



Fast, Private and Regulated Payments in Asynchronous Networks

Maxence Brugerres  

LTCl, Telecom Paris, Institut Polytechnique de Paris, France
Forvis Mazars, Levallois-Perret, France

Victor Languille  

LTCl, Telecom Paris, Institut Polytechnique de Paris, France
EDF R&D, Paris, France

Petr Kuznetsov  

LTCl, Telecom Paris, Institut Polytechnique de Paris, France

Hamza Zarfaoui 

LTCl, Telecom Paris, Institut Polytechnique de Paris, France

Abstract

We propose a decentralized asset-transfer system that enjoys *full privacy*: no party can learn the details of a transaction, except for its issuer and its recipient. Furthermore, the recipient is not aware of the sender's identity. Our system does not rely on consensus or synchrony assumptions, and therefore, it is *responsive*, since it runs at the actual network speed. Under the hood, every transaction creates a consumable *coin* equipped with a non-interactive zero-knowledge proof (NIZK) that confirms that the issuer has sufficient funds without revealing any information about her identity, the recipient's identity, or the payment amount. Moreover, we equip our system with a *regulatory enforcement* mechanism that can be used to regulate transfer limits or restrict specific addresses from sending or receiving funds, while preserving the system's privacy guarantees.

Finally, we report on *PaxPay*, our implementation of Fully Private Asset Transfer (FPAT) that uses the Gnark library for the NIZKs. In our benchmark, PaxPay exhibits better performance than earlier proposals that either ensure only partial privacy, require some kind of network synchrony or do not implement regulation features. Our system thus reconciles privacy, responsiveness, regulation enforcement and performance.

2012 ACM Subject Classification Security and privacy; Computer systems organization → Dependable and fault-tolerant systems and networks; Computing methodologies → Distributed computing methodologies

Keywords and phrases Anonymous, Asset Transfer, Asynchronous System, BFT, CBDC, NIZK, Payment System, Privacy, Regulation, Scalability, zk-SNARK

Digital Object Identifier 10.4230/LIPIcs.AFT.2025.3

Related Version *Full Version*: <https://eprint.iacr.org/2025/098.pdf>

Funding *Maxence Brugerres*: Forvis Mazars Group

Victor Languille: EDF R&D

Petr Kuznetsov: CHIST-ERA-22-SPiDDS-05 (REDONDA project) and Agence Nationale de la Recherche (ANR-23-CHR4-0009)

1 Introduction

Payments, with and without consensus. Bitcoin [33] revolutionized the world of finance by proposing a fully decentralized asset-transfer system that allows an open set of users to consistently exchange assets without any mutual trust. Instead of relying on a central authority, users repeatedly engage in a form of *consensus* [22, 18] to maintain a replicated



© Maxence Brugerres, Victor Languille, Petr Kuznetsov, and Hamza Zarfaoui;
licensed under Creative Commons License CC-BY 4.0

7th Conference on Advances in Financial Technologies (AFT 2025).

Editors: Zeta Avarikioti and Nicolas Christin; Article No. 3; pp. 3:1–3:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

ledger – an ever-growing record of all transactions. It was later observed [27, 26] that, strictly speaking, global consensus, a notoriously hard task [16, 17] that requires partial synchrony [22, 12], is not necessary in payments. In most common cases, when every account is maintained by a dedicated user, asset transfer can be implemented in an asynchronous, *responsive* way on top of the *reliable-broadcast* primitive [10] (instead of consensus).

Reliable broadcast is weaker than consensus as it allows the correct users to eventually and consistently agree on the *set* of issued transactions but not on their order. However, it turned out to be strong enough to prevent *double spending*, a major issue in asset-transfer systems. This observation gave rise to a series of purely asynchronous, *consensus-free* asset-transfer systems [14, 5, 31].

Privacy and regulations in payment systems. However, a decentralized asset-transfer system, as any replicated service, still faces a vital challenge of preserving privacy of its users. Indeed, in all the mentioned payment protocols, all transactions are inherently *public* and every user’s activities are traceable. Hiding the amount of a payment offers *confidentiality* and hiding the identities of the sender and receiver offers *anonymity*. Fully *private* payment systems ought to offer both confidentiality and anonymity. Furthermore, many important payment applications, such as Central Bank Digital Currency (CBDC) [4], require mechanisms for *regulatory enforcement*, such as limiting transaction amounts.

Existing distributed asset-transfer systems ([6, 7, 37, 38, 41, 44, 45], to name a few) arbitrate a trade-off between the four following aspects: **(a) Privacy guarantees**, **(b) Model assumptions**, **(c) Regulation enforcement**, and **(d) Performance**. Reconciling these four aspects is a onerous challenge, as strengthening one aspect often weakens another. Table 1 provides a detailed comparison of several payment systems across these dimensions. This table and its content are detailed and discussed in Section 8.

FPAT and PaxPay. In this paper, we address this challenge. We introduce the abstraction of *Fully Private Asset Transfer* (FPAT) that maintains conventional safety and liveness properties of a payment system (informally, no asset is spent twice and every transaction takes effect), while at the same time making sure that no transaction’s detail is leaked to a non-involved party. We then describe the *PaxPay* protocol, a FPAT implementation over an *asynchronous* network. PaxPay has been implemented in Golang using Gnark library [8].

As in UTXO transactions [33], to execute a transfer, the sender *spends* a set of *old coins* it has received earlier and *generates* a set of *new coins* of the same total value. Every coin is uniquely identified by its *serial number*. Every coin should be signed by a sufficiently large set (a $2/3$ *quorum*) of *validators*, i.e., dedicated parties that maintain lists of spent coins and make sure that no user spends a same coin twice.

The protocol operates in an asynchronous network, where any number of users and less than one third of validators can be *Byzantine* (deviating arbitrarily from its algorithm). The rest of the validators are considered *semi-honest* (following their algorithm but possibly sharing their states with the adversary). To ensure that the transfer is legal, several assertions have to be checked: the total value of the spent coins equals the total value of the new coins, the user that calls the transfer is the owner of the spent coins, all the coins are properly signed, etc.

To provide *full privacy* (confidentiality and anonymity), PaxPay leverages several cryptographic primitives. The legality of transactions is verified within a non-interactive zero-knowledge proofs (NIZK). A *blind* signature scheme is used at the validators’ side and these signatures are verified within the NIZK, allowing not to reveal the coins or their signatures

when spending them. To verify these signatures efficiently, we implemented the NIZK using the Groth16 [25] scheme and employed a pairing-based signature scheme [35]. We then selected a pair of elliptic curves that form a chain to instantiate these cryptographic primitives, allowing efficient verification of the signature within the NIZK (we refer to Section 7 for more details).

Regulation. We also describe and implement a regulation enforcement mechanism that can be built on top of our system. Similar to PEReDi and PARscoin [37, 38], our approach allows setting limits on individual transfer amounts or the cumulative amount transferred by a user since they joined the system. However, regulations for CBDCs mandate stringent measures for Anti-Money Laundering (AML) and Countering the Financing of Terrorism (CFT). These measures go beyond simply limiting transaction amounts.

PEReDi and PARscoin [37, 38] address these requirements and enable validators to reveal the details of transfers if needed. This could put at risk the privacy of the users. Their approach relies on additional trust assumptions, requiring non-Byzantine validators to be *honest* (i.e., not to share any information with the adversary) to maintain privacy. We avoid such extra assumptions by equipping the users with a mechanism to generate cryptographic proofs for arbitrary statements about the histories of their transactions. It enables extracting a requested information without disclosing the full content of the transaction histories.

Additionally, we incorporate a novel mechanism, empowering the regulator to impose *sanctions*, precluding specific addresses from interactions with other users. This ensures compliance with traditional financial sanctions emitted by the central banks or intergovernmental organizations (e.g., the UN¹) and explicitly required by the European Central Bank for the digital Euro [4].

The regulatory mechanism is implemented by introducing an additional type of coin, referred to as the *compliance coin*. At a time, each user may possess several regular coins, but only one unique compliance coin. In each transfer, the compliance coin of the sender must be spent together with the regular coins. Then a new compliance coin containing the information of the transfer is generated (i.e., signed by a quorum of validators). The compliance coin thus allows us to keep track records of all the interactions of a user with the system, which can be seen as a commitment to the transaction history and some aggregated values. To generate an initially valid compliance coin, before joining the system, the user must be registered with a *regulatory authority*.

Our regulation framework incurs very low impact on the transaction throughput. It mainly affects the transaction proving time (by doubling it compared to the non-regulated case), which we do not consider a primary performance factor. In contrast, such a mechanism could be difficult and performance-intensive to implement in other protocols that rely on Sigma protocols [15], such as UTT [41], Zef [6], Peredi [37], and Parscoin [38], which use these protocols to construct their NIZKs.

Performance. Finally, we evaluated the performance of PaxPay, in comparison with state-of-the-art payment systems that provide privacy-preserving and/or regulatory-enforcement features (cf Table 1). In our benchmark (detailed in Section 7.2), we witness that PaxPay can process at least 8 times more transactions per second than any of these systems. The benchmark on AWS EC2 instances demonstrates a throughput of 925 transactions per second. The performance gains can be explained by the use of Groth16 as a succinct NIZK with low

¹ <https://main.un.org/securitycouncil/en/sanctions/information>

verification cost. In return, this comes at the cost of heavy computations required to generate a proof. The time required to create a transfer request with limited computational resources in PaxPay on a one-core CPU is up to 7s for example. However, given the throughput gains, this appears to be a good trade-off. Moreover, PaxPay’s design allows for pushing the throughput even further by parallelizing the validator’s load on several machines.

Full version. The full version of this paper [11] incorporates more details including the algorithms that implement PaxPay, the security proofs, the link toward PaxPay implementation and additional details on the benchmarks.

Summary. To sum up our contributions:

- We formalize the problem by introducing the *Fully Private Asset Transfer* abstraction (FPAT).
- We propose PaxPay, a protocol that implements a FPAT object, and prove its correctness.
- We implement PaxPay in Golang, mainly leveraging on the Gnark [8] library.
- PaxPay achieves a true reconciliation of **(a) Privacy**, **(b) Model Assumptions** (network synchrony and trust assumptions for validators), **(c) Regulation Enforcement**, and **(d) Performance**. Unlike earlier works that propose trade-offs, PaxPay excels in each of these aspects simultaneously, bringing together the best of all these dimensions:
 - Regarding **Privacy**, transactions in PaxPay are *confidential*, *fully anonymous*, and *unlinkable*, ensuring the most robust privacy guarantees.
 - Regarding **Model Assumptions**, the system tolerates up to one-third Byzantine validators, with the remaining validators only required to be *semi-honest*, and does not rely on the network synchrony.
 - Regarding **Regulation**, PaxPay provides a comprehensive regulation framework by enabling users to generate cryptographic proofs about their transaction history, setting spending limits, and incorporating a sanction mechanism preventing sanctioned addresses to receive or spend funds.
 - Finally, our implementation outperforms state-of-the-art protocols. Also, we show that the transaction throughput of PaxPay scales with the computational power of validators. The system can further increase its supported throughput by distributing validators across additional machines.
- PaxPay can find a variety of use cases such as a standalone payment system, a *layer 2* solution on top of a blockchain (brings scalability, compliance, or privacy) or a CBDC.

2 Related work

We overview payment systems related to our work, the summary of our comparative analysis is presented in Table 1 compares this work with the related works. Additional details about the comparison can be found in the full version of the paper.

Consensus-free payment systems. Gupta [27] and Guerraoui *et al.* [26] proposed a payment system that replaces the conventional consensus-based synchronization with the secure broadcast primitive [10, 32]. Systems like Astro [14], FastPay [5], Zef [6], PARScoin [38] and UTT [41] extend this approach, operating without consensus in asynchronous networks.

² Fully anonymous among all the users (Full) or anonymous among an anonymity set (AS).

³ Semi-Honest (SH) or Honest (H).

■ **Table 1** Comparison of PaxPay vs. Zcash [19], Lelantus [29], Quisquis [21], Zef [6], UTT [41], PRCash [44], PEReDi [37] and PARScoin [38].

	Zcash	Lelantus	Quisquis	Zef	UTT	PRCash	PEReDi	PARScoin	PaxPay
PRIVACY PROPERTIES									
Confidential transfers	●	●	●	●	●	●	●	●	●
Sender-anonymous transfers	●	●	●	○	●	●	●	●	●
Receiver-anonymous transfers	●	●	●	●	●	●	●	●	●
Unlinkable transfers	●	●	●	○	●	●	●	●	●
Anonymity strategy ²	Full	AS	AS	Full	Full	Full	Full	Full	Full
MODEL ASSUMPTIONS									
Asynchronous network	○	○	○	●	●	○	○	●	●
Correct validators ³	SH	SH	SH	SH	SH	SH	H	H	SH
REGULATION FEATURES									
Limited held amount per user	○	○	○	○	○	○	●	○	○
Limited spendable amount per tx	○	○	○	○	○	●	●	●	●
Limited spendable amount in total	○	○	○	○	●	○	●	●	●
Full asset tracing	○	○	○	○	○	○	●	●	○
Sanctions	○	○	○	○	○	○	○	○	●
Provable transaction history	○	○	○	○	○	○	○	○	●
PERFORMANCE									
Transaction throughput (tx/s)	25	0	0	88	235	0	0	0	925
Transaction latency (s)	1000	0	0	< 1	< 1	0	0	0	< 1
NIZK Proving time (ms)	21K	2378	2110	438	60	100	3100	392	6959
NIZK Verification time (ms)	9	40	251	142	49	96	518	159	5

Private payment systems without regulation. Zerocash [7], Zexe [9], Monero [34], Quisquis [21] and Lelantus [29] are payment systems built on top of blockchains and consensus. They provide *privacy* via NIZKs or ring signatures. Zerocash and Zexe [7, 9] use *commitments* stored in Merkle trees and nullifiers to prevent double spending. Each transaction is accompanied by an NIZK to provide privacy but these systems offer inherently low transaction throughput. Monero [34] and Quisquis [21] relies on ring signatures to provide anonymity but only limits the sender and receiver’s identity to a small subset of the users, known as the *anonymity set*. Compared to Monero, Quisquis offers stronger privacy guarantees in certain scenarios preventing possible de-anonymization.

Lelantus [29, 30] offers privacy via NIZK requires less advanced cryptographic assumptions than Zerocash or Zexe, for instance it does not require any trusted setup. It also proposes batch verification for validators.

All payment systems mentioned so far are built on a blockchain and, thus, require consensus. Baudet et al. [6] describe Zef, a private payment system, that assumes an asynchronous network but only provides *partial* anonymity and confidentiality. To improve performance, the transaction data in Zef is distributed over a set of *authorities*. Each authority can be sharded – i.e., distributed over several machines, as in Astro [14] and in this work.

Private payment systems with regulation. PRCash [44], Platypus [45], PEReDi [37], PARScoin [38] and UTT [41] incorporate certain regulatory features into private payment systems. A common element in their designs is the reliance on NIZK.

Garman et al. [23] were among the first to address regulation by enforcing *policies* such as fixing a spending limit on a system with a similar construction as Zerocash. Their system has not been implemented and, as it is built on top of Zerocash, it shares the same major drawbacks.

PRCash [44], as Zerocash, is blockchain-based and operates with commitments but offers only partial privacy. Regulation features proposed by PRCash are limited: users cannot spend more than a predefined amount set by the regulator within a certain time window.

UTT [41], a concurrent work to our paper, proposes an asynchronous private payment system with privacy guarantees, a low transaction latency and a high transaction throughput. Both UTT and PaxPay protocols follow a similar approach, using coin commitments and consistent broadcast to create and spend coins in a fully private manner. However, the regulation features offered by UTT are restricted to the *anonymity budget* that limits the amount of anonymous coins that a user can transfer. Compared to UTT, PaxPay offers a more comprehensive set of regulation features and exhibit better performance, due to improved NIZK construction.

Platypus [45] offers regulation features of holding limits, receiving limits, and spending limits. It uses NIZK with low verification costs to increase transaction throughput. However, unlike other payment systems described in this section, Platypus is a centralized payment system rather than a distributed one.

PEReDi [37] is an account-based system. The regulation framework enforces spending and receiving limits for each transaction and imposes restrictions on the maximum amount a user can hold. Additionally, the system allows validators to trace funds and reveal details of certain payments. Users encrypt their transaction details with a threshold encryption scheme so that a quorum of validators can decrypt them. PEReDi needs a synchronous network and *correct* (i.e., non-Byzantine) validators must be *honest* (i.e., strictly follow the protocol and avoid seeking additional information). In contrast, most systems discussed here (as well as ours) are designed to tolerate *semi-honest* correct validators.

Concurrently with this work, the same authors have recently proposed PARScoin [38]. PARScoin enhances PEReDi by allowing asynchronous transactions yet still requires honest validators and faces limited throughput. PARScoin's protocol is similar to the one employed by UTT and our work, relying on Byzantine consistent broadcast. Our work differs from PARScoin in the NIZK construction which offers better performance, fewer assumptions on the validators and additional regulations features. Also, neither PRCash, PARScoin nor PEReDi have conducted latency tests to evaluate the maximum transaction throughput of their systems, they only provide benchmarks of their primitives (such as the time to prove or verify a transaction).

Other decentralised privacy-preserving payment systems adopt this auditability approach such as [3, 36, 13].

Specifications for a responsive and private payment system. In this paper, we introduce a new specification (inspired by the formalism recently proposed by Albouy et al. [1]), as existing ones do not address our particular problem. They describe either private payment systems that cannot be implemented in asynchronous networks [19], asynchronous payment systems that lack anonymity [26, 14, 31], or non-sequential specifications that do not align with how specifications are typically defined in the distributed systems literature [37, 38, 41].

3 Model

3.1 Participants and adversary

The system is composed of two types of participants:

- A set \mathcal{U} of U users of the payment system.
- A set \mathcal{V} of $N = 3f + 1$ validators.

Here f denotes the maximum number of validators that can go *Byzantine*, i.e., deviate from the protocol. A non-Byzantine (faithful following the protocol) is called *correct*. Correct validators may, however, be *semi-honest* [20]: a correct party may try to learn as much as possible from the messages they receive from other parties, which may involve colluding and pooling their views together. In contrast, *honest* validators that are not trying to learn any information. The correct participants are thus following the protocol but they may exchange information with any other participant (Byzantine or not) to learn as much as possible. As we shall see, Byzantine participants pose challenges to all correctness aspects of our system (safety, liveness and privacy), while semi-honest participants affect only privacy.

The adversary \mathcal{A} thus controls any number of users and all validators, but at most f validators can deviate from the protocol. \mathcal{A} can be seen as a hybrid adversary between honest and semi-honest. We assume that \mathcal{A} is *static* and *network-ignorant*: the set of Byzantine participants controlled by \mathcal{A} is chosen a priori, and \mathcal{A} has no information about message delay between the validators and the other participants. It can delay the messages but must eventually deliver them. Conventionally, the adversary \mathcal{A} is computationally bounded. More precisely, \mathcal{A} is *probabilistic polynomial-time* (PPT).

Every participant is provided with a pair of distinct public and secret addresses denoted *apk* and *ask*. These addresses are used to identify users and validators. The public addresses of the validators are known to every participant. Otherwise, a user only needs to know the identifiers of the users she engages in transfers with.

3.2 Network and communications

Concerning the network, we assume *asynchronous* but *reliable* communication. The participants can communicate via *anonymous*, *secure*, and *asynchronous network channels*. The channels do not modify or create messages. If a correct participants sends a message to a correct one, the message is eventually received, though we assume no bounds on the communication delays. The sender's identity is not known to the receiver, but the receiver can still respond to the sender. No other participant can tell who is the sender or the receiver, or tamper with the message content. We could use a network based on Syverson et al. [40] work. We consider that the latency distribution is the same for all the channels.

3.3 Cryptographic tools

Our protocol makes use of several cryptographic tools that we list below. Due to the space constraints, we only give informal definitions and refer to [24] for more details. Also, in the algorithms mentioned below, we omit the security parameters taken as inputs. Some primitives require a setup phase, which will be handled by a trusted third party \mathcal{T} . Note that these setups could be carried out via MPC [20] between the validators.

(k, N) -Threshold Blind Signature Scheme. Tuple of algorithms $\Pi_{\text{sig}} = (\text{SETUP}_{\text{sig}}, \text{BLIND}, \text{SIGN}, \text{UNBLIND}, \text{AGGREGATE}, \text{VERIFY}_{\text{sig}})$ allowing to divide a secret key sk in n fragments $[sk_i]_{i=1}^N$ between n signers, such that valid signatures from any subset of k signers

can be aggregated into a valid signature on behalf of the corresponding public key pk_{agg} . Each fragment sk_i has a corresponding public key pk_i so that a partial signature σ_i generated with sk_i can be verified with respect to pk_i . Moreover, the signers sign a blinding version \tilde{m} of a message m such that no information on m can be derived from \tilde{m} . A valid signature with respect to m can then be computed. Signature are unforgeable, which means that no PPT adversary can forge a valid partial signature σ_i of a message m that correctly verifies against pk_i without the knowledge of sk_i . As a blind signature, this should not also be possible for a signer to eventually make the link between the signature is has issued, and the final signature (after UNBLIND and AGGREGATE).

- $\text{SETUP}_{\text{Sig}}(k, N) \rightarrow ([sk_i]_{i=1}^N, [pk_i]_{i=1}^N, pk_{agg})$: Randomized algorithm run by a trusted party. Takes as input threshold parameters (k, N) and returns a list of N signing keys $[sk_i]_{i=1}^N$ and one corresponding verification key pk_{agg} .
- $\text{BLIND}(m, b) \rightarrow (\tilde{m})$: Takes as input a message m and a blinding factor b , returns a blinded message \tilde{m} .
- $\text{SIGN}(\tilde{m}, sk_i) \rightarrow \tilde{\sigma}_i$: Takes as input a blinded message \tilde{m} and a secret key sk_i , and returns a partial blinded signature $\tilde{\sigma}_i$.
- $\text{UNBLIND}(\tilde{m}, \tilde{\sigma}_i, b) \rightarrow \sigma_i$: Takes as input a blinded message \tilde{m} , the blinding factor b used to blind m and the corresponding partial blinded signature $\tilde{\sigma}_i$, and returns a valid partial signature σ_i for the original message m .
- $\text{AGGREGATE}([\sigma_i]_{i=1}^k) \rightarrow \sigma$: Takes as input a list of k signatures $[\sigma_i]_{i=1}^k$ and produce a signature σ such that σ is a valid signature on behalf of the public key pk_{agg} if and only if all σ_i are valid signatures for distinct public keys pk_i .
- $\text{VERIFY}_{\text{Sig}}(m, \sigma, pk) \rightarrow b$: Takes as input a partial or aggregated signature σ of a message m and returns a bit b of value 1 if σ is valid with respect to pk and 0 otherwise. The signature can be a partial or aggregated signature.

Collision-Resistant and Preimage-Resistant Pseudorandom Function Family [24]. A family of functions $\text{PRF} = \{\text{PRF}_s : \{0, 1\}^* \rightarrow \{0, 1\}^{O(|s|)}\}_s$, where s denotes a seed, computationally indistinguishable from a random function family. *Collision-Resistance* means here that it is computationally infeasible to find couples $(s, x) \neq (s', x')$ such that $\text{PRF}_s(x) = \text{PRF}_{s'}(x')$. *Preimage-Resistance* means that it is computationally infeasible, given y , to find (s, x) such that $\text{PRF}_s(x) = y$.

Non-Interactive Zero-Knowledge Proof (NIZK) [24]. A tuple of algorithms $\Pi_{\text{proof}} = (\text{SETUP}_{\text{NIZK}}, \text{PROVE}, \text{VERIFY}_{\text{NIZK}})$ allowing a *prover* to prove to a *verifier* that, given a statement defined by an NP relation $\mathcal{R}(a, b)$ and an instance `public_input`, she knows a witness `private_input` such that $\mathcal{R}(\text{public_input}, \text{private_input})$.

- $\text{SETUP}_{\text{NIZK}}(\mathcal{R}) \rightarrow \text{pp}_{\text{NIZK}}$: Randomised algorithm run by a trusted party. Takes as input a relation \mathcal{R} and outputs public parameters pp_{NIZK} (also known as common reference string). These public parameters are taken as input by the two following algorithms PROVE and VERIFY , but we omit it to lighten the notation.
- $\text{PROVE}(\text{public_input}, \text{private_input}, \text{pp}_{\text{NIZK}}) \rightarrow \pi$: Randomised algorithm. Takes as inputs instance and witness. Outputs a proof π such that :
 $\mathcal{R}(\text{public_input}, \text{private_input})$.
- $\text{VERIFY}_{\text{NIZK}}(\text{public_input}, \text{pp}_{\text{NIZK}}, \pi) \rightarrow b$: Takes as input an instance `public_input` and a proof π . Outputs 1 if the proof is valid and 0 otherwise.

Incremental commitment scheme. A tuple of algorithms $\Pi_{com} = (\text{COM}, \text{INCR})$ that allows us, given a sequence of messages $[m_i]_{i=1}^n$ and a sequence of randomness $[r_i]_{i=1}^n$, to produce a commitment c . The commitment is *hiding*, i.e., no information on $[m_i]_{i=1}^n$ can be derived from c without prior knowledge on $[r_i]_{i=1}^n$. The commitment is also *binding*, meaning that given a sequence of couples $[(m_i, r_i)]_{i=1}^n$ that commits to a value c , it should be computationally infeasible for a PPT adversary to compute $[(m'_i, r'_i)]_{i=1}^n$ such that $[(m'_i, r'_i)]_{i=1}^n \neq [(m_i, r_i)]_{i=1}^n$ also commits to c . The commitment scheme is also *incremental*: let us consider a sequence $[(m_i, r_i)]_{i=1}^n$, its commitment c and a couple (m', r') . Given the knowledge of c, m' and r' , it is possible to compute c' that commits to $[(m_1, r_1), (m_2, r_2), \dots, (m_n, r_n), (m', r')]$:

- $\text{COM}([(m_i, r_i)]_{i=1}^n) \rightarrow c$ takes as input a sequence of messages and randomness and returns the corresponding commitment. COM is hiding and binding.
- $\text{INCR}(c, m', r') \rightarrow c'$ takes as input a commitment c of a sequence $[(m_i, r_i)]_{i=1}^n$, a new message m' and a randomness r' . It returns c' , the commitment to the sequence $[(m_1, r_1), (m_2, r_2), \dots, (m_n, r_n), (m', r')]$.

4 Fully private asset transfer (FPAT)

State and interface. At the abstract level, the state of the FPAT object is represented as an array of U positive integer values $[v_k]_{k=1}^U$, one for each user, interpreted as the current balances of the users' accounts. Let $[v_k^{init}]_{k=1}^U$ be the *initial state* of the object. The object exports two operations: *transfer* and *balance*. Assuming that a user u invokes an operation, they are defined as follows:

- $\text{transfer}_u(v, w)/r$ takes as inputs a value v and a user identifier w . It transfers the amount v from u to w : updates a state $[v_k]_{k=1}^U$ to the state $[v'_k]_{k=1}^U$ where $v'_k = v_k + v$ if $k = w \wedge u \neq w$, $v'_k = v_k - v$ if $k = u \wedge u \neq w$ and $v'_k = v_k$ otherwise. It returns a binary response $r = \text{confirm}$ if it succeeds and $r = \text{fail}$ otherwise.
- $\text{balance}_u()/v_u$ takes no inputs and returns the value v_u stored at location u of the state $[v_k]_{k=1}^U$.

Let \mathcal{O} be a set of FPAT *operations* – invocations of *transfer* and *balances* provided with matching responses, each associated with a distinct user. Let $\text{transfer}_u(*, *)/C$ (resp., $\text{transfer}_u(*, *)/F$) denotes a transfer operation invoked by a user u that returns *confirm* (resp. *fail*). For each user u , we define a function $\text{total}_u(\mathcal{O})$ as follows:

$$\text{total}_u(\mathcal{O}) = v_u^{init} + \sum_{\text{transfer}_*(v, u)/C \in \mathcal{O}} v - \sum_{\text{transfer}_u(v, *)/C \in \mathcal{O}} v$$

$\text{total}_u(\mathcal{O})$ is thus defined as the current *balance* of u after all successful transfers in \mathcal{O} complete: the initial amount owned by u , minus all the funds sent by u plus all the funds received by u in the set of operation \mathcal{O} .

A *sequential history* S (of FPAT) is a totally ordered set of FPAT operations, let \prec_S be this order. Let $S|_u$ denotes the sub-sequence of S consisting of the events of user u . We say that S is *legal* if:

1. $\forall o = \text{transfer}_u(v, w)/C \in S \quad v \leq \text{total}_u(\{o' \in S : o' \prec_S o\})$
2. $\forall o = \text{balance}_u()/v \in S \quad v \leq \text{total}_u(\{o' \in S : o' \prec_S o\})$
3. If an operation $\text{balance}_u()$ returning v_1 directly precedes a $\text{transfer}_u(v_2, w)$ operation in $S|_u$, and $v_2 \leq v_1$, then the transfer operation cannot return fail.

Histories and serializations. Consider an *execution* of a FPAT algorithm: a sequence of *events* produced by the algorithm, such as invocations and responses of FPAT operations, sending and receiving messages, etc. A *local history* L_i of a user i is the sequence of *operations* invoked by i in that execution. We assume that correct users are *well-formed* – they never invoke a new operation before the previous one returns, and thus if u is correct, then L_u is sequential. In case the last operation of L_u is *incomplete* (not followed by a response), we can add any matching response and get a *completion* of L_u . Now a *history* H is a vector $(L_1, L_2, \dots, L_{|\mathcal{U}|})$ of local histories, one for each user. Notice that if the history is produced by an execution of a FPAT algorithm, only the local histories of correct users are of interest for us: a Byzantine user is not obliged to follow the protocol.

A sequential (FPAT) history S is a *serialization* of a history $H = (L_1, L_2, \dots, L_{|\mathcal{U}|})$ if for each *correct* user u , there exists \hat{L}_u , a completion of L_u , such that $S|_u = \hat{L}_u$ (\prec_S respects the local order $\prec_{\hat{L}_u}$). S can be seen as a *global interpretation* of the local histories of correct users. Notice that we allow any operations to be executed by Byzantine users in S , as long as it “makes sense” to the correct ones.

FPAT-Safety. An implementation of a FPAT-object is *FPAT-Safe* if and only if, for any finite history of execution H it produces, there exists a serialization S of H which is legal.

FPAT-Liveness. Liveness ensures that (1) every operation invoked by a correct user eventually completes and (2) considering two correct users u and w , if w transfers money to u , then u eventually receives this money. Let us consider the existence of a global time during the execution of an implemented FPAT object. An implementation of a FPAT-object is *FPAT-Live* if:

1. All *transfer* and *balance* operations terminate for correct users.
2. Consider a correct user u . For each operation $transfer_*(*, u)$ completed during the execution at time t , there exists a time $t' \geq t$ such that any operation *balance* inserted at time $t'' \geq t'$ will return a value v such that:

$$v \geq \sum_{\substack{transfer_*(v_+, u) \\ \text{invoked by correct users} \\ \text{completed before time } t}} v_+ - \sum_{\substack{transfer_u(v_-, *) \\ \text{completed before time } t''}} v_-$$

FPAT-privacy. We capture the privacy guarantees of FPAT through a *distinguishing game* \mathcal{G}^{priv} described in details the full version of the paper. *FPAT-privacy* has important implications that can be expressed as a collection of properties we use here. Let ϵ be a negligible function and λ the security parameter. Consider a protocol execution, let H be its history and \mathcal{U}_{Byz} the set of users controlled by the adversary, among the \mathcal{U} users of the system. Let $o = transfer_u(v, w) \in H$ be any *transfer* operation such that u is honest. Then for each guess (u', v', w') made by an adversary with no prior knowledge, where u' is the payment sender, v' is the value, and w' is the payment recipient, the following properties hold:

1. **Sender-anonymity:** $\mathbb{P}(u' = u) \leq \frac{1}{|\mathcal{U} - \mathcal{U}_{Byz}|} + \epsilon(\lambda)$, i.e., the adversary only knows that the sender is not itself.
2. **Receiver-anonymity:** If w is honest, then $\mathbb{P}(w' = w) \leq \frac{1}{|\mathcal{U} - \mathcal{U}_{Byz}|} + \epsilon(\lambda)$, i.e., the adversary only knows that the receiver is not itself.
3. **Confidentiality:** If w is honest, then $\mathbb{P}(v' = v) \leq \epsilon(\lambda)$, i.e., the adversary cannot guess the amount of the transaction.

These three properties together constitute **full privacy**. Additionally, *FPAT-privacy*, as captured by the *distinguishing game* \mathcal{G}^{priv} , implies **unlinkability**: given two transfers, no adversary can determine whether the sender, receiver, or amount is the same in both transfers.

5 PaxPay Protocol

We overview our FPAT implementation below and delegate detailed algorithms and proofs of correctness to the full version of the paper.

Setup. Every user is identified by her public address apk that is derived from a secret address ask as follows: $apk = PRF_{ask}(0)$ ⁴

The protocol uses a $(2f + 1, N)$ -threshold blind signature scheme, as defined in Section 3.3. The secret signing keys $[sk_i]_{i=1}^N$ are held by each of the validators. The aggregated public key is denoted by pk_{agg} . The protocol will also make use of NIZK, as defined in Section 3.3. The setup algorithms for these primitives are again described in [11].

Coin structure. A *coin* is a tuple $\mathbf{c} = (v, apk, \rho)$ with:

- v the (integer) value of the coin.
- apk the public address of the owner of the coin.
- ρ the seed of the coin, from which the serial number is derived.

Coins are similar to unspent transactions (UTXO) in Bitcoin: to make a payment, a user “spends” *old* coins and creates *new* coins whose owners are the payment recipients.

Coin validity. Using a quorum of validators signatures (inspired by Byzantine Consistent Broadcast), we decide whether a coin is valid or not. In concrete terms, a coin is valid if it has been signed by $2f + 1$ validators. During a transfer, several old coins $[\mathbf{c}_i^{old}]_{i \in [1, n]}$ are spent to create new coins $[\mathbf{c}_j^{new}]_{j \in [1, m]}$. To make sure that a coin is not spent twice, a unique *serial number* is derived from the coin’s seed ρ as $sn = PRF_{ask}(\rho)$. Each validator maintains a list `snList` of the old coins’ serial numbers. The coins $[\mathbf{c}_i^{old}]_{i \in [1, n]}$ are considered spent when a quorum of $2f + 1$ validators have appended the corresponding serial numbers $[sn_i^{old}]_{i \in [1, n]}$ to their `snList`.

Intuitively, as the computation of sn_i^{old} is using the PRF and the secret address of the payer, no party can link sn_i^{old} to ρ_i^{old} or apk_i^{old} , and thus the sender-anonymity property is preserved. The full version of the paper details an algorithm that initializes the balances for the initial set of users, by creating coins for each of them. ⁵

Local variables. User storage consists of:

- apk/ask : its public/secret address pair
- $\text{coinList} = [(\mathbf{c}_i, \sigma_i)]_i$: the set of coins owned by the user associated with the aggregated signature of each coin.
- rhoList : the list of all seed ρ corresponding to coins received by the user

⁴ The address generation is detailed in [11] and shows the address generation process by sampling a random ask and deriving the apk .

⁵ As discussed in Section 8, depending on the use cases of the system, this algorithm might not be executed. Instead, users might freely join the system and call a `Mint` algorithm that provides them with new coins.

Validator storage consists of:

- sk : its secret signing key
- $snList$: the list of serial numbers identifying spent coins.
- $signedCoinList$: the list of all blinded coins signed by the replica.

Transfer. Assuming a user has enough funds, she can issue payments to several recipients $[apk_j^{new}]_{j=1}^m$ in one transfer by sending v_j^{new} to apk_j^{new} . The full algorithm is described in [11]. The user first chooses n (old) coins denoted $[c_i^{old}]_{i=1}^n$ and their signatures $[\sigma_i^{old}]_{i=1}^n$ from her $coinList$, where $c_i^{old} = (v_i^{old}, apk_i^{old}, \rho_i^{old})$, such that the total value of the old coins does not fall below the total value paid in the transfer: $\sum_{i=1}^n v_i^{old} \geq \sum_{j=1}^m v_j^{new}$. (Otherwise, the transfer operation returns fail.) The user then computes the serial number $[sn_i^{old}]_{i=1}^n$ of every old coin as follows: $sn_i^{old} = \text{PRF}_{ask_i^{old}}(\rho_i^{old})$.

The j new coins are then produced by the user according to the values and addresses to pay: $\forall j \in [1, m]$, $c_j^{new} = (v_j^{new}, apk_j^{new}, \rho_j^{new})$. ρ_j^{new} is derived as $\rho_j^{new} = \text{PRF}_{\rho_{seed}}(sn_1^{old} || \dots || sn_n^{old} || j)$, where ρ_{seed} is sampled randomly. This binds each new coin to the old coins spent for its creation. This way, we make sure that the validators signing c^{new} mark the same old coins as spent. Since the value of the chosen old coins might exceed the value of the new coins, the user also produces a *redeem coin* $c_{m+1}^{new} = (v_{m+1}^{new}, apk_{m+1}^{new}, \rho_{m+1}^{new})$, with $apk_{m+1}^{new} = apk$ and v_{m+1}^{new} the exceeding value. The coins $[c_j^{new}]_{j=1}^{m+1}$ are then *blinded*. The blinding of c_j^{new} is denoted $\tilde{c}_j^{new} = \text{BLIND}(c_j^{new}, b_j)$, where b_j is a randomly sampled blinding factor.

The sender of the payment then computes an NIZK with:

- **public_input** : $([sn_i^{old}]_{i=1}^n, [\tilde{c}_j^{new}]_{j=1}^{m+1})$
- **private_input** : $([c_i^{old}]_{i=1}^n, [\sigma_i^{old}]_{i=1}^n, [ask_i^{old}]_{i=1}^n, [c_j^{new}]_{j=1}^{m+1}, [b_j]_{j=1}^{m+1}, \rho_{seed})$

This NIZK $\pi_{transfer}$ proves the following relations:

1. Serial numbers $[sn_i^{old}]_{i=1}^n$ are correctly derived from the old coins $[c_i^{old}]_{i=1}^n$.
2. New coins $[c_j^{new}]_{j=1}^{m+1}$ are correctly derived from the old coins $[c_i^{old}]_{i=1}^n$.
3. Blinded coins $[\tilde{c}_j^{new}]_{j=1}^{m+1}$ are correctly derived from $[c_j^{new}]_{j=1}^{m+1}$ and $[b_j]_{j=1}^{m+1}$.
4. Signatures $[\sigma_i^{old}]_{i=1}^n$ of the coin $[c_i^{old}]_{i=1}^n$ are correct.
5. Private addresses $[ask_i^{old}]_{i=1}^n$ match the public address $[apk_i^{old}]_{i=1}^n$.
6. The sum of the values of the old coins $[c_i^{old}]_{i=1}^n$ equals the sum of the values of the new coins $[c_j^{new}]_{j=1}^{m+1}$.

The user can now send the NIZK $\pi_{transfer}$ along with the public inputs to all the validators. The detailed algorithm run by the validators is detailed in [11]. Once a validator receives the proof:

1. It checks that none of the $[sn_i^{old}]_{i=1}^n$ appears in its $snList$;
2. It checks that the proof π is correct;
3. If the last two conditions are fulfilled, it adds all $[sn_i^{old}]_{i=1}^n$ to $snList$, add all $[\tilde{c}_j^{new}]_{j=1}^{m+1}$ to $signedCoinList$, it signs all the coins $[\tilde{c}_j^{new}]_{j=1}^{m+1}$ and returns the blinded partial signatures.
4. If any of the previous conditions is not fulfilled, the validator checks if the blinded coins $[\tilde{c}_j^{new}]_{j=1}^{m+1}$ appear in $signedCoinList$. If they all do, then the validator still sends back the signatures. This is done because a Byzantine validator p might receive a transfer request from a user u and send it to other validators using a private channel. As a result, other validators might answer p before answering u , preventing u from receiving the signatures while the old coins are actually spent (their serial numbers would have been added to $snList$ already).

Once the user receives $2f + 1$ valid partial signatures for the coins $[\mathbf{c}_j^{new}]_{j=1}^{m+1}$ from $2f + 1$ validators, she can unblind and aggregate them to form $m + 1$ signatures $[\sigma_j]_{j=1}^{m+1}$ that are valid signatures for $[\mathbf{c}_j^{new}]_{j=1}^{m+1}$. She can now send each couple $(\mathbf{c}_j^{new}, \sigma_j^{new})$ to apk_j^{new} . The user apk_j^{new} only accepts the payment if $\rho_j^{new} \notin \text{rhoList}_{apk_j^{new}}$ and $\text{VERIFY}_{\text{Sig}}(\mathbf{c}_j^{new}, \sigma_j^{new}, pk_{agg}) = 1$.⁶

Double spending. Consider a coin \mathbf{c}^{old} that has been spent. Hence, a set \mathcal{V}_1 of $2f + 1$ validators have claimed to have added sn^{old} to their $snList$, and at least $f + 1$ out of them are correct. If a Byzantine user tries to spend \mathbf{c}^{old} again, it should collect confirmations from a set \mathcal{V}_2 of $2f + 1$ validators that will have to add sn^{old} in their $snList$ again. Since there are $3f + 1$ validators in total, \mathcal{V}_1 and \mathcal{V}_2 have at least one correct validator in common. It must refuse to sign the second transaction because sn^{old} is already in its $snList$, so double spending is prevented and FPAT-Safety is guaranteed. The full proof is presented in [11].

Example. Figure 1 depicts a transfer:

1. Alice wants to pay 3 recipients with 2 coins. She generates 3 blinded coins and a matching NIZK, and sends them to the validators.
2. Each validator checks the validity of the NIZK and that the old coins are not in its $snList$, and, if so, sign the blinded coins and add the serial numbers to their $snList$.
3. Alice receives the signatures.
4. For each coin: once $2f + 1$ blinded partial signatures are received, Alice unblinds and then aggregates them into one aggregated signature.
5. Alice sends the coin tuples and their signatures to their recipients.

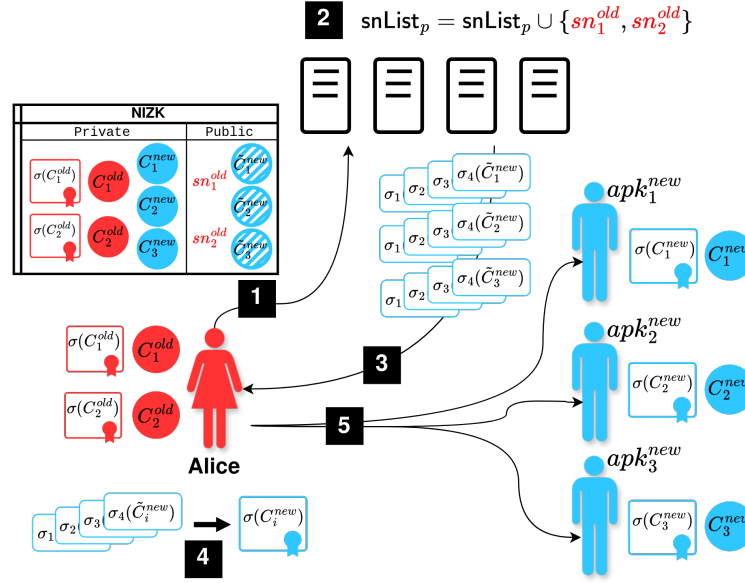
6 Regulatory enforcement

PaxPay is designed as a private payment system. As in PEReDi [37], PRCash [44] or Platypus [45], the protocol can be enhanced to be *regulation-compliant*. The use of succinct proofs limits the impact of regulatory enforcement on the system's performance. This section describes how to build such a regulatory enforcement on top of the protocol described above.

Our regulatory enforcement is achieved via *compliance coins*. Each user owns its unique compliance coin. The coin commits to some data related to the user and all his transactions. The user must attach the compliance coin to each of its transactions as an input and the validators will sign the updated compliance coin. Well-formedness of the updated compliance coin is proved through NIZK. As for a classical coin, the old compliance coin is spent by revealing its serial number. As a result, a user has one valid compliance coin at a time and this coin acts as an append-only tracing mechanism. We introduce the following regulation features, shown in Table 1:

- *Limited held amount per user:* Users cannot hold more than $V_{\text{held}}^{\text{max}}$.
- *Full asset tracing:* A trusted authority (centralized or decentralized) can reveal the content of transaction to trace back assets and users activities.
- *Limited spendable amount per tx:* Users cannot spend more than $V_{\text{sent}}^{\text{max}}$ in one transfer.
- *Limited spendable amount in total:* Users cannot spend more than $V_{\text{total}}^{\text{max}}$ in total.

⁶ For convenience, we allow the **TRANSFER** algorithm to send money to several recipients in a single call, in contrast with the specification, which allows only one recipient at a time. As explained in the security proof in [11], this more generic transfer can be seen as a batch of successive transfer calls, each destined to a single user.



■ **Figure 1** PaxPay example: Alice pays 3 recipients with 2 old coins.

- *Sanctions*: Users in the sanction list cannot send or receive funds.
- *Provable transaction history*: A user can prove arbitrary statements about her transactions. For instance, reveal all of her transactions and prove that the revealed set is complete (no transactions are missing).

Our regulatory construction enforces *Limited spendable amount per tx*, *Limited spendable amount in total*, *Sanction list* and *Provable transaction history*. V_{sent}^{max} , V_{total}^{max} and $SancList$ are public parameters chosen and potentially updated by the regulator. The sanction list $SancList$ is represented as a sorted Merkle tree, which allows efficient proof of non-membership.

Compliance coin. The compliance coin of a user u is defined as a tuple $cc_u = (apk, \rho, v, com)$ where:

- apk is u 's sole public address
- ρ the coin seed, as for a normal coin
- v is the total amount of money sent so far by u
- com is the commitment to the list of all the transfer done by u so far. For each coin created, it commits the tuple (apk, v) with apk the public key of the receiver and v the value sent. It uses the incremental commitment introduced in Section 3.3

com allows to trace u 's activities and allows u to prove additional statements on her transaction history in case of an update of the regulation.

Registration. Adding a user registration process is essential for regulatory enforcement, to comply with the *know-your-customer* (KYC) and *customer-due-diligence* (CDD) checks. Additionally, it ensures the uniqueness of the compliance coin for each physical user.

To this end, the validator storage is provided with a new list **registeredList**, the list of the enrolled users associated with their public key. **registeredList** is used to make sure that a physical user can only register once on the system. The registration process is as follows for the user u :

1. Generate a random ρ and computes $cc_u = (apk_u, \rho, 0, 0)$ and its blinding \tilde{cc}_u (with blinding factor b).
 2. Computes an NIZK $\pi_{register}$ that takes (\tilde{cc}_u, apk_u) as public input and (ask_u, cc_u, b) as private input. It proves that:
 - a. $\tilde{cc}_u = \text{BLIND}((apk_u, \rho, 0, 0), b)$.
 - b. ask_u and apk_u form a correct secret/public address pair.
 3. Send to each validator $\pi_{register}$ and (\tilde{cc}_u, apk_u) .
 4. Upon correct KYC and CDD proving u 's identity, each validator check if $\pi_{register}$ is correct and that: $(u, *) \notin \text{registeredList}$ (u has registered no public address yet). If so, they sign \tilde{cc}_u , send it back to u , and add (u, apk_u) to **registeredList**.
 5. Once u has received $2f + 1$ partial signature, she can unblind and aggregates the partial signatures to form a valid signature for cc_u .
- The detailed algorithm is described in [11].

Transfer. Section 5 described how a transfer is handled by a user u , by spending n old coins $[c_i^{old}]_{i=1}^n$ and creating $m + 1$ new coins $[c_j^{new}]_{j=1}^{m+1}$ to pay m recipients (m coins are created to pay the recipient and a redeem coin is also created with index $m + 1$ to get the excess of money back to u). Now the user also has to update a compliance coin at each payment and prove that their transfer is compliant. The following additional steps are thus required. She generates m random values: $[r_j]_{j=1}^m$. The user provides an additional proof π_{comply} . This proof takes the following public inputs:

- \tilde{cc}_u^{new} the blinding of the new compliance coin
- sn_{cc}^{old} the serial number of the old compliance
- V_{sent}^{max} the maximum amount that can be transferred in one transfer
- V_{total}^{max} the total amount of money that u can transfer in her use of the system
- $SancList$ the list of addresses under sanctions

π_{comply} takes the following as private input:

- $cc^{old} = (apk_{cc}^{old}, \rho_{cc}^{old}, v_{cc}^{old}, com^{old})$ the old compliance coin
- $cc^{new} = (apk_{cc}^{new}, \rho_{cc}^{new}, v_{cc}^{new}, com^{new})$ the new compliance coin
- ask_{cc}^{old} the secret address corresponding to apk_{cc}^{old}

π_{comply} should verify the following clauses:

1. $\forall i \in [1, n], apk_i^{old} = apk_{cc}^{old}$
2. $apk_{cc}^{old} = \text{PRF}_{ask_{cc}^{old}}(0)$
3. $apk_{m+1}^{new} = apk_{cc}^{old}$
4. $apk_{cc}^{new} = apk_{cc}^{old}$
5. $\rho_{cc}^{new} = \text{PRF}_{\rho_{seed}}(sn_1^{old} || \dots || sn_n^{old} || m + 2)$
6. $v_{cc}^{new} = v_{cc}^{old} + \sum_{1 \leq j \leq m} v_j^{new}$
7. $com_0 = com^{old} \wedge \forall j \in [1, m], com_j = \text{INCR}(com_{j-1}, apk_j^{new}, v_j^{new}, r_j) \wedge com^{new} = com_m$
8. $\sum_{1 \leq j \leq m} v_j^{new} \leq V_{sent}^{max}$
9. $v_{cc}^{new} \leq V_{total}^{max}$
10. $apk_{cc}^{old} \notin SancList$
11. $\forall j \in [1, m], apk_j^{new} \notin SancList$

Conditions (1) ensures that all the input coins belong to the owner of the input compliance coin. Condition (2) ensures that u , the user that invokes *transfer*, is the owner of the compliance coin. Condition (3) ensures that the redeem coin will belong to u . Conditions (4), (5), (6) and (7) ensure that the new compliance coin is well-formed. Condition (8), (9) and (10) ensure that the transfer is compliant. (8) verifies that the amount of the transfer does not exceed that maximum allowed for a single transfer. (9) verifies that the total amount of

money sent by u does not exceed that maximum amount. (10) verifies that the sender is not a user under sanctions. (11) verifies that none of the receivers of the transfer is under sanction.

The user storage is also augmented with a list `comList`, in which she stores the tuples (apk, v, r) that she has used at each transfer. At the end of the transfer described above, she thus append to `comList` the following: $\forall j, j \in [1, m], (apk_j^{new}, v_j^{new}, r_j)$. She can later use `comList` and her compliance coin to prove any statement about her transaction.

Upon receiving correct proofs $(\pi_{comply}, \pi_{transfer})$ and their public parameters, the validators execute the same algorithm as for a non regulated transfer, except that they verify both π_{comply} and $\pi_{transfer}$. They can then sign the news coins and then send the signature back. The detailed algorithm is described in [11].

7 PaxPay implementation and performance

PaxPay implementation in Golang is available on Github. We developed with two implementations, with and without regulation support. Below we overview the cryptographic tools we employed and report on the performance of our protocols.

7.1 Cryptographic building blocks

NIZK. NIZK proofs are implemented using the Groth16 [25] protocol. Groth16 allows for constructing succinct non-interactive arguments of knowledge (SNARKs) with constant verification complexity, regardless of the complexity of the relation \mathcal{R} it attests to. The relation \mathcal{R} is represented by an arithmetic circuit. We instantiate this scheme with the BW6-761 [28] curve, designed to cycle with the BL12-377 [28] curve, allowing for efficient pairing verification over BL12-377 within the SNARK.

In the regulated version, each transfer requires two NIZK. For performance reasons, these two proofs $\pi_{transfer}$ and π_{comply} are merged into a single proof with a unique circuit. The proof of non-membership in the merkle tree is implemented as follows: the sanction list has the following form: $SancList = [s_1, s_2, \dots, s_\xi]$. $SancList$ is supposed to be published sorted by the regulator. To verify the non-inclusion of an address apk in the list, we prove that there exist s_i and s_j such that:

- s_i and s_j are elements of $SancList$.
- $j = i + 1$ (s_i and s_j are consecutive elements of $SancList$).
- $s_i < apk < s_j$ (s_1 and s_ξ are set to $0x00..00$ and $0xFF..FF$ to make sure this constraint can always be fulfilled).

Hash functions and PRF. To optimize the computational cost of the proving algorithm in SNARK, we rely on MiMC [2]. MiMC is an arithmetization-oriented, collision-resistant and preimage-resistant hash function that is efficiently represented as arithmetic circuits. We use it to derive the seeds ρ of created coins from the serial numbers sn of spent coins. The pseudorandom function family $\text{PRF}_x()$ is derived from a MiMC function H as $\text{PRF}_x(m) := H(m||x)$.⁷

⁷ The natural way to implement a PRF from a hash function is to consider $\text{PRF}_x := H(m||x)$, however, replacing concatenation by an addition in the field does not compromise security in our case while reducing the computational cost.

Incremental commitment scheme. The incremental commitment scheme is also derived from the MiMC function H . Given a commitment com^{old} and a message m to be added to the commitment with a randomness r , $INCR(com^{old}, m, r)$ is defined as follows: $INCR(com^{old}, m, r) = H(com^{old} || m || r)$. The commitment function COM is then an iteration of increments $INCR$ over the sequence of messages and randomnesses provided.

(k, N)-threshold blind signature. Our signature scheme is based on a slightly modified version of the Coconut scheme [39], which itself is a variant of the Pointcheval and Sanders signature scheme [35]. Indeed, certain checks originally performed using a sigma protocol in Coconut are instead handled in the Groth16 proof in our implementation, reducing the overall verification complexity. We instantiate this scheme with the BLS12-377 curve so that the verification algorithm, which involves pairings, is efficient in the Groth16 proof. More details about the construction are given in [11].

7.2 Performance analysis

We now report on our comparative performance analysis.

NIZK Benchmark. We benchmarked our NIZK, implemented using a Groth16 SNARK, on an Intel i7 @ 2.6 GHz CPU running Ubuntu 22.04. The results are summarized in Table 2. Each value is based on the average of 10 proving and verification runs. A comparison with related protocols is provided in Table 1. The data in this table was gathered from the respective research papers, as the source codes of these protocols were not always publicly available. However, the CPU used for our benchmark has comparable performance to the machines reported in those works: they all use Intel Core i7 CPUs except for Zef and UTT. For Zef and UTT, AWS EC2 instances were used for the benchmark. These two types of instance have similar CPUs (Intel Xeon Platinum 8175 and 8124).

PaxPay has two implementations: with and without the regulatory feature. The benchmark was conducted on both. The regulated version, running on a single core, serves as a reference point. The results indicate that our SNARK verification time is notably short, taking only 5.6 ms. As shown in Table 1, this verification time is between 8 and 100 times shorter than the NIZK verification times reported in other works (except for Zcash, which takes approximately the same time but suffers from other major issues discussed in Section 8). The verification time is a key metric since slow verification limits the transaction throughput of the system.

This fast verification comes at the cost of relatively slow proving, requiring nearly 7 s to generate a proof for a transfer on a single core. However, with 6 cores, the proving time drops to 2 s.

■ **Table 2** SNARK Proving and Verification Times on an Intel i7 2.6 GHz CPU, running Ubuntu 22.04.

	Regulated		Not Regulated	
	1 Core	6 Cores	1 Core	6 Cores
Proving time (ms)	6959 (± 13)	2001 (± 13)	3080 (± 6)	882 (± 5)
Verification time (ms)	5.61 (± 0.20)	5.38 (± 0.28)	5.22 (± 0.15)	4.71 (± 0.85)

Latency test and transaction throughput. Apart from Zcash, which is already deployed in real-world use, the only payment systems in the related work to report on their latency and throughput are Zef [6] and UTT [41]. The performance analyses of UTT and Zef were performed on different types of AWS EC2 instances (c5.4xlarge for UTT and m5.8xlarge for Zef). Since the micro-benchmarks show better performance for PaxPay on UTT’s EC2 instance than on the one used by Zef, we chose to use the same instance type as in the Zef paper for consistency.⁸

We thus run our PaxPay implementation on multiple m5.8xlarge instances, assigning one validator per instance. For each version of PaxPay (regulated and non-regulated), we conducted two tests: one with each validator running on a single CPU core, and another with each validator using 16 CPU cores. The results are presented in Figures 2a and 2b. We consider that the system support a given throughput if the corresponding transaction latency is smaller than 500 ms. Our regulated implementation achieves the throughput of 100 **tx/s** (transactions per second) in the single-core setup and 925 **tx/s** in the 16-core setup. Under the same conditions, Zef processes 5 **tx/s** and 88 **tx/s**, respectively. UTT processes 235 **tx/s** using 16 cores.⁹ For the non-regulated PaxPay implementation, the throughput reaches 115 **tx/s** in the single-core setup and 1200 **tx/s** in the 16-core setup.

These promising results can be attributed to the low verification time of the SNARK presented in the full version of the paper. In the reference setting (regulated implementation with 1-core validators), the verification of the SNARK takes 9 **ms** in our implementation while it takes 142 **ms** in Zef with the same setting [11].

The difference in performance between the regulated and non-regulated versions is due to the additional public parameters in the regulated SNARK. This requires validators to perform more point multiplications on the elliptic curve used in the Groth16 construction.

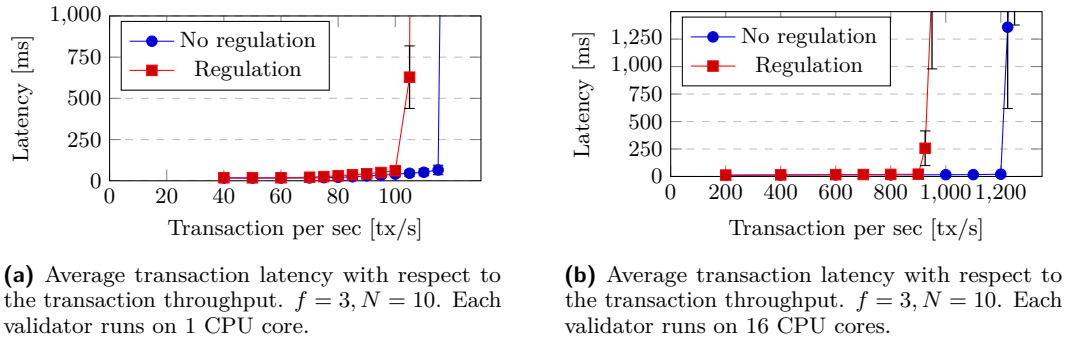
To improve scalability, validators can be *sharded* (run on multiple parallel machines) as in Zef and UTT. The same approach could be applied to our system, though it has not been implemented. A load balancer could efficiently distribute transfer requests among the validators by assigning a specific range of serial numbers to each shard of a validator. Upon receiving a request, the load balancer would determine the appropriate shard to handle it based on these predefined ranges. Our experiments thus show that PaxPay is scalable. The throughput can be improved even further by allocating more computational power to the validators, either by using more CPU cores per validator or by distributing the workload across additional machines.

8 Discussion and comparison with related protocols

Comparative analysis. Table 1 provides a comparison between PaxPay and the related protocols. Lelantus appears to outperform Monero across all metrics in the table and offers larger anonymity sets. We therefore decided to omit Monero from the table. Similarly, Platypus was not included in the table, as we focus exclusively on decentralized payment systems. Details and explanations regarding the data presented in the table can be found in [11].

⁸ We were unable to build and test the UTT code due to dependencies that broke over time. Since the UTT authors chose an EC2 instance type on which PaxPay performs better compared to the m5.8xlarge instance, we consider the results reported in the UTT paper to be suitable for direct comparison.

⁹ As stated in Zef [6], each authority is distributed over multiple shards, each running in a single core. The 16-core result is extrapolated from the linear relationship between throughput and the number of shards per validator, as shown in their paper.



■ **Figure 2** Latency tests on AWS machines within the same region.

Privacy. As shown in Table 1, PaxPay provides the strongest privacy guarantees, comparable to those provided by Zcash, whereas some protocols offer reduced privacy guarantees. This is especially the case for Zef and PRCash.

To ensure full privacy, it is important to hide coin details when creating them with blind signatures. Yet, when the signatures are verified by validators, some users may still learn information on the payment.¹⁰ UTT circumvents this problem by using a re-randomizable signature scheme (where both the signature and the signed message can be randomized), but this solution increases the workload for verifying transactions. Our solution is to verify the signature within the NIZK, which also decreases the validators workload.

Model assumptions. PaxPay is responsive, meaning it can be implemented in an asynchronous network. Among the related protocols only Zef, UTT, and PARScoin also provide responsiveness. All the four protocols are built on some form of Byzantine consistent broadcast. However, PARScoin requires correct validators (those who follow the protocol) to be honest, ensuring they do not share any information with the adversary. In contrast, PaxPay tolerates correct validators being semi-honest, meaning they can share any information they receive or transmit during the protocol's execution without compromising the system's privacy guarantees.

Regulation. As displayed in Table 1, PaxPay does not support the *Limited held amount per user* feature. PEReDi [37] implements this feature, requiring synchrony between sender and receiver. PARScoin [38] claims a different approach: the sender initiates a transaction to reduce their balance, and the receiver can asynchronously claim funds later, provided their balance remains within the imposed limit. However, this method merely limits the amount a user can hold at a given time, as receivers can bypass it by spending excess funds before claiming the pending funds they have received. In practice, this mechanism has the same effect as setting a limit on the amount that can be spent in a single transaction. We believe that implementing a *Limited held amount per user* feature is particularly challenging in a private and asynchronous payment system. Such a feature would require that transfers atomically change the balances of both the sender and the receiver, to prove that receiver

¹⁰ Suppose that, as in Zef [6], we resort to only using blind signatures. Let a user A creates a coin c for a user B during a transfer o_1 . A requests the validators to sign the blinded coin \tilde{c} , and then sends the aggregated signature to B . B now spends c during a transfer o_2 , and in process reveals c and the (randomized) signature. Even though validators cannot reveal that the coin spent in o_2 was created in o_1 , A knows c and, thus, can deduce the sender of o_2 and a lower bound on the amount spent during o_2 .

balance does not exceed a certain amount after executing the transfer. However, to affect the balance of a user, privacy typically requires this user to interact with another user or the validators, in order to provide some secret data known only to him (such as coin data or account commitments). Yet, asynchrony precludes live protocols from relying on interaction between senders and receivers. Indeed, in an asynchronous setting, the sender who gets no response from the receiver cannot distinguish whether this absence of response is caused by network delays or by the receiver refusing to respond. Implementing *Limited Held Amount per User* seems hard when ensuring both asynchrony and privacy.

PaxPay also avoids *Full Asset Tracing* due to privacy risks. Instead, it allows users to generate privacy-preserving proofs about their transactions. Let us give a real life example. A user can prove that her account received funds from multiple addresses, each transfer being under a certain threshold, demonstrating legitimate income patterns. This proof could demonstrate that the income pattern of a business is legitimate and consistent with regular business operations, precluding money laundering (which might involve receiving a large sum from a single address).

Performance. As mentioned in Section 7.2, PaxPay outperforms all other protocols in terms of transaction throughput. This performance advantage is directly attributed to its fast NIZK verification, which surpasses all existing protocols. The only exception is Zcash, which also benefits from a fast NIZK verification but is constrained by the limitations of its underlying consensus protocol.

Additionally, the system’s scalability, that is achieved by increasing the computational power allocated to each validator, further strengthens its practical applicability. The Visa payment system processed an average of 7,388 transactions per second in 2024 [43], which could be supported by PaxPay by distributing each validator across 8 m5.8xlarge EC2 instances each¹¹.

As mentioned in Section 7.2, this high throughput comes at the cost of a relatively slow transfer proving, requiring between 2 and 7 seconds depending on the user computational power to prove a transaction. We believe that this delay is acceptable from the user’s perspective, especially since multiple payments can be aggregated into a single transfer. As a result, even if a user needs to make numerous payments, she can prove them all within a single transaction, requiring only one proof. It is worth noting that the statement being proven in such a case would be larger than for a proof of a single payment, which would still increase the proving time, though less significantly compared to generating a separate proof for each payment.

Use cases. Our system is versatile and adaptable to various use cases. It can function as a standalone payment system, a CBDC, or a scaling solution for an existing blockchain (a so-called Layer 2). It can be enriched with a Mint and Redeem operations that user would call to put or retrieve money in/from the system. Depending on the use case, Mint can be equipped with a proof showing that the user has the right to mint a new coin. For example, if PaxPay were used as a scaling solution for Bitcoin, the Mint operation would take some SPV (Simplified payment verification) [33] as input to prove to the validators that some value has been locked on Bitcoin. The implementation of the Mint and Redeem operations

¹¹ This estimation assumes a direct linear relationship between throughput and the computational power of the validators, as shown in Zef [6].

would vary depending on whether the system is deployed as a CBDC or as a Layer 2 solution for a blockchain. The performance of PaxPay (Section 7.2) and its scaling capacity are key factors that enable our system to effectively address these use cases.

Bad user behavior. If a user tries to pass two transactions spending the same input coin, her account might get blocked because she will not get enough signatures for either of the transactions. In this case, consensus is required to unlock the user’s account [26, 42].

Future work. Our protocol can be reused with some slight modifications to implement the same functionalities as in Zexe [9], hence not only offering a payment system but *private computation*. We can also considerably improve performance of our system by introducing sharding in the implementation, or improve the usability by moving to a different type of NIZK with a universal trusted setup.

9 Conclusion

In this paper, we consider the problem of fully-private asset transfer (FPAT). We propose PaxPay, an asynchronous FPAT protocol, prove its correctness, detail its implementation in Golang, and analyze its performance. PaxPay leverages succinct non-interactive zero-knowledge proofs and threshold blind signatures. Our system demonstrates significant improvements over existing systems, processing transactions at a high rate, while maintaining full anonymity for the users and responsiveness. PaxPay also ensures regulatory enforcement, including some features that no other private payment system has managed to provide so far. The flexibility offered by succinct NIZKs opens avenues for incorporating other regulatory compliance features with minimal impact on the system’s performance. These elements make PaxPay suitable for a wide range of applications, from a scaling solution of blockchain to a standalone private payment system that can be used as a technical layer for CBDCs.

References

- 1 Timoth   Albouy, Emmanuelle Anceaume, Davide Frey, Mathieu Gestin, Arthur Rauch, Michel Raynal, and Fran  ois Taiani. Asynchronous BFT asset transfer: Quasi-anonymous, light, and consensus-free, 2024. doi:10.48550/arXiv.2405.18072.
- 2 Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. Mimc: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I*, volume 10031 of *Lecture Notes in Computer Science*, pages 191–219, 2016. doi:10.1007/978-3-662-53887-6_7.
- 3 Elli Androulaki, Jan Camenisch, Angelo De Caro, Maria Dubovitskaya, Kaoutar Elkhayaoui, and Bj  rn Tackmann. Privacy-preserving auditable token payments in a permissioned blockchain system. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, AFT ’20, pages 255–267, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3419614.3423259.
- 4 European Central Bank. Report on a digital euro, 2020. URL: https://www.ecb.europa.eu/pub/pdf/other/Report_on_a_digital_euro~4d7268b458.en.pdf.
- 5 Mathieu Baudet, George Danezis, and Alberto Sonnino. Fastpay: High-performance byzantine fault tolerant settlement. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 163–177, 2020. doi:10.1145/3419614.3423249.

- 6 Mathieu Baudet, Alberto Sonnino, Mahimna Kelkar, and George Danezis. Zef: Low-latency, scalable, private payments. In Bart P. Knijnenburg and Panos Papadimitratos, editors, *Proceedings of the 22nd Workshop on Privacy in the Electronic Society, WPES 2023, Copenhagen, Denmark, 26 November 2023*, pages 1–16. ACM, 2023. doi:10.1145/3603216.3624952.
- 7 Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474, 2014. doi:10.1109/SP.2014.36.
- 8 Gautam Botrel, Thomas Piellard, Youssef El Housni, Ivo Kubjas, and Arya Tabaie. Consensus/gnark: v0.11.0, September 2024.
- 9 Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. Zexe: Enabling decentralized private computation. *IACR Cryptol. ePrint Arch.*, page 962, 2018. URL: <https://eprint.iacr.org/2018/962>.
- 10 Gabriel Bracha. Asynchronous Byzantine agreement protocols. *Information and Computation*, 75(2):130–143, 1987. doi:10.1016/0890-5401(87)90054-X.
- 11 Maxence Brugerres, Victor Languille, Petr Kuznetsov, and Hamza Zarfaoui. Fast, private and regulated payments in asynchronous networks. Cryptology ePrint Archive, Paper 2025/098, 2025. URL: <https://eprint.iacr.org/2025/098>.
- 12 Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. The weakest failure detector for solving consensus. *Journal of the ACM (JACM)*, 43(4):685–722, 1996. doi:10.1145/234533.234549.
- 13 Panagiotis Chatzigiannis and Foteini Baldimtsi. Miniledger: Compact-sized anonymous and auditable distributed payments. In *European Symposium on Research in Computer Security*, pages 407–429. Springer, 2021. doi:10.1007/978-3-030-88418-5_20.
- 14 Daniel Collins, Rachid Guerraoui, Jovan Komatovic, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, Andrei Tonkikh, and Athanasios Xytkis. Online payments by merely broadcasting messages. In *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2020, Valencia, Spain, June 29 - July 2, 2020*, pages 26–38. IEEE, 2020. doi:10.1109/DSN48063.2020.00023.
- 15 Ronald Cramer. *Modular design of secure yet practical cryptographic protocols*. PhD thesis, U. of Amsterdam, 1996.
- 16 Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for byzantine agreement. *Journal of the ACM (JACM)*, 32(1):191–204, 1985. doi:10.1145/2455.214112.
- 17 Danny Dolev and H. Raymond Strong. Authenticated algorithms for byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983. doi:10.1137/0212045.
- 18 Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony (preliminary version). In Tiko Kameda, Jayadev Misra, Joseph G. Peters, and Nicola Santoro, editors, *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing, Vancouver, B. C., Canada, August 27-29, 1984*, pages 103–118. ACM, 1984. URL: <https://dl.acm.org/citation.cfm?id=1599406>.
- 19 Electric Coin Company and Zcash Contributors. Zcash: Protocol and reference implementation. <https://github.com/zcash/zcash>, 2023. Accessed: 01/09/2024.
- 20 David Evans, Vladimir Kolesnikov, and Mike Rosulek. *A Pragmatic Introduction to Secure Multi-Party Computation*. Now Publishers Inc, 2018. doi:10.1561/33000000019.
- 21 Prastudy Fauzi, Sarah Meiklejohn, Rebekah Mercer, and Claudio Orlandi. Quisquis: A new design for anonymous cryptocurrencies. Cryptology ePrint Archive, Paper 2018/990, 2018. URL: <https://eprint.iacr.org/2018/990>.
- 22 Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. Impossibility of distributed consensus with one faulty process. In Ronald Fagin and Philip A. Bernstein, editors, *Proceedings of the Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, March 21-23, 1983, Colony Square Hotel, Atlanta, Georgia, USA*, pages 1–7. ACM, 1983. doi:10.1145/588058.588060.

- 23 Christina Garman, Matthew Green, and Ian Miers. Accountable privacy for decentralized anonymous payments. In Jens Grossklags and Bart Preneel, editors, *Financial Cryptography and Data Security - 20th International Conference, FC 2016, Christ Church, Barbados, February 22-26, 2016, Revised Selected Papers*, volume 9603 of *Lecture Notes in Computer Science*, pages 81–98. Springer, 2016. doi:10.1007/978-3-662-54970-4_5.
- 24 Oded Goldreich. *Foundations of Cryptography*, volume 1. Cambridge University Press, 2001. doi:10.1017/CB09780511546891.
- 25 Jens Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology—EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35*, pages 305–326. Springer, 2016. doi:10.1007/978-3-662-49896-5_11.
- 26 Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovič, and Dragos-Adrian Seredinschi. The consensus number of a cryptocurrency. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC '19. ACM, July 2019. doi:10.1145/3293611.3331589.
- 27 Saurabh Gupta. A Non-Consensus Based Decentralized Financial Transaction Processing Model with Support for Efficient Auditing. Master's thesis, Arizona State University, USA, 2016.
- 28 Youssef El Housni and Aurore Guillevic. Optimized and secure pairing-friendly elliptic curves suitable for one layer proof composition. Cryptology ePrint Archive, Paper 2020/351, 2020. URL: <https://eprint.iacr.org/2020/351>.
- 29 Aram Jivanyan. Lelantus: A new design for anonymous and confidential cryptocurrencies. Cryptology ePrint Archive, Paper 2019/373, 2019. URL: <https://eprint.iacr.org/2019/373>.
- 30 Aram Jivanyan and Aaron Feickert. Lelantus spark: Secure and flexible private transactions. Cryptology ePrint Archive, Paper 2021/1173, 2021. doi:10.1007/978-3-031-32415-4_28.
- 31 Petr Kuznetsov, Yvonne-Anne Pignolet, Pavel Ponomarev, and Andrei Tonkikh. Permissionless and asynchronous asset transfer. *Distributed Comput.*, 36(3):349–371, 2023. doi:10.1007/S00446-023-00449-X.
- 32 Dahlia Malkhi and Michael K. Reiter. A high-throughput secure reliable multicast protocol. *Journal of Computer Security*, 5(2):113–128, 1997. doi:10.3233/JCS-1997-5203.
- 33 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. URL: <http://bitcoin.org/bitcoin.pdf>.
- 34 Shen Noether, Adam Mackenzie, and The Monero Research Lab. Ring Confidential Transactions. *Ledger*, 1:1–18, December 2016. doi:10.5195/ledger.2016.34.
- 35 David Pointcheval and Olivier Sanders. Short randomizable signatures. Cryptology ePrint Archive, Paper 2015/525, 2015. URL: <https://eprint.iacr.org/2015/525>.
- 36 Guillaume Quispe, Pierre Jouvelot, and Gérard Memmi. Nickpay, an auditable, privacy-preserving, nickname-based payment system. *CoRR*, abs/2503.19872, 2025. doi:10.48550/arXiv.2503.19872.
- 37 Amirreza Sarencheh, Aggelos Kiayias, and Markulf Kohlweiss. PEReDi: Privacy-enhanced, regulated and distributed central bank digital currencies. Cryptology ePrint Archive, Paper 2022/974, 2022. URL: <https://eprint.iacr.org/2022/974>.
- 38 Amirreza Sarencheh, Aggelos Kiayias, and Markulf Kohlweiss. Parscoin: A privacy-preserving, auditable, and regulation-friendly stablecoin. In *Cryptology and Network Security: 23rd International Conference, CANS 2024, Cambridge, UK, September 24–27, 2024, Proceedings, Part I*, pages 289–313, Berlin, Heidelberg, 2024. Springer-Verlag. doi:10.1007/978-981-97-8013-6_13.
- 39 Alberto Sonnino, Mustafa Al-Bassam, Shehar Bano, Sarah Meiklejohn, and George Danezis. Coconut: Threshold issuance selective disclosure credentials with applications to distributed ledgers. *arXiv preprint*, 2018. arXiv:1802.07344.

- 40 Paul F. Syverson, David M. Goldschlag, and Michael G. Reed. Anonymous connections and onion routing. In *1997 IEEE Symposium on Security and Privacy, May 4-7, 1997, Oakland, CA, USA*, pages 44–54. IEEE Computer Society, 1997. doi:10.1109/SECPRI.1997.601314.
- 41 Alin Tomescu, Adithya Bhat, Benny Applebaum, Ittai Abraham, Guy Gueta, Benny Pinkas, and Avishay Yanai. UTT: decentralized ecash with accountable privacy. *IACR Cryptol. ePrint Arch.*, 2022. URL: <https://eprint.iacr.org/2022/452>.
- 42 Andrei Tonkikh, Pavel Ponomarev, Petr Kuznetsov, and Yvonne-Anne Pignolet. Cryptoconcurrency: (almost) consensusless asset transfer with shared accounts. In *CCS*, pages 1556–1570. ACM, 2023. doi:10.1145/3576915.3616587.
- 43 Visa. Visa annual report 2024, 2024. URL: https://s29.q4cdn.com/385744025/files/doc_downloads/2024/Visa-Fiscal-2024-Annual-Report.pdf.
- 44 Karl Wüst, Kari Kostiainen, Vedran Capkun, and Srdjan Capkun. PRCash: Fast, private and regulated transactions for digital currencies. *Cryptology ePrint Archive*, Paper 2018/412, 2018. URL: <https://eprint.iacr.org/2018/412>.
- 45 Karl Wüst, Kari Kostiainen, Noah Delius, and Srdjan Capkun. Platypus: A central bank digital currency with unlinkable transactions and privacy preserving regulation. *Cryptology ePrint Archive*, Paper 2021/1443, 2021. doi:10.1145/3548606.3560617.