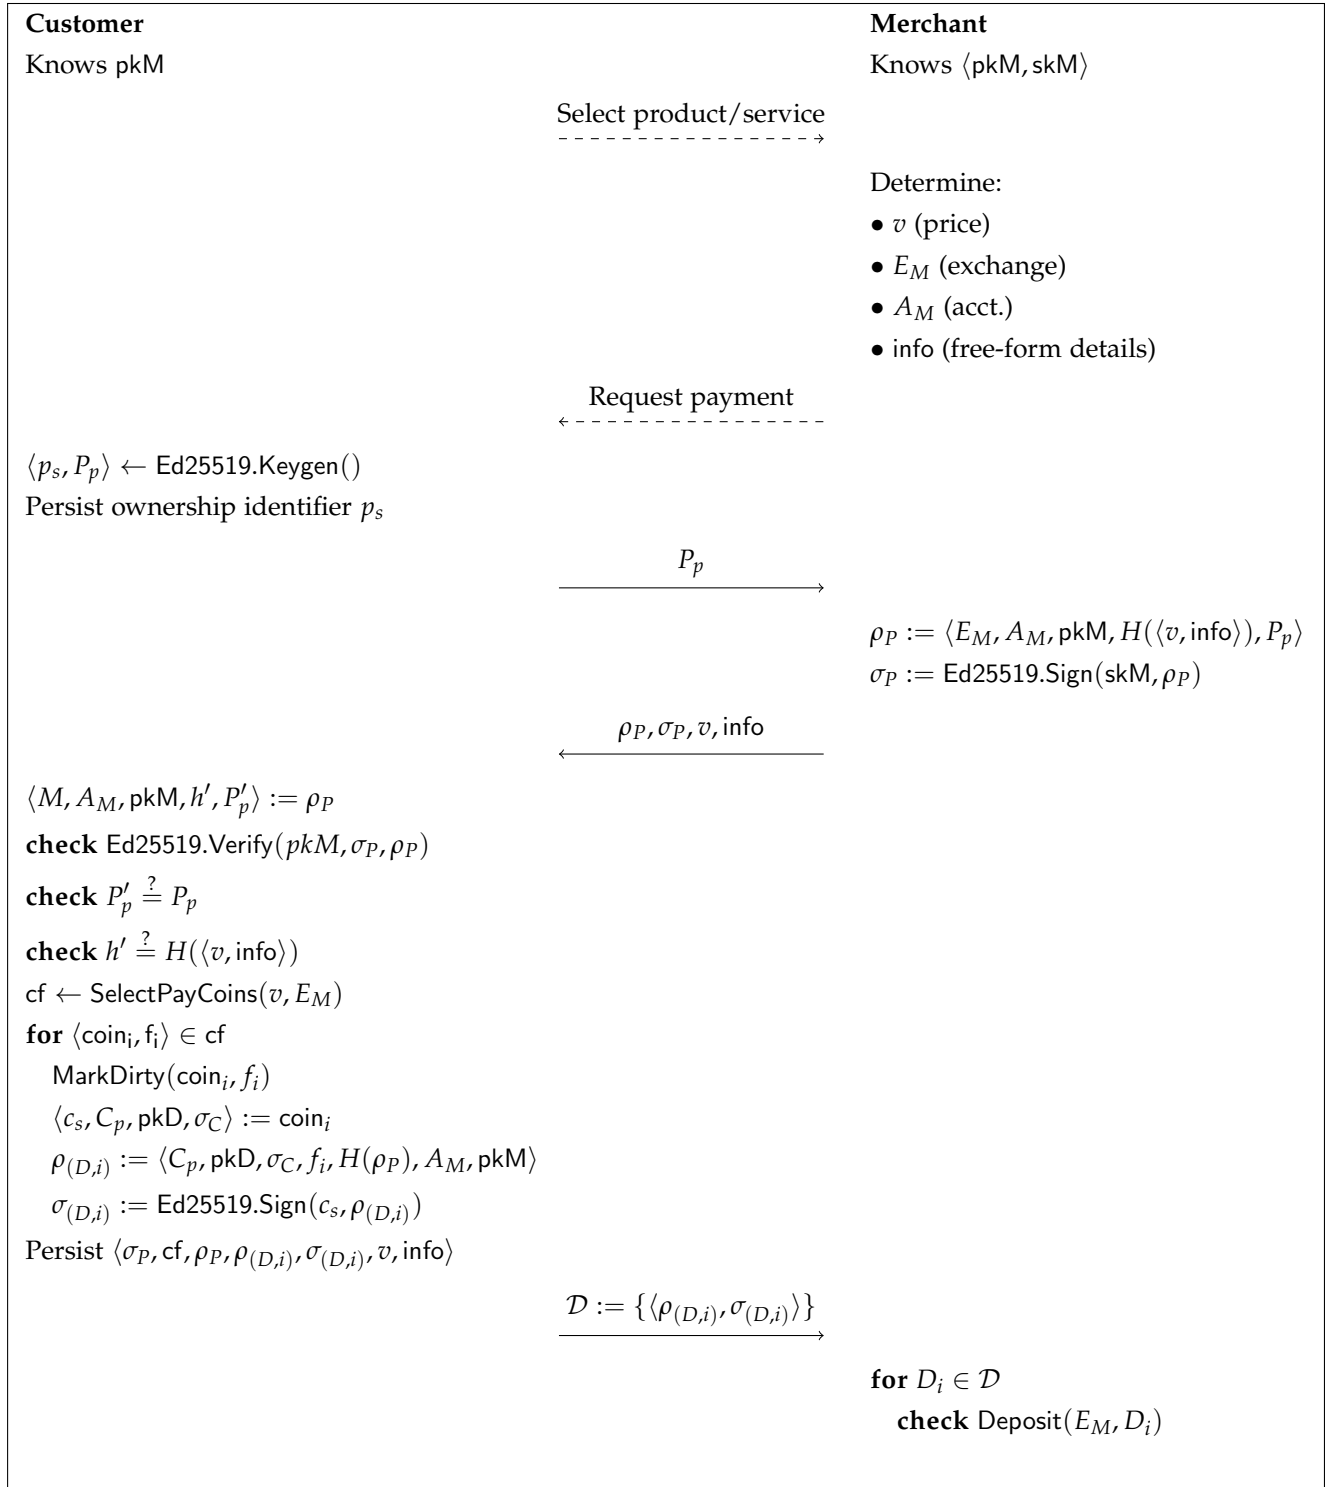


Figure 4.13.: Spend Protocol executed between customer and merchant for the purchase of an article of price v using coins from exchange E_M . The merchant has provided his account details to the exchange under an identifier A_M . The customer can identify themselves as the one who received the offer using p_s .

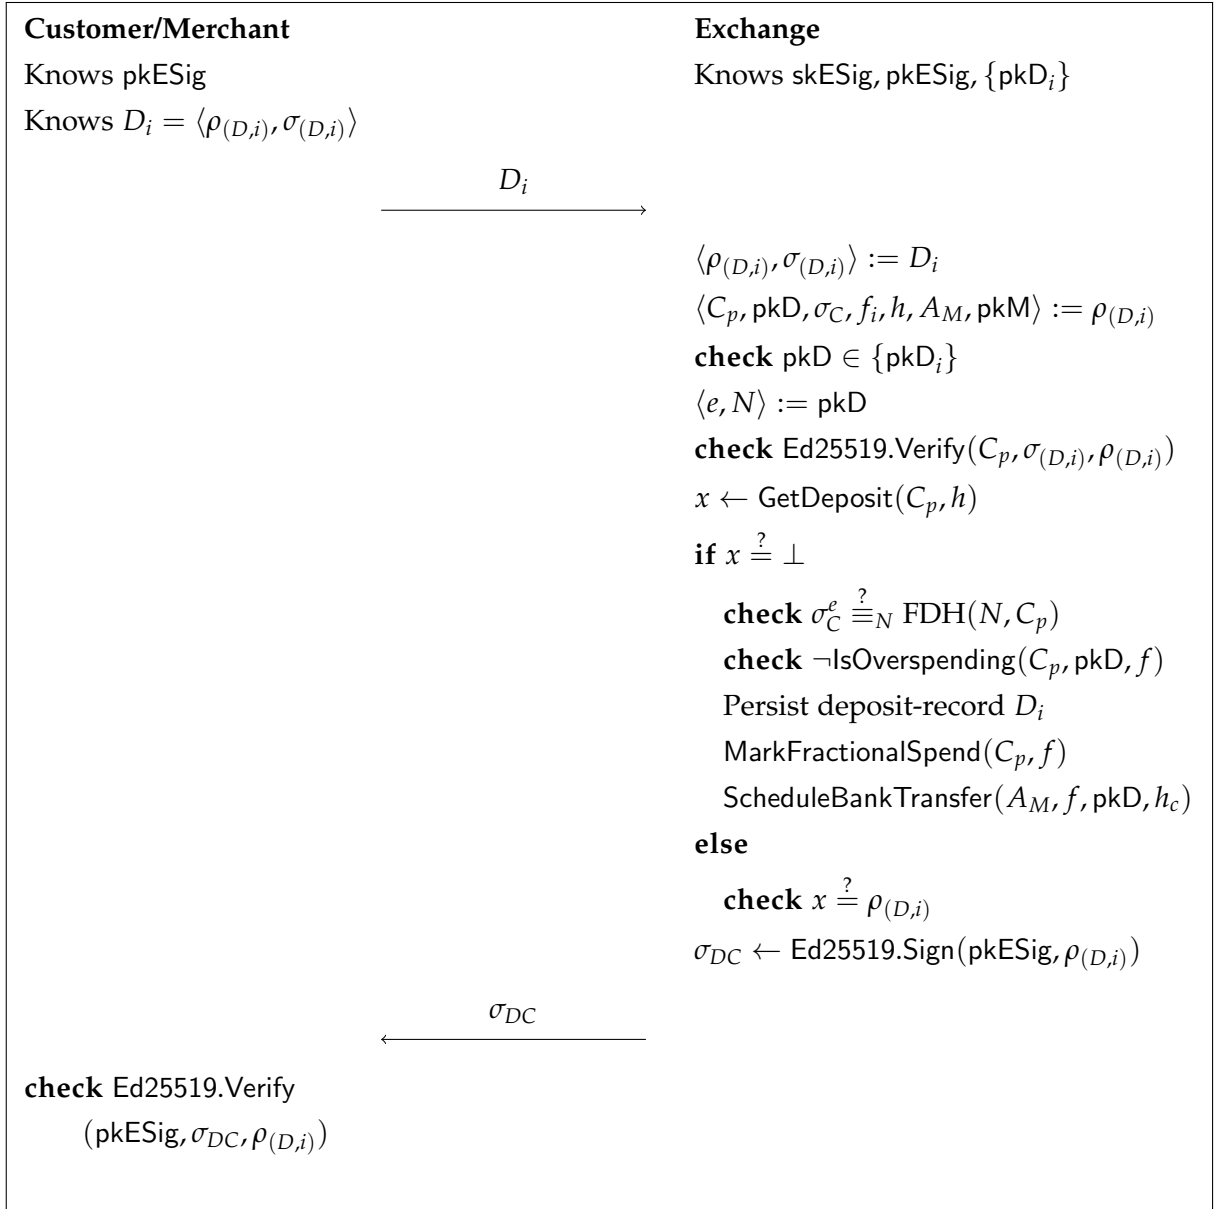


Figure 4.14.: Deposit Protocol run for each deposited coin $D_i \in \mathcal{D}$ with the exchange that signed the coin.

4. Implementation of GNU Taler

RefreshDerive($s, \langle e, N \rangle, C_p$)

```
 $t := \text{HKDF}(256, s, \text{"t"})$   
 $T := \text{Curve25519.GetPub}(t)$   
 $x := \text{ECDH-EC}(t, C_p)$   
 $r := \text{SelectSeeded}(x, \mathbb{Z}_N^*)$   
 $c_s := \text{HKDF}(256, x, \text{"c"})$   
 $C_p := \text{Ed25519.GetPub}(c_s)$   
 $\bar{m} := r^e \cdot C_p \pmod{N}$   
return  $\langle t, T, x, c_s, C_p, \bar{m} \rangle$ 
```

Figure 4.15.: The RefreshDerive algorithm running with the seed s on dirty coin C_p to generate a fresh coin to be later signed with denomination key $pkD := \langle e, N \rangle$.

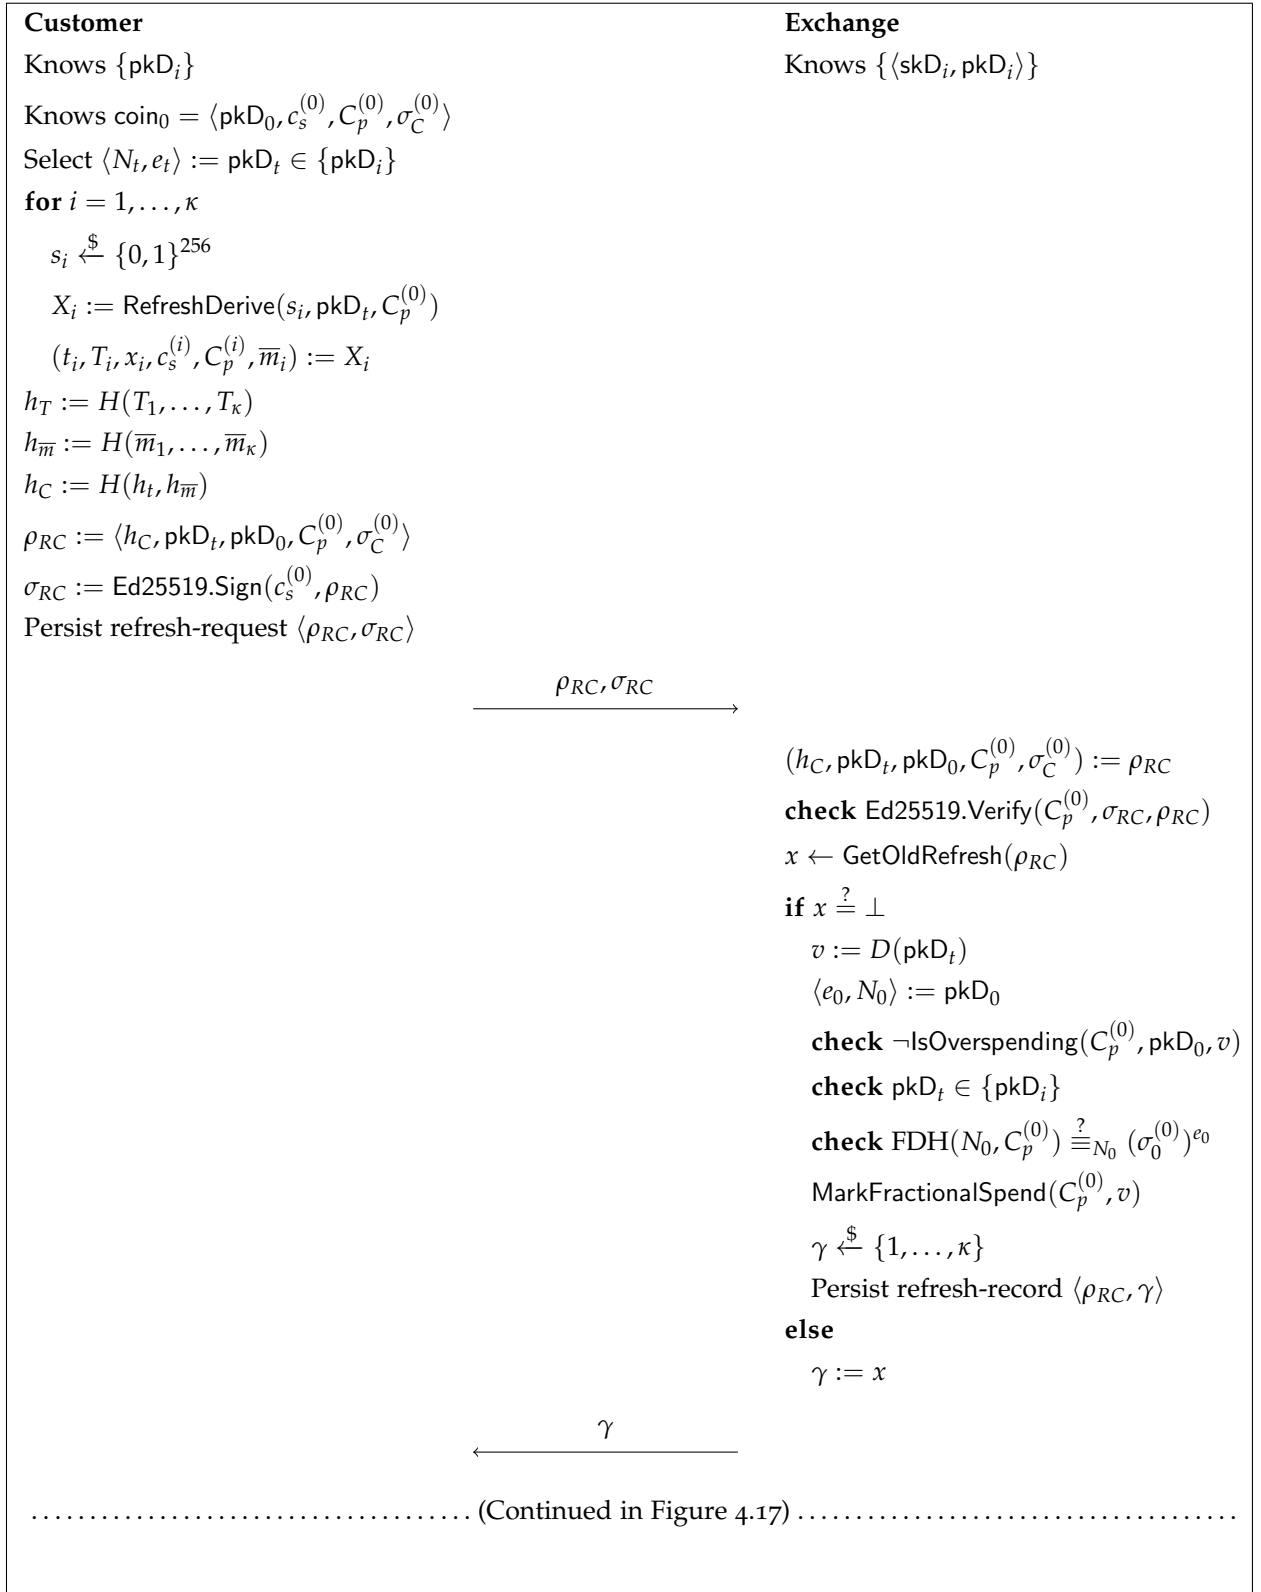


Figure 4.16.: Refresh Protocol (Commit Phase)

4. Implementation of GNU Taler




Figure 4.17.: Refresh Protocol (Reveal Phase)

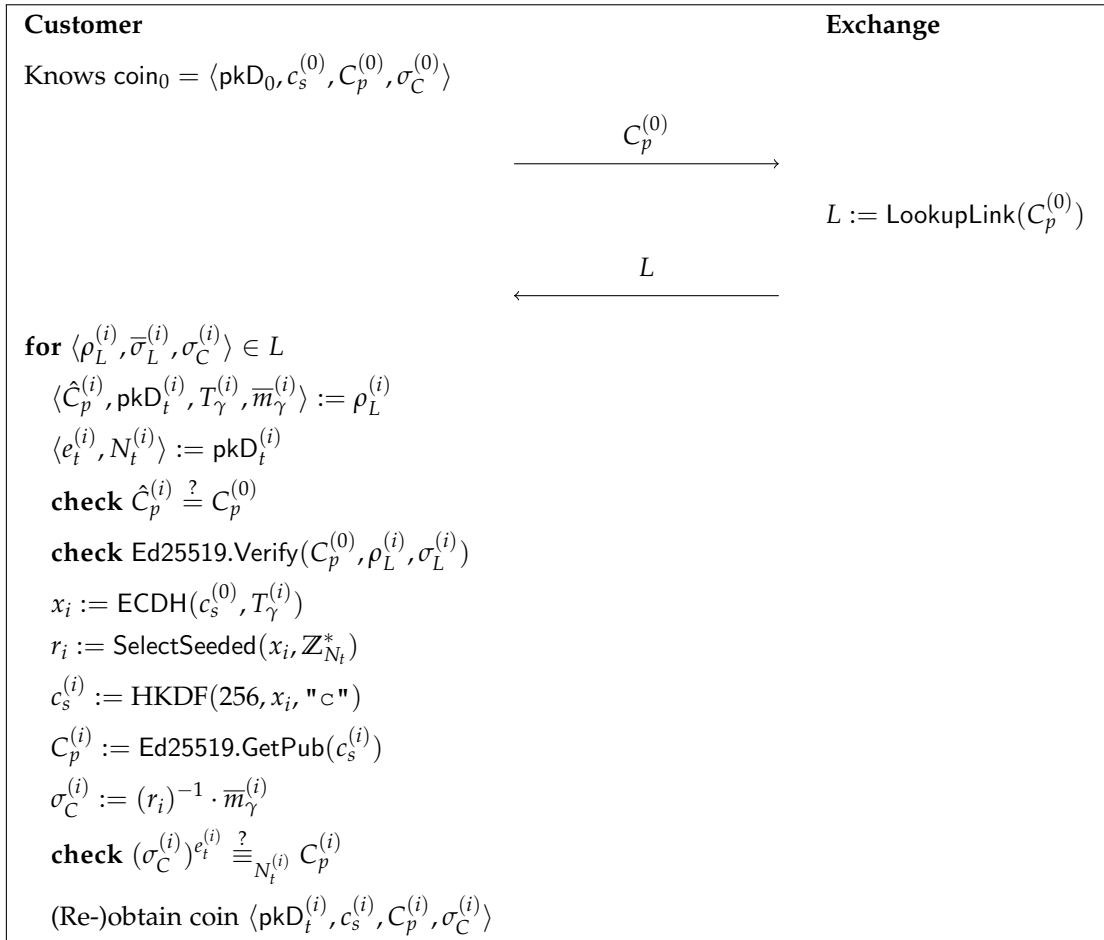


Figure 4.18.: Linking protocol

4.7.5. Refunds

The refund protocol is defined in Figure 4.19. The customer requests from the merchant that a deposit should be “reversed”, and if the merchant allows the refund, it authorizes the exchange to apply the refund and sends the refund confirmation back to the customer. Note that in practice, refunds are only possible before the refund deadline, which is not considered here.

We use the following algorithms, defined informally here:

- $\text{ShouldRefund}(\rho_P, m) \mapsto \{\top, \perp\}$ is used by the merchant to check whether a refund with reason m should be given for the purchase identified by the contract terms ρ_P . The decision is made according to the merchant’s business rules.
- $\text{LookupDeposits}(\rho_P, m) \mapsto \{\langle \rho_{(D,i)}, \sigma_{(D,i)} \rangle\}$ is used by the merchant to retrieve deposit permissions that were previously sent by the customer and already deposited with the exchange.
- $\text{RefundDeposit}(C_p, h, f, \text{pkM})$ is used by the exchange to modify its database. It (partially) reverses the amount f of a deposit of coin C_p to the merchant pkM for the contract identified by h . The procedure is idempotent, and subsequent invocations with a larger f increase the refund.

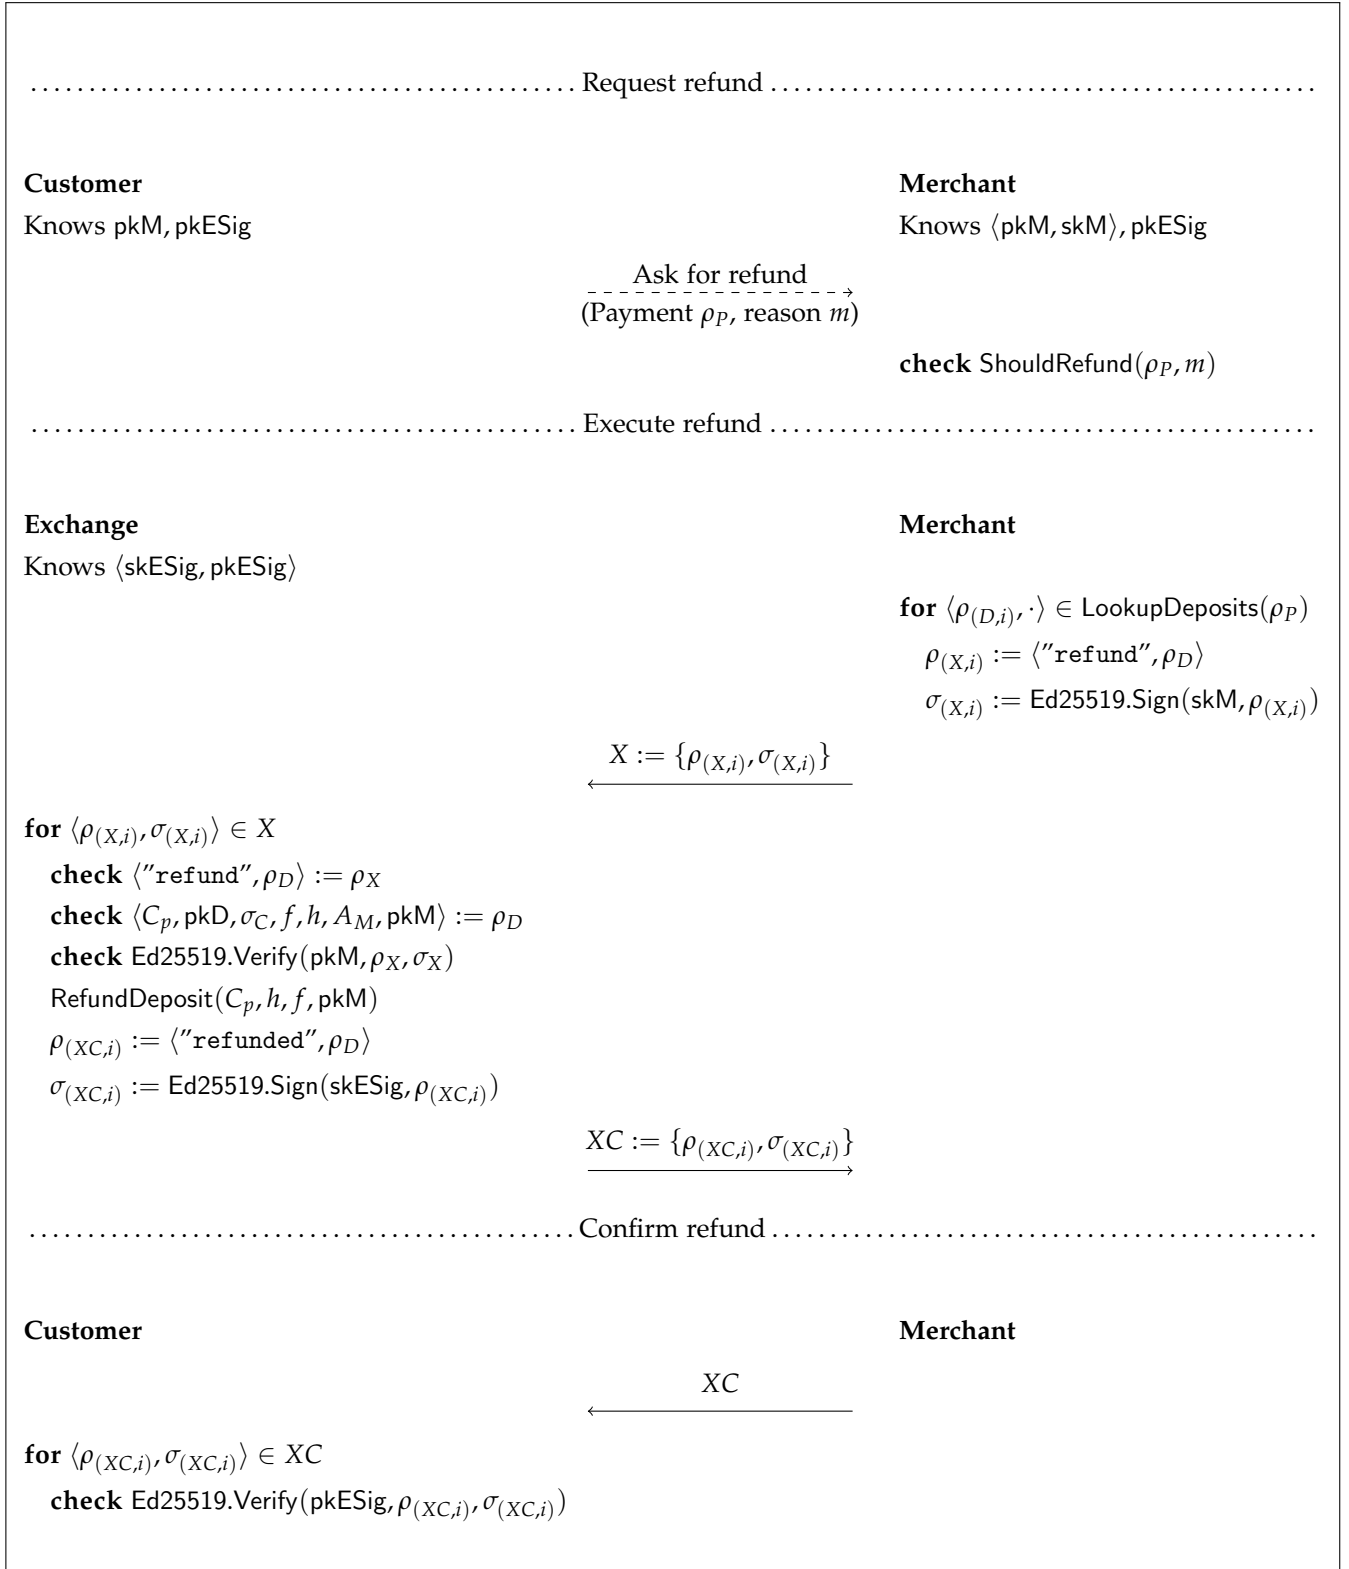


Figure 4.19.: Refund protocol

4. Implementation of GNU Taler

Operation	Time (ms)	Operation	Time (ms)
eddsa create	9.69	eddsa create	34.80
eddsa sign	22.31	eddsa sign	78.55
eddsa verify	19.28	eddsa verify	72.50
hash big	0.05	hash big	0.51
hash small	0.13	hash small	1.37
rsa 2048 blind	3.35	rsa 2048 blind	14.35
rsa 2048 unblind	4.94	rsa 2048 unblind	19.78
rsa 2048 verify	1.97	rsa 2048 verify	9.10
rsa 4096 blind	10.38	rsa 4096 blind	47.86
rsa 4096 unblind	16.13	rsa 4096 unblind	69.91
rsa 4096 verify	6.57	rsa 4096 verify	29.02

(a) Wallet microbenchmark on a Laptop (Intel i7-4600U) with Firefox

(b) Wallet microbenchmark on Android Moto G3 with Firefox

Table 4.1.: Wallet microbenchmarks

4.8. Experimental results

We now evaluate the performance of the core components of the reference implementation of GNU Taler. No separate benchmarks are provided for the merchant backend, as the work done by the merchant per transaction is relatively negligible compared to the work done by the exchange, and one exchange needs to provide service many merchants and all of their customers. Thus, the exchange is the bottleneck for the performance of the system.

We provide a microbenchmark for the performance of cryptographic operations in the wallet (Table 4.1. Even on a low-end smartphone device, the most expensive cryptographic operations remain well under $150ms$, a threshold for user-interface latency under which user happiness and productivity is considered to be unaffected [TASo6].

We implemented a benchmarking tool that starts a single (multi-threaded) exchange and a bank process for the taler-test wire transfer protocol. It then generates workload on the exchange with a configurable number of concurrent clients and operations. The benchmarking tool is able to run the exchange on a different machine (via SSH¹⁴) than the benchmark driver, mock bank and clients. At the end, the benchmark outputs the number of deposited coins per second and latency statistics.

¹⁴<https://www.openssh.com/>

4.8.1. Hardware Setup

We used two server machines (`firefly` and `gv`) with the following hardware specifications for our tests:

- `firefly` has a 96-core AMD EPYC 7451 CPU and 256GiB DDR4 2667 MHz RAM.
- `gv` has a 16-core Intel(R) Xeon X5550 (2.67GHz) CPU and 128GiB DDR3 1333 MHz RAM.

We used 2048-bit RSA denomination keys for all of our exchange benchmarks. We used a development version of the exchange (with git commit hash `5fbda29b76c24d...`). PostgreSQL version 11.3 was used as the database. As our server machines have only slower hard-disk drives instead of faster solid-state drives, we ran the benchmarks with an in-memory database.

4.8.2. Coins Per Transaction

The transaction rate is an important characteristic of a payment system. Since GNU Taler internally operates on the level of coins instead of transactions, we need to define what actually constitutes one transaction in our measurements. This includes both how many coins are used per transaction on average, as well as how often refresh operations are run.

We ran a simple simulation to determine rather conservative upper bounds for the parameters that characterize the average transaction. The source code for the simulation can be found in Appendix C.

In the simulation, thirteen denominations of values $2^0, \dots, 2^{12}$ are available. Customers repeatedly select a random value to be spent between 4 and 5000. When customers do not have enough coins for a transaction, they withdraw a uniform random amount between the minimum amount to complete the transaction and 10000. The denominations selected for withdrawal are chosen by greedy selection of the largest possible denomination. When spending, the customer first tries to use one coin, namely the smallest coin larger than the requested amount. If no such coin exists in the customer's wallet, the customer pays with multiple coins, spending smallest coins first.

Choosing a random uniform amount for withdrawal could be considered unrealistic, as customers in practice likely would select from a fixed list of common withdrawal amounts, just like most ATMs operate. Thus, we also implemented a variation of the simulation that withdraws a constant amount of 1250 (i.e., 1/4 of the maximum transaction value) if it is sufficient for the transaction, and the exact required amount otherwise.

We obtained the following results for the number of average operations executed for one “business transaction”:

4. Implementation of GNU Taler

	random withdraw	constant withdraw
#spend operations	8.3	7.0
#refresh operations	1.3	0.51
#refresh output coins	4.2	3.62

Based on these results, we chose the parameters for our benchmark: for every spend operation we run a refresh operation with probability $1/10$, where each refresh operation produces 4 output coins. In order to arrive at the transaction rate, the rate of spend operations should be divided by 10.

Note that this is a rather conservative analysis. In practice, the coin selection for withdrawal/spending can use more sophisticated optimization algorithms, rather than using greedy selection. Furthermore, we expect that the amounts paid in real-world transactions will have more predictable distributions, and thus the offered denominations can be adjusted to typical amounts.

Baseline Sequential Resource Usage

To obtain a baseline for the resource usage of the exchange, we ran the benchmark on `firefly` with a single client that executes sequential requests to withdraw and spend 10000 coins, with 10% refresh probability.

Table 4.2 shows the time used for cryptographic operations, together with the number of times they are executed by the clients (plus the mock bank and benchmark setup) and exchange, respectively. Note that while we measured the wall-clock time for these operations, the averages should correspond to the actual CPU time required for the respective operations, as the benchmark with one client runs significantly fewer processes/threads than the number of available CPUs on our machine.

The benchmark completed in 15.10 minutes on `firefly`. We obtained the total CPU usage of the benchmark testbed and exchange. The refresh operations are rather slow in comparison to spends and deposits, as the benchmark with a refresh probability of 0% only took 8.84 minutes to complete.

The size of the exchange’s database after the experiment (starting from an empty database) is shown in Table 4.3. We measured the size of tables and indexes using the `pg_relation_size / pg_indexes_size` functions of PostgreSQL.

We observe that even though the refresh operations account for about half of the time taken by the benchmark, they contribute to only $\approx 16\%$ of the database’s size. The computational costs for refresh are higher than the storage costs (compared to other operations), as the database stores only needs to store one commitment, one transfer key and the blinded coins that are actually signed.

In our sequential baseline benchmark run, only one reserve was used to withdraw coins, and thus the tables that store the reserves are very small. In practice, information for multiple reserves would be tracked for each active customers.

Operation	Time/Op (μ s)	Count (exchange)	Count (clients)
ecdh eddsa	1338.62	2430	3645
ecdhe key create	1825.38	0	3645
ecdhe key get public	1272.64	2430	4860
eddsa ecdh	1301.78	0	4860
eddsa key create	1896.27	0	12 180
eddsa key get public	1729.69	9720	80 340
eddsa sign	5182.33	13 393	25 608
eddsa verify	3976.96	25 586	25 627
hash	1.41	165 608	169 445
hash context finish	0.28	1215	1227
hash context read	0.81	25 515	25 655
hash context start	11.38	1215	1227
hkdf	40.61	65 057	193 506
rsa blind	695.25	9720	31 633
rsa private key get public	5.30	0	40
rsa sign blinded	5284.88	17 053	0
rsa unblind	1348.62	0	21 898
rsa verify	421.19	13 393	29 216

Table 4.2.: Cryptographic operations in the benchmark with one client and 10000 operations.

Relation	Table (MiB)	Indexes (MiB)	Total (MiB)
denominations	0.02	0.03	0.05
reserves_in	0.01	0.08	0.09
reserves	0.02	0.25	0.27
refresh_commitments	0.36	0.28	0.64
refresh_transfer_keys	0.38	0.34	0.73
refresh_revealed_coins	4.19	0.91	5.14
known_coins	7.37	0.70	8.07
deposits	4.85	6.80	11.66
reserves_out	8.95	4.48	13.43
<i>Sum</i>	26.14	13.88	40.02

Table 4.3.: Space usage by database table for 10000 deposits with 10% refresh probability.

4. Implementation of GNU Taler

The TCP/IP network traffic between the exchange, clients and the mock bank was 57.95 MiB, measured by the Linux kernel's statistics for transmitted/received bytes on the relevant network interface. As expected, the traffic is larger than the size of the database, since some data (such as signatures) is only verified/generated and not stored in the database.

4.8.3. Transaction Rate and Scalability

Figure 4.20 shows the mean time taken to process one coin for different numbers of parallel clients. With increasing parallelism, the throughput continues to rise roughly until after the number of parallel clients saturates the number of available CPU cores (96). There is no significant decrease in throughput even when the system is under rather high load, as clients whose requests cannot be handled in parallel are either waiting in the exchange's listen backlog or waiting in a retry timeout (with randomized, truncated, exponential back-off) after being refused when the exchange's listen backlog is full.

Figure 4.21 shows the CPU time (sum of user and system time) of both the exchange and the whole benchmark testbed (including the exchange) in relation to the wall-clock time the benchmark took to complete. We can see that the gap between the wall-clock time and CPU time used by the benchmark grows with an increase in the number of parallel clients. This can be explained by the CPU usage of the database (whose CPU usage is not measured as part of the benchmark). With a growing number of parallel transactions, the database runs into an increasing number of failed commits due to read/write conflicts, leading to retries of the corresponding transactions.

To estimate the time taken up by cryptographic operations in the exchange, we first measured a baseline with a single client, where the wall-clock time for cryptographic operations is very close to the actual CPU time, as virtually no context switching occurs. We then extrapolated these timings to experiment runs with parallelism by counting the number of times each operation is executed and multiplying with the baseline. As seen in the dot-and-dash line in Figure 4.21, by our extrapolation slightly more than half of the time is spent in cryptographic routines.

We furthermore observe in Figure 4.21 that under full load, less than 1/3 of the CPU time is spent by the exchange. A majority of the CPU time in the benchmark is used by the simulation of clients. As we did not have a machine available that is powerful enough to generate traffic that can saturate a single exchange running on *firefly*, we estimate the throughput that would be possible if the machine only ran the exchange. The highest rate of spends was 780 per second. Thus, the theoretically achievable transaction rate on our single test machine (and a dedicated machine for the database) would be $780 \cdot 3/10 = 234$ transactions per second under the relatively pessimistic assumptions we made about what constitutes a transaction.

If a GNU Taler deployment was used to pay for items of fixed price (e.g.,




Figure 4.20.: Coin throughput in relation to number of parallel clients, with 1000 coins per client per experiment run.

online news articles), the overhead of multiple coins and refresh operations (which accounts for $\approx 50\%$ of spent time as measured earlier) and multiple coins per payment would vanish, giving an estimated maximum transaction rate of $742 \cdot 2 = 1484$ transactions per second.

4.8.4. Latency

We connected `firefly` and `gv` directly with a patch cable, and introduced artificial network latency by configuring the Linux packet scheduler with the `tc` tool. The goal of this experiment was to observe the network latency characteristics of the implementation. Note that we do not consider the overhead of TLS in our experiments, as we assume that TLS traffic is already terminated before it reaches the exchange service, and exchanges can be operated securely even without TLS.

The comparison between no additional delay and a 100 ms delay is shown in Table 4.4. TCP Fast Open [Che+14] was enabled on both `gv` and `firefly`. Since for all operations except `/refresh/reveal`, both request and response fit into one TCP segment, these operations complete within one round-trip time. This explains the additional delay of ≈ 200 ms when the artificial delay is introduced. Without TCP Fast Open, we would observe an extra round trip for the SYN and SYN/ACK packages without any payload. The `/refresh/reveal` operation takes an extra roundtrip due to the relatively large size of the request (as show in Table 4.5), which exceeds the MTU of 1500 for the link between `gv` and `firefly`, and thus does not fit into the first TCP Fast Open packet.

4. Implementation of GNU Taler




Figure 4.21.: Comparison of real time, the CPU time for the exchange and the whole benchmark.

Figure 4.22 shows the latency for the exchange’s HTTP endpoints in relation to different network delays. As expected, the additional delay grows linearly for a single client. We note that in larger benchmarks with multiple parallel clients, the effect of additional delay would likely not just be linear, due to timeouts raised by clients.

4.9. Current Limitations and Future Improvements

Currently the auditor does not support taking samples of deposit confirmations that merchants receive. The API and user interface to receive and process

Endpoint	Base latency (ms)	Latency with 100 ms delay (ms)
/keys	1.14	201.25
/reserve/withdraw	22.68	222.46
/deposit	22.36	223.22
/refresh/melt	20.71	223.9
/refresh/reveal	63.64	466.30

Table 4.4.: Effects of 100 ms symmetric network delay on total latency.

Endpoint	Request size 2048-bit RSA (kB)	Response size 2048-bit RSA (kB)	Request size 1024-bit RSA (kB)	Response size 1024-bit RSA (kB)
/keys	0.14	3.75	0.14	3.43
/reserve/withdraw	0.73	0.71	0.60	0.49
/deposit	1.40	0.34	1.14	0.24
/refresh/melt	1.06	0.35	0.85	0.35
/refresh/reveal	1.67	2.11	1.16	1.23

Table 4.5.: Request and response sizes for the exchange’s API. In addition to the sizes for 2048-bit RSA keys (used throughout the benchmark), the sizes for 1024-bit RSA keys are also provided.

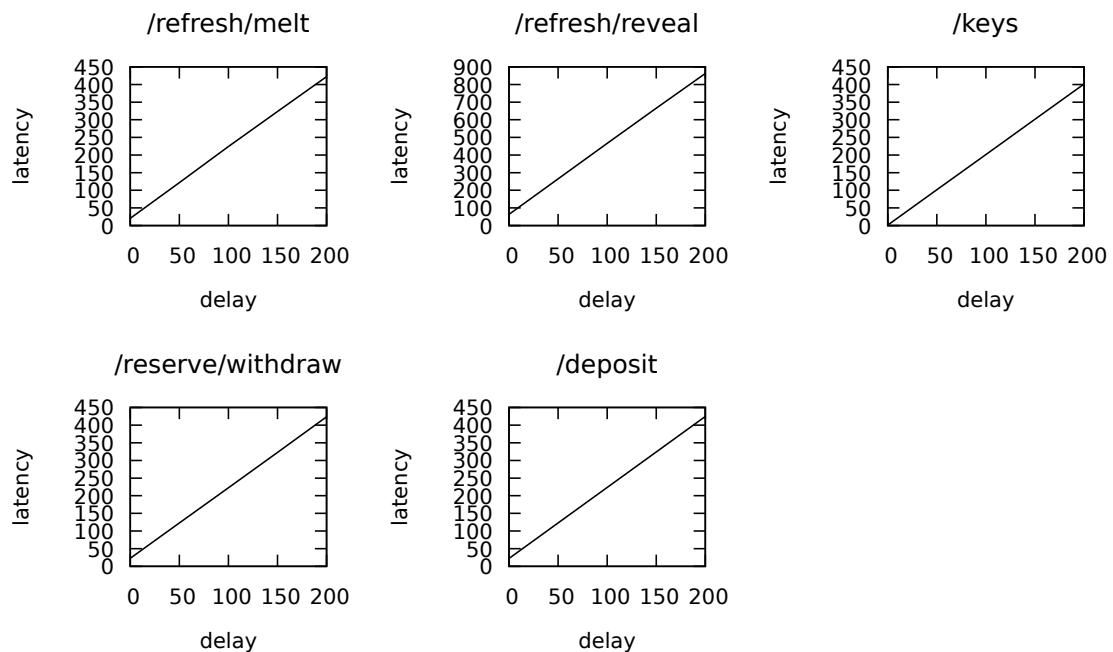


Figure 4.22.: Effect of artificial network delay on exchange’s latency.

4. *Implementation of GNU Taler*

proofs of misbehavior of the exchange/merchant generated by the wallet is not implemented yet.

As a real-world deployment of electronic cash has rather high requirements for the operational security, the usage of hardware security modules for generation of signatures should be considered. Currently, a single process has access to all key material. For a lower-cost improvement that decreases the performance of the system, a threshold signature scheme could be used.

The current implementation is focused on web payments. To use GNU Taler for payments in brick-and-mortar stores, hardware wallets and smartphone apps for devices with near-field-communication (NFC) must be developed. In some scenarios, either the customer or the merchant might not have an Internet connection, and this must be considered in the protocol design. In typical western brick-and-mortar stores, it is currently more likely that the merchant has Internet connectivity, and thus the protocol must allow operations of the wallet (such as refreshing) to be securely routed over the merchant's connection. In other scenarios, typically in developing countries, the merchant (for example, a street vendor) might not have Internet connection. If the vendor has a smartphone, the connection to the merchant can be routed through the customer. In other cases, street vendors only have a "dumb phone" that can receive text messages, and the payment goes through a provider trusted by the merchant that sends text messages as confirmation for payments. All these possibilities must be considered both from the perspective of the protocol and APIs as well as the user experience.

Our experiments were only done with single exchange process and a single database on the same machine. There are various ways to horizontally scale the exchange:

- Multiple exchange processes can be run on multiple machines and access the database that runs a separate machine. Requests are directed to the machines running the exchange process via a load balancer. In this scenario, the throughput of the database is likely to be the bottleneck.
- To avoid having the database as a bottleneck, the contents can be partitioned into shards. For this technique to be effective, data in the shards should not have any dependencies in other shards. A natural way to do sharding for the Taler exchange is to give each shard the sole responsibility for a subset of all available denominations.
- If the transaction volume on one denomination is too high to handle for a single shard, transactions can be further partitioned based on the coin's public key. Each would maintain the database of spent/refreshed coins for a subset of all possible coin public keys. This approach has been suggested for a centrally-banked cryptocurrency by Danezis and Meiklejohn [DM16].

5. Byzantine Set-Union Consensus *

5.1. Introduction

Byzantine consensus is a fundamental building block for fault-tolerant distributed systems. It allows a group of peers to reach agreement on some value, even if a fraction of the peers are controlled by an active adversary. Theory-oriented work on Byzantine consensus often focuses on finding a single agreement on a binary flag or bit string [FHo6]. More recent approaches for practical applications are mainly based on state machine replication (SMR), wherein peers agree on a sequence of state machine transitions. State machine replication makes it relatively easy to lift existing, non-fault-tolerant services to a Byzantine fault-tolerant implementation [CL99]. Each request from a client triggers a state transition in the replicated state machine that provides the service.

A major shortcoming of SMR is that all requests to the service need to be individually agreed upon in sequence by the replica peers of the state machine. This is undesirable since in unoptimized SMR protocols, such as PBFT [CL99], a single transition requires $O(n^2)$ messages to be exchanged for n replicas. Some implementations [Kot+07] try to address this inefficiency by optimistically processing requests and falling back to individual Byzantine agreements only when Byzantine behavior is detected. In practice, this leads to very complex implementations whose correctness is hard to verify and that have weak progress guarantees [Cle+09].

The canonical example for a service where this inefficiency becomes apparent is the aggregation of values submitted by clients into a set. This scenario is relevant for the implementation of secure multiparty computation protocols such as electronic voting [CGS97], where ballots must be collected, and auctions [Bog+09], where bids must be collected. A direct implementation that reaches agreement on a set of m elements with SMR requires m sequential agreements, each consisting of $O(n^2)$ messages.

We introduce Byzantine set-union consensus (BSC) as an alternative communication primitive that allows this aggregation to be implemented more efficiently. In order to implement the set aggregation service described above, the peers first reconcile their sets using an efficient set reconciliation protocol that is not fault-tolerant but where the complexity is bounded even in the case of failures. Then, they use a variant of ByzConsensus [BDH10] to reach Byzantine agreement

*The content of this chapter has been previously published in the EURASIP Journal on Wireless Communications and Networking [DG17].

on the union.

We assume a partially synchronous communication model, where non-faulty peers are guaranteed to successfully receive values transmitted by other non-faulty peers within an existing but unknown finite time bound [DLS88]. Peers communicate over pairwise channels that are authenticated. Message delivery is reliable (i.e., messages arrive uncorrupted and in the right order) but the receipt of messages may be delayed. We make the same assumption as Castro and Liskov [CL99; CL02] about this delay, namely that it does not grow faster than some (usually exponential) function of wall-clock time. We assume a computationally unbounded adversary that can corrupt at most $t = \lceil n/3 \rceil - 1$ peers creating Byzantine faults. The adversary is static, that is the set of corrupted peers is fixed before the protocol starts, but this set is not available to the correct peers. The actual number of faulty peers is denoted by f , with $f \leq t$.

The BSC protocol has message complexity $O(mn + n^2)$ when no peers show Byzantine behavior. When f peers show Byzantine behavior, the message complexity is $O(mnf + kfn^2)$, where k is the number of valid set elements exclusively available to the adversary. We will show how k can be bounded for common practical applications, since in the general case k is only bounded by the bandwidth available to the adversary. In practice, we expect kf to be significantly smaller than m . Thus, $O(mnf + kfn^2)$ is an improvement over using SMR-PBFT which would have complexity $O(mn^2)$.

We have created an implementation of the BSC protocol by combining Ben-Or's protocol for Byzantine consensus [BDH10] with a bounded variant of Eppstein's protocol for efficient set reconciliation [Epp+11]. We demonstrate the practical applicability of our resulting abstraction by using BSC to implement distributed key generation, ballot collection and cooperative decryption from the Cramer-Gennaro-Schoenmakers remote electronic voting scheme [CGS97] in the GUNet framework. Our experimental results show that the behavior of the implementation matches our predictions from theory.

In summary, we make the following contributions in this chapter:

- The introduction of Byzantine Set-Union Consensus (BSC) with Byzantine Eppstein Set Reconciliation.
- The analysis and proof of correctness of Byzantine Set Union Consensus.
- An implementation and experimental evaluation of the protocol.
- A discussion of practical applications to Secure Multiparty Computation.

5.2. Background

The Byzantine consensus problem [LSP82] is a generalization of the consensus problem where the peers that need to reach agreement on some value might also exhibit Byzantine faults.

Many specific variants of the agreement problem (such as interactive consistency [FL81], k -set consensus [DMR01], or leader election [MWVoo] and many others [FLM86]) exist. We will focus on the consensus problem, wherein each peer in a set of peers $\{P_1, \dots, P_n\}$ starts with an initial value $v_i \in M$ for an arbitrary fixed set M . At some point during the execution of the consensus protocol, each peer irrevocably decides on some output value $\hat{v}_i \in M$. Informally, a protocol that solves the consensus problem must fulfill the following properties:²

- *Agreement*: If peers P_i, P_j are correct, then $\hat{v}_i = \hat{v}_j$.
- *Termination*: The protocol terminates in a finite number of steps.
- *Validity*: If all correct peers have the same input value v , then all correct peers decide on $\hat{v} = v$.

Some definitions of the consensus problem also include *strong validity*, which requires the value that is agreed upon to be the initial value of some correct peer [Nei94]. The consensus protocol presented in this chapter does not offer strong validity; in fact, for a set union operation this is not exactly desirable as the goal is to have all peers agree a union of all of the original sets, not on some peer's initial subset.

5.2.1. The FLP Impossibility Result

A fundamental theoretical result (often called FLP impossibility for the initials of the authors) states, informally, that no deterministic protocol can solve the consensus problem in the asynchronous communication model, even in the presence of only one crash-fault [FLP85].

While this result initially seems discouraging, the conditions under which FLP impossibility holds are quite specific and subtle [Agu10]. There are a number of techniques to avoid these conditions while still resulting in a useful protocol. For example:

- *Common coins*: Some protocols introduce a shared source of randomness that the adversary cannot predict or bias. This avoids the FLP impossibility result, since the protocol is not deterministic anymore. In practice, these protocols are very complex and often use variants of secret-sharing and weaker forms of Byzantine agreement to implement the common coin [Fel88; FM88; MMR14]. Implementing a common coin oracle resilient against an active adversary is non-trivial and usually required extra assumptions such as a trusted dealer in the startup phase [CKS05] or shared memory [Asp98]. Recent designs to implement a Byzantine fault-tolerant bias-resistant public random generator only scale to hundreds of participants and still have relatively high failure rates (reported at 0.08% for and adversary power bounded at $\frac{1}{3}$ and 32 participants) [Syt+16].

²Different variations and names can be found in the literature. We have chosen a definition that extends to our generalization to sets later on.

5. Byzantine Set-Union Consensus

- *Failure oracles:* Approaches based on unreliable failure detectors [Gue+00] augment the model with oracles for the detection of faulty nodes. Much care has to be taken not to violate correctness of the protocol by classifying too many correct peers as faulty; this is a problem present in early systems such as Rampart [Rei95] and SecureRing [KMM98] as noted by Castro and Liskov [CL99; CL02]. While the theory of failure detectors is quite established for the non-Byzantine case, it is not clear whether they are still useful in the presence of Byzantine faults.
- *Partial synchrony:* A model where a bound on the message delay or clock shift exists but is unknown or is known but only holds from an unknown future point in time is called partial synchrony. The FLP result does not hold in this model [DLS88].
- *Minimal synchrony:* The definition of synchrony used by the FLP impossibility result can be split into three types of synchrony: Processor synchrony, communication synchrony and message ordering synchrony. Dolev et al. [DDS87] show that consensus is still possible if only certain subsets of these three synchrony assumptions are fulfilled.

This work follows the path of [DLS88] in relaxing the full asynchrony assumption behind the FLP impossibility result.

5.2.2. Byzantine Consensus in the Partially Synchronous Model

The protocols presented in this chapter operate within the constraints of the partially synchronous model, where participants have some approximate information about time.

A fundamental result is that no Byzantine consensus protocol with n peers can support $\lceil n/3 \rceil$ or more Byzantine faults in the partially synchronous model [DLS88].

Early attempts at implementing Byzantine consensus with state machine replication are SecureRing [KMM98] and Rampart [Rei95]. A popular design in the partially synchronous model is Castro and Liskov's Practical Byzantine Fault Tolerance (PBFT) [CL99; CL02]. PBFT uses a leader to coordinate peers (called *replicas* in PBFT terminology). When replicas detect that the leader is faulty, they run a leader-election protocol to appoint a new leader.

PBFT guarantees progress as long as the message delay does not grow indefinitely for some fixed growth function³. The approach taken by PBFT (and several derived protocols) has several problems [Cle+09]: In practice, malicious participants are able to slow down the system significantly. When facing an adversarial scheduler that violates PBFT's weak synchrony assumption, PBFT can fail to make progress entirely [Mil+16].

Some more recent Byzantine state machine replication protocols such as Q/U [Abd+05] or Zyzzyva [Kot+07] have less overhead per request since they

³In practice, exponential back-off is used.

optimize for the non-Byzantine case. This comes, however, often at the expense of robustness in the presence of Byzantine faults [Cle+09], not to mention that correctness proofs for the respective protocols and the implementation of state machine replication are notoriously difficult [Aub+15].

5.2.3. Gradecast

A key building block for our protocol is Feldman’s Gradecast protocol [FM88]. In contrast to an unreliable broadcast, Gradecast provides correctness properties to the receivers, even if the leader is exhibiting Byzantine faults.

In a Gradecast, a leader P_L broadcasts a message m among a fixed set $\mathcal{P} = \{P_1, \dots, P_n\}$ of peers. For notational convenience, we assume that $P_L \in \mathcal{P}$. These are the communication steps for peer P_i :

1. LEAD: If $i = L$, send the input value v_L to \mathcal{P}
2. ECHO: Send the value received in LEAD to \mathcal{P} .
3. CONFIRM: If a common value \bar{v} was received at least $n - t$ times in round ECHO, send \bar{v} to \mathcal{P} . Otherwise, send nothing.

Afterwards, each peer assigns a confidence value $c_i \in \{0, 1, 2\}$ that “grades” the correctness of the broadcast. The result is a graded result tuple $\langle \hat{v}_i, c_i \rangle$ containing the output value \hat{v}_i and the confidence c_i . The grading is done with the following rules:

- If some \hat{v} was received at least $n - t$ times in CONFIRM, output $\langle \hat{v}, 2 \rangle$.
- Otherwise, if some \hat{v} was received at least $t + 1$ times in CONFIRM, output $\langle \hat{v}, 1 \rangle$.
- Otherwise, output $\langle \perp, 0 \rangle$. Here, \perp denotes a special value that indicates the absence of a meaningful value.

For the c_i , the following correctness properties must hold:

1. If $c_i \geq 1$ then $\hat{v}_i = \hat{v}_j$ for correct P_i and P_j
2. If P_L is correct, then $c_i = 2$ and $\hat{v}_i = v_L$ for correct P_i .
3. $|c_i - c_j| \leq 1$ for correct P_i and P_j .

When a correct peer P_i receives a Gradecast with confidence 2, it can deduce that all other peers received the same message, but some other peers might have only received it with a confidence of 1. Receiving a Gradecast with confidence 1 also guarantees that all other correct peers received the same message. However, it indicates that the leader behaved incorrectly. No assumption can be made about the confidence of other peers. Receiving a Gradecast with confidence 0

5. Byzantine Set-Union Consensus

indicates that the leader behaved incorrectly and, crucially, that all other correct peers *know* that the leader behaved incorrectly.

A simple counting argument proves that the above protocol satisfies the three Gradecast properties [FM88].

5.2.4. ByzConsensus

ByzConsensus [BDH10] uses Gradecast to implement a consensus protocol for simple values. Each peer begins with a starting value $s_i^{(1)}$ and the list of all n participants \mathcal{P} . Each peer also starts with an empty blacklist of corrupted peers. If a peer is ever blacklisted, it is henceforth excluded from the protocol. In ByzConsensus, Gradecast is used to force corrupt peers to either expose themselves as faulty—and consequently be excluded—by gradecasting a value with low confidence, or to follow the protocol and allow all peers to reach agreement.

ByzConsensus consists of at most $f + 1$ sequentially executed super-rounds $r \in 1 \dots f + 1$ where $f \leq t$. In each super-round, each peer leads a Gradecast using their candidate value $s_i^{(r)}$; these n Gradecasts can be executed in parallel. Leaders where the Gradecast results in a confidence of less than 2 are put on the blacklist. Recall that different correct peers might receive a Gradecast with different confidence; thus, peers do not necessarily agree on the blacklist.

At the end of each super-round, each peer computes a new candidate value $s_i^{(r+1)}$ using the value that was received most often from the Gradecasts with a confidence of at least 1. If $s_i^{(r)}$ was received more than $n - t$ times, then $r = f$ and the next round is the last round.

If the final candidate value does not receive a majority of at least $2t + 1$ among the n Gradecasts, or if the blacklist has more than t entries, then the protocol failed: either more than t faults happened or, in the partially synchronous model, correct peers did not receive a message within the designated round due to the delayed delivery.

ByzConsensus has message complexity $O(fn^3)$. While the asymptotic message complexity is obviously worse than the $O(n^2)$ of PBFT, there is a way to use set reconciliation to benefit from the parallelism of the Gradecast rounds and thereby reduce the complexity to $O(fn^2)$.

5.2.5. Set Reconciliation

The goal of set reconciliation is to identify the differences between two large sets, say S_a and S_b , that are stored on two different machines in a network. A simple but inefficient solution would be to transmit the smaller of the two sets, and let the receiver compute and announce the difference. Research has thus focused on protocols that are more efficient than this naïve approach with respect to the amount of data that needs to be communicated when the sets S_a and S_b are large,

but their symmetric difference $S_a \oplus S_b$ is small.

An early attempt to efficiently reconcile sets [MTZ03] was based on representing sets by their characteristic polynomial over a finite field. Conceptually, dividing the characteristic polynomials of two sets cancels out the common elements, leaving only the set difference. The characteristic polynomials are transmitted as a sequence of sampling points, where the number of sampling points is proportional to the size of the symmetric difference of the sets S_a and S_b . The number of sampling points can be approximated with an upper bound, or increased on the fly should a peer be unable to interpolate a polynomial. While theoretically elegant, the protocol is not efficient in practice. The computational complexity of the polynomial interpolation grows as $O(|S_a \oplus S_b|^3)$ and uses rather expensive arithmetic operations over large finite fields.

A practical protocol was first proposed by Eppstein et al. in 2011 [Epp+11]. It is based on invertible Bloom filters (IBFs), a probabilistic data structure that is related to Bloom filters [Blo70], and stratas for difference estimation.

Invertible Bloom Filters

An IBF is a constant-size data structure that supports four basic operations, *insert*, *delete*, *decode* and *subtract*.

Insert and *delete* operations are commutative operations encoding a *key* that uniquely identifies a set element, typically derived from the element via a hash function.

The *decode* operation can be used to extract some or all of the updates, returning the key and the sign of the operation, that is either *insert* or *delete*. Since the data structure uses constant space, decoding cannot always succeed. Decoding is a probabilistic operation that is more likely to succeed when the IBF is sparse, that is the number of encoded operations (*excluding* the operations that canceled each other out) is small. The decoding process can also be partially successful, if some elements could be extracted but the remaining IBF is non-empty. Extracting an update by decoding an IBF is only possible if the key was recorded only once in the IBF. However, storing a deletion or insertion of the same key twice or more (not counting operations that canceled each other out) makes both updates impossible to decode.

IBFs of the same size can also be *subtracted* from each other. When subtracting IBF_b from IBF_a , the resulting structure $\text{IBF}_c := \text{IBF}_a - \text{IBF}_b$ contains all insertions and deletions from IBF_a , and insertions from IBF_b are recorded as deletions in IBF_c and deletions from IBF_b are recorded as insertions in IBF_c . Effectively, the IBF subtraction allows to compute the difference between two sets simply by encoding each set as an IBF using only insertion operations.

Under the hood, an IBF of size n is an array of n buckets. Each bucket holds three values:

- A signed counter that handles overflow via wrap-around. Recording an insertion or deletion adds -1 or $+1$ to the counter, respectively.

5. Byzantine Set-Union Consensus

- An \oplus -sum⁴, called the `keySum`, over the keys that identify set the elements that were recorded in the bucket.
- An \oplus -sum, called the `keyHashSum`, over a the hash $h(\cdot)$ of each key that was recorded in the bucket.

As with ordinary Bloom filters, encoding an update in an IBF records the update in k different buckets of the IBF. The indices of buckets that record the update are derived via a k independent hash functions from the key of the element that is subject of the update. We write $\text{Pos}(x)$ for the set of array positions that correspond to the element key x .

Before we describe the decoding process, we introduce some terminology. A bucket is called a *candidate bucket* if its counter is -1 or $+1$, which might indicate that the `keySum` field contains the key of an element that was the subject of an update. Candidate buckets that contain the key of an element that was previously updated are called *pure buckets*. Candidate buckets are not necessarily pure buckets, since a candidate bucket could also result from, for example, first inserting an element key e_1 and then deleting e_2 when $\text{Pos}(e_1) \cap \text{Pos}(e_2) \neq \emptyset$ and $\text{Pos}(e_1) \neq \text{Pos}(e_2)$.

The `keyHashSum` provides a way to detect if a candidate bucket is not a pure bucket, namely when $h(\text{keySum}) \neq \text{keyHashSum}$. The probability of classifying an impure bucket as pure with this method is dependent on the probability of a hash collision. Another method to check for an impure candidate bucket with index i is to check whether $i \notin \text{Pos}(\text{keySum})$.

The decoding process then simply searches for buckets that are, with high probability, pure. When the `count` field of the bucket is 1 , the key decoding procedure reports the key as “inserted” and executes a deletion operation with that key. When the `count` field is -1 , the key is reported as “deleted” and subsequently an insertion operation is executed.

With a probability that increases with sparser IBFs, decoding one element may cause one or more other buckets to become pure, allowing the decoding to be repeated. If none of the buckets is pure, the IBF is undecodable, and a larger IBF must be used, or the reconciliation could fall back to the naïve approach of sending the whole set.

The IBF decoding process is particularly suitable for reconciling large sets with small differences. When the symmetric difference between the sets is small enough compared to the size of the IBFs, the result IBF_c of the subtraction can be decoded, since the common elements encoded in IBF_a and IBF_b cancel each other out. This makes it possible to obtain the elements of the symmetric difference, even when the IBFs that represent the full sets cannot be decoded.

As long as the symmetric difference between the original sets S_a and S_b can be approximated well enough, IBFs can be used for set reconciliation by encoding S_a in IBF_a and S_b in IBF_b . One of the IBFs is sent over the network, the $\text{IBF}_c =$

⁴The \oplus denotes bit-wise exclusive or.

$\text{IBF}_a - \text{IBF}_b$ is computed and decoded. Should the decoding (partially) fail, the same procedure is repeated with larger IBFs.

Difference Estimation with Stratas

In order to select the initial size of the IBF appropriately for the set reconciliation protocol, one needs an estimate of the symmetric difference between the sets that are being reconciled. Eppstein et al. [Epp+11] describe a simple technique, called strata estimation, that is accurate for small differences. While Eppstein et al. suggest combining the strata estimator, with a min-wise estimator, which is more accurate for large differences, our work only requires the strata estimators.

A strata estimator is an array of fixed-size IBFs. These fixed-size IBFs are called *strata* since each of them contains a sample of the whole set, with increased sampling probability towards inner strata. Similar to how two IBFs can be subtracted, strata estimators are subtracted by pairwise subtraction of the IBFs they consist of.

The set difference is estimated by having both peers encode their set in a strata estimator. One of the strata estimators is then sent over to the other peer, which subtracts the strata estimators from each other. With every IBF of the strata estimator that results from the subtraction, a decoding attempt is made. The number of successfully decoded elements in each stratum allows an estimate to be made on the set difference, which is then used to determine the size of the IBF for the actual set reconciliation.

5.3. Our Approach

We now describe how to combine the previous approaches into a protocol for Byzantine fault-tolerant set consensus. The goal of the adversary is to sabotage timely consensus among correct peers, e.g., by increasing message complexity or forcing timeouts.

A major difficulty with agreeing on a set of elements as a whole is that malicious peers can initially withhold elements from the correct peers and later send them only to a subset of the correct peers. This could possibly happen at a time when it is too late to reconcile the remaining difference caused by distributing these elements. We assume that the number of these elements that are initially known to the adversary but not to all correct peers is bounded by k , where k exists but is not necessarily known to the correct participants.

5.3.1. Definition

We now give a definition of set-union consensus that is motivated by practical applications to secure multiparty computation protocols such as electronic voting, which are discussed in more detail in Section 5.7.

5. Byzantine Set-Union Consensus

Consider a set of n peers $\mathcal{P} = \{P_1, \dots, P_n\}$. Fix some (possibly infinite) universe M of elements that can be represented by a bit string. Each peer P_i has an initial set $S_i^{(0)} \subseteq M$.

Let $R : \mathcal{P}(M) \rightarrow \mathcal{P}(M)$ be an idempotent function that canonicalizes subsets of M by replacing multiple conflicting elements with the lexicographically smallest element in the conflict set and removes invalid elements. What is considered conflicting or invalid is application-specific. During the execution of the set-union consensus protocol, after finite time each peer P_i irrevocably commits to a set S_i such that:

1. For any pair of correct peers P_i, P_j it holds that $S_i = S_j$.
2. If P_i is correct and $e \in S_i^0$ then $e \in S_i$.
3. The set S_i is canonical, that is $S_i = R(S_i)$.

The canonicalization function allows us to set an upper bound on the number of elements that can simultaneously be in a set. For example, in electronic voting, canonicalization would remove malformed ballots and combine multiple different (encrypted) ballots submitted by the same voter into a single “invalid” ballot for that voter.

5.3.2. Byzantine Set-Union Consensus (BSC) Protocol

Recall that every peer P_i , $0 < i \leq n$ starts with a set $S_i^{(0)}$. The BSC protocol incorporates two subprotocols, bounded set reconciliation and lower bound agreement, and uses those to realize an efficient Byzantine fault-tolerant variant of ByzConsensus. An existing generalization of IBFs to multi-party set reconciliation [MP13] based on network coding is not applicable to this problem, as it requires trusted intermediaries.

The basic problem solved by the two subprotocols is bounding the cost of Eppstein’s set reconciliation. Given a set size difference between two peers of k , the expected cost of Eppstein’s set reconciliation is $O(k)$ if both participants are honest. However, we need to ensure that malicious peers cannot generally raise the complexity to $O(m)$ where m is the size of the union.

For this, we use a bounded variant of Eppstein’s set reconciliation protocol, which is given a lower bound \mathcal{L} on the size of the set of elements shared by all honest participants. Given such a lower bound, the bounded set reconciliation protocol must detect faulty participants in $O(k + (m - \mathcal{L}))$. We note that for $\mathcal{L} = 0$, the bounded set reconciliation is still allowed to cost $O(m)$.

Bounded Set Reconciliation

In *bounded* set reconciliation we are thus concerned with modifications that ensure that a set reconciliation step between an honest and a faulty peer either succeeds after $O(k)$ traffic, or aborts notifying the honest peer that the other peer is faulty. While we use probabilities to detect faulty behavior, we note that suitable

parameters can be used to ensure that false-positives are rare, say $1 : 2^{128}$, and thus as unlikely as successful brute-force attacks against canonical cryptographic primitives, which BSC also assumes to be safe.

To begin with, to bound the complexity of Eppstein set reconciliation one needs to bound the number of iterations the protocol performs. Assuming honest peers, the initial strata estimation ensures that the IBFs will decode with high probability, resulting in Eppstein’s claim of single-round complexity. Given aggressive choices of the parameters to improve the balance between round-trips and bandwidth consumption, decoding failures can happen with non-negligible probability in practice. In this case, the process can simply be restarted using a different set of hash functions and an IBF doubled in size. This addresses issues caused by conservative choices for IBF sizes that optimize for the average case. What is critical is that the probability of such failures remains small enough that after if the number of rounds exceeds some constant, we can assert faulty behavior and overall remain within the $O(k)$ bound assuming individual rounds are bounded by $O(k)$.

Another problem with Eppstein’s original protocol related to aggressive parameter choices is that iterative decoding does not always have to end with an empty or an undecodable IBF. Specifically, the decoding step can sometimes decode a key that was never added to the IBF, simply because the two purity checks are also probabilistic. This is usually not an issue, as when a decoder requests the transmission of the element corresponding to improperly decoded key, the presumed element’s owner can indicate a decoding failure at that time. Here, another round of the protocol is unlikely to produce the same error and would again fix the problem. However, given reasonably short strings for the `hashKeySum`, it is actually even possible to obtain a looping IBF that spawns an infinite series of “successfully” decoded keys. Here, the implementor has to be careful to ensure that the iterated hash decoding algorithm terminates. Instead of mandating an excessively long `hashKeySum` to prevent this, it is in practice more efficient to handle this case by stopping the iteration and reporting the IBF as undecodable when the number of decoded keys exceeds a threshold proportional to the size of the IBF.

We also need to consider the bandwidth consumption of an individual round. To cause more than $O(k)$ traffic, a malicious peer could produce strata that result in a huge initial symmetric difference. In this case, the initial size of the IBF may exceed $O(k)$. We address this problem by not permitting the use of Eppstein’s method if the symmetric difference definitively exceeds $\frac{n-L}{2}$, where n is the smaller of the two set sizes.⁵ Instead, once the estimate of the symmetric difference substantially exceeds this threshold, the reconciliation algorithm falls back to sending the complete set. As this creates $O(m)$ traffic, it must only be allowed under certain conditions.

⁵The optimal formula here depends on the size ratio of IBF element to the transmission size of an individual element and the estimated size of the set overlap. However, to simplify the exposition, we will assume a simple 50% threshold henceforth.

5. Byzantine Set-Union Consensus

First, we consider the case where the honest peer has the larger set. Here, the honest peer P_i will only send its full set if the set difference is no larger than $|S_i| - \mathcal{L}$, and otherwise report a fault. This ensures that a malicious peer cannot arbitrarily request the full set from honest peers.

Second, we consider the case where the honest peer P_i is facing a faulty peer that claims to have a huge set. This can happen either directly from the strata estimator, or after P_i observes a constant number of successive IBF decoding failures.⁶ At this point, instead of passively accepting the transmission of elements, the receiver P_i checks that a sufficient number of the elements received are not in S_i . Let R be the stream of elements e received at any point in time. We assume that the sender is required to transmit the elements in randomized order. Thus, if $|R \cap S_i| - |R \setminus S_i| \geq 128$, P_i can determine that the sender is faulty with probability $2^{128} : 1$, as the $\frac{n}{2}$ -threshold for converting to complete set transmission ensures that for an honest sender less than half of the elements would be in S_i .

Finally, we note that the individual *insert*, *delete*, *decode* and *subtract* operations on the IBF are all constant time and that IBFs are also constant size. Thus, given a constant number of rounds and a bound on the bandwidth per round, we have implicitly assured that memory and CPU consumption of the bounded set reconciliation is also $O(k + (m - \mathcal{L}))$.

Lower Bound Agreement

To provide a lower bound on the permissible set size for set reconciliation, BSC first executes a protocol for *lower bound agreement* (LBA). In this first step, every correct peer P_i learns a superset $S_i^{(1)}$ of the union of all correct peers' initial sets, as well as a lower bound ℓ_i for the minimum number of elements shared by all correct peers where $n - \ell_i \leq k$. Note that neither $S_i^{(1)} = S_j^{(1)}$ nor $\ell_i = \ell_j$ necessarily hold even for correct peers P_i and P_j . Our LBA protocol proceeds in three steps:

- (i) All peers reconcile their initial set with each other, using pairwise bounded set reconciliation using a lower bound of $\mathcal{L} = 0$.
- (ii) All peers send their current set size to each other, and each peer P_i sets ℓ_i to the $(t + 1)$ -smallest set size that P_i received.
- (iii) All peers again reconcile their sets with each other, using pairwise bounded set reconciliation.

The third step is necessary to ensure that every correct P_i has at least ℓ_i elements, since malicious peers could use the k elements initially withheld to force an honest

⁶Each failure causes the IBF size to double and thus corresponds to a doubling of the set difference estimate. Thus, the number of decoding failures could remain the threshold that causes an abort, while the set difference estimate substantially exceeds $2(|S_i| - \mathcal{L})$.

peer's set size below the $(t + 1)$ -smallest set size. Thanks to the repetition even if ℓ_i is different for each peer, it is guaranteed that P_i has at least ℓ_i elements in common with every other good peer.

In subsequent set reconciliations, ℓ_i can be used to bound the traffic that malicious peers are able to cause by falsely claiming to have a large number of elements missing. LBA itself has complexity $O(nmf)$: initially all malicious peers can *once* claim to have empty sets with all other peers. LBA ensures that for the remainder of the protocol, a correct peer with m_i elements can stop sending elements to malicious peer P_M after P_M requested $m_i - \ell_i \leq k$ elements by reducing the complexity of bounded set reconciliation with peer m_i to $O(k)$ using $\mathcal{L} = \ell_i$.

Exact Set Agreement

After LBA, an *exact set agreement* is executed, where all peers reach Byzantine agreement for a super-set of the set reached in LBA. The exact set agreement is implemented by executing a variant of ByzConsensus which instead of sending values reconciles sets.

The Gradecast is adapted as follows:

- (i) LEAD: If $i = L$, reconcile the input set V_L with \mathcal{P} .
- (ii) ECHO: Reconcile the set received in LEAD with \mathcal{P} .
- (iii) CONFIRM: Let \mathcal{U}_E be the union of all sets received in the ECHO round, and $N_E(e)$ the number of times a single set element e was received.
If $\forall e \in \mathcal{U}_E \ t < N_E(e) < n - t$, send \perp (where $\perp \neq \emptyset$). Otherwise send $\mathcal{U}_E - \{e \mid N_E(e) \leq t\}$ to \mathcal{P} .

The grading rules are also adapted to sets. Let \mathcal{U}_C be the union of sets received in CONFIRM, $N_C^+(e)$ the number of times a single element $e \in \mathcal{U}_C$ was received, and $N_C^-(e)$ the number of sets (not \perp) received in CONFIRM that excluded e .

- If $\bigwedge_{e \in \mathcal{U}} N_C^+(e) \geq n - t \vee N_C^-(e) \geq n - t$,
output $\langle \{e \mid N_C^+(e) \geq n - t\}, 2 \rangle$.
- Otherwise if $\bigwedge_{e \in \mathcal{U}_C} N_C^+(e) > t \wedge N_C^+(e) \geq N_C^-(e)$
or $\bigwedge_{e \in \mathcal{U}_C} N_C^-(e) > t \wedge N_C^-(e) > N_C^+(e)$,
output $\langle \{e \mid N_C^+(e) > t \wedge N_C^+(e) \geq N_C^-(e)\}, 1 \rangle$.
- Otherwise, output $\langle \perp, 0 \rangle$.

Similar to ByzConsensus, the BSC consists of at most $f + 1$ super-rounds, where $f \leq t$. Each peer P_i starts with $S_i^{(1)}$ as its current set. In sequential super-rounds, all peers lead a Gradecast for their candidate set. Like in ByzConsensus, if P_i receives a Gradecast with a confidence value that is not 2, then P_i puts the leader

5. Byzantine Set-Union Consensus

of the Gradecast on its blacklist, and correct peers stop all communication with peers on their blacklist.

At the end of each super-round, peers update their candidate set as follows. Let n' be the number of leaders that gradecasted a set with a non-zero confidence. The new candidate set contains all set elements that were included in at least $\lceil n'/2 \rceil$ sets that were gradecasted with a non-zero confidence value. If all elements occur with a $(n - t)$ -majority, then the next round is the last round. The output of the consensus protocol is the candidate set after the last round—or failure if $f > t$.

We give a correctness proof that generalizes Feldman's proof for Gradecast of single values [Fel88, Section 4.1].

Lemma 1. *If two correct peers send sets $A \neq \perp$ and $B \neq \perp$ respectively in CONFIRM, then $A = B$.*

Proof. Proof by contradiction and counting argument. Assume w.l.o.g. that $e \in A$ and $e \notin B$. At least $n - t$ peers must have echoed a set that includes e to the first peer. Suppose f of these peers were faulty, then at least $n - t - f > t$ good peers included e in the ECHO transmission to the second peer. If $e \notin B$, then $t < N_E(e) < n - t$. In this case, an honest second peer must output $B = \perp$. Contradiction. \square

Theorem 5. *The generalization of Gradecast to sets satisfies the three Gradecast properties.*

Proof. We show that each property holds:

- Property 1 (If $c_i, c_j \geq 1$ then $\hat{V}_i = \hat{V}_j$ for correct P_i and P_j): Assume w.l.o.g. that $e \in \hat{V}_i \setminus \hat{V}_j$.

For $e \in \hat{V}_i$, P_i must have received e at least $N_C^+(e) > t$ times in CONFIRM. Given $f \leq t$ failures, at least one honest peer must thus have included e in CONFIRM. According to Lemma 1, then all $n - f$ honest peers must either include e in CONFIRM or send \perp .

Because \perp is not a set, this leaves at most all $f \leq t$ faulty peers that can send a set without e . But for $e \notin \hat{V}_j$ we need $N_C^-(e) \geq t + 1$. Contradiction.

- Property 2 (If P_L is correct, then $c_i = 2$ and $\hat{V}_i = \hat{V}_L$ for correct P_i): All $n - f \geq n - t$ good peers ECHO and CONFIRM the same set. By the grading rules, they must output a confidence of 2.
- Property 3 ($|c_i - c_j| \leq 1$ for correct P_i and P_j): Proof by contradiction. Assume w.l.o.g. $c_i = 2$ and $c_j = 0$. $c_i = 2$ implies that for each $x \in \hat{V}_i$ at least $n - t$ peers (and thus $(n - t) - f \geq t + 1$ correct peers) must have sent a set in CONFIRM that includes x . For any $y \notin \hat{V}_i$, $n - t$ peers (and thus $(n - t) - f \geq t + 1$ correct peers) must have sent a non- \perp set in CONFIRM that excludes y .

Given $c_j = 0$, there must have been an element e such that, $N_C^+(e) \leq t$ and $N_C^-(e) \leq t$ for P_j . However, we just derived that for all elements either $N_C^+(e) > t$ or $N_C^-(e) > t$. Contradiction. \square

\square

Given the Gradecast properties for sets, the correctness argument given by Ben-Or [BDH10] for the Byzantine consensus applies to BSC's generalization to sets.

As described, the protocol has complexity $O(mnf + fkn^3)$. However, the n parallel set reconciliation rounds in each super-round can be combined by tagging the set elements that are being reconciled in the LEAD, ECHO and CONFIRM rounds with the respective leader L . Because LBA (via $n - \ell_i \leq k$) and bounded set reconciliation limit mischief for the combined super-round, each malicious peer can, as leader, *once* cause bounded set reconciliation during the ECHO round to all-to-all transmit at most k extra elements, resulting in a total of $O(fkn^2)$ extra traffic over all $f + 1$ rounds. Before exposing themselves this way, non-leading malicious peers can only cause $O(f^2kn)$ additional traffic during all ECHO rounds. Finally, malicious peers can also cause at most $O(fkn^2)$ traffic in the CONFIRM round. Thus, BSC has overall message complexity of $O(mnf + fkn^2)$.

5.4. Implementation

We implemented the BSC protocol in the SET and CONSENSUS services of GNUnet [GNUNET].

5.4.1. The GNUnet Framework

GNUnet is composed of various components that run in separate operating system processes and communicate via message passing. Components that expose an interface to other components are called *services* in GNUnet. The main service used by our implementation is the CADET service, which offers pairwise authenticated end-to-end encryption between all participants. CADET uses a variation of the Axolotl public key ratcheting scheme and double-encrypts using both TwoFish and AES [PG14]. The resulting encryption is relatively expensive compared to the other operations, and thus dominates in terms of CPU consumption for the experiments.

5.4.2. Set Reconciliation

Bounded set reconciliation is implemented in the SET service. The SET service provides a generic interface for set operations between two peers; the operations currently implemented are the IBF-based set reconciliation and set intersection [TRL12].

5. Byzantine Set-Union Consensus

In addition to the operation-specific protocols, the following aspects are handled generically (i.e., independent of the specific remote set operation) in the SET service:

Local set operations

Applications need to create sets and perform actions (iteration, insertion, deletion) on them locally.

Concurrent modifications

While a local set is in use in a network operation, the application may still continue to mutate that set. To allow this without interfering with concurrent the network operations, changes are versioned. A network operation only sees the state of a set at the time the operation was started.

Lazy copying

Some applications building on the SET service—especially the CONSENSUS service described in the next section—manage many local sets that are large but only differ in a few elements. We optimize for this case by providing a lazy copy operation that returns a logical copy of the set without duplicating the sets in memory.

Negotiating remote operations

In a network operation, the involved peers have one of two roles: The acceptor, which waits for remote operation requests and accepts or rejects them, as well as the initiator, which sends the request.

Our implementation estimates the initial difference between sets only using *strata estimators* as described by Eppstein [Epp+11]. However, we compress the strata estimator—which is 60KB uncompressed—using `gzip`. The compression is highly effective at reducing bandwidth consumption due to the high probability of long runs of zeros or ones in the most sparse or most dense strata respectively.

We also use a *salt* when deriving the bucket indices from the element keys. When the decoding of an IBF fails, the IBF size is doubled and the salt is changed. This prevents decoding failures in scenarios where keys map to the same bucket indices even modulo a power of two, where doubling the size of the IBF does not remove the collision.

5.4.3. Set-Union Consensus

To keep the description of the set-union consensus protocol in the previous section succinct, we merely stated that peers efficiently transmit sets using the reconciliation protocol. However, given that the receiving peer has usually many sets to reconcile against, an implementation needs to be careful to ensure that it scales to large sets as intended.

The key goal is to avoid duplicating full sets and to instead focus on the differences. New sets usually differ in only a few elements, thus our implementation

avoids copying entire sets. Instead, in the leader round we just store the set of differences with a reference to the original set. In the ECHO and CONFIRM round, we also reconcile with respect to the set we received from the leader, and not a peer's current set. In the ECHO round, we only store one set and annotate each element to indicate which peer included or excluded that element. This also allows for a rather efficient computation of the set to determine the \perp -result in the CONFIRM round.

5.4.4. Evaluating Malicious Behavior

For the evaluation, our CONSENSUS service can be configured to exhibit the following types of adversarial behavior:

- *SpamAlways*: A malicious peer adds a constant number of additional elements in every reconciliation.
- *SpamLeader*: A malicious peer adds a constant number of additional elements in reconciliations where the peer is the leader.
- *SpamEcho*: A malicious peer adds a constant number of additional elements in echo rounds.
- *Idle*: Malicious peers do not participate actively in the protocol, which amounts to a crash fault from the start of the protocol. This type of behavior is not interesting for the evaluation, but used to test the implementation with regards to timeouts and majority counting.

For the *Spam*-* behaviors, two different variations are implemented. One of them ("*-replace") always generates new elements for every reconciliation. This is not typical for real applications where the number of storable elements ought to be limited by set canonicalization. However, this shows the performance impact in the worst case. The other variation ("*-noreplace") reuses the same set of additional elements for all reconciliations, which is more realistic for most cases. We did not implement adversarial behaviour where elements are elided, since the resulting traffic is the same as for additional elements, and memory usage would only be reduced.

5.5. Experimental Results

All of the experiments were run on a single machine with a 24-core 2.30GHz Intel Xeon E5-2630 CPU, and GUNet SVN revision 36765. We used the `gnunet-consensus-profiler` tool, which is based on GUNet's TESTBED service [Tot13], to configure and launch multiple peers on the target system. We configured the profiler to emulate a network of peers connected in a clique topology (via loop-back, without artificial latency). Elements for the set operations are randomly generated and always 64 bytes large.




Figure 5.1.: CPU system time for the SET service in relation to total set size. Average over 50 executions.

Bandwidth consumption was measured using the statistics that GUNet’s CADET service [PG14] provides. Processor time was measured using GUNet’s resource reporting functionality, which uses the `wait3` system call for that purpose.

5.5.1. Bounded Set Reconciliation

We now summarize the experimental results for the bounded set reconciliation protocol between two peers. We first measured the behavior of the set reconciliation if identical sets were given to both peers (Figure 5.1 and 5.2). Figure 5.1 shows that total CPU utilization generally grows slowly as the set size increases. The sudden jump in processing time that is visible at around 7,000 elements can most likely be explained by cache effects. The effect could not be observed when we ran the experiment under profiling tools.

Figure 5.2 shows that bandwidth consumption does not grow linearly with the total set size, as long as the set size difference between the two peers is small. The logarithmic increase of the traffic with larger sets can be explained by the compression of strata estimators: The k -th strata samples the set with probability 2^{-k} , and for small input sets the strata tends to contain long runs of zeros that are more easily compressed.

We also measured the behavior of the set reconciliation implementation if the sets differed. Figure 5.3 and 5.4 show that—as expected—CPU time and bandwidth do grow linearly with the symmetric difference between the two sets.




Figure 5.2.: CADET traffic for the SET service in relation to total set size. Average over 50 executions.




Figure 5.3.: CPU system time for the SET service in relation to symmetric set difference. Average over 50 executions.




Figure 5.4.: CADET traffic for the SET service in relation to symmetric difference. Average over 50 executions.

Finally, we analyzed what happens when the algorithm switches from transmitting set differences to full sets. Figure 5.5 shows the bandwidth in relation to the symmetric set difference, for different total numbers of elements in the shared set. Up to the threshold where the algorithm switches from IBFs to full set transmission, we expect the transmission size to grow steeply, and then afterwards continue linearly at a lower rate again. If the handover threshold is chosen well, the two lines should meet. This is the case in the dashed curve in Figure 5.5. The small bump at a set difference of ≈ 800 is due to an unlucky size estimate by the strata estimator causing the algorithm to initially attempt set reconciliation, before switching to full set transmission. If the threshold between IBF and full set transmission is picked a bit too high and IBFs are sent slightly beyond the point where they are beneficial, the curve from the IBF transmission will peak above the one that represents the full set transmission. This is the case in the solid curve in Figure 5.5. Finally, the dotted curve shows the case where the threshold is picked too low, causing expensive full set transmission to occur when IBFs would have been more useful. Here, we also see a lucky case of underestimating the size of the difference. We note that given the size of an IBF entry, the average size of a set element and an estimate of the size overlap, near-perfect thresholds (instead of the 50%-heuristic we described earlier) can be trivially computed.




Figure 5.5.: CADET traffic for the SET service in relation to symmetric difference at the boundary between IBF and full set transmission. Note that we did cherry-pick runs for this graph. Our goal is to illustrate how the curves evolve with regard to different thresholds between IBF and full set transmission. We also wanted to show how significant deviations in set difference estimates generated by the strata estimator can have a minor impact on performance.

5. Byzantine Set-Union Consensus




Figure 5.6.: CADET traffic for BSC per peer for 100 elements and only correct peers. Average over 50 executions.

5.5.2. Byzantine Set Consensus

For our experiments with the BSC implementation, all ordinary peers start with the same set of elements; different sets would only affect the all-to-all union phase of the protocol which does pairwise set reconciliation, resulting in increased bandwidth and CPU consumption proportional to the set difference as shown in the previous section.

As expected, traffic increases cubically with the number of peers when no malicious peers are present (Figure 5.6). Most of the CPU time (Figure 5.7) is taken up by CADET, which uses expensive cryptographic operations [PG14]. Since we ran the experiments on a multicore machine, the total runtime follows the same pattern as the traffic (Figure 5.8).

We now consider the performance implications from the presence of malicious peers. Figures 5.10 and Figure 5.11 show that bandwidth and runtime increase linearly with the additional elements malicious peers can exclusively supply, in contrast to the sub-linear growth for the non-Byzantine case (Figure 5.2).

Figure 5.11 highlights how the different attack strategies impact the number of additional elements that were received during set reconciliations: The number of stuffed elements for the “SpamEcho” behavior is significantly larger than for “SpamLead”, since multiple ECHO rounds are executed for one LEAD round, and the number of stuffed elements is fixed per reconciliation. When malicious peers add extra elements during the LEAD round, the effect of that is amplified, since all correct receivers have to re-distribute the additional elements in the




Figure 5.7.: CPU of BSC for 100 elements of 64 bytes and only correct peers. Average over 50 executions.




Figure 5.8.: Runtime of BSC for 100 elements of 64 bytes and only correct peers. Average over 50 executions.

5. Byzantine Set-Union Consensus




Figure 5.9.: CADET traffic for BSC on 100 elements of 64 bytes and one malicious peer with the indicated mode. Average over 50 executions.




Figure 5.10.: Latency for BSC with 4 peers on 100 elements of 64 bytes and one malicious peer with the indicated mode. Average over 50 executions.

ECHO/CONFIRM round. Even though adding elements in the LEAD round requires the least bandwidth from the leader, the effect on traffic and latency is the largest (see Figures 5.9 and 5.10).

As expected, when the number of stuffed elements is limited to a fixed set, the effect on the performance is limited (“SpamAll-noreplace” in Figures 5.9, 5.10, 5.11).

5.6. Opportunities for Further Improving BSC

We now discuss some of the key limitations of the current implementation and, how it could be optimized further.

5.6.1. Extension to Partial Synchrony

The prototype used in the evaluation only works in the synchronous model. It would be trivial to extend it to the partially synchronous model with synchronous clocks by using the same construction as PBFT [CL99], namely retrying the protocol with larger round timeouts (usually doubled on each retry) when it did not succeed.

It might be worthwhile to further investigate the Byzantine round synchronization protocols discovered independently by Attya and Dolev [ADG84] as well




Figure 5.11.: Total number of extra elements received by each peer for BSC on 100 elements of 64 bytes and one malicious peer with the indicated mode. Average over 50 executions.

5. Byzantine Set-Union Consensus

as Dwork, Lynch and Stockmeyer [DLS88]. Running a Byzantine clock synchronization protocol interleaved with consensus protocol might lead to a protocol with lower latency, since the timeouts are dynamically adjusted instead of being increased for each failed iteration.

5.6.2. Persistent Data Structures

Both the SET and CONSENSUS service have to store many variations of the same set when faulty peers elide or add elements. While the SET service API already supports lazy copying, the underlying implementation is inefficient and based on a log of changes per element with an associated version number. It might be possible to reduce memory usage and increase performance of the element storage by using data structures that are more well suited, such as the persistent data structures described by Okasaki [Oka98].

5.6.3. Fast Dissemination

Recall that in order to be included in the final set, an element must be sent to at least $t + 1$ peers, so that at least one correct peer will receive the element. In applications of set-union consensus such as electronic voting, the effort to the client should be minimized, and thus in practice, elements might be sent only to $t + 1$ peers, which would lead to large initial symmetric differences between peers.

A possible optimization would be to add another dissemination round that only requires $n \log_2 n$ reconciliations to achieve perfect element distribution when only correct peers are present. The n^2 reconciliations that follow will consequently be more efficient, since no difference has to be reconciled when all peers are correct. In the presence of faulty peers, the optimization adds more overhead due to the additional dissemination round.

More concretely, in the additional dissemination round the peers reconcile with their 2^ℓ -th neighbour (for some arbitrary, fixed order on the peers) in the ℓ -th subround of the dissemination round. After $\lceil \log_2 \rceil$ of these subrounds, the elements are perfectly distributed as long as every peer passed along their current set correctly.

5.7. Application to SMC

Secure multiparty computation (SMC) is an area of cryptography that is concerned with protocols that allow a group of peers $\mathcal{P} = P_1, \dots, P_n$ to jointly compute a function $y = f(x_1, \dots, x_n)$ over private input values x_1, \dots, x_n without using a trusted third party [GL05]. Each peer P_i contributes its own input value x_i , and during the course of the SMC protocol, P_i ideally only learns the output y , but no additional information about the other peers' input values. Applications of SMC include electronic voting, secure auctions and privacy-preserving data mining.

SMC protocols often assume a threshold $t < n$ on the amount of peers controlled by an adversary, which is typically either *honest-but-curious* (i.e., tries to learn as much information as possible but follows the protocol) or *actively malicious*. The actively malicious case mandates the availability of Byzantine consensus as a building block [SZ15].⁷

In practical applications, the inputs typically consist of sets of values that were given to the peers \mathcal{P} by external clients: In electronic voting protocols the peers need to agree on the set of votes; with secure auctions, the peers need to agree on bids, and so on.

In this section, we focus on one practical problem, namely electronic voting. We show how BSC is useful at multiple stages of the protocol, and discuss how our approach differs from existing solutions found in the literature.

5.7.1. Bulletin Board for Electronic Voting

The *bulletin board* is a communication abstraction commonly used for electronic voting [Ben87; Peto5]. It is a stateful, append-only channel that participants of the election can post messages to. Participants of the election identify themselves with a public signing key and must sign all messages that they post to the bulletin board. The posted messages are publicly available to facilitate independent auditing of elections.

Existing work on electronic voting either does not provide a Byzantine fault-tolerant bulletin board in the first place [Adio8] and instead relies on trusted third parties, or suggests the use of state machine replication [CGS97].

Some of the bulletin board protocols surveyed by Peters [Peto5] use threshold signatures to certify to the voter that the vote was accepted by a sufficiently large fraction of the peers that jointly provide the bulletin board service. With a naïve approach, the message that certifies acceptance by t peers is the concatenation of the peers' individual signatures and thus $O(t)$ bits large. Threshold signature schemes allow smaller signatures, but at the expense of a more complex protocol. Since the number of peers is typically not very large, a linear growth in t is acceptable, which makes the simple scheme sufficient for practical implementations.

It is easy to implement a variant of the bulletin board with set-union consensus. In contrast to traditional bulletin boards, this variant has *phases*, where posted messages are only visible after the group of peers have agreed that a phase is concluded. The concept of phases maps well to the requirements of existing voting protocols. Every phase is implemented with one set-union consensus execution. To guarantee that a message is posted to the bulletin board, it must be sent to at least one correct peer from the group of peers that jointly implements

⁷An attempt has been made to relax the definition of SMC to alleviate this requirement, resulting in a weaker definition that includes non-unanimous *aborts* as a possible result [GL05]. This definition is mainly useful in scenarios without an non-faulty 2/3 majority, where Byzantine consensus is not possible in the asynchronous model [DLS88].




Figure 5.12.: Relation of different SMC protocols and communication primitives in GNUnet. Dashed arrows indicate optional dependencies.

the bulletin board.

5.7.2. Distributed Threshold Key Generation and Cooperative Decryption

Voting schemes as well as other secure multiparty computation protocols often rely on threshold cryptography [Des94]. The basic intuition behind threshold cryptography is that some operations—such as signing a message or decrypting a ciphertext—should only succeed if a large enough fraction of some group of peers cooperate. Typically, the public key of the threshold cryptosystem is publicly known, while the private key is not known by any entity but reconstructible from the shares that are distributed among the participants, for example, with Shamir’s secret sharing scheme [Sha79].

Generating this shared secret key either requires a trusted third party, or a protocol for distributed key generation [FS01; Ped91]. The former is undesirable for most practical applications since it creates a single point of failure.

In a distributed key generation protocol, each peer contributes a number of *pre-shares*. The peers agree on the set of pre-shares and each peer re-combines

them in a different way, yielding the shares of the private threshold key.

In the key generation protocol used for the Cramer et al. voting scheme, the number of pre-shares that need to be agreed upon is quadratic in the number of peers. Every peer needs to know every pre-share, even if it is not required by the individual peer for reconstructing the share, since the pre-shares are accompanied by non-interactive proofs of correctness. Thus, the number of values that need to be agreed upon is quadratic in the number of peers, which makes the use of set-union consensus attractive compared to individual agreement.

Even though the pre-shares can be checked for correctness, Byzantine consensus on the set of shares is still necessary for the case when a malicious peer submits a incorrect share to only some peers. Without Byzantine consensus, different correct recipients might exclude different peers, resulting in inconsistent shares.

Similarly, when a message that was encrypted with the threshold public key shall be decrypted, every peer contributes a *partial decryption* with a proof of correctness. While the set of partial decryptions is typically linear in the number of peers, set-union consensus is still a reasonable choice here, this way the whole system only needs one agreement primitive.

5.7.3. Electronic Voting with Homomorphic Encryption

Various conceptually different voting schemes use homomorphic encryption; we look at the scheme by Cramer et al. [CGS97] as a modern and practical representative. A fundamental mechanism of the voting scheme is that a set of voting authorities A_1, \dots, A_n establish a threshold key pair that allows any entity that knows the public part of the key to encrypt a message that can only be decrypted when a threshold of the voting authorities cooperate. The homomorphism in the cryptosystem enables the computation of an encrypted tally with only the ciphertext of the submitted ballots. Ballots represent a choice of one candidate from a list of candidate options. The validity of encrypted ballot is ensured by equipping them with a non-interactive zero-knowledge proof of their validity.

It is assumed that the adversary is not able to corrupt more than $1/3$ of the authorities. The voting process itself is then facilitated by all voters encrypting their vote and submitting it to the authorities. The encrypted tally is computed by every authority and then cooperatively decrypted by the authorities and published. Since correct authorities will only agree to decrypt the final tally and not individual ballots, the anonymity of the voter is preserved. For the voting scheme to work correctly, all correct peers must agree on exactly the same set of ballots before the cooperative decryption process starts, otherwise the decryption of the tally will fail.

Using BSC for this final step to agree on a set of ballots again makes sense, as the number of ballots is typically much larger than the number of authorities. Figure 5.12 summarizes the various ways how BSC and is used in our implementation [Dol14] of Cramer-style [CGS97] electronic voting.

5.7.4. Other Applications of BSC

Bitcoin [Nako8] has gained immense popularity over the past few years. Bitcoin solves a slight variation of Byzantine consensus without strong validity [ML14; GKL15]. Given that a block in Bitcoin is basically just a set of (valid) transactions, BSC could be used to efficiently achieve agreement between participants about the next transaction group. Here, the most natural application would be to use BSC to improve the efficiency of proof-of-stake incentivized peers running BFT consensus in Cosmos [KB16].

5.8. Conclusions

Given m ballots, n authorities, f Byzantine faults and k ballots exclusively available to the adversary, voting with BSC achieves a complexity of $O(mn + (f + k)n^3)$, which in practice is better than the $O(mn^2)$ complexity of using SMR as m is usually significantly larger than n . Equivalent arguments hold for other applications requiring consensus over large sets. Furthermore, BSC remains advantageous in the absence of Byzantine failures, and the bounded set reconciliation makes it particularly efficient at handling various non-Byzantine faults.

To ensure these performance bounds, BSC requires a bounded variant of Epstein's set reconciliation protocol that ensures that individual steps in the protocol cannot consume unbounded amounts of bandwidth. We are currently applying bounded set reconciliation in related domains, as any set reconciliation can be made more robust if the complexity of the operation is bounded. For example, the GNU Name System [WSG14] can use bounded set reconciliation when gossiping sets of key revocation sets. Here, the use of bounded set reconciliation protects the key revocation protocol against denial-of-service attacks where an attacker might have previously sent excessively large IBFs or retransmitted known revocation messages already known to the recipient. The result is an efficient and resilient method for disseminating key revocation data.

In future work, it would be interesting to apply bounded set reconciliation to Byzantine consensus protocols that are more efficient than the simple gradecast consensus. It would also be interesting to experimentally compare bulletin boards using BSC with those using traditional replicated state machines.

6. Future Work

We now discuss future work that builds upon the results presented so far.

Standard Model

Our current instantiation of the Taler protocol relies heavily on hash functions. Since the result by Canetti and others [CGH04] about the theoretical impossibility of securely instantiating protocols that rely on the random oracle assumption for their security, a vast amount of literature has been devoted to find instantiations of interesting protocols in the standard model [KM15]. The Taler protocol syntax could likely be also instantiated securely in the standard model, based existing on blind signature schemes in the standard model. The trade-off however is that while removing the random oracle assumption, typically other less well known assumptions must be made.

Post-Quantum security

The possibility of post-quantum computers breaking the security of established cryptographic primitives has lately received a lot of attention from cryptographers. While currently most schemes with post-quantum security are impractical, it might be worthwhile to further investigate their application to e-cash, based on existing work such as [Zha+18].

Applications to network incentives

Some peer-to-peer networking protocols (such as onion routing [DMS04]) do not have inherent incentives and rely on volunteers to provide infrastructure. In future work, we want to look at adding incentives in the form of Taler payments to a peer-to-peer networking platform such as GUNet.

Smart(er) Contracts and Auctions

Contract terms in Taler are relatively limited. There are some interesting secure multiparty computations, such as privacy-preserving auctions [Bra06] that could be offered by exchanges as a fixed smart contract. This would allow a full privacy-preserving auction platform, as current auction protocols only output the winner of a privacy-preserving auction but do not address the required anonymous payments.

Backup and Sync

Synchronization of wallets between multiple devices is a useful feature, but a naïve implementation endangers privacy. A carefully designed protocol for backup and synchronization must make sure that the hosting service for the wallet's data cannot collaborate with the exchange and merchants to deanonymize users or transactions. Thus when spending coins for a payment, devices should not have to synchronously talk to their backup/sync provider. This creates the challenge of allocating the total available balance to individual devices in a way that fits the customer's spending pattern, and only require synchronous communication at fixed intervals or when really necessary to re-allocate coins.

Another possible approach might be to use Private Information Retrieval (PIR) [Golo7] to access backup and synchronization information.

Machine-Verified Proofs

We currently model only a subset of the GNU Taler protocol formally, and proofs are handwritten and verified by humans. A tool such as CryptoVerif [Blao7] can allow a higher coverage and computer-checked proofs, and would allow protocol changes to be validated in shorter time.

Coin Restrictions / “Taler for Children”

By designating certain denominations for different purposes, GNU Taler could be used to implement a very simple form of anonymous credentials [PZ11; CLo4], which then could be used to implement a Taler wallet specifically aimed at children, in order to teach them responsible and autonomous spending behavior, while granting them privacy and at the same time preventing them from making age-inappropriate purchases online, as the discretion of parents.

7. Conclusion

This thesis presented efficient protocols for both register- and value-based electronic payment systems with focus on security and privacy. While we believe our approach to be socially and economically beneficial, a technological impact analysis is in order prior to adopting new systems that have broad economic and socio-political implications.

Currencies serve three key functions in society: [Man10]

1. As a unit for measurement of value,
2. a medium of exchange, and
3. a store of value.

How do the various methods measure up to these requirements?

7.1. Cryptocurrencies vs. Central-Bank-Issued Currencies

Cryptocurrencies generally fail to achieve the required stability to serve as a reasonable unit of measurement (Figure 7.1). The volatility of cryptocurrencies is caused by a combination of a lack of institutions that could intervene to dampen fluctuations and a comparatively limited liquidity in the respective markets. The latter is exacerbated by the limited ability of decentralized cryptocurrencies to handle large transaction volumes, despite their extreme levels of resource consumption. As a result, the utility of decentralized cryptocurrencies is limited to highly speculative investments and to the facilitation of criminal transactions.

With respect to privacy, completely decentralized cryptocurrencies provide either too much or too little anonymity. Transparent cryptocurrencies create the spectre of discriminatory pricing, while especially for privacy-enhanced cryptocurrencies the lack of regulation creates an attractive environment for fraud and criminal activity from tax evasion to financing of terrorism.

These problems are easily addressed by combining the register (or ledger) with a central bank providing a regulatory framework and monetary policy, including anti-money-laundering and know-your-customer enforcement. Such central-bank-issued currencies

may be able to improve the availability, integrity and performance of their register using our Byzantine Set Consensus protocol in lieu of simplistic provisioned blockchains.

7. Conclusion




Figure 7.1.: Historical market price (in USD) of Bitcoin across major exchanges
(Source: <https://blockchain.com>).

7.2. Electronic Payments

Day-to-day payments using registers are expensive and inconvenient. Using a register requires users to *identify* themselves to *authorize* transactions, and the use of register-based banking systems tends to be more expensive than the direct exchange of physical cash. However, with the ongoing digitalization of daily life where a significant number of transactions is realized over networks, some form of electronic payments remain inevitable.

The current alternative to (centrally banked) electronic cash are a payment systems under full control of oligopoly companies such as Google, Apple, Facebook or Visa. The resulting oligopolies are anti-competitive. In addition to excessive fees, they sometimes even refuse to process payments with certain types of legal businesses, which then are often ruined due to lack of alternatives. Combining payment services with companies where the core business model is advertising is also particularly damaging for privacy. Finally, the sheer size of these companies creates systemic risks, just as their global scale creates challenges for regulation.

As GNU Taler is free software, even without backing by a central bank, Taler would not suffer from these drawbacks arising from the use of proprietary technology.

Furthermore, Taler-style electronic cash comes with some unique benefits:

- improved income transparency compared to cash and traditional Chaum-style e-cash,
- anonymity for payers,
- avoidance of enticement towards consumer debt — especially compared to credit cards, and
- support of new business models and Internet security mechanisms which require (anonymous) micro-transactions.

Central banks are carefully considering what might be the right technology to implement an electronic version of their centrally banked currency, and with Taler we hope to address most of their concerns. Nevertheless, all electronic payment systems, including Taler even when backed by central-bank-issued currencies, come with their own inherent set of risks: [Rik17]

- increased risk of a bank run: in a banking crisis, as it is easier to withdraw large amounts of digital cash quickly — even from remote locations;
- increased volatility due to foreign holdings that would not be as easily possible with physical cash;
- increased risk of theft and disruption: while physical cash can also be stolen (and likely with much less effort), it is difficult to transport in volume [FEF15], the risk is increased with computers because attacks scale [Ham18], and generally many small incidents are socially preferable over a tiny number of very large-scale incidents; and
- unavailability in crisis situations without electricity and Internet connectivity.

We believe that in the case of Taler, some of the risks mentioned above can be mitigated:

- Volatility due to foreign holdings and the resulting increased risk of bank runs can be reduced by putting limits on the amount of electronic coins that customers can withdraw. Limiting the validity periods of coins is another method that can help disincentivize the use of Taler as a value store.
- The use of open standards and reference implementations enables white-hat security research around GNU Taler, which together with good operational security procedures and the possibility of competing providers should reduce the risks from attacks.
- GNU Taler can co-exist with physical cash, and might even help revive the use of cash if it succeeds in reducing credit card use online thereby eliminating a key reason for people to have credit cards.

Unlike cryptocurrencies, Taler does not prescribe a solution for monetary policy or just taxation, as we believe these issues need to be subject to continuous political debate and cannot be “solved” by simplistic algorithms. What we offer to society is an open and free (as in free speech) system with mechanisms to audit merchants’ income, instead of proprietary systems controlled by a few oligopoly companies.

Bibliography

- [Abd+05] Michael Abd-El-Malek, Gregory R Ganger, Garth R Goodson, Michael K Reiter, and Jay J Wylie. “Fault-scalable Byzantine fault-tolerant services”. In: *ACM SIGOPS Operating Systems Review* 39.5 (2005), pp. 59–74 (cit. on p. 128).
- [ADG84] Chagit Attiya, Danny Dolev, and Joseph Gil. “Asynchronous Byzantine Consensus”. In: *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing*. PODC ’84. Vancouver, British Columbia, Canada: ACM, 1984, pp. 119–133. URL: <http://doi.acm.org/10.1145/800222.806740> (cit. on p. 149).
- [Adio8] Ben Adida. “Helios: Web-based Open-audit Voting”. In: *Proceedings of the 17th Conference on Security Symposium*. SS’08. San Jose, CA: USENIX Association, 2008, pp. 335–348. URL: <http://dl.acm.org/citation.cfm?id=1496711.1496734> (cit. on p. 151).
- [Ady16] Adyen. *The Global E-Commerce Payments Guide*. 2016 (cit. on p. 2).
- [Agu10] Marcos K. Aguilera. “Replication”. In: ed. by Bernadette Charron-Bost, Fernando Pedone, and André Schiper. Berlin, Heidelberg: Springer-Verlag, 2010. Chap. Stumbling over Consensus Research: Misunderstandings and Issues, pp. 59–72. URL: <http://dl.acm.org/citation.cfm?id=2172338.2172342> (cit. on p. 127).
- [And+18] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, et al. “Hyperledger fabric: a distributed operating system for permissioned blockchains”. In: *Proceedings of the Thirteenth EuroSys Conference*. ACM. 2018, p. 30 (cit. on p. 34).
- [Ano99] Anonymous. *How DigiCash Blew Everything*. 1999 (cit. on p. 31).
- [AO00] Masayuki Abe and Tatsuoaki Okamoto. “Provably secure partially blind signatures”. In: *Annual International Cryptology Conference*. Springer. 2000, pp. 271–286 (cit. on pp. 31, 59).
- [Arn+18] Douglas W Arner, Dirk A Zetsche, Ross P Buckley, and Janos Nathan Barberis. “The Identity Challenge in Finance: From Analogue Identity to Digitized Identification to Digital KYC Utilities”. In: *European Banking Institute* (2018) (cit. on p. 80).

- [ASM11] Man Ho Au, Willy Susilo, and Yi Mu. “Electronic cash with anonymous user suspension”. In: *Australasian Conference on Information Security and Privacy*. Springer. 2011, pp. 172–188 (cit. on pp. 6, 30).
- [Asp98] James Aspnes. “Lower bounds for distributed coin-flipping and randomized consensus”. In: *Journal of the ACM (JACM)* 45.3 (1998), pp. 415–450 (cit. on p. 127).
- [Aub+15] Pierre-Louis Aublin, Rachid Guerraoui, Nikola Knežević, Vivien Quéma, and Marko Vukolić. “The Next 700 BFT Protocols”. In: *ACM Trans. Comput. Syst.* 32.4 (Jan. 2015), 12:1–12:45. URL: <http://doi.acm.org/10.1145/2658994> (cit. on p. 129).
- [Bad15] Heinz-Peter Bader. *France steps up monitoring of cash payments to fight low-cost terrorism*. <http://www.reuters.com/article/2015/03/18/us-france-security-financing-idUSKBN0ME14720150318>. Mar. 2015 (cit. on p. 18).
- [Bar11] Adam Barth. *The Web Origin Concept*. RFC 6454. Dec. 2011. URL: <https://rfc-editor.org/rfc/rfc6454.txt> (cit. on p. 85).
- [BDH10] Michael Ben-Or, Danny Dolev, and Ezra N Hoch. “Simple gradecast based algorithms”. In: *arXiv preprint arXiv:1007.1049* (2010) (cit. on pp. 125, 126, 130, 139).
- [Bel+03] Bellare, Namprempre, Pointcheval, and Semanko. “The One-More-RSA-Inversion Problems and the Security of Chaum’s Blind Signature Scheme”. In: *Journal of Cryptology* 16.3 (June 2003), pp. 185–215. URL: <https://doi.org/10.1007/s00145-002-0120-1> (cit. on pp. 64, 76).
- [Bel+98] Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. “Relations among notions of security for public-key encryption schemes”. In: *Annual International Cryptology Conference*. Springer. 1998, pp. 26–45 (cit. on p. 45).
- [Ben+14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. “Zerocash: Decentralized Anonymous Payments from Bitcoin”. In: *IEEE Symposium on Security & Privacy*. 2014 (cit. on pp. 13, 15, 35).
- [Ben+18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. “Scalable, transparent, and post-quantum secure computational integrity”. In: *Cryptol. ePrint Arch., Tech. Rep* 46 (2018), p. 2018 (cit. on p. 35).
- [Ben87] Josh Daniel Cohen Benaloh. *Verifiable secret-ballot elections*. Yale University. Department of Computer Science, 1987 (cit. on p. 151).
- [Ber+12] Daniel J Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. “High-speed high-security signatures”. In: *Journal of Cryptographic Engineering* 2.2 (2012), pp. 77–89 (cit. on p. 64).

- [Bero6] Daniel J Bernstein. “Curve25519: new Diffie-Hellman speed records”. In: *International Workshop on Public Key Cryptography*. Springer. 2006, pp. 207–228 (cit. on pp. 64, 76).
- [BGK95] Ernest F Brickell, Peter Gemmell, and David W Kravitz. “Trustee-based Tracing Extensions to Anonymous Cash and the Making of Anonymous Change.” In: *SODA*. Vol. 95. 1995, pp. 457–466 (cit. on pp. 6, 30).
- [BGM16] Iddo Bentov, Ariel Gabizon, and Alex Mizrahi. “Cryptocurrencies without proof of work”. In: *International Conference on Financial Cryptography and Data Security*. Springer. 2016, pp. 142–157 (cit. on p. 34).
- [BH13] Carsten Bormann and Paul E. Hoffman. *Concise Binary Object Representation (CBOR)*. RFC 7049. Oct. 2013. URL: <https://rfc-editor.org/rfc/rfc7049.txt> (cit. on p. 76).
- [Bla07] Bruno Blanchet. “CryptoVerif: Computationally sound mechanized prover for cryptographic protocols”. In: *Dagstuhl seminar “Formal Protocol Verification Applied*. Vol. 117. 2007 (cit. on p. 156).
- [Blo70] Burton H Bloom. “Space/time trade-offs in hash coding with allowable errors”. In: *Communications of the ACM* 13.7 (1970), pp. 422–426 (cit. on p. 131).
- [Bog+09] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael Schwartzbach, and Tomas Toft. “Financial Cryptography and Data Security”. In: ed. by Roger Dingledine and Philippe Golle. Berlin, Heidelberg: Springer-Verlag, 2009. Chap. Secure Multiparty Computation Goes Live, pp. 325–343. URL: http://dx.doi.org/10.1007/978-3-642-03549-4_20 (cit. on p. 125).
- [Böh+15] Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. “Bitcoin: Economics, technology, and governance”. In: *Journal of Economic Perspectives* 29.2 (2015), pp. 213–38 (cit. on p. 13).
- [Bolo3] Alexandra Boldyreva. “Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme”. In: *International Workshop on Public Key Cryptography*. Springer. 2003, pp. 31–46 (cit. on p. 77).
- [Bon+14] Joseph Bonneau, Arvind Narayanan, Andrew Miller, Jeremy Clark, Joshua A Kroll, and Edward W Felten. “Mixcoin: Anonymity for Bitcoin with accountable mixes”. In: *International Conference on Financial Cryptography and Data Security*. Springer. 2014, pp. 486–504 (cit. on p. 35).

- [Bon98] Dan Boneh. “The decision diffie-hellman problem”. In: *International Algorithmic Number Theory Symposium*. Springer. 1998, pp. 48–63 (cit. on p. 45).
- [BPo6] Bruno Blanchet and David Pointcheval. “Automated security proofs with sequences of games”. In: *Annual International Cryptology Conference*. Springer. 2006, pp. 537–554 (cit. on p. 46).
- [BRo6] Mihir Bellare and Phillip Rogaway. “Code-based game-playing proofs and the security of triple encryption”. In: *Advances in Cryptology—EUROCRYPT*. Vol. 4004. 2006, p. 10 (cit. on pp. 42, 46, 54).
- [BR93] Mihir Bellare and Phillip Rogaway. “Random oracles are practical: A paradigm for designing efficient protocols”. In: *Proceedings of the 1st ACM conference on Computer and communications security*. ACM. 1993, pp. 62–73 (cit. on pp. 41, 42).
- [BR96] Mihir Bellare and Phillip Rogaway. “The exact security of digital signatures—How to sign with RSA and Rabin”. In: *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 1996, pp. 399–416 (cit. on p. 64).
- [Brao6] Felix Brandt. “How to obtain full privacy in auctions”. In: *International Journal of Information Security* 5.4 (2006), pp. 201–216 (cit. on p. 155).
- [Bra17] Tim Bray. “The JavaScript Object Notation (JSON) Data Interchange Format”. In: *RFC 8259* (2017), pp. 1–16. URL: <https://doi.org/10.17487/RFC8259> (cit. on p. 76).
- [BS15] Alireza Beikverdi and JooSeok Song. “Trend of centralization in Bitcoin’s distributed network”. In: *Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD), 2015 16th IEEE/ACIS International Conference on*. IEEE. 2015, pp. 1–6 (cit. on p. 13).
- [BY18] Lynn Batten and Xun Yi. “Off-line digital cash schemes providing untraceability, anonymity and change”. In: *Electronic Commerce Research* (Jan. 2018). URL: <https://doi.org/10.1007/s10660-018-9289-8> (cit. on p. 30).
- [CFN90] David Chaum, Amos Fiat, and Moni Naor. “Untraceable Electronic Cash”. In: *Advances in Cryptology — CRYPTO’ 88: Proceedings*. Ed. by Shafi Goldwasser. New York, NY: Springer New York, 1990, pp. 319–327. URL: https://doi.org/10.1007/0-387-34799-2_25 (cit. on pp. 29, 32).
- [CGo7] Sébastien Canard and Aline Gouget. “Divisible e-cash systems can be truly anonymous”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2007, pp. 482–497 (cit. on p. 29).

- [CGHo4] Ran Canetti, Oded Goldreich, and Shai Halevi. “The random oracle methodology, revisited”. In: *Journal of the ACM (JACM)* 51.4 (2004), pp. 557–594 (cit. on pp. 42, 155).
- [CGHo6] Sébastien Canard, Aline Gouget, and Emeline Hufschmitt. “A handy multi-coupon system”. In: *International Conference on Applied Cryptography and Network Security*. Springer. 2006, pp. 66–81 (cit. on p. 30).
- [CGS97] Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. “A secure and optimally efficient multi-authority election scheme”. In: *European transactions on Telecommunications* 8.5 (1997), pp. 481–490 (cit. on pp. 125, 126, 151, 153).
- [Cha+89] David Chaum, Bert den Boer, Eugène van Heyst, Stig Mjølsnes, and Adri Steenbeek. “Efficient offline electronic checks”. In: *Workshop on the theory and application of cryptographic techniques*. Springer. 1989, pp. 294–301 (cit. on p. 29).
- [Cha83] David Chaum. “Blind Signatures for Untraceable Payments”. In: *Advances in Cryptology: Proceedings of Crypto 82*. Ed. by David Chaum, Ronald L. Rivest, and Alan T. Sherman. Boston, MA: Springer US, 1983, pp. 199–203. URL: https://doi.org/10.1007/978-1-4757-0602-4_18 (cit. on pp. 4, 8, 17, 29, 32, 64).
- [Che+14] Yuchung Cheng, Jerry Chu, Sivasankar Radhakrishnan, and Arvind Jain. *TCP Fast Open*. RFC 7413. Dec. 2014. URL: <https://rfc-editor.org/rfc/rfc7413.txt> (cit. on p. 121).
- [CHLo5] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. “Compact E-Cash”. In: *Advances in Cryptology – EUROCRYPT 2005: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005. Proceedings*. Ed. by Ronald Cramer. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 302–321. URL: https://doi.org/10.1007/11426639_18 (cit. on pp. 29, 32, 59).
- [CKSo5] Christian Cachin, Klaus Kursawe, and Victor Shoup. “Random oracles in Constantinople: Practical asynchronous Byzantine agreement using cryptography”. In: *Journal of Cryptology* 18.3 (2005), pp. 219–246 (cit. on p. 127).
- [CLo2] Miguel Castro and Barbara Liskov. “Practical Byzantine fault tolerance and proactive recovery”. In: *ACM Transactions on Computer Systems (TOCS)* 20.4 (2002), pp. 398–461 (cit. on pp. 126, 128).
- [CLo4] Jan Camenisch and Anna Lysyanskaya. “Signature schemes and anonymous credentials from bilinear maps”. In: *Annual International Cryptology Conference*. Springer. 2004, pp. 56–72 (cit. on p. 156).

- [CL99] Miguel Castro and Barbara Liskov. “Practical Byzantine Fault Tolerance”. In: *Third Symposium on Operating Systems Design and Implementation (OSDI)*. Vol. 99. New Orleans, Louisiana: USENIX Association, Co-sponsored by IEEE TCOS and ACM SIGOPS, Feb. 1999, pp. 173–186 (cit. on pp. 125, 126, 128, 149).
- [Cle+09] Allen Clement, Edmund Wong, Lorenzo Alvisi, Mike Dahlin, and Mirco Marchetti. “Making Byzantine Fault Tolerant Systems Tolerate Byzantine Faults”. In: *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*. NSDI’09. Boston, Massachusetts: USENIX Association, 2009, pp. 153–168. URL: <http://dl.acm.org/citation.cfm?id=1558977.1558988> (cit. on pp. 125, 128, 129).
- [CLM07] J. Camenisch, A. Lysyanskaya, and M. Meyerovich. “Endorsed E-Cash”. In: *2007 IEEE Symposium on Security and Privacy (SP ’07)*. May 2007, pp. 101–115 (cit. on pp. 7, 46, 70).
- [Cor00] Jean-Sébastien Coron. “On the exact security of full domain hash”. In: *Annual International Cryptology Conference*. Springer. 2000, pp. 229–235 (cit. on pp. 41, 45).
- [CP92] David Chaum and Torben Pryds Pedersen. “Wallet databases with observers”. In: *Annual International Cryptology Conference*. Springer. 1992, pp. 89–105 (cit. on p. 30).
- [Cro] Douglas Crockford. *Base32 Encoding*. URL: <https://www.crockford.com/wrmg/base32.html> (cit. on p. 76).
- [CSS11] Bert Bos, ed. *Cascading Style Sheets Level 2 Revision 1 (CSS 2.1) Specification*. 2011 (cit. on p. 101).
- [Dal16] Therése Dalebrant. “The Monetary Policy Effects of Sweden’s Transition Towards a Cashless Society: An Econometric Analysis”. In: (2016) (cit. on p. 2).
- [Dam07] Ivan Damgård. “A “proof-reading” of some issues in cryptography”. In: *International Colloquium on Automata, Languages, and Programming*. Springer. 2007, pp. 2–11 (cit. on pp. 30, 41).
- [Dav+97] George Davida, Yair Frankel, Yiannis Tsiounis, and Moti Yung. “Anonymity control in e-cash systems”. In: *International Conference on Financial Cryptography*. Springer. 1997, pp. 1–16 (cit. on p. 30).
- [DDS87] Danny Dolev, Cynthia Dwork, and Larry Stockmeyer. “On the minimal synchronism needed for distributed consensus”. In: *Journal of the ACM (JACM)* 34.1 (1987), pp. 77–97 (cit. on p. 128).
- [Des94] Yvo G Desmedt. “Threshold cryptography”. In: *European Transactions on Telecommunications* 5.4 (1994), pp. 449–458 (cit. on p. 152).

- [DG17] Florian Dold and Christian Grothoff. “Byzantine set-union consensus using efficient set reconciliation”. In: *EURASIP Journal on Information Security* 2017.1 (July 2017), p. 14. URL: <https://doi.org/10.1186/s13635-017-0066-3> (cit. on p. 125).
- [DKL15] Jannik Dreier, Ali Kassem, and Pascal Lafourcade. “Formal analysis of e-cash protocols”. In: *e-Business and Telecommunications (ICETE), 2015 12th International Joint Conference on*. Vol. 4. IEEE. 2015, pp. 65–75 (cit. on p. 30).
- [DLS88] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. “Consensus in the presence of partial synchrony”. In: *Journal of the ACM (JACM)* 35.2 (1988), pp. 288–323 (cit. on pp. 126, 128, 150, 151).
- [DM16] George Danezis and Sarah Meiklejohn. “Centrally Banked Cryptocurrencies”. In: *23rd Annual Network and Distributed System Security Symposium, NDSS 2016, San Diego, California, USA, February 21-24, 2016*. The Internet Society, 2016 (cit. on p. 124).
- [DMR01] Roberto De Prisco, Dahlia Malkhi, and Michael Reiter. “On k-set consensus problems in asynchronous systems”. In: *Parallel and Distributed Systems, IEEE Transactions on* 12.1 (2001), pp. 7–21 (cit. on p. 127).
- [DMS04] Roger Dingledine, Nick Mathewson, and Paul Syverson. “Tor: The Second-Generation Onion Router”. In: *Proceedings of the 13th USENIX Security Symposium*. Aug. 2004 (cit. on pp. 18, 77, 155).
- [Dol14] Florian Dold. “Cryptographically Secure, Distributed Electronic Voting”. Bachelor’s Thesis. Technische Universität München, 2014 (cit. on p. 153).
- [DPWo8] AW Dent, KG Paterson, and PR Wild. “Extensions to Chaum’s Blind Signature Scheme and OpenCoin Requirements”. In: (2008) (cit. on p. 31).
- [Epp+11] David Eppstein, Michael T. Goodrich, Frank Uyeda, and George Varghese. “What’s the Difference?: Efficient Set Reconciliation Without Prior Context”. In: *SIGCOMM Comput. Commun. Rev.* 41.4 (Aug. 2011), pp. 218–229. URL: <http://doi.acm.org/10.1145/2043164.2018462> (cit. on pp. 126, 131, 133, 140).
- [ES18] Ittay Eyal and Emin Gün Sirer. “Majority is not enough: Bitcoin mining is vulnerable”. In: *Communications of the ACM* 61.7 (2018), pp. 95–102 (cit. on pp. 14, 34).
- [Eya+16] Ittay Eyal, Adem Efe Gencer, Emin Gün Sirer, and Robbert Van Renesse. “Bitcoin-NG: A Scalable Blockchain Protocol.” In: *NSDI*. 2016, pp. 45–59 (cit. on p. 34).

- [FEF15] Financial Action Task Force, Middle East, and North Africa Financial Action Task Force. *Money laundering through the physical transportation of cash*. 2015 (cit. on p. 159).
- [Fel88] Paul Neil Feldman. “Optimal algorithms for Byzantine agreement”. PhD thesis. Massachusetts Institute of Technology, 1988 (cit. on pp. 127, 138).
- [FH06] Matthias Fitzi and Martin Hirt. “Optimally Efficient Multi-valued Byzantine Agreement”. In: *Proceedings of the Twenty-fifth Annual ACM Symposium on Principles of Distributed Computing*. PODC ’06. Denver, Colorado, USA: ACM, 2006, pp. 163–168. URL: <http://doi.acm.org/10.1145/1146381.1146407> (cit. on p. 125).
- [FKH13] Ned Freed, Dr. John C. Klensin, and Tony Hansen. *Media Type Specifications and Registration Procedures*. RFC 6838. Jan. 2013. URL: <https://rfc-editor.org/rfc/rfc6838.txt> (cit. on p. 91).
- [FL81] Michael J Fischer and Nancy A Lynch. *A lower bound for the time to assure interactive consistency*. Tech. rep. DTIC Document, 1981 (cit. on p. 127).
- [FLM86] Michael J Fischer, Nancy A Lynch, and Michael Merritt. “Easy impossibility proofs for distributed consensus problems”. In: *Distributed Computing* 1.1 (1986), pp. 26–39 (cit. on p. 127).
- [FLP85] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. “Impossibility of distributed consensus with one faulty process”. In: *Journal of the ACM (JACM)* 32.2 (1985), pp. 374–382 (cit. on p. 127).
- [FM88] Paul Feldman and Silvio Micali. “Optimal Algorithms for Byzantine Agreement”. In: *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*. STOC ’88. Chicago, Illinois, USA: ACM, 1988, pp. 148–161. URL: <http://doi.acm.org/10.1145/62212.62225> (cit. on pp. 127, 129, 130).
- [FPV09] Georg Fuchsbauer, David Pointcheval, and Damien Vergnaud. “Transferable constant-size fair e-cash”. In: *International Conference on Cryptology and Network Security*. Springer. 2009, pp. 226–247 (cit. on pp. 6, 30).
- [FS01] Pierre-Alain Fouque and Jacques Stern. “One Round Threshold Discrete-Log Key Generation without Private Channels”. In: *Public Key Cryptography: 4th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2001 Cheju Island, Korea, February 13–15, 2001 Proceedings*. Ed. by Kwangjo Kim. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 300–316. URL: https://doi.org/10.1007/3-540-44586-2_22 (cit. on p. 152).

- [FS09] Marc Fischlin and Dominique Schröder. “Security of blind signatures under aborts”. In: *International Workshop on Public Key Cryptography*. Springer. 2009, pp. 297–316 (cit. on p. 59).
- [FT00] Roy T Fielding and Richard N Taylor. *Architectural styles and the design of network-based software architectures*. Vol. 7. University of California, Irvine Doctoral dissertation, 2000 (cit. on p. 76).
- [Gar+07] Sujata Garera, Niels Provos, Monica Chew, and Aviel D Rubin. “A framework for detection and measurement of phishing attacks”. In: *Proceedings of the 2007 ACM workshop on Recurring malware*. ACM. 2007, pp. 1–8 (cit. on p. 3).
- [GGM16] Christina Garman, Matthew Green, and Ian Miers. “Accountable privacy for decentralized anonymous payments”. In: *International Conference on Financial Cryptography and Data Security*. Springer. 2016, pp. 81–98 (cit. on p. 35).
- [Gil+17] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. “Algorand: Scaling byzantine agreements for cryptocurrencies”. In: *Proceedings of the 26th Symposium on Operating Systems Principles*. ACM. 2017, pp. 51–68 (cit. on p. 34).
- [GKL15] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. “The Bitcoin Backbone Protocol: Analysis and Applications”. In: *Advances in Cryptology - EUROCRYPT 2015: 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26–30, 2015, Proceedings, Part II*. Ed. by Elisabeth Oswald and Marc Fischlin. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 281–310. URL: https://doi.org/10.1007/978-3-662-46803-6_10 (cit. on p. 154).
- [GL05] Shafi Goldwasser and Yehuda Lindell. “Secure multi-party computation without agreement”. In: *Journal of Cryptology* 18.3 (2005), pp. 247–287 (cit. on pp. 150, 151).
- [GM16] Matthew Green and Ian Miers. *Bolt: Anonymous Payment Channels for Decentralized Currencies*. Cryptology ePrint Archive, Report 2016/701. <http://eprint.iacr.org/2016/701>. 2016 (cit. on pp. 13, 36).
- [GM82] Shafi Goldwasser and Silvio Micali. “Probabilistic encryption & how to play mental poker keeping secret all partial information”. In: *Proceedings of the fourteenth annual ACM symposium on Theory of computing*. ACM. 1982, pp. 365–377 (cit. on p. 41).
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L Rivest. “A digital signature scheme secure against adaptive chosen-message attacks”. In: *SIAM Journal on Computing* 17.2 (1988), pp. 281–308 (cit. on p. 45).

- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. “The knowledge complexity of interactive proof systems”. In: *SIAM Journal on computing* 18.1 (1989), pp. 186–208 (cit. on p. 42).
- [GNUNET] *The GNUnet Project*. <https://gnunet.org/>. Accessed 28 Feb 2017 (cit. on p. 139).
- [Golo7] Ian Goldberg. “Improving the robustness of private information retrieval”. In: *Security and Privacy, 2007. SP’07. IEEE Symposium on*. IEEE. 2007, pp. 131–148 (cit. on p. 156).
- [GSM18] Fuchun Guo, Willy Susilo, and Yi Mu. *Introduction to Security Reduction*. Springer, 2018 (cit. on p. 42).
- [Gue+00] Rachid Guerraoui, Michel Hurfinn, Achour Mostefaoui, Riucarlos Oliveira, Michel Raynal, and Andre Schiper. “Consensus in Asynchronous Distributed Systems: A Concise Guided Tour”. In: *Advances in Distributed Systems: Advanced Distributed Computing: From Algorithms to Systems*. Ed. by Sacha Krakowiak and Santosh Shrivastava. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 33–47. URL: https://doi.org/10.1007/3-540-46475-1_2 (cit. on p. 128).
- [H306] Tony Hansen and Donald E. Eastlake 3rd. *US Secure Hash Algorithms (SHA and HMAC-SHA)*. RFC 4634. Aug. 2006. URL: <https://rfc-editor.org/rfc/rfc4634.txt> (cit. on pp. 64, 76).
- [Ham18] Joshua Hammer. “The Billion-Dollar Bank Job”. In: *The New York Times Magazine* (2018) (cit. on p. 159).
- [Hei+17] Ethan Heilman, Leen Alshenibr, Foteini Baldimtsi, Alessandra Scafuro, and Sharon Goldberg. “TumbleBit: An untrusted Bitcoin-compatible anonymous payment hub”. In: *Network and Distributed System Security Symposium*. 2017 (cit. on p. 35).
- [HS90] Stuart Haber and W Scott Stornetta. “How to time-stamp a digital document”. In: *Conference on the Theory and Application of Cryptography*. Springer. 1990, pp. 437–455 (cit. on p. 33).
- [HTI97] Mei-Chen Hsueh, Timothy K Tsai, and Ravishankar K Iyer. “Fault injection techniques and tools”. In: *Computer* 30.4 (1997), pp. 75–82 (cit. on p. 42).
- [ibi14] ibi research. *Digitalisierung der Gesellschaft 2014 — Aktuelle Einschätzungen und Trends*. <http://www.ecommerce-leitfaden.de/digitalisierung-der-gesellschaft-2014.html>. 2014 (cit. on p. 39).

- [IL13] Malika Izabachène and Benoît Libert. “Divisible E-Cash in the Standard Model”. In: *Pairing-Based Cryptography – Pairing 2012: 5th International Conference, Cologne, Germany, May 16-18, 2012, Revised Selected Papers*. Ed. by Michel Abdalla and Tanja Lange. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 314–332. URL: https://doi.org/10.1007/978-3-642-36334-4_20 (cit. on p. 30).
- [Jaw+18] Husam Al Jawaheri, Mashaël Al Sabah, Yazan Boshmaf, and Aimen Erbad. “When A Small Leak Sinks A Great Ship: Deanonymizing Tor Hidden Service Users Through Bitcoin Transactions Analysis”. In: *arXiv preprint arXiv:1801.07501* (2018) (cit. on p. 15).
- [Joh+13] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. “Users get routed: Traffic correlation on Tor by realistic adversaries”. In: *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 337–348 (cit. on p. 18).
- [Jon15] Rupert Jones. *Cap on card fees could lead to lower prices for consumers*. <http://www.theguardian.com/money/2015/jul/27/cap-on-card-fees-retailers>. July 2015 (cit. on p. 37).
- [KB16] Jae Kwon and Ethan Buchman. *Cosmos: A Network of Distributed Ledgers*. <https://cosmos.network/whitepaper>. Accessed 22 Feb 2017. 2016 (cit. on p. 154).
- [KE10] Dr. Hugo Krawczyk and Pasi Eronen. *HMAC-based Extract-and-Expand Key Derivation Function (HKDF)*. RFC 5869. May 2010. URL: <https://rfc-editor.org/rfc/rfc5869.txt> (cit. on pp. 64, 76, 104).
- [KM07] Neal Koblitz and Alfred J Menezes. “Another look at” provable security”. In: *Journal of Cryptology* 20.1 (2007), pp. 3–37 (cit. on p. 41).
- [KM10] Neal Koblitz and Alfred Menezes. “The brave new world of bodacious assumptions in cryptography”. In: *Notices of the American Mathematical Society* 57.3 (2010), pp. 357–365 (cit. on p. 41).
- [KM15] Neal Koblitz and Alfred J Menezes. “The random oracle model: a twenty-year retrospective”. In: *Designs, Codes and Cryptography* 77.2-3 (2015), pp. 587–610 (cit. on pp. 42, 155).
- [KMM98] Kim Potter Kihlstrom, L. E. Moser, and P. M. Melliar-Smith. “The SecureRing Protocols for Securing Group Communication”. In: *Proceedings of the Thirty-First Annual Hawaii International Conference on System Sciences - Volume 3*. HICSS ’98. Washington, DC, USA: IEEE Computer Society, 1998, pp. 317–. URL: <http://dx.doi.org/10.1109/HICSS.1998.656294> (cit. on p. 128).

- [Kot+07] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. “Zyzyva: Speculative Byzantine Fault Tolerance”. In: *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating Systems Principles*. SOSP ’07. Stevenson, Washington, USA: ACM, 2007, pp. 45–58. URL: <http://doi.acm.org/10.1145/1294261.1294267> (cit. on pp. 125, 128).
- [Kwo14] Jae Kwon. *Tendermint: Consensus without mining*. Draft v. 0.6, fall. 2014 (cit. on p. 34).
- [Lev17] Karen EC Levy. “Book-smart, not street-smart: blockchain-based smart contracts and the social workings of law”. In: *Engaging Science, Technology, and Society* 3 (2017), pp. 1–15 (cit. on p. 35).
- [LI16] Jason Luu and Edward J Imwinkelried. “The challenge of Bitcoin pseudo-anonymity to computer forensics”. In: *Criminal Law Bulletin* 52.1 (2016) (cit. on p. 15).
- [Lin17] Yehuda Lindell. “How to simulate it—a tutorial on the simulation proof technique”. In: *Tutorials on the Foundations of Cryptography*. Springer, 2017, pp. 277–346 (cit. on p. 42).
- [LMS16] Sebastian Lupu, Melisande Mual, and Mees van Stiphout. “Ecommerce Payment Methods Report 2016”. In: (2016) (cit. on p. 2).
- [Lom+11] Victor Lomne, A Dehaboui, Philippe Maurine, L Torres, and M Robert. “Side channel attacks”. In: *Security trends for FPGAs*. Springer, 2011, pp. 47–72 (cit. on p. 42).
- [LSP82] Leslie Lamport, Robert Shostak, and Marshall Pease. “The Byzantine generals problem”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 4.3 (1982), pp. 382–401 (cit. on p. 126).
- [MA10] Steven J. Murdoch and Ross Anderson. “Verified by Visa and Mastercard Securecode: Or, How Not to Design Authentication”. In: *Proceedings of the 14th International Conference on Financial Cryptography and Data Security*. FC’10. Tenerife, Spain: Springer-Verlag, 2010, pp. 336–342. URL: <https://www.cl.cam.ac.uk/~rja14/Papers/fc10vbwsecurecode.pdf> (cit. on p. 38).
- [MA14] Stephen Murdoch and Ross Anderson. “Security Protocols and Evidence: Where Many Payment Systems Fail”. In: *Financial Cryptography and Data Security*. 2014 (cit. on p. 28).
- [Man10] N.G. Mankiw. *Macroeconomics, 7th Edition*. Worth Publishers, 2010 (cit. on p. 157).
- [Mär15] Patrick Märten. *Practical Compact E-Cash with Arbitrary Wallet Size*. Cryptology ePrint Archive, Report 2015/086. <http://eprint.iacr.org/2015/086>. 2015 (cit. on p. 31).

- [Mei+13] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. “A fistful of bitcoins: characterizing payments among men with no names”. In: *Proceedings of the 2013 conference on Internet measurement conference*. ACM. 2013, pp. 127–140 (cit. on p. 15).
- [Mil+16] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. “The Honey Badger of BFT Protocols”. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’16. Vienna, Austria: ACM, 2016, pp. 31–42. URL: <http://doi.acm.org/10.1145/2976749.2978399> (cit. on p. 128).
- [ML14] Andrew Miller and Joseph J LaViola Jr. *Anonymous byzantine consensus from moderately-hard puzzles: A model for bitcoin*. Tech. rep. CS-TR-14-01. University of Central Florida, Apr. 2014 (cit. on p. 154).
- [MMR14] Achour Mostefaoui, Hamouma Moumen, and Michel Raynal. “Signature-free Asynchronous Byzantine Consensus with $t < n/3$ and $O(n^2)$ Messages”. In: *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*. PODC ’14. Paris, France: ACM, 2014, pp. 2–9. URL: <http://doi.acm.org/10.1145/2611462.2611468> (cit. on p. 127).
- [MP13] Michael Mitzenmacher and Rasmus Pagh. “Simple Multi-Party Set Reconciliation”. In: *arXiv preprint arXiv:1311.2037* (2013) (cit. on p. 134).
- [MTZ03] Yaron Minsky, Ari Trachtenberg, and Richard Zippel. “Set reconciliation with nearly optimal communication complexity”. In: *Information Theory, IEEE Transactions on* 49.9 (2003), pp. 2213–2218 (cit. on p. 131).
- [Mul+13] Martin Mulazzani, Philipp Reschl, Markus Huber, Manuel Leithner, Sebastian Schrittwieser, Edgar Weippl, and FC Wien. “Fast and reliable browser identification with javascript engine fingerprinting”. In: *Web 2.0 Workshop on Security and Privacy (W2SP)*. Vol. 5. Citeseer. 2013 (cit. on p. 101).
- [MWV00] Navneet Malpani, Jennifer L. Welch, and Nitin Vaidya. “Leader Election Algorithms for Mobile Ad Hoc Networks”. In: *Proceedings of the 4th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*. DIALM ’00. Boston, Massachusetts, USA: ACM, 2000, pp. 96–103. URL: <http://doi.acm.org/10.1145/345848.345871> (cit. on p. 127).
- [Nako8] Satoshi Nakamoto. “Bitcoin: A peer-to-peer electronic cash system”. In: *Consulted 1.2012* (2008), p. 28 (cit. on pp. 33, 154).
- [Nei94] Gil Neiger. “Distributed consensus revisited”. In: *Information Processing Letters* 49.4 (1994), pp. 195–201 (cit. on p. 127).

- [Oka95] Tatsuaki Okamoto. “An Efficient Divisible Electronic Cash Scheme”. In: *Advances in Cryptology — CRYPTO’ 95: 15th Annual International Cryptology Conference Santa Barbara, California, USA, August 27–31, 1995 Proceedings*. Ed. by Don Coppersmith. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 438–451. URL: https://doi.org/10.1007/3-540-44750-4_35 (cit. on p. 29).
- [Oka98] Chris Okasaki. *Purely Functional Data Structures*. New York, NY, USA: Cambridge University Press, 1998 (cit. on p. 150).
- [PB17] Joseph Poon and Vitalik Buterin. *Plasma: Scalable autonomous smart contracts*. White paper. 2017 (cit. on p. 37).
- [PD16] Joseph Poon and Thaddeus Dryja. “The bitcoin lightning network: Scalable off-chain instant payments”. In: *draft version 0.5* (2016), p. 14 (cit. on p. 36).
- [Ped91] Torben Pryds Pedersen. “A threshold cryptosystem without a trusted party”. In: *Advances in Cryptology—EUROCRYPT’91*. Springer. 1991, pp. 522–526 (cit. on p. 152).
- [Ped96] Torben P Pedersen. “Electronic payments of small amounts”. In: *International Workshop on Security Protocols*. Springer. 1996, pp. 59–68 (cit. on p. 36).
- [Per17] Nathaniel Persily. “The 2016 US Election: Can democracy survive the internet?” In: *Journal of democracy* 28.2 (2017), pp. 63–76 (cit. on p. 2).
- [Pet05] RA Peters. “A Secure Bulletin Board”. Master’s Thesis. Technische Universiteit Eindhoven, 2005 (cit. on p. 151).
- [PG14] B. Polot and C. Grothoff. “CADET: Confidential ad-hoc decentralized end-to-end transport”. In: *2014 13th Annual Mediterranean Ad Hoc Networking Workshop (MED-HOC-NET)*. June 2014, pp. 71–78 (cit. on pp. 139, 142, 146).
- [Poi05] David Pointcheval. “Provable security for public key schemes”. In: *Contemporary cryptology*. Springer, 2005, pp. 133–190 (cit. on pp. 41, 42).
- [Pro+07] Niels Provos, Dean McNamee, Panayiotis Mavrommatis, Ke Wang, Nagendra Modadugu, et al. “The Ghost in the Browser: Analysis of Web-based Malware.” In: *HotBots* 7 (2007), pp. 4–4 (cit. on p. 2).
- [PS00] David Pointcheval and Jacques Stern. “Security arguments for digital signatures and blind signatures”. In: *Journal of cryptology* 13.3 (2000), pp. 361–396 (cit. on p. 30).

- [PS96] David Pointcheval and Jacques Stern. “Provably secure blind signature schemes”. In: *Advances in Cryptology — ASIACRYPT ’96: International Conference on the Theory and Applications of Cryptology and Information Security Kyongju, Korea, November 3–7, 1996 Proceedings*. Ed. by Kwangjo Kim and Tsutomu Matsumoto. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, pp. 252–265. URL: <https://doi.org/10.1007/BFb0034852> (cit. on p. 30).
- [PST17] David Pointcheval, Olivier Sanders, and Jacques Traoré. “Cut Down the Tree to Achieve Constant Complexity in Divisible E-cash”. In: *Public-Key Cryptography – PKC 2017: 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Amsterdam, The Netherlands, March 28–31, 2017, Proceedings, Part I*. Ed. by Serge Fehr. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, pp. 61–90. URL: https://doi.org/10.1007/978-3-662-54365-8_4 (cit. on pp. 5, 30, 32, 46, 55).
- [PZ11] Christian Paquin and Greg Zaverucha. “U-prove cryptographic specification v1. 1”. In: *Technical Report, Microsoft Corporation* (2011) (cit. on p. 156).
- [RDo8] Eric Rescorla and Tim Dierks. *The Transport Layer Security (TLS) Protocol Version 1.2*. RFC 5246. Aug. 2008. URL: <https://rfc-editor.org/rfc/rfc5246.txt> (cit. on p. 77).
- [Rei95] Michael K. Reiter. “The Rampart Toolkit for Building High-Integrity Services”. In: *Selected Papers from the International Workshop on Theory and Practice in Distributed Systems*. London, UK, UK: Springer-Verlag, 1995, pp. 99–110. URL: <http://dl.acm.org/citation.cfm?id=647369.723763> (cit. on p. 128).
- [Reu15] Reuters. *Greek council recommends 60 euro limit on ATM withdrawals from Tuesday*. <http://www.reuters.com/article/2015/06/28/eurozone-greece-limits-idUSA8N0Z302P20150628>. June 2015 (cit. on p. 18).
- [RH13] Fergal Reid and Martin Harrigan. “An analysis of anonymity in the bitcoin system”. In: *Security and privacy in social networks*. Springer, 2013, pp. 197–223 (cit. on p. 35).
- [Ric16] Jean-Loup Richet. “Extortion on the internet: the rise of crypto-ransomware”. In: *Harvard* (2016) (cit. on p. 6).
- [Rik17] Sveriges Riksbank. *The Riksbank’s e-krona project*. 2017 (cit. on pp. 1, 159).
- [Rivo4] Ronald L Rivest. “Peppercorn micropayments”. In: *Financial Cryptography*. Springer. 2004, pp. 2–8 (cit. on p. 35).

- [ROH16] Wessel Reijers, Fiachra O’Brolcháin, and Paul Haynes. “Governance in blockchain technologies & social contract theories”. In: *Ledger* 1 (2016), pp. 134–151 (cit. on p. 35).
- [Run11] Guy Rundle. *The humble credit card is now a political tool*. <http://www.crikey.com.au/2011/10/25/rundle-humble-credit-card-now-a-political-tool-just-ask-wikileaks/>. Oct. 2011 (cit. on p. 37).
- [Rup+13] Andy Rupp, Gesine Hinterwälder, Foteini Baldimtsi, and Christof Paar. “P4R: Privacy-preserving pre-payments with refunds for transportation systems”. In: *International Conference on Financial Cryptography and Data Security*. Springer. 2013, pp. 205–212 (cit. on p. 30).
- [SB17] Sunny Kumar Singh and Kaushik Bhattacharya. “Does easy availability of cash affect corruption? Evidence from a panel of countries”. In: *Economic Systems* 41.2 (2017), pp. 236–247 (cit. on p. 2).
- [Sch98] Berry Schoenmakers. “Security Aspects of the Ecash(TM) Payment System”. In: *State of the Art in Applied Cryptography, Course on Computer Security and Industrial Cryptography - Revised Lectures*. Leuven, Belgium: Springer-Verlag, 1998, pp. 338–352. URL: <http://dl.acm.org/citation.cfm?id=647443.726912> (cit. on pp. 31, 32).
- [SD10] Y Sahin and E Duman. “An overview of business domains where fraud can take place, and a survey of various fraud detection techniques”. In: *Proceedings of the 1st international symposium on computing in science and engineering, Aydin, Turkey*. 2010 (cit. on p. 3).
- [Sha79] Adi Shamir. “How to share a secret”. In: *Communications of the ACM* 22.11 (1979), pp. 612–613 (cit. on p. 152).
- [Sho04] Victor Shoup. “Sequences of games: a tool for taming complexity in security proofs.” In: *IACR Cryptology ePrint Archive* 2004 (2004), p. 332 (cit. on pp. 41, 42, 46).
- [SPC95] Markus Stadler, Jean-Marc Piveteau, and Jan Camenisch. “Fair blind signatures”. In: *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 1995, pp. 209–219 (cit. on pp. 6, 30).
- [ST99] Tomas Sander and Amnon Ta-Shma. “On Anonymous Electronic Cash and Crime”. In: *ISW’99*. LNCS 1729. 1999, pp. 202–206 (cit. on p. 6).
- [Stao2] Richard Stallman. *Free software, free society: Selected essays of Richard M. Stallman*. Lulu.com, 2002 (cit. on pp. 2, 7).
- [SU17] Dominique Schröder and Dominique Unruh. “Security of blind signatures revisited”. In: *Journal of Cryptology* 30.2 (2017), pp. 470–494 (cit. on p. 59).

- [Sun+17] Shi-Feng Sun, Man Ho Au, Joseph K Liu, and Tsz Hon Yuen. “RingCT 2.0: a compact accumulator-based (linkable ring signature) protocol for blockchain cryptocurrency monero”. In: *European Symposium on Research in Computer Security*. Springer. 2017, pp. 456–474 (cit. on p. 35).
- [SWP16] David Shrier, Weige Wu, and Alex Pentland. “Blockchain & infrastructure (identity, data security)”. In: *Massachusetts Institute of Technology-Connection Science* 1.3 (2016) (cit. on p. 15).
- [Syt+16] Ewa Syta, Philipp Jovanovic, Eleftherios Kokoris Kogias, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Michael J. Fischer, and Bryan Ford. *Scalable Bias-Resistant Distributed Randomness*. Cryptology ePrint Archive, Report 2016/1067. <http://eprint.iacr.org/2016/1067>, Accessed 22 Feb 2017. 2016 (cit. on p. 127).
- [SZ15] Jared Saia and Mahdi Zamani. “Recent Results in Scalable Multi-Party Computation”. In: *SOFSEM 2015: Theory and Practice of Computer Science: 41st International Conference on Current Trends in Theory and Practice of Computer Science, Pec pod Sněžkou, Czech Republic, January 24-29, 2015. Proceedings*. Ed. by Giuseppe F. Italiano, Tiziana Margaria-Steffen, Jaroslav Pokorný, Jean-Jacques Quisquater, and Roger Wattenhofer. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 24–44. URL: https://doi.org/10.1007/978-3-662-46078-8_3 (cit. on p. 151).
- [TASo6] Niraj Tolia, David G Andersen, and Mahadev Satyanarayanan. “Quantifying interactive user experience on thin clients”. In: *Computer* 3 (2006), pp. 46–52 (cit. on p. 116).
- [Tea18] Team Rocket. *Snowflake to Avalanche: A Novel Metastable Consensus Protocol Family for Cryptocurrencies*. IPFS. 2018 (cit. on p. 34).
- [Tot13] Sree Harsha Totakura. “Large Scale Distributed Evaluation of Peer-to-Peer Protocols”. Master’s Thesis. Garching bei München: Technische Universität München, June 2013, p. 76 (cit. on p. 141).
- [TRL12] Sasu Tarkoma, Christian Esteve Rothenberg, and Eemil Lagerspetz. “Theory and practice of bloom filters for distributed systems”. In: *Communications Surveys & Tutorials, IEEE* 14.1 (2012), pp. 131–155 (cit. on p. 139).
- [TWo1] Robert Tracz and Konrad Wrona. “Fair Electronic Cash Withdrawal and Change Return for Wireless Networks”. In: *Proceedings of the 1st International Workshop on Mobile Commerce*. WMC ’01. Rome, Italy: ACM, 2001, pp. 14–19. URL: <http://doi.acm.org/10.1145/381461.381464> (cit. on pp. 6, 30, 32).

- [Vuk15] Marko Vukolić. “The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication”. In: *International Workshop on Open Problems in Network Security*. Springer. 2015, pp. 112–125 (cit. on p. 34).
- [VV17] Paul Voigt and Axel Von dem Bussche. *The EU General Data Protection Regulation (GDPR)*. Vol. 18. Springer, 2017 (cit. on p. 3).
- [Wah+18] Riad S Wahby, Ioanna Tzialla, Abhi Shelat, Justin Thaler, and Michael Walfish. “Doubly-efficient zkSNARKs without trusted setup”. In: *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2018, pp. 926–943 (cit. on p. 35).
- [Wal19] Angela Walch. “Deconstructing ‘Decentralization’: Exploring the Core Claim of Crypto Systems”. In: *Crypto Assets: Legal and Monetary Perspectives (OUP, forthcoming 2019)* (2019) (cit. on p. 33).
- [WG17] Karl Wüst and Arthur Gervais. “Do you need a Blockchain?” In: *IACR Cryptology ePrint Archive 2017* (2017), p. 375 (cit. on p. 35).
- [Woo14] Gavin Wood. “Ethereum: A secure decentralised generalised transaction ledger”. In: *Ethereum project yellow paper 151* (2014), pp. 1–32 (cit. on p. 35).
- [WSG14] Matthias Wachs, Martin Schanzenbach, and Christian Grothoff. “A Censorship-Resistant, Privacy-Enhancing and Fully Decentralized Name System”. In: *Proceedings of the 13th International Conference on Cryptology and Network Security - Volume 8813*. New York, NY, USA: Springer-Verlag New York, Inc., 2014, pp. 127–142. URL: http://dx.doi.org/10.1007/978-3-319-12280-9_9 (cit. on p. 154).
- [Yee13] P. Yee. *Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*. RFC 6818 (Proposed Standard). RFC. Fremont, CA, USA: RFC Editor, Jan. 2013. URL: <https://www.rfc-editor.org/rfc/rfc6818.txt> (cit. on p. 19).
- [Zak11] Alon Zakai. “Emscripten: an LLVM-to-JavaScript compiler”. In: *Proceedings of the ACM international conference companion on Object oriented programming systems languages and applications companion*. ACM. 2011, pp. 301–312 (cit. on p. 99).
- [Zha+18] Pingyuan Zhang, Han Jiang, Zhihua Zheng, Peichu Hu, and Qiuliang Xu. “A New Post-Quantum Blind Signature From Lattice Assumptions”. In: *IEEE Access* 6 (2018), pp. 27251–27258 (cit. on p. 155).
- [ZSI13] Mark Zandi, Virendra Singh, and Justin Irving. “The impact of electronic payments on economic growth”. In: *Moody’s Analytics: Economic and Consumer Credit Analytics* 217 (2013) (cit. on p. 2).

A. Résumé en Français

Les nouveaux protocoles de réseautage et cryptographiques peuvent considérablement améliorer les systèmes de paiement électroniques en ligne. Le présent mémoire porte sur la conception, la mise en œuvre et l'analyse sécuritaire du GNU Taler, un système de paiement respectueux de la vie privée conçu pour être pratique pour l'utilisation en ligne comme méthode de (micro-)paiement, et en même temps socialement et moralement responsable.

La base technique du GNU Taler peut être dû à l'e-cash de David Chaum. Notre travail va au-delà de l'e-cash de Chaum avec un changement efficace, et la nouvelle notion de transparence des revenus garantissant que les marchands ne peuvent recevoir de manière fiable un paiement d'un payeur non fiable que lorsque leurs revenus du paiement est visible aux autorités fiscales.

La transparence des revenus est obtenue grâce à l'introduction d'un protocole d'actualisation donnant lieu à un changement anonyme pour un jeton partiellement dépensé sans avoir besoin de l'introduction d'une évasion fiscale échappatoire. En outre, le protocole d'actualisation peut être utilisé pour la mise en œuvre des swaps atomiques de style Camenisch, et pour la préservation de l'anonymat en présence d'annulation du protocole et d'erreurs de crash avec perte de données par les participants. De plus, nous démontrons la sécurité prouvable de la transparence anonyme de nos revenus e-cash, qui concerne en plus l'anonymat habituel et les propriétés infalsifiables de l'e-cash, ainsi que la conservation formelle des fonds et la transparence des revenus.

Notre mise en œuvre du GNU Taler est utilisable par des utilisateurs non-experts et s'intègre à l'architecture du web moderne. Notre plateforme de paiement aborde une série de questions pratiques telles que la prodigue des conseils aux clients, le mode de remboursement, l'intégration avec les banques et les chèques "know-your-customer (KYC)", ainsi que les exigences de sécurité et de fiabilité de la plateforme web. Sur une seule machine, nous réalisons des taux d'opérations qui rivalisent avec ceux des processeurs de cartes de crédit commerciaux globaux. Nous améliorons la robustesse des échanges - la composante qui détient l'argent de banque en mains tierces en échange de l'e-cash - en ajoutant une composante d'auditeur qui vérifie le fonctionnement correct du système, et permet une détection tôt d'un compromis ou d'un mauvais comportement de l'échange.

Tout comme les comptes bancaires ont raisons d'exister de même que les billets de banque, l'e-cash ne sert que dans le cadre d'un empilement de système de paiement. Les registres distribués ont récemment gagné une immense popularité en tant que substituant potentiel des parties de l'industrie financière

traditionnelle.

Pendant que les crypto-monnaies basées sur la preuve de travail à l'instar de Bitcoin doivent encore être mises à l'échelle pour servir de substituant aux systèmes de paiement établis, d'autres systèmes plus efficaces basés sur les Blockchains avec des algorithmes de consensus plus classiques pourraient avoir des applications prometteurs dans le secteur financier. Nous faisons dans la conception, la mise en œuvre et l'analyse de la Byzantine Set Union Consensus, un protocole de Byzantine consensus qui s'accorde sur un (Super-)ensemble d'éléments à la fois, au lieu d'accepter en séquence les éléments individuels sur un ensemble. Byzantine Set consensus peut être utilisé comme composante de base pour des chaînes de blocs de permissions, où (à l'instar du style Nakamoto consensus) des blocs entiers d'opérations sont convenus à la fois d'augmenter le taux d'opération.

Objectifs du GNU Taler

Nous avons entamé la conception du GNU Taler avec un ensemble d'objectifs de conception de haut niveau. Ces objectifs sont classés par ordre de leur importance, et lorsqu'un compromis doit être fait, celui qui apporte son soutien à l'objectif le plus haut classé est préféré :

1. **GNU Taler doit être mise en œuvre comme un logiciel libre.** Le vocable libre ici renvoie à "liberté, comme dans liberté d'expression" et non "gratuité, comme dans bière gratuite". Plus particulièrement, les quatre libertés essentielles du logiciel libre doivent être respectées, notamment que les utilisateurs doivent avoir la liberté (1) d'exécuter le logiciel, (2) l'étudier et le modifier, (3) redistribuer des copies, et (4) distribuer des copies de la version modifiée.

En ce qui concerne les marchands, cela empêche le verrouillage par le fournisseur, car un autre fournisseur de paiement peut prendre le relais, si l'actuel offre une qualité de service insuffisante. Comme le logiciel de prestataire de paiement lui-même est libre, les plus petits pays ou organisations défavorisés peuvent exécuter le système de paiement sans besoin d'être contrôlés par une société étrangère. Les clients bénéficient de cette liberté, car le portefeuille électronique peut être conçu pour fonctionner sur une variété de plateformes, et des fonctionnalités hostiles à l'utilisateur telles que le tracking ou la télémétrie pourrait facilement être supprimé du portefeuille électronique.

Cela exclut l'utilisation obligatoire du matériel informatique spécialisé tel que les cartes électroniques ou d'autres hardware security modules, car le logiciel qu'ils exécutent ne peut être modifié par l'utilisateur.

Cependant, ces composantes peuvent être volontairement utilisées par des marchands, clients ou organismes de paiement dans le but d'améliorer leur

sécurité opérationnelle.

2. GNU Taler doit protéger la vie privée des acheteurs.

La protection de la vie privée devrait être garantie par des mesures techniques, aux antipodes des simples politiques. Ceci particulièrement à l'aide des micropaiements pour des contenus en ligne, une quantité disproportionnée de données plutôt privées sur les acheteurs serait révélée, si le système de paiement n'a pas de mécanisme de protection de la vie privée.

Dans les législations relatives à la protection des données (à l'instar du GDPR récemment introduit en Europe), les marchands en profitent également, car aucune violation des données des clients ne peut se produire si ces informations ne sont pas collectées premièrement par conception. De toute évidence, certaines données privées, telles l'adresse de livraison pour une livraison physique, doivent toujours être collectées en fonction des besoins de l'entreprise.

La sécurité des systèmes de paiement en profite également, car le modèle passe de l'authentification des clients à la simple autorisation de paiement. Cette approche exclut les classes entières d'attaques telles que l'hameçonnage ou la fraude par carte de crédit.

3. GNU Taler doit permettre à l'État de procéder à l'imposition des revenus et la répréhension des activités commerciales illégales.

Comme un système de paiement doit toujours être légal pour son fonctionnement et utilisation, il doit se conformer aux exigences susmentionnées. En outre, nous considérons que la perception des impôts est bénéfique pour la société.

4. GNU Taler doit prévenir la fraude.

Cela impose des exigences sur la sécurité du système, ainsi que sur la conception générale, car la fraude de paiement peut également se produire par la conception illusoire de l'interface utilisateur, ou le manque de preuves cryptographiques pour certains processus.

5. GNU Taler ne doit permettre uniquement que la divulgation de la quantité minimale d'informations nécessaires.

La raison derrière cet objectif est semblable à (2). La vie privée des acheteurs est prioritaire, mais d'autres parties à l'instar des marchands en profitent encore, par exemple en conservant des détails sur leurs finances cachés aux concurrents.

6. GNU Taler doit être utilisable.

Il doit particulièrement être utilisable pour des clients non-experts. Cette utilité s'applique également à l'intégration avec les marchands, et renseigne sur les choix concernant l'architecture, telles que les procédures

d'encapsulation requièrent des opérations cryptographiques dans une composante isolée avec une simple API.

7. GNU Taler doit être efficace.

Les approches telles que la preuve de travail sont exclues par cette exigence. L'efficacité est nécessaire pour que GNU Taler soit utilisé pour les micropaielements.

8. GNU Taler doit éviter les points de défaillance uniques.

Alors que la conception à présenter plus tard est plutôt centralisée, son objectif demeure d'éviter les points uniques de défaillance. Cela se manifeste dans les choix architecturaux tels que l'isolement de certaines composantes et procédures d'audit.

9. GNU Taler doit promouvoir la concurrence.

Le processus d'intégration des concurrents aux systèmes doit relativement être facile. Alors que les obstacles à ces systèmes financiers traditionnels sont assez nombreux, le fardeau technique d'adhésion aux nouveaux concurrents doit être minimisé. Un autre choix de conception qui soutient ceci est de diviser l'ensemble du système en de plus petites composantes pouvant être exploitées, développées et améliorées indépendamment, au lieu d'avoir un système complètement monolithique.

Byzantine Set Union Consensus

Le protocole de Byzantine Set Union Consensus que nous procédons à la conception, la mise en œuvre et l'évaluation, offre une amélioration asymptotique au-delà d'une mise en œuvre naïve à l'aide de la machine de duplication de l'État.

Pour des pairs n et un ensemble d'éléments m , la messagerie de notre protocole est d'une complexité $O(mn + n^2)$ lorsqu'aucun pair ne montre le comportement Byzantine. Lorsque les pairs f montrent un comportement Byzantine, la complexité du message est $O(mnf + kfn^2)$, où k est le nombre d'éléments valides exclusivement disponibles pour l'adversaire.

Nous démontrons comment k peut être délimité pour des applications pratiques communes, puisqu'en général k est seulement délimité par la largeur disponible à l'adversaire. En pratique, on s'attend à ce que $k f$ soit significativement plus petit que m . Ainsi, $O(mnf + kfn^2)$ est une amélioration par rapport à l'utilisation de SMR-PBFT (duplication de la machine d'État avec la tolérance de faille Byzantine pratique) qui aurait la complexité $O(mn^2)$ sous ces hypothèses.

Nous parvenons à ce résultat en combinant un protocole de Byzantine Consensus existant au protocole de réconciliation efficace d'Eppstein. La même construction s'applique également aux autres protocoles de consensus.

Contributions

Nous revendiquons les contributions clés ci-dessous dans le cadre ce mémoire :

- Nous procédons à la conception, la mise en œuvre et une analyse efficace du Byzantine consensus protocol sur des structures définies permettant une mise en œuvre optimisée des registres d'opérations distribuées.
- Nous introduisons la notion de transparence des revenus pour l'e-cash, avec une instanciation en e-cash et des épreuves du style Chaum.
- Nous procédons à la conception du système de paiement GNU Taler en tenant compte des aspects pratiques de l'e-cash, notamment les annulations, les défaillances de réseau, les remboursements, les paiements multi-coin, les défauts de synchronisation des portefeuilles et leurs effets sur l'anonymat; montrant la nécessité d'une opération d'actualisation.
- Nous proposons une modification de notre protocole offrant une protection contre certains scénarios de chantage et d'enlèvement.
- Nous procédons à la conception et la mise en œuvre d'une intégration homogène et native de l'e-cash dans l'architecture du web, et discutons des aspects de sécurité et de confidentialité de cette intégration.
- Nous avons mis en œuvre le système de paiement GNU Taler et avons évalué ses performances.

B. dold-draft-payto *

Independent Stream
Internet-Draft
Intended status: Informational
Expires: October 19, 2019

F. Dold
Taler Systems SA
C. Grothoff
BFH
April 17, 2019

The 'payto' URI scheme for payments
draft-dold-payto-06

Abstract

This document defines the 'payto' Uniform Resource Identifier (URI) scheme for designating targets for payments.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 19, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

*The most recent version of the draft can be found at <https://datatracker.ietf.org/doc/draft-dold-payto/>.

B. dold-draft-payto

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
 2. Syntax of a 'payto' URL
 3. Semantics
 4. Examples
 5. Generic Options
 6. Internationalization and Character Encoding
 7. Security Considerations
 8. IANA Considerations
 - 8.1. URI Scheme Registration
 - 8.2. Payment Target Type Registry
 - 8.2.1. ACH Bank Account
 - 8.2.2. Business Identifier Code
 - 8.2.3. International Bank Account Number
 - 8.2.4. Unified Payments Interface
 - 8.2.5. Bitcoin Address
 - 8.2.6. Interledger Protocol Address
 9. References
 - 9.1. Normative References
 - 9.2. Informational References
- Authors' Addresses

1. Introduction

This document defines the 'payto' Uniform Resource Identifier (URI) [RFC3986] scheme for designating transfer form data for payments. In particular, it always identifies the target of a payment. A 'payto' URL consists of a payment target type, a target identifier and optional parameters such as an amount or a payment reference.

The interpretation of the target identifier is defined by the payment target type, and typically represents either a bank account or an (unsettled) transaction.

A unified URI scheme for all payment target types allows applications to offer user interactions with URIs that represent payment targets,

simplifying the introduction of new payment systems and applications.

2. Syntax of a 'payto' URL

This document uses the Augmented Backus-Naur Form (ABNF) of [RFC5234].

```
payto-URI = "payto" "://" authority path-abempty [ "?" opts ]
opts = opt *( "&" opt )
opt = (generic-opt / authority-specific-opt) "=" *( pchar )
generic-opt = "amount" / "receiver-name" / "sender-name" /
              "message" / "instruction"
authority = ALPHA *( ALPHA / DIGIT / "-" / "." )
path-abempty = <path-abempty, see [RFC3986], Section 3.3>
pchar = <pchar, see [RFC3986], Appendix A.>
```

3. Semantics

The authority component of a payment URI identifies the payment target type. The payment target types are defined in the "Payment Target Types" registry, see Section 8.2. The path component of the URI identifies the target for a payment as interpreted by the respective payment target type. The query component of the URI can provide additional parameters for a payment. Every payment method SHOULD accept the options defined in generic-opt. The default operation of applications that invoke a URI with the payto scheme SHOULD be to launch an application (if available) associated with the payment target type that can initiate a payment. If multiple handlers are registered for the same payment target type, the user SHOULD be able to choose which application to launch. This allows users with multiple bank accounts (each accessed the respective bank's banking application) to choose which account to pay with. An application SHOULD allow dereferencing a payto URI even if the payment target type of that URI is not registered in the "Payment Target Types" registry. Details of the payment MUST be taken from the path and options given in the URI. The user SHOULD be allowed to modify these details before confirming a payment.

4. Examples

```
payto://iban/DE75512108001245126199?amount=EUR:200.0&message=hello
```

```
INVALID (authority missing): payto:iban/12345
```

5. Generic Options

B. *dold-draft-payto*

Applications MUST accept URIs with options in any order. The "amount" option MUST only occur at most once. Other options MAY be allowed multiple times, with further restrictions depending on the payment method. The following options SHOULD be understood by every payment method.

amount: The amount to transfer, including currency information if applicable. The format MUST be:

```
amount = [ currency ":" ] unit [ "." fraction ]
currency = 1*ALPHA
unit = 1*(DIGIT / ",")
fraction = 1*(DIGIT / ",")
```

The unit value MUST be smaller than 2^{53} . If present, the fraction MUST consist of no more than 8 decimal digits. The use of commas is optional for readability and they MUST be ignored.

receiver-name: Name of the entity that receives the payment (creditor).

sender-name: Name of the entity that makes the payment (debtor).

message: A short message to identify the purpose of the payment, which MAY be subject to lossy conversions (for example, due to character set encoding limitations).

instruction: A short message giving instructions to the recipient, which MUST NOT be subject to lossy conversions. Character set limitations allowed for such instructions depend on the payment method.

6. Internationalization and Character Encoding

Various payment systems use restricted character sets. An application that processes 'payto' URIs MUST convert characters that are not allowed by the respective payment systems into allowable character using either an encoding or a replacement table. This conversion process MAY be lossy, except for the instruction field.

To avoid special encoding rules for the payment target identifier, the userinfo component [RFC3986] is disallowed in payto URIs. Instead, the payment target identifier is given as an option, where encoding rules are uniform for all options.

7. Security Considerations

Interactive applications handling the payto URI scheme MUST NOT initiate any financial transactions without prior review and confirmation from the user, and MUST take measures to prevent clickjacking [HMW12].

Unless a payto URI is received over a trusted, authenticated channel, a user might not be able to identify the target of a payment. In particular due to homographs [unicode-tr36], a payment target type SHOULD NOT use human-readable names in combination with unicode in the target account specification, as it could give the user the illusion of being able to identify the target account from the URL.

To avoid unnecessary data collection, payment target types SHOULD NOT include personally identifying information about the sender of a payment that is not essential for an application to conduct a payment.

8. IANA Considerations

8.1. URI Scheme Registration

The "payto" URI scheme is already registered in the "Provisional URI Schemes" registry.

Scheme name: payto

Status: permanent

URI scheme syntax: See Section 2.

URI scheme semantics: See Section 3.

Applications/protocols that use this scheme name: payto URIs are mainly used by financial software, as well as by interactive applications (e.g. email clients, chat applications) that detect payto URIs and allow the user to interact with them (e.g. make them clickable)

Contact: grothoff@gnu.org

Change controller: grothoff@gnu.org

References: See References section of this document.

8.2. Payment Target Type Registry

B. dold-draft-payto

This document defines a registry for payment methods. The name of the registry is "Payment Target Types".

The registry shall record for each entry:

- o Name: The name of the payment target type (case insensitive ASCII string, restricted to alphanumeric characters, dots and dashes)
- o Description: A description of the payment target type, including the semantics of the path in the URI if applicable.
- o Example: At least one example URI to illustrate the payment target type.
- o Contact: The contact information of a person to contact for further information
- o References: Optionally, references describing the payment method (such as an RFC) and method-specific options, or references describing the payment system underlying the payment target type.

The registration policy for this registry is "expert review", as described in [RFC5226]. The expert is appointed by the IETF Independent Stream Editor. The expert's review SHOULD consider the following criteria:

1. The proposed registry entry contains all mandatory information.
2. The description clearly defines the syntax and semantics of the payment target and optional parameters if applicable.
3. Relevant references are provided if they are available.
4. The chosen name is appropriate for the payment target type, does not conflict with well-known payment systems, and avoids potential to confuse users.
5. The payment system underlying the payment target type is not fundamentally incompatible with the general options (such as positive decimal amounts) in this specification.
6. The payment target type is not a vendor-specific version of a payment target type that could be described more generally by a vendor-neutral payment target type.
7. The specification of the new payment target type remains within

the scope of payment transfer form data. In particular specifying complete invoices is not in scope. Neither are processing instructions to the payment processor or bank beyond a simple payment.

8. The payment target and the options do not contain the payment sender's account details.

8.2.1. ACH Bank Account

- o Name: ach
- o Description: Automated Clearing House. The path consist of two components, the routing number and the account number.
- o Example: payto://ach/122000661/1234
- o Contact: N/A
- o References: [NACHA]

8.2.2. Business Identifier Code

- o Name: bic
- o Description: Business Identifier Code. The path consist of just a BIC. This is used for wire transfers between banks. The registry for BICs is provided by SWIFT. The path does not allow specifying a bank account number.
- o Example: payto://bic/SOGEDEFFXXX
- o Contact: N/A
- o References: [BIC]

8.2.3. International Bank Account Number

- o Name: iban
- o Description: International Bank Account Number (IBAN). Generally the IBAN allows to unambiguously derive the the associated Business Identifier Code (BIC). However, some legacy applications process payments to the same IBAN differently based on the specified BIC. Thus the path can either consist of a single component (the IBAN) or two components (BIC and IBAN).

B. dold-draft-payto

- o Example: `payto://iban/DE75512108001245126199`
`payto://iban/SOGEDEFFXXX/DE75512108001245126199`
- o Contact: N/A
- o References: [ISO20022]

8.2.4. Unified Payments Interface

- o Name: `upi`
- o Description: Unified Payment Interface. The path is an account alias. The amount and receiver-name options are mandatory for this payment target.
- o Example: `payto://upi/alice@example.com?receiver-name=Alice&amount=INR:200`
- o Contact: N/A
- o References: [UPILinking]

8.2.5. Bitcoin Address

- o Name: `bitcoin`
- o Description: Bitcoin protocol. The path is a "bitcoinaddress" as per [BIP0021].
- o Example: `payto://bitcoin/12A1MyfXbW6RhdRAZEqofac5jCQQjwEPBu`
- o Contact: N/A
- o References: [BIP0021]

8.2.6. Interledger Protocol Address

- o Name: `ilp`
- o Description: Interledger protocol. The path is an ILP address as per [ILP-ADDR].
- o Example: `payto://ilp/g.acme.bob`
- o Contact: N/A
- o References: [ILP-ADDR]

9. References

9.1. Normative References

- [ISO20022] International Organization for Standardization, "ISO 20022 Financial Services – Universal financial industry message scheme", May 2013.
- [NACHA] NACHA, "NACHA Operating Rules & Guidelines", January 2017.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <<https://www.rfc-editor.org/info/rfc5226>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [unicode-tr36] Davis, M., Ed. and M. Suignard, "Unicode Technical Report #36: Unicode Security Considerations", September 2014.

9.2. Informational References

- [BIC] International Organization for Standardization, "ISO 9362:2014 Business Identifier Code (BIC)", March 2019, <<https://www.iso.org/standard/60390.html>>.
- [BIP0021] Schneider, N. and M. Corallo, "Bitcoin Improvement Proposal 21", January 2012, <https://en.bitcoin.it/wiki/BIP_0021>.
- [HWM12] Huang, L., Moshchuk, A., Wang, H., Schecter, S., and C. Jackson, "Clickjacking: Attacks and Defenses", January 2012, <<https://www.usenix.org/system/files/conference/usenixsecurity12/sec12-final39.pdf>>.
- [ILP-ADDR] Interledger Team, "ILP Addresses – v2.0.0", September 2018, <<https://interledger.org/rfcs/0015-ilp-addresses/>>.

B. dold-draft-payto

[UPILinking]

National Payment Corporation of India, "Unified Payment Interface - Common URL Specifications For Deep Linking And Proximity Integration", May 2016,
<<http://www.npci.org.in/documents/UPILinkingSpecificationsVersion10draft.pdf>>.

Authors' Addresses

Florian Dold
Taler Systems SA
7, rue de Mondorf
Erpeldange L-5421
LU

Email: dold@taler.net

Christian Grothoff
BFH
Hoeheweg 80
Biel/Bienne CH-2501
CH

Email: christian.grothoff@bfh.ch

C. Coin Spending Simulation

The most recent version of this TypeScript program can be found in the repository of the wallet reference implementation (<https://git.taler.net/wallet-webex.git/tree/contrib/coinsim.ts>).

```
/*
  This file is part of GNU Taler
  (C) 2018 GNUnet e.V.

  GNU Taler is free software; you can redistribute it and/or modify it under the
  terms of the GNU General Public License as published by the Free Software
  Foundation; either version 3, or (at your option) any later version.

  GNU Taler is distributed in the hope that it will be useful, but WITHOUT ANY
  WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR
  A PARTICULAR PURPOSE. See the GNU General Public License for more details.

  You should have received a copy of the GNU General Public License along with
  GNU Taler; see the file COPYING. If not, see <http://www.gnu.org/licenses/>
  */

function getRandomInt(min, max) {
  return Math.floor(Math.random() * (max - min + 1)) + min;
}

const denoms = [8096, 4096, 2048, 1024, 512, 256, 128, 64, 32, 16, 8, 4, 2, 1];

// mapping from denomination index to count
const wallet = denoms.map(() => 0);

const trans_max = 5000;
const trans_min = 4;

const withdraw_max = 10000;

const num_transactions = parseInt(process.argv[2]);

// Refresh or withdraw operations
let ops = 0;
let ops_refresh = 0;
let ops_withdraw = 0;
let ops_spend = 0;
let refresh_output = 0;

function withdraw(amount, is_refresh) {
```


C. Coin Spending Simulation

```
while (amount != 0) {
  for (let i = 0; i < denoms.length; i++) {
    let d = denoms[i];
    if (d <= amount) {
      amount -= d;
      wallet[i]++;
      ops++;
      if (!is_refresh) {
        ops_withdraw++;
      } else {
        refresh_output++;
      }
      break;
    }
  }
}

function spendSmallestFirst(cost) {
  while (cost != 0) {
    for (let j = 0; j < denoms.length; j++) {
      const k = denoms.length - j - 1;
      const d = denoms[k];
      const w = wallet[k];
      if (w == 0) {
        continue;
      }
      if (d <= cost) {
        // spend
        wallet[k]--;
        cost -= d;
        ops++;
        ops_spend++;
        break;
      }
      // partially spend and then refresh
      ops++;
      ops_spend++;
      let r = d - cost;
      ops_refresh++;
      wallet[k]--;
      withdraw(r, true);
      cost = 0;
    }
  }
}

function spendLargestFirst(cost) {
  while (cost != 0) {
    for (let j = 0; j < denoms.length; j++) {
      const d = denoms[j];
      const w = wallet[j];
      if (w == 0) {
```

```

        continue;
    }
    if (d <= cost) {
        // spend
        wallet[j]--;
        cost -= d;
        ops++;
        ops_spend++;
        break;
    }
    // partially spend and then refresh
    ops++;
    ops_spend++;
    let r = d - cost;
    ops_refresh++;
    wallet[j]--;
    withdraw(r, true);
    cost = 0;
}
}
}

function spendHybrid(cost) {
    for (let j = 0; j < denoms.length; j++) {
        const k = denoms.length - j - 1;
        const d = denoms[k];
        const w = wallet[k];
        if (w == 0) {
            continue;
        }
        if (d < cost) {
            continue;
        }
        // partially spend and then refresh
        ops++;
        ops_spend++;
        let r = d - cost;
        ops_refresh++;
        wallet[k]--;
        withdraw(r, true);
        cost = 0;
    }

    spendSmallestFirst(cost);
}

for (let i = 0; i < num_transactions; i++) {
    // check existing wallet balance
    let balance = 0;
    for (let j = 0; j < denoms.length; j++) {
        balance += wallet[j] * denoms[j]
    }
    // choose how much we want to spend

```

C. Coin Spending Simulation

```
let cost = getRandomInt(trans_min, trans_max);
if (balance < cost) {
  // we need to withdraw
  let amount = getRandomInt(cost - balance, withdraw_max);
  withdraw(amount, false);
}

// check that we now have enough balance
balance = 0;
for (let j = 0; j < denoms.length; j++) {
  balance += wallet[j] * denoms[j]
}

if (balance < cost) {
  throw Error("not enough balance");
}

// now we spend
spendHybrid(cost);
}

console.log("total ops", ops / num_transactions);
console.log("spend ops", ops_spend / num_transactions);
console.log("pure withdraw ops", ops_withdraw / num_transactions);
console.log("refresh (multi output) ops", ops_refresh / num_transactions);
console.log("refresh output", refresh_output / ops_refresh);
```