

Black-Box Accumulation Based on Lattices^{*} ^{**}

Sebastian H. Faller¹, Pascal Baumer², Michael Klooß³, Alexander Koch³, Astrid Ottenhues³, and Markus Raiber³

Institute of Information Security and Dependability (KASTEL)
Karlsruhe Institute of Technology
Karlsruhe, Germany

¹sebastian.faller@mailbox.org, ²ueeap@student.kit.edu,
³firstname.lastname@kit.edu

Abstract. Black-box accumulation (BBA) is a cryptographic protocol that allows users to accumulate and redeem points, e.g. in payment systems, and offers provable security and privacy guarantees. Loosely speaking, the transactions of users remain unlinkable, while adversaries cannot claim a false amount of points or use points from other users. Attempts to spend the same points multiple times (double spending) reveal the identity of the misbehaving user and an undeniable proof of guilt. Known instantiations of BBA rely on classical number-theoretic assumptions, which are not post-quantum secure. In this work, we propose the first lattice-based instantiation of BBA, which is plausibly post-quantum secure. It relies on the hardness of the Learning with Errors (LWE) and Short Integer Solution (SIS) assumptions and is secure in the Random Oracle Model (ROM).

Our work shows that a lattice-based instantiation of BBA can be realized with a communication cost per transaction of about 199 MB if built on the zero-knowledge protocol by Yang et al. (CRYPTO 2019) and the CL-type signature of Libert et al. (ASIACRYPT 2017). Without any zero-knowledge overhead, our protocol requires 1.8 MB communication.

Keywords: Lattice-based Cryptography · Black-box Accumulation (BBA) · Electronic Funds Transfer · Security and Privacy · Learning with Errors (LWE) · Short Integer Solution (SIS)

1 Introduction

Black-box accumulation (BBA), introduced in [24], allows the anonymous collection and redemption of points. BBA protocols feature two roles: users and operators. The users can accumulate and spend points on a cryptographic token issued by the operators, via the respective interactive protocols. In real-world

^{*} This is the full version of a paper accepted at the 18th IMA International Conference on Cryptography and Coding. https://doi.org/10.1007/978-3-030-92641-0_11

^{**} This work was supported by funding of the Helmholtz Association (HGF) through the Competence Center for Applied Security Technology (KASTEL).

scenarios like loyalty programs in shops or prepayment systems for public transport, users can collect incentives or bonus points. For the operators, the secure transfer of points is of paramount importance, whereas users want to protect their privacy. BBA offers a provably secure solution to both concerns. It allows users to collect and redeem points in an unlinkable manner and it protects operators from malicious users trying to claim more points than collected.

Several works have extended the framework of BBA. BBA+ [22] added stronger notions of both security and privacy as well as offline-transactions (in the sense that no permanent connection to a central database is required). More recently, [6, 7, 23] improved several aspects of BBA+. However, all of the proposed instantiations are based on classical cryptographic building blocks whose security guarantees rely on number-theoretic assumptions which are broken by Shor’s algorithms [33], rendering them insecure against quantum adversaries.

In contrast, lattice-based hardness assumptions have so-far withstood attempts to break them with quantum algorithms and allow to construct an extensive variety of cryptographic primitives, including commitments, public-key encryption [20, 25, 31, 32] and fully homomorphic encryption (FHE) [19], and are hence considered an ideal candidate to achieve post-quantum (PQ) security. Moreover, lattice-based protocols usually feature good asymptotic efficiency, parallelism, and security under worst-case intractability assumptions. The downside is an increase in communication costs for certain important building blocks, such as zero-knowledge (ZK) proofs. As all known BBA constructions are heavily based on ZK proofs, it gives rise to the difficult question of how to instantiate BBA from lattice-based assumptions, while remaining relatively efficient.

Contribution. In this work, we propose the first lattice-based instantiation of BBA, called BABL (**B**lack-**B**ox **A**ccumulation **B**ased on **L**attices). It relies on the LWE and SIS problems, and is proven secure in the ROM. We follow the security framework of [23], referred to as BBW in the following.

Moreover, we give a concrete instantiation, together with a suitable choice of lattice parameters, and evaluate the scheme’s communication complexity. Without any zero-knowledge overhead, our protocol requires 1.8 MB communication, which shows that the efficiency baseline of our general approach/construction is low. When using the popular ZK proof system by Yang et al. [36], an optimized version requires 199 MB communication, too much to be practically usable. However, lattice-based ZK proofs are improving rapidly, see e.g. [35] for some performance comparisons. Thus, it is plausible that the added computation and communication cost shrinks to an actually practical level in the near future. Our construction is the most efficient lattice-based payment system (BBA or E-Cash, cf. Section 1) to date. The closest competitor, E-Cash, needs 262 MB per transaction, using the same ZK protocol. While this does not yet make our protocol fit for practice, it places lattice-based BBA schemes into the range of practicality, where a further round of optimizations could likely allow its real-world use.

Our Construction in a Nutshell. On a high level, our construction follows the approach of BBW [23], but it requires care to translate it to lattices, without

reaching a giga-/terabyte range of communication cost per transaction. In BBW, the user holds a token which is basically a commitment whose contents are signed by the operator. This commitment contains a serial number for double-spending detection, a secret key uniquely identifying the user, and the amount of points. (For simplicity, we omit double-spending tags for now.) An update of the number of points (e.g. in a purchase) works as follows. The user sends a fresh (rerandomized) commitment to the operator, reveals the serial number, and proves in zero-knowledge that the commitment’s contents are signed and that the new balance lies within admitted bounds. Using a property of the commitment scheme, the operator updates the balance while keeping the committed content intact. A serial number is chosen by a two-party coin-toss, to ensure it cannot be used to track users. Finally, the operator provides a signature for the new token.

To implement this strategy, [23] uses group-based commitments, ZK proofs, and so-called CL-type signatures [10, 11], which have practically efficient ZK proofs for proving possession of a signature on a commitment or committed value. To replace Multi-Pedersen commitments, we use the lattice-based multi-block commitment scheme of [25], called KTX commitment in the following. We make use of a structural property of these commitments, which allows to “add blocks” to the committed message later, without knowing the messages in the other blocks, similar to Multi-Pedersen commitments. Finally, to replace the group-based CL-type signatures and ZK proofs, we rely on the ZK protocol of Yang et al. [36] combined with the CL-type signatures of Libert et al. [28].

Related Work. The previous BBA protocols by Blömer et al. [6], Bobolz et al. [7], Hartung et al. [22], Hoffmann et al. [23], and Jager and Rupp [24] are all based on number-theoretic hardness assumptions. While this allows them to be much more efficient, it also makes them insecure when quantum computers become available. We think that future privacy issues regarding payments made today, and the security of a users’ collected points in the future are reason enough to switch to post-quantum payment systems in the long term. The closest relatives to BBA are Electronic Cash (E-Cash) cryptosystems [14]. The first *compact* E-Cash scheme was given in [9], where compact means that the complexity of withdrawal and spending is logarithmic in the size of an (electronic) wallet. In (compact) E-Cash, there exist three parties, namely a bank, a user, and a merchant. The bank allows withdrawing a wallet containing coins and depositing coins. The wallet is signed by the bank, to make it possible for a user to prove the legitimacy of their wallet. Further, the user can spend the coins from their withdrawn wallet in a privacy-preserving way at a merchant. This is achieved by proving in zero-knowledge the legitimacy of the origin of the coin. The merchant can then deposit the received coin at the bank. The bank can detect double-spenders and prove their guilt, if and only if they are guilty. There are two lattice-based versions of compact E-Cash in the literature: The work of Libert et al. [28] – which propose (implicit) CL-type signatures, and an abstraction of Stern-type ZK protocols [34] – and Yang et al. [36]’s system, which applies their ZK argument system and further optimizations to construct a more efficient system similar to

[28]. [8] showed some major issues with the double-spending in compact E-Cash. [15] solved these problems with their lattice-based version of E-Cash.

We stress that E-Cash and BBA are quite different. Most importantly, BBA allows payments in both directions, i.e., points can be accumulated and spent. E-Cash does allow to deposit points but only at the bank, which means that a wallet cannot be “refilled”. Instead, a new wallet with fresh coins has to be generated. Additionally, E-Cash and BBA have different assumptions on the involved parties. On the one hand, it is not possible in BBA to separate issuer and accumulator – such as when the issuer is an E-Cash bank, and the accumulator a merchant. Issuer and accumulator have the same secret key. On the other hand, a merchant and bank must not collude in E-Cash, as this can break privacy. A BBA issuer and an accumulator can collude without breaking privacy. This is necessary due to an impossibility result, cf. [12]. Further, E-Cash only allows the transferal of a single coin per transaction. The transaction value in BBA is an integer in a certain range (e.g. 32-bit integers). Hence, when payments for products with different prices are made, BBA requires just one transaction, whereas many E-Cash transactions would be necessary.

Acknowledgements. We thank the anonymous reviewers for their feedback. The work presented in this paper has been funded by the German Federal Ministry of Education and Research (BMBF) under the project “PQC4MED” (ID 16KIS1044) and the topic Engineering Secure Systems of the Helmholtz Association (HGF) and by KASTEL Security Research Labs.

2 Preliminaries

Notation. We use $\lambda \in \mathbb{N}$ as security parameter. Vectors and matrices are in bold. For $n \in \mathbb{N}$ we write \mathbf{I}_n for the identity matrix of dimension n . We denote by \log the binary logarithm. $\|\cdot\|$ denotes the Euclidean norm and $\|\cdot\|_\infty$ the maximum norm. For $q \in \mathbb{N}$ we denote by $\mathbb{Z}_q = \{-\lfloor (q-1)/2 \rfloor, \dots, \lfloor (q-1)/2 \rfloor\}$ the ring of congruence classes of integers modulo q . We denote by $x \leftarrow S$ that x is drawn uniformly at random from set S and by $y \leftarrow D$ that y is drawn according to distribution D . We denote by $\cdot\| \cdot$ the concatenation of vectors, i.e., for $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_q^n$ we have $(\mathbf{x}\| \mathbf{y}) \in \mathbb{Z}_q^{2n}$. For $\mathbf{x} \in \mathbb{Z}_q^m$, we denote by $\text{bin}(\mathbf{x}) \in \mathbb{Z}_2^{m \lceil \log q \rceil}$ the binary decomposite of $\mathbf{x} \in \mathbb{Z}_q^m$, i.e. $\mathbf{x}_j = \sum_{i=0}^{\lceil \log q \rceil - 1} \text{bin}(\mathbf{x})_{\lceil \log q \rceil \cdot j + i} \cdot 2^i$. Inversely, for a $\mathbf{y} \in \mathbb{Z}_2^{\lceil \log q \rceil}$, we denote by $\text{tolnt}(\mathbf{y}) := \sum_{i=0}^{\lceil \log q \rceil - 1} \mathbf{y}_i \cdot 2^i \in \mathbb{Z}_q$ the integer (modulo q) represented by \mathbf{y} . For a full-rank matrix $\mathbf{M} \in \mathbb{Z}_q^{n \times m}$, we denote by $\overline{\mathbf{M}}$ the Gram-Schmidt orthogonalization of \mathbf{M} 's columns.

2.1 Black-Box Accumulation

In this section, we give an overview of the BBW framework defined in [23], which we base this work on. It allows a user to anonymously collect (and redeem) points from the operators, which cover the following three roles: i) the *issuer* issues new

tokens to the users of the system, ii) the *accumulator* adds points to a token, and iii) the *verifier* subtracts points from a token and verifies that a user’s balance is large enough to perform that transaction. As these roles share the same key pair, we do not distinguish them within the paper and refer to them as the *operator*.

The protocols are *offline*, meaning transactions can be performed without a permanent connection of the operator to a database. Nonetheless, regular access to a shared database to store the double-spending tags is needed. We require a common reference string which is set-up by a Trusted Third Party (TTP).¹ Next, let us give the formal definition of a BBW scheme.

Definition 2.1 (BBW Scheme). A black-box wallet *scheme* $\text{BBW} = (\text{Setup}, \text{OGen}, \text{UGen}, \text{Issue}, \text{Update}, \text{UVer}, \text{IdentDS}, \text{VerifyGuilt})$ consists of probabilistic polynomial time (PPT) algorithms Setup , OGen and UGen , interactive protocols Issue and Update and deterministic polynomial time algorithms UVer , IdentDS and VerifyGuilt :

- $\text{CRS} \leftarrow \text{Setup}(1^\lambda)$: On input 1^λ , returns a common reference string CRS . All following algorithms always receive CRS (implicitly) as input.
- $(\text{pk}_\mathcal{O}, \text{sk}_\mathcal{O}) \leftarrow \text{OGen}(\text{CRS})$: Returns a public and secret key for operator \mathcal{O} .
- $(\text{pk}_\mathcal{U}, \text{sk}_\mathcal{U}) \leftarrow \text{UGen}(\text{CRS})$: Returns a public and secret key for user \mathcal{U} .
- $((\mathcal{T}, b_\mathcal{U}), b_\mathcal{O}) \leftarrow \text{Issue}(\mathcal{U}(\text{pk}_\mathcal{O}, \text{pk}_\mathcal{U}, \text{sk}_\mathcal{U}), \mathcal{O}(\text{pk}_\mathcal{O}, \text{sk}_\mathcal{O}, \text{pk}_\mathcal{U}))$: User \mathcal{U} communicates with operator \mathcal{O} , who produces a new token \mathcal{T} for \mathcal{U} with balance 0. The user’s input is their key pair $(\text{pk}_\mathcal{U}, \text{sk}_\mathcal{U})$, and \mathcal{O} ’s public key $\text{pk}_\mathcal{O}$, while \mathcal{O} ’s input is its key pair $(\text{pk}_\mathcal{O}, \text{sk}_\mathcal{O})$, and the user’s public key $\text{pk}_\mathcal{U}$. The bits $b_\mathcal{O}$ and $b_\mathcal{U}$ indicate whether \mathcal{O} and \mathcal{U} “accept” the protocol run, respectively.
- $((\mathcal{T}^*, b_\mathcal{U}), (\text{dstag}, b_\mathcal{O})) \leftarrow \text{Update}(\mathcal{U}(\text{pk}_\mathcal{O}, \text{pk}_\mathcal{U}, \text{sk}_\mathcal{U}, \mathcal{T}, v), \mathcal{O}(\text{pk}_\mathcal{U}, \text{sk}_\mathcal{O}, v))$: User \mathcal{U} updates the token by interacting with the operator \mathcal{O} . Both get as inputs the public keys $\text{pk}_\mathcal{O}$ and $\text{pk}_\mathcal{U}$ and their respective secret key, and the (possibly negative) value v to be added to the token’s balance. \mathcal{U} additionally gets their token \mathcal{T} (with balance w) as input. In the end, \mathcal{U} outputs an updated token \mathcal{T}^* with balance $w + v$, and a bit $b_\mathcal{U}$ indicating acceptance of the execution. The operator outputs an acceptance bit $b_\mathcal{O}$ and a so-called double-spending tag dstag (which later allows detection of reuses of the same token).²
- $b \leftarrow \text{UVer}(\text{pk}_\mathcal{O}, \text{pk}_\mathcal{U}, \text{sk}_\mathcal{U}, \mathcal{T}, w)$: User \mathcal{U} verifies a token \mathcal{T} , given the operator’s public key $\text{pk}_\mathcal{O}$, the user’s key pair $(\text{pk}_\mathcal{U}, \text{sk}_\mathcal{U})$, and a value w , and outputs 1 if \mathcal{T} is a valid token of \mathcal{U} with balance w , or 0 otherwise.
- $(\text{pk}_\mathcal{U}, \Pi) \leftarrow \text{IdentDS}(\text{pk}_\mathcal{O}, \text{dstag}_1, \text{dstag}_2)$: Takes as input the operator’s public key $\text{pk}_\mathcal{O}$ and two double-spending tags $\text{dstag}_1, \text{dstag}_2$. If $\text{dstag}_1, \text{dstag}_2$ come from a transaction with the same token, then IdentDS outputs the public key $\text{pk}_\mathcal{U}$ of the user \mathcal{U} that “double-spent” their token and a proof of guilt Π . (Π

¹ Our setup only requires a *uniform random string (URS)*, also called *transparent setup*. In practice, it can be heuristically chosen, e.g. as a hash image.

² Note, that [22–24] distinguish between an *Add* and a *Sub* transaction for updating the token, where the first one hides the user’s balance and the latter one reveals it (or hides it via expensive range proofs). As we will discuss in [Section 3](#) there is no need for us to distinguish those cases, as the balance is always hidden in our construction.

can later be verified by a third party, using the `VerifyGuilt` algorithm described next.) Otherwise, it outputs an error symbol \perp .

- $b \leftarrow \text{VerifyGuilt}(\text{pk}_{\mathcal{O}}, \text{pk}_{\mathcal{U}}, \Pi)$: Given a proof of guilt Π , \mathcal{O} 's and \mathcal{U} 's public keys $\text{pk}_{\mathcal{O}}$, and $\text{pk}_{\mathcal{U}}$, it outputs 1 if \mathcal{U} is guilty of double-spending, 0 otherwise.

We say a BBW scheme is *correct* if the two interactive protocols `Issue` and `Update` and the algorithms `UVer`, `IdentDS`, and `VerifyGuilt` are correct. For `Issue` this means, if both parties follow the protocol, `Issue` outputs a valid token \mathcal{T} (as verified by `UVer`) and both parties accept the execution. Similarly, `Update` is correct if both parties accept the execution and the output is a valid (as above) updated token (with new balance $w + v$) if the parties follow the protocol. Correctness of `UVer`, `IdentDS`, and `VerifyGuilt` are defined in the canonical way.

Privacy and Security Properties. We give an informal description of the security properties and refer to [App. B](#) for the full definitions.

On the system side we formalize security by three properties: i) a scheme is *owner-binding* if a token is bound to a unique user, and can only be used by it, ii) a scheme is *balance-binding* if no false balance can be claimed, i.e., one can only claim a certain (overall) balance for a token if this balance equals the exact amount of points that have been legitimately collected with this token up to this point in time, and iii) a scheme features *double-spending detection* if a user that presents an already used token in a transaction can be (provably) identified.

For the privacy of the user, we demand the following properties: i) the scheme is *privacy preserving*, i.e., an adversary is not able to link any transactions of the user, even with corrupt operators, ii) the scheme offers *false-accusation protection*, if no malicious operator can falsely produce a proof of guilt for an honest user, and iii) a scheme should feature *post-compromise security*, i.e., that after a temporary compromise of the user, the unlinkability (but not the false-accusation property) can be recovered (by introducing new randomness into the token).

A difference in the description of BBW and our framework is that BBW allows embedding attributes in the token, i.e., the token's expiration date or data for age verification. Including such attributes is direct, but omitted for simplicity.

2.2 Lattices

We recall the basics of lattice-based cryptography required for our construction.

Definition 2.2. A lattice \mathcal{L} is the group of all integer linear combinations of k linearly independent vectors $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\} \subseteq \mathbb{R}^n$, for $k \in \mathbb{N}$: $\mathcal{L} = \mathcal{L}(\mathbf{B}) := \left\{ \sum_{i=1}^k z_i \cdot \mathbf{b}_i \mid z_i \in \mathbb{Z} \right\}$. Let $m \geq n \geq 1$, a prime number $q > 2$, $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{u} \in \mathbb{Z}_q^n$. We write:

$$\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{e} = \mathbf{0}^n \pmod{q}\}, \quad \Lambda_q^{\mathbf{u}}(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{e} = \mathbf{u} \pmod{q}\}.$$

Definition 2.3 (Discrete Gaussian Distribution). For a lattice \mathcal{L} , a vector $\mathbf{c} \in \mathbb{R}^m$, and a real number $\sigma > 0$, define $\rho_{\sigma, \mathbf{c}}(\mathbf{x}) = \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / \sigma^2)$. The

discrete Gaussian distribution of support \mathcal{L} , center \mathbf{c} and parameter σ is defined as $D_{\mathcal{L},\sigma,\mathbf{c}}(\mathbf{y}) = \rho_{\sigma,\mathbf{c}}(\mathbf{y})/\rho_{\sigma,\mathbf{c}}(\mathcal{L})$ for any $\mathbf{y} \in \mathcal{L}$, where $\rho_{\sigma,\mathbf{c}}(\mathcal{L}) = \sum_{x \in \mathcal{L}} \rho_{\sigma,\mathbf{c}}(x)$. We denote by $D_{\mathcal{L},\sigma}(\mathbf{y})$ the distribution centered in $\mathbf{c} = \mathbf{0}^m$ and exploit the fact that samples from $D_{\mathcal{L},\sigma}$ have small maximum norm with high probability.

Lemma 1 ([5, Lemma 1.5]). *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ and positive real number $\sigma > 0$, we have $\Pr_{\mathbf{b} \leftarrow D_{\mathcal{L},\sigma}}[\|\mathbf{b}\| \leq \sqrt{n}\sigma] \geq 1 - 2^{-\Omega(n)}$.*

The following lemmas specify how one can sample an (almost) random lattice basis of $\Lambda_q^\perp(\mathbf{A})$, together with a short trapdoor basis, and how to extend a basis:

Lemma 2 ([4, Theorem 3.2]). *There is a PPT algorithm TrapGen, that takes as input $1^n, 1^m$ and an integer $q > 2$ with $m \geq \Omega(n \log q)$, and outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a basis $\mathbf{T}_\mathbf{A}$ of $\Lambda_q^\perp(\mathbf{A})$ such that \mathbf{A} is within statistical distance $2^{-\Omega(n)}$ to the uniform distribution over $\mathbb{Z}_q^{n \times m}$ and $\|\widetilde{\mathbf{T}}_\mathbf{A}\| \leq \mathcal{O}(\sqrt{n \log q})$.*

Lemma 3 ([13, Lemma 3.2]). *For $m' > m$, there exists a PPT algorithm ExtBasis that takes as inputs a matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m'}$ whose first m columns span \mathbb{Z}_q^n , and a basis $\mathbf{T}_\mathbf{A}$ of $\Lambda_q^\perp(\mathbf{A})$ where \mathbf{A} is the left $n \times m$ submatrix of \mathbf{B} , and outputs a basis $\mathbf{T}_\mathbf{B}$ of $\Lambda_q^\perp(\mathbf{B})$ with $\|\widetilde{\mathbf{T}}_\mathbf{B}\| \leq \|\widetilde{\mathbf{T}}_\mathbf{A}\|$.*

2.3 Instantiation of Building Blocks

KTX-Commitments. In our construction, we use the commitment scheme of [25]. Let $n \in \mathcal{O}(\lambda)$, $q \in \mathcal{O}(n^4)$, $m_0, m_1 \in \Theta(n \log q)$, $0 < \sigma_{\text{Com}} \in \mathbb{R}$, where m_0 is the size of the randomness vector \mathbf{r} , m_1 is the size of the message vector \mathbf{m} and σ_{Com} is the parameter of the Gaussian distribution for the randomness. In the simplest case, one can commit to one message block $\mathbf{m} \in \mathbb{Z}_2^{m_1}$ by computing:

$$\begin{aligned} \text{Gen}(1^\lambda): \mathbf{D}_0 &\leftarrow \mathbb{Z}_q^{n \times m_0}, \mathbf{D}_1 \leftarrow \mathbb{Z}_q^{n \times m_1}, \text{ output } (\mathbf{D}_0, \mathbf{D}_1, \sigma_{\text{Com}}). \\ \text{Com}(\text{params}, \mathbf{m}; \mathbf{r}) &:= \mathbf{D}_0 \cdot \mathbf{r} + \mathbf{D}_1 \cdot \mathbf{m} \in \mathbb{Z}_q^n, \end{aligned}$$

where $\text{params} = (\mathbf{D}_0, \mathbf{D}_1, \sigma_{\text{Com}})$ are the public parameters.

The matrices $\mathbf{D}_0, \mathbf{D}_1$ are drawn uniformly at random. For each new commitment the randomness $\mathbf{r} \leftarrow D_{\mathbb{Z}_q^{m_0}, \sigma_{\text{Com}}}$ is chosen according to the discrete Gaussian distribution $D_{\mathbb{Z}_q^{m_0}, \sigma_{\text{Com}}}$. Usually, we set $m_0 = 2m_1$. A commitment \mathbf{c} can be opened by showing \mathbf{m} and \mathbf{r} . If it holds that $\mathbf{m} \in \mathbb{Z}_2^{m_1}$, $\|\mathbf{r}\| \leq \sigma_{\text{Com}}\sqrt{m_0}$ and $\mathbf{c} = \mathbf{D}_0 \cdot \mathbf{r} + \mathbf{D}_1 \cdot \mathbf{m}$, the commitment is valid. This scheme is statistically hiding. It is computationally binding, which can be seen by a straightforward reduction on $SIS_{n,q,2\sigma_{\text{Com}}\sqrt{m_0}, m_0+m_1}$. For $N \in \mathbb{N}$, the scheme can be extended to a commitment scheme on N messages by using N matrices. For the security proof of our construction, we require our commitment scheme to be equivocal. The scheme presented above can easily be turned into an equivocal commitment scheme, by using lattice trapdoor gadgets, as discussed in the full version [17] of this paper. It is necessary in the construction of the trapdoor to have $m_0 > n \lceil \log q \rceil$.

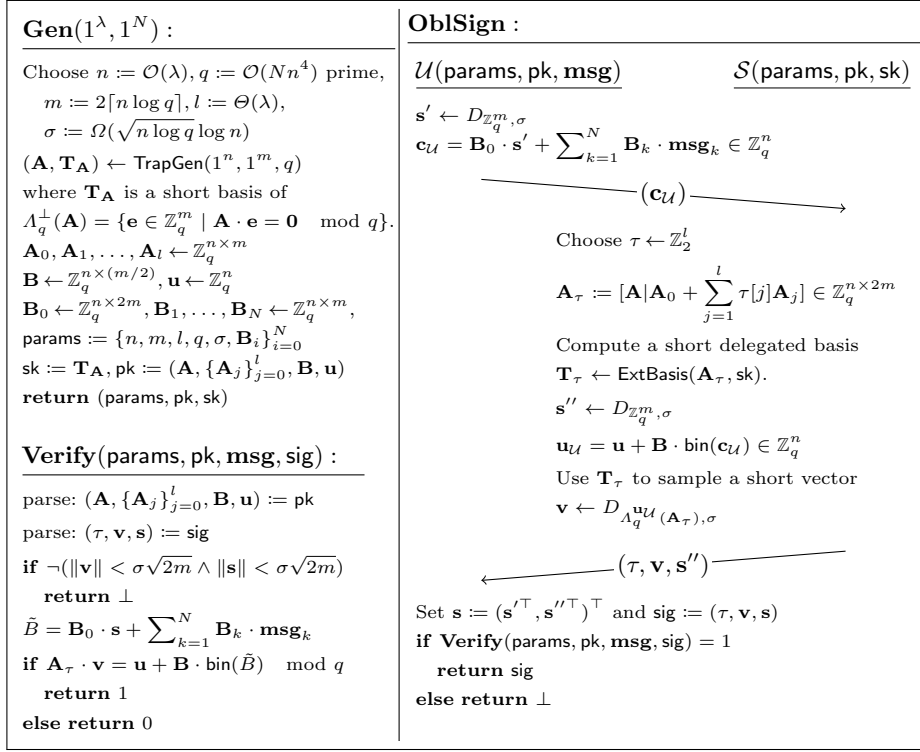


Fig. 1: Gen, Verify and OblSign algorithms of the signature scheme.

Signature Scheme by Libert et al. The scheme for obliviously signing committed messages by Libert et al. [27] consists of the two algorithms **Gen**, **Vfy**, and the interactive protocol **OblSign** (described in Fig. 1). It allows the signing of N -block messages $\mathbf{msg} = (\mathbf{m}_1, \dots, \mathbf{m}_N)$, for $N = \text{poly}(\lambda)$. In our construction we will use the notation $\text{OblSign}.\mathcal{S}(\text{pk}, \text{sk}, \mathbf{c}_\mathcal{U}) \rightarrow (\tau, \mathbf{v}, \mathbf{s}'')$ to denote the part of the protocol, which is executed by the signer, where **params** is derived from the relevant parts of the implicitly given CRS. The algorithm takes a key pair (pk, sk) and a commitment $\mathbf{c}_\mathcal{U}$ and outputs the signer's part of the signature $(\tau, \mathbf{v}, \mathbf{s}'')$ on the content of the commitment $\mathbf{c}_\mathcal{U}$. See the full version [17] of this paper for a proof sketch or [28, Theorem 2] for a full proof of Lemma 4. In the full-version, we additionally give a formal definition of the security, which was only implicit in [28][27].

Lemma 4 ([28, Theorem 2]). *Let $\beta' := \sigma^2 m \sqrt{2m}(l+2) + \sigma m \sqrt{m}$ and $\beta'' = \sigma^2 m \sqrt{2m} + \sqrt{2m} + 4\sigma m \sqrt{2m \log q}$. Then the above scheme is secure if the $\text{SIS}_{n,q,\beta',m}$ and $\text{SIS}_{n,q,\beta'',m}$ assumptions hold.*

OGen(CRS)	UGen(CRS)
Generate $(\text{params}', \text{pk}_S, \text{sk}_S) \leftarrow \text{S.Gen}(1^\lambda, 1^5)$	Draw $\text{sk}_U \leftarrow \mathbb{Z}_2^{m_{\text{sk}}}$
return $(\text{sk}_O := \text{sk}_S, \text{pk}_O := \text{pk}_S)$	return $(\text{pk}_U := \mathbf{F} \cdot \text{sk}_U, \text{sk}_U)$

Fig. 2: Generation algorithms **OGen**, **UGen** for operators and user, respectively

3 Our Construction of BABL

We denote by **S** the signature scheme of [28] (cf. Section 2.3) and by **C** the commitment scheme of [25] used in **Issue** and **Update** (cf. Section 2.3). The two zero-knowledge proof systems P1 and P2 are instantiations of the ZK scheme from [36] and are used in **Issue** and **Update**, respectively. (For a general description of the building blocks and their security notions, see App. A.2).

System Setup and Key Generation. We describe the choice of moduli, (matrix) dimensions, and their relation for the building blocks (as defined in Section 2.3). The parameter generation for both parties is described in Fig. 2.

Setup(1^λ):

- Choose a prime modulus $q_0 = \mathcal{O}(\lambda^4)$.
- For the signature scheme, set the modulus to $q = q_0^e$ for some $e > 0$. Let dimensions $n_S = \mathcal{O}(\lambda)$ and $m_S = 2n_S \lceil \log q \rceil$, and tag length $l_\tau = \Theta(\lambda)$.
- The Gaussian parameter is set to $\sigma = \Omega(\sqrt{n_S \log q \log n_S})$.
- Choose $n_{\text{sk}} = \mathcal{O}(\lambda)$, $m_{\text{sk}} = n_{\text{sk}} \log q_0$ and draw $\mathbf{F} \leftarrow \mathbb{Z}_{q_0}^{n_{\text{sk}} \times m_{\text{sk}}}$ (for pk_U later).
- For the commitment scheme, let dimension $n_C = n_S$. Let $m_r = 2n_C \lceil \log q \rceil$ be the size of commitment randomness, let $m_{\text{nr}} = \mathcal{O}(\lambda)$ be the size of serial numbers, and let $m_b = \mathcal{O}(\lambda)$ be the size of balance vectors. We require $2^{m_b} < q/4$ for the balance space \mathbb{V} , and choose $\mathbb{V} = \{0, \dots, 2^{m_b} - 1\}$. Thus, for all $x, y \in \mathbb{V}$, there is no wrap-around for $x + y$. Draw $\mathbf{D}_0 = (\mathbf{D}_0^0, \mathbf{D}_0^1) \leftarrow \mathbb{Z}_q^{n_C \times 2m_r}$, $\mathbf{D}_1 \leftarrow \mathbb{Z}_q^{n_C \times m_{\text{sk}}}$, $\mathbf{D}_2 \leftarrow \mathbb{Z}_q^{n_C \times m_b}$, $\mathbf{D}_3, \mathbf{D}_4 \leftarrow \mathbb{Z}_q^{n_C \times m_{\text{nr}} \lceil \log q_0 \rceil}$, $\mathbf{D}_5 \leftarrow \mathbb{Z}_q^{n_C \times m_{\text{sk}} \lceil \log q_0 \rceil}$ and set $N = 5$.³ Choose a Gaussian parameter $\sigma_{\text{Com}} > 0$ and use the same modulus q , as for the signature scheme.
- Let $H_{FRD} : \mathbb{Z}_{q_0}^{m_{\text{sk}}} \rightarrow \mathbb{Z}_{q_0}^{m_{\text{sk}} \times m_{\text{sk}}}$ be a full-rank difference function, (see App. A.2).
- **return** $\text{CRS} := (1^\lambda, q_0, q, n_S, m_S, n_C, m_r, n_{\text{sk}}, m_{\text{sk}}, m_b, m_{\text{nr}}, \sigma, \sigma_{\text{Com}}, \mathbf{F}, H_{FRD}, \{\mathbf{D}_i\}_{i=0}^5)$.

Issuing a New Token. In this protocol, the user \mathcal{U} interacts with the operator \mathcal{O} (in the issuer role) to get a fresh token with balance $\mathbf{b} = \mathbf{0}$. The token is a tuple of the form $\mathcal{T} = (\mathbf{c}, \mathbf{r}, \text{sk}_U, \mathbf{b}, \mathbf{s}_U, \mathbf{s}_O, \mathbf{u}_U, \text{sig})$, where \mathbf{c} is a multi-block commitment to the values $\text{sk}_U, \mathbf{b}, \mathbf{s}_U, \mathbf{s}_O$, and \mathbf{u}_U with randomness \mathbf{r} . Here, sk_U is the user’s secret key, and the vectors \mathbf{s}_U and \mathbf{s}_O are the two shares of the token’s serial number, chosen by \mathcal{U} and \mathcal{O} , respectively. The vector \mathbf{u}_U is randomly drawn by the user and is used in the generation of the \mathbf{t} -part of the double-spending tag (to be explained below). Finally, sig is a signature on the commitment \mathbf{c} .

³ We will use these matrices for the signature, too. We ignore params' , output by **S.Gen**.

The `Issue` protocol is the only protocol in which the operator sees pk_U , the public key of the user. In subsequent transactions of `Update`, possession of sk_U (and thus, of pk_U) is proven via ZK proof.




Fig. 3: Protocols for issuing (left) and updating (right) a token \mathcal{T} .

Fig. 3 (left) shows the `Issue` protocol in detail. First, the user \mathcal{U} draws their part of the serial number $\mathbf{s}_U \leftarrow \mathbb{Z}_{q_0}^{m_{nr}}$, the vector $\mathbf{u}_U \leftarrow \mathbb{Z}_{q_0}^{m_{sk}}$ for the computation of the double-spending tag, and a random vector for the commitment $\mathbf{r}' \leftarrow D_{\mathbb{Z}_q^{m_r}, \sigma_{\text{Com}}}$. The other half of the randomness vector $\mathbf{r} \in \mathbb{Z}_q^{2m_r}$ is set to 0, so the randomness chosen by \mathcal{O} can later be added (after \mathcal{U} received \mathbf{r}'' from \mathcal{O}). \mathcal{U} then commits (using randomness \mathbf{r}') on a five-block message, containing the secret key sk_U , \mathbf{s}_U and \mathbf{u}_U , with the second and fourth message block of the commitment being initialized to $\mathbf{0}$. This is because the second block represents the balance of the token, and is supposed to be $\mathbf{0}$ after issuance of the token, and the fourth block is zero for the operator to later add their share \mathbf{s}_O of the serial number to the block. Afterwards, the user sends the commitment to \mathcal{O} , together with a ZK proof π that ensures that the commitment contains the secret key which belongs to the user's public key.

After verification of π , \mathcal{O} proceeds by adding their share $\mathbf{s}_{\mathcal{O}}$ of the serial number ($\mathbf{s}_{\mathcal{U}}, \mathbf{s}_{\mathcal{O}}$) to the commitment. Then, the operator signs the committed message obliviously and sends the final commitment and the signature back to the user. The user verifies if the token is correct and accepts if this is the case.

We denote the used ZK proof system by P1. With P1, the user proves the following relations to the operator:

1. $\text{pk}_{\mathcal{U}} = \mathbf{F} \cdot \text{sk}_{\mathcal{U}} \pmod{q_0}$
2. $\mathbf{c}_{\text{new}} = \mathbf{D}_0^0 \cdot \mathbf{r}' + \mathbf{D}_0^1 \cdot \mathbf{0} + \mathbf{D}_1 \cdot \text{sk}_{\mathcal{U}} + \mathbf{D}_2 \cdot \mathbf{0} + \mathbf{D}_3 \cdot \bar{\mathbf{s}}_{\mathcal{U}} + \mathbf{D}_4 \cdot \mathbf{0} + \mathbf{D}_5 \cdot \bar{\mathbf{u}}_{\mathcal{U}}$
3. $\|\mathbf{r}'\| \leq \sigma_{\text{Com}} \sqrt{m_{\mathbf{r}}}$
4. $\text{sk}_{\mathcal{U}}, \bar{\mathbf{s}}_{\mathcal{U}}, \bar{\mathbf{u}}_{\mathcal{U}}$ are binary and $\text{bin}(\mathbf{s}_{\mathcal{U}}) = \bar{\mathbf{s}}_{\mathcal{U}}, \text{bin}(\mathbf{u}_{\mathcal{U}}) = \bar{\mathbf{u}}_{\mathcal{U}}$

In the full-version [17] of this paper, we show how these equations can be proven via the ZK protocol of [36]. We denote this proof by $\pi := \text{P1.Prove}()$. We note here, that the “actual” serial number of the token will be $\mathbf{s} = \mathbf{s}_{\mathcal{U}} + \mathbf{s}_{\mathcal{O}} \pmod{q_0} \in \mathbb{Z}_{q_0}^{m_{\text{nr}}}$. Thus, \mathbf{s} is uniformly random in $\mathbb{Z}_{q_0}^{m_{\text{nr}}}$ if one of the parties was honest. Hence, the collision probability is negligible if $q_0^{-m_{\text{nr}}}$ is negligible, which must be ensured by parameter choices.

Updating the Balance of a Token. We start with an overview of the update protocol, and then a detailed explanation. See Fig. 3 (right). The user starts with a token from a run of `Issue` or `Update`. Showing the token to the operator in the plain would make transactions linkable – we hence use ZK proofs for this.

First, the operator sends a “challenge” $\mathbf{u}_{\mathcal{O}}$. From this, the user generates a double-spending tag $\text{dstag} = (\mathbf{s}, \mathbf{t}, \mathbf{u}_{\mathcal{O}})$, where \mathbf{s} is the serial number, computed as $\mathbf{s} = \mathbf{s}_{\mathcal{U}} + \mathbf{s}_{\mathcal{O}}$, and \mathbf{t} is a masking of $\text{sk}_{\mathcal{U}}$, effectively a one-time pad encryption of $\text{sk}_{\mathcal{U}}$ with $\mathbf{u}_{\mathcal{U}}$. Given two values \mathbf{t}, \mathbf{t}' with distinct “challenges” $\mathbf{u}_{\mathcal{O}}, \mathbf{u}'_{\mathcal{O}}$, one can easily compute $\text{sk}_{\mathcal{U}}$. Since the user is bound to \mathbf{s} and $\text{sk}_{\mathcal{U}}$ by the token (and the commitment), this ensures that a double-spending user must reuse a serial number \mathbf{s} , and, since with overwhelming probability $\mathbf{u}_{\mathcal{O}} \neq \mathbf{u}'_{\mathcal{O}}$, is caught and deanonymized when doing so. This implements double spending detection. (Note, that a benign user can not be deanonymized).

The user also sets up a new (partial) token \mathbf{c}_{new} , analogous to the `Issue` protocol, but with balance set $\text{tolnt}(\mathbf{b}) + v$, according to the transaction. Then it proves that \mathbf{c}_{new} is a valid token (analogous to `Issue`) with correct balance, its connection and the validity of the “old” token (and its balance), and also, that the double spending tag $\text{dstag} = (\mathbf{s}, \mathbf{t}, \mathbf{u}_{\mathcal{O}})$ was correctly computed.

As in `Issue`, the operator first verifies the proof. It then chooses its part $\mathbf{s}'_{\mathcal{O}}$ of the new serial number and obliviously signs the adapted commitment \mathbf{c} . Also as in `Issue`, the user verifies the commitment and signature. After the transaction, the user has a freshly updated token, and the operator a double-spending tag.

Now, we describe the relevant parts in more detail. `Update` is defined as an interactive protocol between user \mathcal{U} and operator \mathcal{O} . Both parties take as inputs their key pairs, and the transaction value $v \in \mathbb{V}$. Additionally, the user gets as input the operator’s public key and the old token \mathcal{T} containing their current balance $\text{tolnt}(\mathbf{b}) \in \mathbb{V}$. The protocol outputs for \mathcal{U} and \mathcal{O} consist of a new token

\mathcal{T}^* and a double-spending tag dstag , respectively, as well as the output bits $b_{\mathcal{U}}, b_{\mathcal{O}}$, respectively, indicating a party accepts the execution of the protocol.

The user's token is $\mathcal{T} = (\mathbf{c}_{\text{old}}, \mathbf{r}, \text{sk}_{\mathcal{U}}, \mathbf{b}, \mathbf{s}_{\mathcal{U}}, \mathbf{s}_{\mathcal{O}}, \mathbf{u}_{\mathcal{U}}, \text{sig})$. As explained above, the protocol generates a new token, but with a different balance and additional consistency proofs and double-spending detection. This is reflected by the variables with an additional prime, corresponding to those of `Issue`. The updated balance is $\mathbf{b}_{\text{new}} := \text{bin}(\text{tolnt}(\mathbf{b}) + v \bmod q)$. The partial commitment is of the form

$$\mathbf{c}_{\text{new}} = \mathbf{D}_0^0 \cdot \mathbf{r}' + \mathbf{D}_1 \cdot \text{sk}_{\mathcal{U}} + \mathbf{D}_2 \cdot \mathbf{b}_{\text{new}} + \mathbf{D}_3 \cdot \text{bin}(\mathbf{s}'_{\mathcal{U}}) + \mathbf{D}_5 \cdot \text{bin}(\mathbf{u}'_{\mathcal{U}}), \quad (1)$$

which differs from `Issue` in the term $\mathbf{D}_2 \cdot \mathbf{b}_{\text{new}}$, where $\mathbf{b}_{\text{new}} = \mathbf{0}$. As in `Issue`, the new serial number is calculated from $(\mathbf{s}'_{\mathcal{U}}, \mathbf{s}'_{\mathcal{O}})$ chosen by user and operator, respectively. Also note the user only adds by $\mathbf{D}_0^0 \cdot \mathbf{r}'$ the first half of the randomness to the commitment. The second half will be chosen by the operator. For double-spending detection, \mathcal{O} sends $\mathbf{u}_{\mathcal{O}} \leftarrow \mathbb{Z}_{q_0}^{m_{\text{sk}}}$. The user calculates the vector $\mathbf{s} := \mathbf{s}_{\mathcal{U}} + \mathbf{s}_{\mathcal{O}} \bmod q_0$, and $\mathbf{t} := H_{FRD}(\mathbf{u}_{\mathcal{O}}) \cdot \text{sk}_{\mathcal{U}} + \mathbf{u}_{\mathcal{U}}$. Recall that H_{FRD} denotes the full-rank difference function from [App. A.2](#). Also note, that $\mathbf{u}_{\mathcal{U}}$ perfectly masks $H_{FRD}(\mathbf{u}_{\mathcal{O}}) \cdot \text{sk}_{\mathcal{U}}$, as it is a uniformly random value (chosen when \mathbf{c}_{old} was issued). Hence, \mathbf{t} reveals nothing about $\text{sk}_{\mathcal{U}}$. However, if the user double-spends by reusing an old token, given $\mathbf{u}_{\mathcal{O}} \neq \mathbf{u}'_{\mathcal{O}}$ in these two executions (which happens with overwhelming probability), then $(H_{FRD}(\mathbf{u}_{\mathcal{O}}) - H_{FRD}(\mathbf{u}'_{\mathcal{O}}))^{-1}(\mathbf{t}' - \mathbf{t}'') = \text{sk}_{\mathcal{U}}$. Thus, the identity of the misbehaving user is revealed. We denote by `P2` a non-interactive ZK proof system, and by $\pi := \text{P2.Prove}()$ its output. `P2` proves following equations:

1. $\text{Vfy}(\text{pk}_{\mathcal{O}}, \text{sig}, (\text{sk}_{\mathcal{U}}, \mathbf{b}, \bar{\mathbf{s}}_{\mathcal{U}}, \bar{\mathbf{s}}_{\mathcal{O}}, \bar{\mathbf{u}}_{\mathcal{U}})) = 1$
2. $\mathbf{s} = \mathbf{s}_{\mathcal{U}} + \mathbf{s}_{\mathcal{O}} \bmod q_0$
3. $\mathbf{t} = H_{FRD}(\mathbf{u}_{\mathcal{O}}) \cdot \text{sk}_{\mathcal{U}} + \mathbf{u}_{\mathcal{U}} \bmod q_0$
4. $\mathbf{c}_{\text{new}} = \mathbf{D}_0^0 \cdot \mathbf{r}'' + \mathbf{D}_0^1 \cdot \mathbf{0} + \mathbf{D}_1 \cdot \text{sk}_{\mathcal{U}} + \mathbf{D}_2 \cdot \mathbf{b}_{\text{new}} + \mathbf{D}_3 \cdot \bar{\mathbf{s}}'_{\mathcal{U}} + \mathbf{D}_5 \cdot \bar{\mathbf{u}}'_{\mathcal{U}}$
5. $\mathbf{b}_{\text{new}} = \text{bin}(\text{tolnt}(\mathbf{b}) + v \bmod q)$
6. $\|\mathbf{r}'\| \leq \sigma_{\text{Com}} \sqrt{m_{\mathbf{r}}}$
7. $\text{sk}_{\mathcal{U}}, \bar{\mathbf{s}}_{\mathcal{O}}, \bar{\mathbf{u}}_{\mathcal{U}}, \bar{\mathbf{u}}'_{\mathcal{U}}, \mathbf{b}_{\text{new}}, \bar{\mathbf{s}}'_{\mathcal{U}}$ are binary, $\text{bin}(\mathbf{s}_{\mathcal{U}}) = \bar{\mathbf{s}}_{\mathcal{U}}, \text{bin}(\mathbf{s}_{\mathcal{O}}) = \bar{\mathbf{s}}_{\mathcal{O}}, \text{bin}(\mathbf{u}_{\mathcal{U}}) = \bar{\mathbf{u}}_{\mathcal{U}}, \text{bin}(\mathbf{u}'_{\mathcal{U}}) = \bar{\mathbf{u}}'_{\mathcal{U}},$ and $\text{bin}(\mathbf{s}'_{\mathcal{U}}) = \bar{\mathbf{s}}'_{\mathcal{U}}$

The equations prove that the old token was valid, and its contents were used to compute $\mathbf{b}_{\text{new}}, \text{dstag}$ and \mathbf{c}_{new} . [Items 2](#) and [3](#) are for showing that the serial number \mathbf{s} and tag \mathbf{t} were computed correctly (from these values). The remaining equations prove the well-formedness of the new token, similar to `Issue`. In the full-version [\[17\]](#) of this paper, we show how the equations can be transformed into the generic form $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \bmod q$, where \mathbf{A} is a public matrix, \mathbf{y} is a public vector and \mathbf{x} is the secret witness. Once the equations are in this form, the ZK protocol from [\[36\]](#) can be leveraged.

The user sends $(\mathbf{c}_{\text{new}}, \mathbf{s}, \mathbf{t}, \pi)$ to \mathcal{O} , who checks the validity of the proof. The remainder of the protocol is essentially as in `Issue`, i.e. \mathcal{O} picks their share $\mathbf{s}_{\mathcal{O}}$ of the serial number, obviously signs the extended commitment \mathbf{c} , and sends the respective values to \mathcal{U} . The double-spending tag $\text{dstag} = (\mathbf{s}, \mathbf{t}, \mathbf{u}_{\mathcal{O}})$ is stored in a database of the operator, after the transaction ended successfully. If an entry $(\mathbf{s}, \mathbf{t}')$ is already recorded, `IdentDS` can be used to identify the offending user.

$\text{UVer}(\text{pk}_{\mathcal{O}}, \text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}, \mathcal{T}, \mathbf{b})$	$\text{IdentDS}(\text{pk}_{\mathcal{O}}, (\mathbf{s}, z_1), (\mathbf{s}', z_2))$
<pre> parse: $(\mathbf{c}, \mathbf{r}, \text{sk}_{\mathcal{U}}, \mathbf{b}, \mathbf{s}_{\mathcal{U}}, \mathbf{s}_{\mathcal{O}}, \mathbf{u}_{\mathcal{U}}, \text{sig}) := \mathcal{T}$ parse: $(\tau_{\text{sig}}, \mathbf{v}_{\text{sig}}, \mathbf{r}) := \text{sig}$ msg := $\text{sk}_{\mathcal{U}}, \mathbf{b}, \text{bin}(\mathbf{s}_{\mathcal{U}}), \text{bin}(\mathbf{s}_{\mathcal{O}}), \text{bin}(\mathbf{u}_{\mathcal{U}})$ if $\mathbf{c} = \text{C.Com}(\text{msg}; \mathbf{r})$ $\wedge \text{S.Verify}(\text{pk}_{\mathcal{O}}, \text{sig}, \text{msg}) = 1$ return 1 else return 0 </pre>	<pre> parse: $(\mathbf{t}, \mathbf{u}_{\mathcal{O}}) := z_1, (\mathbf{t}', \mathbf{u}'_{\mathcal{O}}) := z_2$ if $\mathbf{s} \neq \mathbf{s}' \vee \mathbf{u}_{\mathcal{O}} = \mathbf{u}'_{\mathcal{O}}$ return \perp else $\text{sk}_{\mathcal{U}} := (\mathbf{t} - \mathbf{t}')$ $\cdot (H_{FRD}(\mathbf{u}_{\mathcal{O}}) - H_{FRD}(\mathbf{u}'_{\mathcal{O}}))^{-1} \bmod q_0$ $\text{pk}_{\mathcal{U}} := \mathbf{F} \cdot \text{sk}_{\mathcal{U}}$ return $(\text{pk}_{\mathcal{U}}, \Pi = \text{sk}_{\mathcal{U}})$ </pre>

Fig. 4: UVer for token verification, and IdentDS to handle double-spending.

Security of the Construction. We give an intuition on why the protocol is secure (see App. B.3 for the formal version of the argument): the commitment binds the user to the values in a token. The ZK property of the proof system protects the user’s actions from being linked between executions. Its soundness ensures that the user cannot cheat. Thus, at the end of the protocol, the user has a new token with an updated balance. The operator only learns that the old and new token are valid, the value v of the transaction, and the double-spending tag.

Discussion. We point out a difference between our construction of BBA and previous ones [22–24]. There, range proofs were expensive and therefore optional. However, going without range proofs was only possible if the user revealed the balance when spending points, leaking a lot of information. Due to the lattice-based setting, our ZK proofs implicitly ensure that the balance is within the allowed range \mathbb{V} . This is because our proofs rely on bit decomposition. That is, we prove $\text{tolnt}(\mathbf{x}) = \sum_{i=0}^{m_b-1} \mathbf{x}_i \cdot 2^i \in \{0, \dots, 2^{m_b} - 1\} = \mathbb{V}$ and $\mathbf{x} \in \{0, 1\}^{m_b}$. Consequently, $\text{tolnt}(\mathbf{x})$ is a positive integer. (Recall that $2^{m_b} < q/4$, so the unique representative of the congruence class in \mathbb{Z}_q is positive.) More precisely, we prove that \mathbf{b}_{new} is of the form $\mathbf{b}_{\text{new}} = \text{bin}(\text{tolnt}(\mathbf{b}) + v \bmod q)$. Now, if $\text{tolnt}(\mathbf{b}) + v \bmod q$ would be negative or bigger than $2^{m_b} - 1$, the user could not generate a ZK proof which would be accepted by the operator.

Detecting Double-Spending. To identify a double-spender, in other words, a user who tries to spend the same token twice, the operator runs the IdentDS algorithm (see Fig. 4). For double-spending detection, the operator requires access to the database of double-spending tags. When the user did double-spend, the operator can calculate the user’s secret-key $\text{sk}_{\mathcal{U}}$ from the double-spending tags.

The VerifyGuilt algorithm takes the users public key $\text{pk}_{\mathcal{U}}$ and a proof Π of double-spending. The algorithm outputs 1 if $\mathbf{F} \cdot \Pi = \text{pk}_{\mathcal{U}} \bmod q_0 \wedge \|\Pi\|_{\infty} \leq 1$. Note, that it is not possible to generate a proof of guilt for benign user, as our construction offers *false-accusation protection*.

User-Verify Algorithm. The UVer algorithm (see Fig. 4) checks if the commitment \mathbf{c} contained in the token is truly a commitment on the messages $\text{sk}_{\mathcal{U}}, \mathbf{b}, \text{bin}(\mathbf{s}_{\mathcal{U}}), \text{bin}(\mathbf{s}_{\mathcal{O}})$ and $\text{bin}(\mathbf{u}_{\mathcal{U}})$ with randomness \mathbf{r} . Further the algorithm

checks if \mathbf{sig} is a valid signature on the commitment \mathbf{c} under the operator's secret key. The algorithm outputs 1 if both conditions are fulfilled, otherwise 0.

4 Efficiency Evaluation

In this section, we evaluate the efficiency of our construction. We concentrate on the communication cost, as this is the main bottleneck for mobile payments. Therefore, we first analyze the communication cost of an **Update** transaction. Next, we briefly explain our choice of parameters and calculate the concrete communication cost of **Update** and **Issue** given those parameters. We also compare the efficiency of our instantiation to similar protocols.

Communication Cost. From Fig. 3 (right) we can derive the cost of **Update**. We denote the communication cost by $\mathcal{C}_{\text{Update}}$ and we write $|\mathbf{vec}|$ for the number of bits needed to represent a vector \mathbf{vec} .

$$\mathcal{C}_{\text{Update}} = |\mathbf{u}_{\mathcal{O}}| + |\mathbf{c}_{\text{new}}| + |\mathbf{s}| + |\mathbf{t}| + |\pi| + |\mathbf{c}| + |(\tau, \mathbf{v}_{\text{sig}}, \mathbf{r}'')| + |\mathbf{s}'_{\mathcal{O}}|$$

The biggest part of this sum is $|\pi|$. Therefore, we will further analyze the size of the proof. According to Yang et al. [36] the size of this proof is

$$|\pi| = (\log(2p + 1) + \kappa + (3l_1 + 2l_2 + 2m_{\text{Update}} + 2l_{\text{Update}}) \cdot \log q) \cdot N + (l_1 + m_{\text{Update}}) \cdot \log q, \quad (2)$$

where p, κ, l_1, l_2, N are parameters of the zero-knowledge protocol, m_{Update} is the length of the witness and l_{Update} is the size of the set \mathcal{M} .

Using the fast mode (cf. [36, Sec. 1.2]) to prove the norm bounds on \mathbf{r} and \mathbf{v} we arrive at the size $m_{\text{Update}} := |\tau| + 2n_S|\tau| + 2m_S + n_C \log q + 2m_{\mathbf{r}} + m_{\text{sk}} + 2m_{\mathbf{b}} + 3m_{\text{nr}} \log q_0 + 2m_{\text{sk}} \log q_0 + m_{\mathbf{r}} + b\lambda(\log(10m_S\beta/b) + 1)$ for the witness and $l_{\text{Update}} := |\tau| + n_S|\tau| + n_C \log q + 2m_{\mathbf{b}} \log q_0 + 3m_{\text{nr}} \log q_0 + 2m_{\text{sk}} + b\lambda(\log(10m_S\beta/b) + 1)$ for \mathcal{M} , where b allows for a trade-off between proof size and tightness of the proven bound.

Next, we look at the signature $|\mathbf{sig}|$. By definition, we have $\mathbf{sig} = (\tau, \mathbf{v}, \mathbf{s})$, where $\tau \in \mathbb{Z}_2^l$, $\mathbf{v} \in \mathbb{Z}_q^{2m_S}$ and $\mathbf{s} \in \mathbb{Z}_q^{2m_S}$. Hence, we get $|\mathbf{sig}| = l + 4m_S \cdot \log(\beta)$ bits.

Choice of Parameters. To provide practical parameters for our scheme, we follow the heuristic approach of setting parameters high enough to withstand best-known attacks instead of deriving them from a reduction to a hard lattice problem such as SVP. For the sake of comparison we choose parameters for 80-bit security. In real-world scenarios, a higher level of security is desirable. To do so, we examine the root Hermite factor (RHF) [18] of our SIS/LWE problems. According to [36], to achieve 80-bit security a RHF of at most 1.0048 is required. We follow [36] and estimate the required RHF as

$$\text{RHF}(\text{SIS}_{n,q,\beta}) \approx 2^{\frac{\log^2 \beta}{4n \log q}} \quad \text{resp.} \quad \text{RHF}(\text{LWE}_{n,q,\alpha}) \approx 2^{\frac{\log^2 \frac{\alpha}{5.31}}{4n \log q}}$$

see [26], resp. [3]. BABL relies on the the following assumptions:

- $SIS_{n_{sk},q,\sqrt{m_{sk}}}$ so it is infeasible to derive the secret key from the public key.
- $SIS_{n_C,q,2\sigma_{Com}\sqrt{2m_r}\frac{m_r}{b}}$ for the commitment scheme to be binding.
- $SIS_{n_S,q,\beta'\frac{m_S}{b}}, SIS_{n_S,q,\beta''\frac{m_S}{b}}$ for the signature to be secure.
- $SIS_{l_1,q,\beta_1}, SIS_{l_2,q,\beta_2}$ and $LWE_{l_2,q,\alpha}$ for the zero-knowledge protocol [36].

where factors $\frac{m_r}{b}, \frac{m_S}{b}$ are due to the soundness loss from the fast mode of [36]. From the ZK argument we have the constraints that $q_0 > p$ and $q > \max(\beta_1, \beta_2)$, where we repeat the proof $\lceil 2^{80}/p \rceil$ times to achieve a soundness error of 2^{-80} .

Param.	Value	Param.	Value	Param.	Value
p	2^{80}	$\log q_0$	100	$\log q$	200
l_1	7050	n_S	880	n_C	880
l_2	7000	m_S	352 000	m_r	352 000
σ_3	67	l_τ	80	σ_{Com}	4195.0
σ_4	5.96×10^{30}	σ	671.0	m_{nr}	1
β_1	1.75×10^{59}	β	8055.0	m_b	32
β_2	8×10^{58}	β'	9.6×10^{20}		
α	1.05×10^{-58}	β''	1.29×10^{19}		
b	16			n_{sk}	9
κ	128			m_{sk}	900
	RHF		RHF		RHF
SIS_{l_1,q,β_1}	1.0048	$SIS_{n_S,q,\beta'}$	1.0048	$SIS_{n_C,q,2\sqrt{2m_r}\sigma_{Com}\frac{m_r}{b}}$	1.0014
SIS_{l_2,q,β_2}	1.0048	$SIS_{n_S,q,\beta''}$	1.004	$SIS_{n_{sk},q,\sqrt{m_{sk}}}$	1.0046
$LWE_{l_2,q,\alpha}$	1.0047				

Table 1: Concrete choices of parameters and resulting values for the underlying assumptions of the zero-knowledge proof (left), the signature scheme (middle), the commitment scheme (right) and for the secret keys (bottom right).

We tested for values of $p = 2^{10}$ up to $p = 2^{80}$ and arrived at the smallest proof size for $p = 2^{80}$. Then we set q_0 such that $q = q_0^2$ is just big enough. Finally we set all dimensions n just high enough to achieve the desired RHF of 1.0048. This resulted in the parameters shown in Table 1. The size of the proof π is 197 MB. Overall, the communication cost for Update is 199 MB and for Issue 70 MB.

Comparison with Similar Protocols. In Table 2, we compare the result for our instantiation with other protocols. The given values for our construction, and E-Cash of [36] and of [28] are theoretical estimations, while the values for BBA+ [22] and BBW [23] are results of empirical experiments on a software implementation (the benchmarks are described in [23]). Therefore, the comparison should be taken with a grain of salt. However, it suffices to illustrate the efficiency gap between the lattice-based constructions and the elliptic curve-based ones. Note, that our construction is slightly more efficient than the construction of [36]. Even though we used the same zero-knowledge proof, we do not need to prove statements about the correct evaluation of weak pseudorandom functions.

5 Future Work

Post-Quantum Security. Despite recent progress in proving Fiat–Shamir transformations of Σ -protocols secure in the quantum random oracle model [16, 30] none of the results seems to apply to our setting. That is, even if we assume that the results apply to [36], the resulting notion of security is not sufficient for our proofs. We essentially require witness-extended emulation (WEE) [21, 29], i.e., except with negligible probability an accepting proof can be extracted. This is a stronger notion than the knowledge soundness, which [16, 30] use. In the classical setting, the difference is small since amplification (via rewinding) can be used to obtain WEE from knowledge soundness [29]. In the quantum setting, this is unclear. A possible remedy would be a transformation which allows online extraction of the witness, e.g. by (additionally) committing to the witness with an extractable commitment scheme or a dual-mode commitment scheme. Knowledge of the extraction trapdoor allows to prove “operator soundness”, while the hiding property still ensures the “privacy notions”. However, this would further increase proof sizes and introduce a global system trapdoor (which is undesirable).

Efficiency. As seen in Section 4, our construction requires a high amount of network traffic. For real-world scenarios this cost is still unacceptably high. Basing a construction on the stronger assumptions of Ring-LWE and Ring-SIS should allow more efficient schemes. However, while more efficient zero-knowledge proofs are known in the ring setting, we are not aware of more efficient CL-type signatures. Thus, this remains the most important open question.

References

- [1] S. Agrawal, D. Boneh, and X. Boyen. “Efficient Lattice (H)IBE in the Standard Model”. In: *EUROCRYPT 2010*. Ed. by H. Gilbert. Vol. 6110. LNCS. Springer, Heidelberg, 2010, pp. 553–572. DOI: [10.1007/978-3-642-13190-5_28](https://doi.org/10.1007/978-3-642-13190-5_28).
- [2] M. Ajtai. “Generating Hard Instances of Lattice Problems (Extended Abstract)”. In: *28th ACM STOC*. ACM Press, 1996, pp. 99–108. DOI: [10.1145/237814.237838](https://doi.org/10.1145/237814.237838).

Protocol	Issuance	Transaction	Token/Wallet	Based on
Our work	70 MB	199 MB	11 MB	Lattices
E-Cash [28]	33 TB	720 TB	4 MB	Lattices
E-Cash [36]	53 MB	262 MB	4 MB	Lattices
BBA+ [22]	1 kB	14 kB	<1 kB	Elliptic Curves
BBW [23]	1 kB	5 kB	<1 kB	Elliptic Curves

Table 2: Comparison of the efficiency of similar protocols with our work

- [3] M. R. Albrecht, R. Player, and S. Scott. *On The Concrete Hardness Of Learning With Errors*. Cryptology ePrint Archive, Report 2015/046. <https://eprint.iacr.org/2015/046>.
- [4] J. Alwen and C. Peikert. *Generating Shorter Bases for Hard Random Lattices*. Cryptology ePrint Archive, Report 2008/521. <https://eprint.iacr.org/2008/521>.
- [5] W. Banaszczyk. “New bounds in some transference theorems in the geometry of numbers”. In: *Mathematische Annalen* 296.1 (1993), pp. 625–635. DOI: [10.1007/BF01445125](https://doi.org/10.1007/BF01445125).
- [6] J. Blömer, J. Bobolz, D. Diemert, and F. Eidens. “Updatable Anonymous Credentials and Applications to Incentive Systems”. In: *ACM CCS 2019*. Ed. by L. Cavallaro, J. Kinder, X. Wang, and J. Katz. ACM Press, 2019, pp. 1671–1685. DOI: [10.1145/3319535.3354223](https://doi.org/10.1145/3319535.3354223).
- [7] J. Bobolz, F. Eidens, S. Krenn, D. Slamanig, and C. Striecks. “Privacy-Preserving Incentive Systems with Highly Efficient Point-Collection”. In: *ASIACCS 20*. Ed. by H.-M. Sun, S.-P. Shieh, G. Gu, and G. Ateniese. ACM Press, 2020, pp. 319–333. DOI: [10.1145/3320269.3384769](https://doi.org/10.1145/3320269.3384769).
- [8] F. Bourse, D. Pointcheval, and O. Sanders. “Divisible E-Cash from Constrained Pseudo-Random Functions”. In: *ASIACRYPT 2019, Part I*. Ed. by S. D. Galbraith and S. Moriai. Vol. 11921. LNCS. Springer, Heidelberg, 2019, pp. 679–708. DOI: [10.1007/978-3-030-34578-5_24](https://doi.org/10.1007/978-3-030-34578-5_24).
- [9] J. Camenisch, S. Hohenberger, and A. Lysyanskaya. “Compact E-Cash”. In: *EUROCRYPT 2005*. Ed. by R. Cramer. Vol. 3494. LNCS. Springer, Heidelberg, 2005, pp. 302–321. DOI: [10.1007/11426639_18](https://doi.org/10.1007/11426639_18).
- [10] J. Camenisch and A. Lysyanskaya. “A Signature Scheme with Efficient Protocols”. In: *SCN 02*. Ed. by S. Cimato, C. Galdi, and G. Persiano. Vol. 2576. LNCS. Springer, Heidelberg, 2003, pp. 268–289. DOI: [10.1007/3-540-36413-7_20](https://doi.org/10.1007/3-540-36413-7_20).
- [11] J. Camenisch and A. Lysyanskaya. “Signature Schemes and Anonymous Credentials from Bilinear Maps”. In: *CRYPTO 2004*. Ed. by M. Franklin. Vol. 3152. LNCS. Springer, Heidelberg, 2004, pp. 56–72. DOI: [10.1007/978-3-540-28628-8_4](https://doi.org/10.1007/978-3-540-28628-8_4).
- [12] S. Canard and A. Gouget. “Anonymity in Transferable E-cash”. In: *ACNS 08*. Ed. by S. M. Bellovin, R. Gennaro, A. D. Keromytis, and M. Yung. Vol. 5037. LNCS. Springer, Heidelberg, 2008, pp. 207–223. DOI: [10.1007/978-3-540-68914-0_13](https://doi.org/10.1007/978-3-540-68914-0_13).
- [13] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. “Bonsai Trees, or How to Delegate a Lattice Basis”. In: *Journal of Cryptology* 25.4 (2012), pp. 601–639. DOI: [10.1007/s00145-011-9105-2](https://doi.org/10.1007/s00145-011-9105-2).
- [14] D. Chaum. “Blind Signatures for Untraceable Payments”. In: *Advances in Cryptology* (1983), pp. 199–203. DOI: [10.1007/978-1-4757-0602-4_18](https://doi.org/10.1007/978-1-4757-0602-4_18).
- [15] A. Deo, B. Libert, K. Nguyen, and O. Sanders. “Lattice-Based E-Cash, Revisited”. In: *ASIACRYPT 2020, Part II*. Ed. by S. Moriai and H. Wang. Vol. 12492. LNCS. Springer, Heidelberg, 2020, pp. 318–348. DOI: [10.1007/978-3-030-64834-3_11](https://doi.org/10.1007/978-3-030-64834-3_11).

- [16] J. Don, S. Fehr, C. Majenz, and C. Schaffner. “Security of the Fiat-Shamir Transformation in the Quantum Random-Oracle Model”. In: *CRYPTO 2019, Part II*. Ed. by A. Boldyreva and D. Micciancio. Vol. 11693. LNCS. Springer, Heidelberg, 2019, pp. 356–383. DOI: [10.1007/978-3-030-26951-7_13](https://doi.org/10.1007/978-3-030-26951-7_13).
- [17] S. H. Faller, P. Baumer, M. Kloß, A. Koch, A. Ottenhues, and M. Raiber. *Black-Box Accumulation Based on Lattices*. Cryptology ePrint Archive, Report 2021/1303. <https://eprint.iacr.org/2021/1303>.
- [18] N. Gama and P. Q. Nguyen. “Predicting Lattice Reduction”. In: *EUROCRYPT 2008*. Ed. by N. P. Smart. Vol. 4965. LNCS. Springer, Heidelberg, 2008, pp. 31–51. DOI: [10.1007/978-3-540-78967-3_3](https://doi.org/10.1007/978-3-540-78967-3_3).
- [19] C. Gentry. “Fully homomorphic encryption using ideal lattices”. In: *41st ACM STOC*. Ed. by M. Mitzenmacher. ACM Press, 2009, pp. 169–178. DOI: [10.1145/1536414.1536440](https://doi.org/10.1145/1536414.1536440).
- [20] C. Gentry, C. Peikert, and V. Vaikuntanathan. “Trapdoors for hard lattices and new cryptographic constructions”. In: *40th ACM STOC*. Ed. by R. E. Ladner and C. Dwork. ACM Press, 2008, pp. 197–206. DOI: [10.1145/1374376.1374407](https://doi.org/10.1145/1374376.1374407).
- [21] J. Groth and Y. Ishai. “Sub-linear Zero-Knowledge Argument for Correctness of a Shuffle”. In: *EUROCRYPT 2008*. Ed. by N. P. Smart. Vol. 4965. LNCS. Springer, Heidelberg, 2008, pp. 379–396. DOI: [10.1007/978-3-540-78967-3_22](https://doi.org/10.1007/978-3-540-78967-3_22).
- [22] G. Hartung, M. Hoffmann, M. Nagel, and A. Rupp. “BBA+: Improving the Security and Applicability of Privacy-Preserving Point Collection”. In: *ACM CCS 2017*. Ed. by B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu. ACM Press, 2017, pp. 1925–1942. DOI: [10.1145/3133956.3134071](https://doi.org/10.1145/3133956.3134071).
- [23] M. Hoffmann, M. Kloß, M. Raiber, and A. Rupp. “Black-Box Wallets: Fast Anonymous Two-Way Payments for Constrained Devices”. In: *PoPETs 2020.1* (2020), pp. 165–194. DOI: [10.2478/popets-2020-0010](https://doi.org/10.2478/popets-2020-0010).
- [24] T. Jager and A. Rupp. “Black-Box Accumulation: Collecting Incentives in a Privacy-Preserving Way”. In: *PoPETs 2016.3* (2016), pp. 62–82. DOI: [10.1515/popets-2016-0016](https://doi.org/10.1515/popets-2016-0016).
- [25] A. Kawachi, K. Tanaka, and K. Xagawa. “Concurrently Secure Identification Schemes Based on the Worst-Case Hardness of Lattice Problems”. In: *ASIACRYPT 2008*. Ed. by J. Pieprzyk. Vol. 5350. LNCS. Springer, Heidelberg, 2008, pp. 372–389. DOI: [10.1007/978-3-540-89255-7_23](https://doi.org/10.1007/978-3-540-89255-7_23).
- [26] A. Kosba, Z. Zhao, A. Miller, Y. Qian, H. Chan, C. Papamanthou, R. Pass, abhi shelat, and E. Shi. *CC0: A Framework for Building Composable Zero-Knowledge Proofs*. Cryptology ePrint Archive, Report 2015/1093. <https://eprint.iacr.org/2015/1093>.
- [27] B. Libert, S. Ling, F. Mouhartem, K. Nguyen, and H. Wang. “Signature Schemes with Efficient Protocols and Dynamic Group Signatures from Lattice Assumptions”. In: *ASIACRYPT 2016, Part II*. Ed. by J. H. Cheon and T. Takagi. Vol. 10032. LNCS. Springer, Heidelberg, 2016, pp. 373–403. DOI: [10.1007/978-3-662-53890-6_13](https://doi.org/10.1007/978-3-662-53890-6_13).

- [28] B. Libert, S. Ling, K. Nguyen, and H. Wang. “Zero-Knowledge Arguments for Lattice-Based PRFs and Applications to E-Cash”. In: *ASIACRYPT 2017, Part III*. Ed. by T. Takagi and T. Peyrin. Vol. 10626. LNCS. Springer, Heidelberg, 2017, pp. 304–335. DOI: [10.1007/978-3-319-70700-6_11](https://doi.org/10.1007/978-3-319-70700-6_11).
- [29] Y. Lindell. “Parallel Coin-Tossing and Constant-Round Secure Two-Party Computation”. In: *CRYPTO 2001*. Ed. by J. Kilian. Vol. 2139. LNCS. Springer, Heidelberg, 2001, pp. 171–189. DOI: [10.1007/3-540-44647-8_10](https://doi.org/10.1007/3-540-44647-8_10).
- [30] Q. Liu and M. Zhandry. “Revisiting Post-quantum Fiat-Shamir”. In: *CRYPTO 2019, Part II*. Ed. by A. Boldyreva and D. Micciancio. Vol. 11693. LNCS. Springer, Heidelberg, 2019, pp. 326–355. DOI: [10.1007/978-3-030-26951-7_12](https://doi.org/10.1007/978-3-030-26951-7_12).
- [31] C. Peikert and B. Waters. “Lossy trapdoor functions and their applications”. In: *40th ACM STOC*. Ed. by R. E. Ladner and C. Dwork. ACM Press, 2008, pp. 187–196. DOI: [10.1145/1374376.1374406](https://doi.org/10.1145/1374376.1374406).
- [32] O. Regev. “On lattices, learning with errors, random linear codes, and cryptography”. In: *37th ACM STOC*. Ed. by H. N. Gabow and R. Fagin. ACM Press, 2005, pp. 84–93. DOI: [10.1145/1060590.1060603](https://doi.org/10.1145/1060590.1060603).
- [33] P. W. Shor. “Algorithms for Quantum Computation: Discrete Logarithms and Factoring”. In: *35th FOCS*. IEEE Computer Society Press, 1994, pp. 124–134. DOI: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700).
- [34] J. Stern. “A new paradigm for public key identification”. In: *IEEE Transactions on Information Theory* 42.6 (1996), pp. 1757–1768. DOI: [10.1109/18.556672](https://doi.org/10.1109/18.556672).
- [35] C. Weng, K. Yang, J. Katz, and X. Wang. *Wolverine: Fast, Scalable, and Communication-Efficient Zero-Knowledge Proofs for Boolean and Arithmetic Circuits*. Cryptology ePrint Archive, Report 2020/925. <https://eprint.iacr.org/2020/925>.
- [36] R. Yang, M. H. Au, Z. Zhang, Q. Xu, Z. Yu, and W. Whyte. “Efficient Lattice-Based Zero-Knowledge Arguments with Standard Soundness: Construction and Applications”. In: *CRYPTO 2019, Part I*. Ed. by A. Boldyreva and D. Micciancio. Vol. 11692. LNCS. Springer, Heidelberg, 2019, pp. 147–175. DOI: [10.1007/978-3-030-26948-7_6](https://doi.org/10.1007/978-3-030-26948-7_6).

A Hardness Assumptions and Cryptographic Building Blocks

A.1 Lattice-based Hardness Assumptions

Definition A.1 (Short Integer Solution). *Given a modulus $q \in \mathbb{N}$, $m \in \mathbb{N}$ uniformly random vectors $\mathbf{a}_i \in \mathbb{Z}_q^n$ (written as a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$), and a uniformly random vector $\mathbf{u} \in \mathbb{Z}_q^n$, the Inhomogeneous Short Integer Solution*

(ISIS) problem is to find a non-zero integer vector $\mathbf{z} \in \mathbb{Z}^m$ of norm $\|\mathbf{z}\| \leq \beta$ such that

$$\mathbf{A}\mathbf{z} = \sum_{i=1}^m \mathbf{a}_i \cdot z_i = \mathbf{u} \in \mathbb{Z}_q^n,$$

where $\beta \in \mathbb{R}$ is a parameter with $\beta < q$.

In the case where \mathbf{u} is not uniform but fixed to $\mathbf{0}$, the problem is called Short Integer Solution (SIS). We write $SIS_{n,q,\beta,m}$, if we want to emphasize the respective parameters.

For typical parameter choices, SIS and ISIS are equivalent. Ajtai showed in his seminal work [2] that the average-case SIS problem can be reduced in polynomial time to the short integer vector problem (SIVP), a worst-case problem on lattices.

Regev [32] introduced the LWE problem and gave a quantum reduction to SIVP. We define the decisional variant of the respective hardness assumption.

Definition A.2 (Learning with Errors). *LWE $_{n,q,\chi,m}$: For a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$ and a probability distribution χ over \mathbb{Z}_q^m , sample a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ uniformly at random and a vector $\mathbf{e} \leftarrow \chi$. Given $(\mathbf{A}, \mathbf{b}^*)$ where \mathbf{b}^* is either \mathbf{b}_0 or \mathbf{b}_1 , where $\mathbf{b}_0^\top = \mathbf{s}^\top \mathbf{A} + \mathbf{e} \pmod q$ and \mathbf{b}_1 is chosen uniformly at random. Decide whether $\mathbf{b}^* = \mathbf{b}_0$ or $\mathbf{b}^* = \mathbf{b}_1$.*

A.2 Building Blocks

We give a brief overview of the used building blocks. In Section 2.3, we give instantiations of these building blocks based on the hardness of SIS and LWE.

Commitments. A commitment scheme allows one to *commit* to (i.e., fix) a value, without revealing it immediately. At a later point, the commitment can be *opened* and the committed value is revealed. A commitment scheme consists of two PPT algorithms: The parameter generation $\text{Gen}(1^\lambda)$ outputs public parameters params , and the commitment algorithm $\text{Com}(\text{params}, m; r) \rightarrow c$ outputs, for a given message m , some explicit randomness r and the public parameters params , a commitment c on that value m . We often omit the input of the public parameters.

To open the commitment c , one can reveal the randomness r , to check, if $\text{Com}(\text{params}, m; r) = c$ holds. Informally, we want a commitment scheme to be *hiding*, i.e. no (efficient) adversary can learn anything about the message m in a commitment, prior to opening. Furthermore, a commitment should be *binding* which means it should be (computationally) infeasible to open a commitment on m to any other value than m . A commitment scheme is *equivocal*, if there exist an additional trapdoor generation algorithm $\text{EqGen}(1^\lambda)$ that outputs the public parameters params , together with a trapdoor td . With the trapdoor td , it is possible to open a commitment c on the value m to another value m' with $m' \neq m$, by running a second algorithm, called $\text{Equiv}(\text{td}, c, m')$ that outputs a randomness value r' for opening c to m' . We require the two setups via Gen and EqGen to be computationally indistinguishable.

Oblivious Signing of Committed Messages. In our construction, we use the signature scheme of Libert et al. [28]; in particular, their protocol for obliviously signing a committed message (see Section 2.3). A *signature scheme for oblivious signing of committed messages* consists of the following algorithms/protocols:

- A key-generation algorithm $\text{Gen}(1^\lambda)$ that outputs $(\text{params}, \text{pk}, \text{sk})$, namely public parameters params , and a pair consisting of a public and a secret key.
- $\text{Ob!Sign}(\mathcal{U}(\text{params}, \text{pk}, m), \mathcal{S}(\text{params}, \text{pk}, \text{sk}))$, an interactive protocol, where a user \mathcal{U} interacts with a signer \mathcal{S} to obtain a signature on a message m inside of a commitment. In this protocol, \mathcal{U} sends a commitment $c \leftarrow \text{Com}(m; r)$ on m to the signer \mathcal{S} and eventually \mathcal{U} outputs a valid signature on m .
- a verification algorithm $\text{Vfy}(\text{params}, \text{pk}, m, \text{sig}) \rightarrow b$ that allows to check, whether sig is a valid signature on message m public key pk .

The signer does not learn anything about m , as the commitment scheme is hiding. This protocol offers a security notion that is almost identical to common EUF-CMA security but takes into account that the user sends commitments and not plain messages. Libert et al. forgo an abstract definition of the signature’s security as they directly apply the signature scheme to their E-Cash. We give a formal definition in the full-version [17] of this paper.

Zero-Knowledge Proofs. A proof system allows a party, called *prover*, to prove to another party, called *verifier*, that some statement is true. It is a *zero-knowledge* (ZK) protocol, if (informally) the verifier gains no additional knowledge, except for the truth of the statement. More precisely, the prover can convince the verifier that a word x belongs to a certain \mathcal{NP} -language L , while even a malicious verifier learns nothing about x except for the truth of $x \in L$. The protocol is a *proof of knowledge* (PoK), or *extractable*, if a convincing prover must know an \mathcal{NP} -witness w . For example, if $x = m$, $w = \sigma$, and the language is “I know a signature σ on message m ”, then a ZK-PoK guarantees that a convincing prover knows a signature σ , yet, the verifier learns nothing about σ . A ZK-PoK is correct, if an honest execution with correct statement always accepts. It has *soundness error* $p \in [0, 1]$, if the probability that the verifier accepts a false statement is at most p .

Full-Rank Difference Function. We define full-rank differences as introduced by Agrawal, Boneh, and Boyen [1], and refer to them for a concrete instantiation. We use this in the calculation of the double-spending tag. Let $q \in \mathbb{N}$ be a prime and $n \in \mathbb{N}$. A *full-rank difference function* is an efficiently computable function $H_{FRD}: \mathbb{Z}_q^n \rightarrow \mathbb{Z}_q^{n \times n}$ satisfying that for all distinct $u, v \in \mathbb{Z}_q^n$, the matrix $H_{FRD}(u) - H_{FRD}(v) \in \mathbb{Z}_q^{n \times n}$ is full rank.

B Security and Privacy Notions

In this section we render the precise definitions of the security and privacy notions as defined by [23].

Definition B.1 (Correctness of BABL). Similar to [22] the BABL scheme is called correct if the following holds: If the system is set up by the Setup algorithm, the keys are generated by UGen and OGen, all parties follow the protocol honestly, then, the following properties hold: (1) Correctness of the Issue protocol: Both parties return as acceptance bit 1. (2) Correctness of the Update protocol: For all valid tokens and balances, after adding a value the user always returns as acceptance bit 1.

Definition B.2 (Oracles, from [23, Def. 3.2]). MalIssue lets the adversary initiate the Issue protocol with an honest issuer \mathcal{O} provided that there is no pending MalIssue call for $\text{pk}_{\mathcal{U}}$ and $\text{pk}_{\mathcal{U}}$ has also not been used in a successful call to MalIssue before. MalUpdate lets the adversary initiate the Update protocol with an honest operator \mathcal{O} for an input value v . We say that a call to an oracle is successful if the honest party represented by the oracle accepts the run.

B.1 System Security

We denote by $\mathcal{T}_{\lambda, \text{CRS}}^{\text{Update}}$ the set of all transcripts of Update transactions, meaning all exchanged messages from the beginning, until both parties terminate.

Definition B.3. A scheme is called simulation-linkable if it satisfies the following conditions:

Completeness: Let $n \in \mathbb{N}$, $\text{CRS} \leftarrow \text{Setup}(1^\lambda)$ and $tr \in \mathcal{T}_{\lambda, \text{CRS}}^{\text{Update}}$ be a transcript. Then there exist inputs $\text{pk}_{\mathcal{U}}, \text{sk}_{\mathcal{U}}, \mathcal{T}, \mathbf{b}$ and random choices for an honest user \mathcal{U} and honest operator \mathcal{O} such that a run of the Update protocol between \mathcal{U} and \mathcal{O} with those inputs leads to the same transcript tr .

Extractability: There exists a PPT algorithm ExtractUID that, given two related transcripts $tr_1, tr_2 \in \mathcal{T}_{\lambda, \text{CRS}}^{\text{Update}}$ produced by the interaction of a honest user \mathcal{U} with public key $\text{pk}_{\mathcal{U}}$ with a honest operator \mathcal{O} outputs the public key $\text{pk}_{\mathcal{U}}$. Two transcripts tr_1, tr_2 are called related if they are identical except for the zero-knowledge challenges, output by the Random Oracle.

Additionally, there exists an expected PPT algorithm GenerateTranscripts that, given access to a transcript oracle $\mathcal{O} = \langle \mathcal{U}, \mathcal{O} \rangle$ which outputs transcripts between a user and an operator, outputs two related transcripts $tr_1, tr_2 \in \mathcal{T}_{\lambda, \text{CRS}}^{\text{Update}}$ with overwhelming probability. GenerateTranscripts is allowed to rewind \mathcal{O} and reprogram the Random Oracle.

Definition B.4. A simulation-linkable scheme is called owner-binding if for any PPT adversary \mathcal{A} in the experiments $\text{Exp}_{\text{BABL}, \mathcal{A}}^{\text{ob-issue}}(\lambda)$ and $\text{Exp}_{\text{BABL}, \mathcal{A}}^{\text{ob-update}}(\lambda)$ from Fig. 5 the advantages of \mathcal{A} defined by

$$\text{Adv}_{\text{BABL}, \mathcal{A}}^{\text{ob-issue}}(\lambda) = \Pr \left[\text{Exp}_{\text{BABL}, \mathcal{A}}^{\text{ob-issue}}(\lambda) = 1 \right]$$

$$\text{Adv}_{\text{BABL}, \mathcal{A}}^{\text{ob-update}}(\lambda) = \Pr \left[\text{Exp}_{\text{BABL}, \mathcal{A}}^{\text{ob-update}}(\lambda) = 1 \right]$$

are negligible in λ .

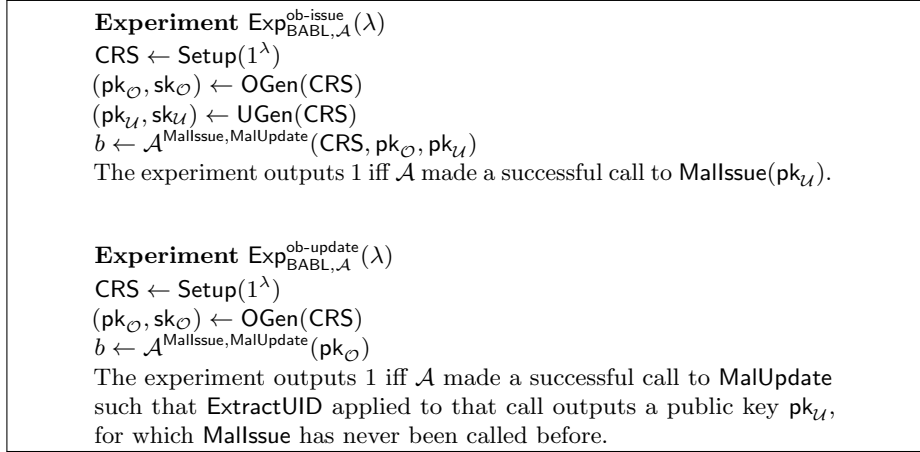


Fig. 5: Owner-binding experiment

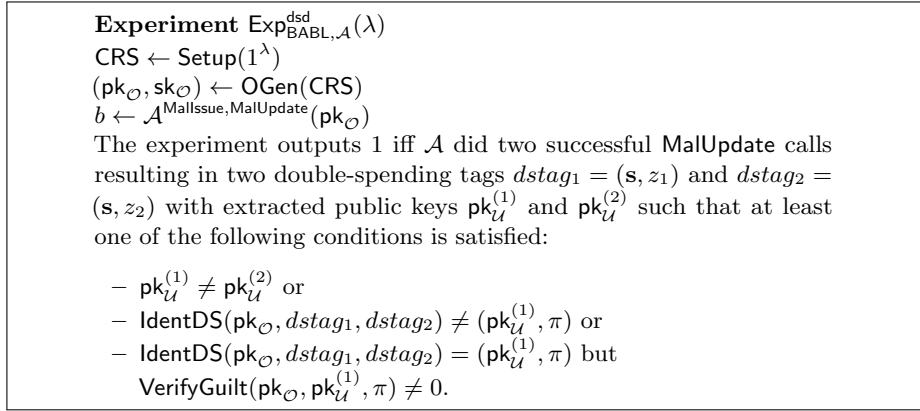


Fig. 6: Double-spending experiment

Definition B.5. A simulation-linkable scheme ensures double-spending detection if for any PPT adversary \mathcal{A} in the experiments $\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{dsd}}(\lambda)$ from Fig. 6 the advantage of \mathcal{A} defined by

$$\text{Adv}_{\text{BABL},\mathcal{A}}^{\text{dsd}}(\lambda) = \Pr \left[\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{dsd}}(\lambda) = 1 \right]$$

is negligible in λ .

Definition B.6. A simulation-linkable scheme is called balance-binding if for any PPT adversary \mathcal{A} in the experiments $\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{bb}}(\lambda)$ from Fig. 7 the advantage of \mathcal{A} defined by

$$\text{Adv}_{\text{BABL},\mathcal{A}}^{\text{bb}}(\lambda) = \Pr \left[\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{bb}}(\lambda) = 1 \right]$$

is negligible in λ .

<p>Experiment $\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{bb}}(\lambda)$ $\text{CRS} \leftarrow \text{Setup}(1^\lambda)$ $(\text{pk}_\mathcal{O}, \text{sk}_\mathcal{O}) \leftarrow \text{OGen}(\text{CRS})$ $b \leftarrow \mathcal{A}^{\text{MalIssue, MalUpdate}}(\text{pk}_\mathcal{O})$ The experiment outputs 1 iff \mathcal{A} made a successful call to MalUpdate with extracted user public-key $\text{pk}_\mathcal{U}$, s.t.</p> <ul style="list-style-type: none"> – all successful MalIssue / MalUpdate calls produce unique token version numbers, – the claimed balance $w \in \mathbb{V}$ does not equal the sum of previously collected accumulation values v for $\text{pk}_\mathcal{U}$, i.e. $w \neq \sum_{v \in V_{\text{pk}_\mathcal{U}}} v$ where $V_{\text{pk}_\mathcal{U}}$ is the list of all accumulation values $v \in \mathbb{V}$ that appeared in previous successful calls to MalUpdate for which $\text{pk}_\mathcal{U}$ could be extracted using ExtractUID.

Fig. 7: Balance-binding experiment

B.2 User Security and Privacy

User security is defined using the real/ideal world paradigm. The adversary can query the **HonUser** oracle to spawn new users. In the real world, the adversary interacts with oracles **RHonIssue** and **RHonUpdate** implementing the real user protocols. In the ideal world, the adversary interacts with a simulator. The simulator has to play the role of the oracles, but without receiving any private user information. We denote this by **SHonIssue**, **SHonUpdate**. In both worlds, the adversary can query **RCorrupt** or **SCorrupt**, respectively, to corrupt a user. By this, they learn all private information of the respective user.

Definition B.7. *A scheme is called privacy-preserving if there exist PPT algorithms **SimSetup** and **SCorrupt** as well as PPT protocols **SHonIssue**, **SHonUpdate** that receive no private user information, such that for all PPT adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the experiment depicted in Fig. 8, the advantage $\text{Adv}_{\text{BABL},\mathcal{A}}^{\text{priv}}(\lambda)$ of \mathcal{A} defined by*

$$\left| \Pr \left[\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{priv-real}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{priv-ideal}}(\lambda) = 1 \right] \right|$$

is negligible in λ .

Definition B.8. *A simulation-linkable scheme ensures false-accusation protection if for any PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the experiments $\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{facp}}(\lambda)$ from Fig. 9 the advantage of \mathcal{A} defined by*

$$\text{Adv}_{\text{BABL},\mathcal{A}}^{\text{facp}}(\lambda) = \Pr \left[\text{Exp}_{\text{BABL},\mathcal{A}}^{\text{facp}}(\lambda) = 1 \right]$$

is negligible in λ . (Note, that this does not guarantee anything, once the user was compromised.)

```

Experiment  $\text{Exp}_{\text{BABL}, \mathcal{A}}^{\text{priv-real}}(\lambda)$ 
CRS  $\leftarrow \text{Setup}(1^\lambda)$ 
 $(\text{pk}_\mathcal{O}, \text{state}) \leftarrow \mathcal{A}_1(\text{CRS})$ 
 $b \leftarrow \mathcal{A}_2^{\text{HonUser, RHonIssue, RHonUpdate, RCorrupt}}(\text{pk}_\mathcal{O}, \text{state})$ 
return  $b$ 

```

```

Experiment  $\text{Exp}_{\text{BABL}, \mathcal{A}}^{\text{priv-ideal}}(\lambda)$ 
 $(\text{CRS}, \text{td}_{\text{sim}}) \leftarrow \text{SimSetup}(1^\lambda)$ 
 $(\text{pk}_\mathcal{O}, \text{state}) \leftarrow \mathcal{A}_1(\text{CRS})$ 
 $b \leftarrow \mathcal{A}_2^{\text{HonUser, SHonIssue, SHonUpdate, SCorrupt}}(\text{pk}_\mathcal{O}, \text{state})$ 
return  $b$ 

```

Fig. 8: Real/Ideal world privacy experiment

```

Experiment  $\text{Exp}_{\text{BABL}, \mathcal{A}}^{\text{facp}}(\lambda)$ 
CRS  $\leftarrow \text{Setup}(1^\lambda)$ 
 $(\text{pk}_\mathcal{O}, \text{sk}_\mathcal{O}) \leftarrow \mathcal{A}_1(\text{CRS})$ 
 $(\text{pk}_\mathcal{U}, \text{sk}_\mathcal{U}) \leftarrow \text{UGen}(\text{CRS})$ 
 $\pi \leftarrow \mathcal{A}_2^{\text{RHonIssue, RHonUpdate}}(\text{pk}_\mathcal{O}, \text{pk}_\mathcal{U})$ 
Return 1 iff  $\text{VerifyGuilt}(\text{pk}_\mathcal{O}, \text{pk}_\mathcal{U}, \pi) = 1$ .

```

Fig. 9: False-accusation experiment

B.3 Security and Privacy

Our construction fulfills the desired security and privacy properties mentioned in Section 2.1. We formulate the theorems and give proof sketches. Note, that the proofs follow closely the proofs from [23]. Only small changes were necessary to adopt the proofs to the lattice-setting:

Theorem B.1 (Simulation-Linkability). *Suppose BABL is correct, S is secure and P1, P2 are sound. Then BABL is simulation-linkable.*

Proof sketch (Simulation-Linkability (Theorem B.1)). As by definition, a scheme is called simulation-linkable when it is complete and extractable, we have to show that BABL fulfills these properties.

Completeness requires, that for every accepted transcript, there is a choice of parameters, such that the transcript is the result of an honest protocol run. This is given in our case, as the sum of the serial number and the \mathbf{t} -part of the double-spending tag are indistinguishable from random values. Further, as the transcript is accepted, the soundness property of the zero-knowledge protocol from [36] guarantees that a commitment is well-formed. It remains to show that the signature is honestly generated, but as the token is accepted by the user and the signature is secure, this is given.

To prove the extractability property we can rely on the fact that the protocol from [36] is extractable, because it is a proof of knowledge.

Theorem B.2 (Owner-Binding w.r.t. Issue). *Suppose the $SIS_{n_{sk},q,\sqrt{m_{sk}}}$ assumption holds and P1 is extractable. Then BABL is owner-binding w.r.t. Issue.*

Proof sketch (Owner-Binding wrt. Issue (Theorem B.2)). Proving this property is a straightforward reduction on $SIS_{n_{sk},q,\sqrt{m_{sk}}}$.

Theorem B.3 (Owner-Binding w.r.t. Update). *Suppose BABL is simulation-linkable, P2 is extractable and S is secure. Then BABL is owner-binding w.r.t. Update.*

Proof sketch (Owner-Binding wrt. Update (Theorem B.3)). To prove the owner-binding property for the Update protocol we define a series of games for a hybrid argument. In these games, we test the required properties for the owner-binding property step by step. Finally, we show that the advantage of the adversary to win in the original game differs only negligibly from the other games.

First, we show that it is indeed possible to extract the user’s secret key by reducing this problem on the already proven simulation-linkability property of our scheme. Then it is possible to extract witnesses for all occurred zero-knowledge proofs as the zero-knowledge argument from [36] is extractable. Finally, as there are extracted witnesses, the only way left for the adversary to win the owner-binding game is to forge a signature. This is prevented by the security of our signature. Note here, that the security of the signature is not exactly the usual EUF-CMA security, as the user sends commitments and the signer signs the committed messages (and not the commitment itself). However, the notion of security given in the full-version [17] of this paper suffices for our proofs.

Theorem B.4 (Double-Spending Detection). *Suppose BABL is simulation-linkable, the $SIS_{n_{sk},q,\sqrt{m_{sk}}}$ assumption holds, P2 is extractable and S is secure. Then BABL ensures double-spending detection.*

Proof sketch (Double-Spending Detection (Theorem B.4)). Similar to the last theorem, we prove this property with a hybrid argument. In particular, we consider all ways in which IdentDS (Fig. 4) could be tricked into not recognizing an actual act of double-spending. We show that an adversary would therefore either be able to find a collision on the serial number, or they were able to manipulate the double-spending tag in a specific way. The former happens only with negligible probability, as the serial number is chosen in a coin-toss-like manner. The latter happens only with negligible probability, as the double-spending certainly includes values that the operators drew at random, and as the zero-knowledge proof is sound.

Theorem B.5 (Balance-Binding). *Suppose P1 and P2 are extractable and sound and C is statistically hiding, and S is secure. Then BABL is balance-binding.*

Proof sketch (Balance-Binding (Theorem B.5)). Similar to the proofs for the previous properties, we prove the balance-binding property by defining several games, where we show step by step, that the probability for an adversary to break the balance-binding property of BABL is negligible. More precisely, we

interpret every transaction as a node in a graph. Two nodes are connected if the output serial number of the first transaction is the input serial number of the last. We ensure through the game hops, that all nodes have an indegree of exactly one (except for the issuance of the token) and outdegrees of at most one. If this was not the case, there would be collisions on the serial number, double-spending, or forged signatures. Additionally, every such chain of transactions must be started by the issuance of a token and the balance must only change according to the transaction values of the nodes.

Theorem B.6 (Privacy-Preserving). *Suppose P1 and P2 are zero-knowledge and C is equivocal. Then BABL is privacy-preserving.*

Proof sketch (Privacy-Preserving (Theorem B.6)). We prove this property by defining several games, where the oracles of the real experiment are step-by-step replaced by oracles that hold no personal information of the user, called the ideal world. By showing that an adversary is only with negligible probability able to tell apart the real from the ideal world, we prove that BABL is indeed privacy-preserving. In more detail, we make use of the fact that the zero-knowledge proof from [36] is indeed zero-knowledge and the commitment scheme is equivocal. We use the equivocality property to replace the real values in the token with random ones, which makes it impossible to extract personal information from the user in the ideal world.

Theorem B.7 (False-Accusation Protection). *Suppose the $SIS_{n_{sk},q,\sqrt{m_{sk}}}$ assumption holds and the scheme ensures double-spending detection. Then BABL ensures false-accusation protection.*

Proof sketch (False-Accusation Protection (Theorem B.7)). Just like in the privacy-preserving proof we use the real/ideal world paradigm. Now, if the adversary is able to output a false proof of guilt for an honest user, one can directly construct an adversary breaking the $SIS_{n_{sk},q,\sqrt{m_{sk}}}$ assumption. If the adversary is not able to output a proof of guilt for a guilty user, this can be leveraged to distinguish the real world from the ideal world.