

# Lattice-Based E-Cash, Revisited

Amit Deo<sup>1,3</sup>, Benoît Libert<sup>2,1</sup>, Khoa Nguyen<sup>4</sup>, and Olivier Sanders<sup>5</sup>

<sup>1</sup> ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, Inria, UCBL), France

<sup>2</sup> CNRS, Laboratoire LIP, France

<sup>3</sup> Inria, France

<sup>4</sup> School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore

<sup>5</sup> Orange Labs, Applied Crypto Group, Cesson-Sévigné, France

**Abstract.** Electronic cash (e-cash) was introduced 40 years ago as the digital analogue of traditional cash. It allows users to withdraw electronic coins that can be spent anonymously with merchants. As advocated by Camenisch *et al.* (Eurocrypt 2005), it should be possible to store the withdrawn coins compactly (i.e., with logarithmic cost in the total number of coins), which has led to the notion of *compact* e-cash. Many solutions were proposed for this problem but the security proofs of most of them were invalidated by a very recent paper by Bourse *et al.* (Asiacrypt 2019). The same paper describes a generic way of fixing existing constructions/proofs but concrete instantiations of this patch are currently unknown in some settings. In particular, compact e-cash is no longer known to exist under quantum-safe assumptions. In this work, we resolve this problem by proposing the first secure compact e-cash system based on lattices following the result from Bourse *et al.* Contrarily to the latter work, our construction is not only generic, but we describe two concrete instantiations. We depart from previous frameworks of e-cash systems by leveraging lossy trapdoor functions to construct our coins. The indistinguishability of lossy and injective keys allows us to avoid the very strong requirements on the involved pseudo-random functions that were necessary to instantiate the generic patch proposed by Bourse *et al.*

**Keywords.** Lattice-based cryptography, e-cash, anonymity, exculpability, provable security.

## 1 Introduction

The last decades have witnessed major changes in consumer habits, with a gradual shift to credit/debit cards for payments. Since 2016, the total amount of card payment transactions worldwide has indeed exceeded that of cash transactions,<sup>6</sup> as card transactions simply make spending easier and enable online purchases.

However, the benefits of electronic payments come at a price. Each transaction indeed leaks very sensitive information (at least to the entity managing the payment system), such as the identity of the recipient, the amount, the location

---

<sup>6</sup> <https://avpsolutions.com/blog/payment-cards-now-set-to-surpass-cash/>

of the spender, etc. For example, a patient paying his cardiologist with his card implicitly reveals to his bank that he probably has a heart condition, which is far from insignificant.

One could argue that, in some cases, the users' or recipients' identities can be masked through pseudonyms, but the concrete privacy benefits of this solution are questionable. Indeed, even for systems without central authority such as Bitcoin, pseudonymity only provides limited anonymity guarantees as shown for example by Ron and Shamir [31]. A natural question in this context is whether we can achieve the best of the two worlds. Namely, can we combine the features of electronic payments together with the anonymity of traditional cash?

**Related Work.** A first answer to this question was provided by Chaum in 1982 [13] when he introduced the notion of electronic cash (e-cash). Concretely, an electronic coin is the digital analogue of a standard coin/banknote that is issued by an authority, called a bank, to users. The authenticity of coins can be checked publicly, which allows users to spend them anonymously with any merchant who knows the bank public key. Unfortunately, the comparison stops there, as there is a major difference between physical and electronic coins. In the first case, the physical support is assumed to be unclonable, unless for extremely powerful adversaries. Obviously, the same assumption does not hold for digital data, and it is thus necessary to deter multiple spendings of the same coin.

However, detecting multiple uses of the same coin without affecting the anonymity of honest users is challenging. Chaum achieved this using blind signatures [13], by associating each coin with a serial number that remains hidden until the coin is spent. At this time, the serial number is added to a register that can be public, precluding crypto-currency ledgers. Using this register, anyone can detect the reuse of a coin, which leads to two families of e-cash systems.

The first one allows detecting frauds but does not enable the identification of perpetrators. In this case, detection must be performed before accepting payments. These systems are thus inherently *online*, as any recipient must be able to check the ledger at any time. This entails incompressible latencies to process payments that can be prohibitive in some situations, such as payments at toll-gates, or at turnstiles for public transport.

The second family allows for the identification of defrauders. In this case, it is no longer necessary to check the coin upfront, as the defrauders know that they will ultimately be identified and then prosecuted. This simplifies the whole payment process as the e-cash system can now work *offline*.

It is hard today to discuss e-cash without mentioning crypto-currencies such as Bitcoin. The distinction between the two families above highlights the first difference between such systems. Crypto-currencies are indeed necessarily online whereas e-cash can be offline. However the main difference between these two systems rather lies in the trust model. The main strength of crypto-currencies is probably the absence of a central authority. This helps them circumvent the traditional reluctance of banks to novelties because a crypto-currency can be launched (almost) from scratch. In contrast, an e-cash system requires the sup-

port of a financial institution. Nevertheless, the absence of a central authority is also the main drawback of crypto-currencies. It indeed means that, in case of theft or loss of secret keys, the users lose everything, which is a major issue that we believe to be unacceptable for the general public. In the e-cash setting, where some authority manages the system, handling these situations is quite easy (corresponding procedures already exist for current payments systems such as debit/credit cards). There are also differences such as compliance with legislation, etc. In all cases, the very different features of both systems mean that they cannot be opposed. The reality is in fact the opposite and we should rather see crypto-currencies and e-cash systems as complementary solutions for privacy-preserving payments. From now on, we will only consider offline e-cash.

Following Chaum’s seminal work, blind signatures were the first cornerstone of e-cash systems. Unfortunately, this design strategy comes with some limitations, such as the need to withdraw and store coins one by one, which quickly becomes cumbersome (see, e.g., [9]). This problem was addressed by Camenisch *et al.* [11] who proposed the notion of *compact* e-cash, where users withdraw and store  $N$  coins (that constitute a wallet) with constant, or at least logarithmic, complexity. The core idea of their construction – which has become the blueprint of most following works – is to associate each wallet with two seeds  $k$  and  $t$  for a pseudo-random function (PRF) family. These two seeds are then used to generate  $N$  pairs of pseudo-random values  $(\text{PRF}_k(i), \text{PRF}_t(i))$ . The former (i.e.,  $\text{PRF}_k(i)$ ) serves as the coin serial number whereas  $\text{PRF}_t(j)$  essentially acts as a one-time pad on the spender’s identity, resulting in a so-called double-spending tag. In case a coin is spent more than once, the same mask  $\text{PRF}_t(j)$  is used twice and can thus be cancelled out to reveal the defrauder’s identity.

This elegant construction underlies many subsequent systems, including a candidate based on lattices [24]. Unfortunately, a recent result by Bourse *et al.* [9] has shown the limitations of this framework. In particular, they highlighted that systems based on it may fail to provably achieve exculpability, i.e., the property that honest users cannot be wrongly accused of double-spending a coin, even when the bank and other users conspire against them. As this issue underlies most of our design choices, we need to recall some details on it.

In the CHL construction [11], the serial number and the double-spending tag are constructed from the PRF outputs mentioned above but also from the spender’s public key and some public data that can be seen as the unique identifier of a transaction. In case of double-spending, it can be shown that the perpetrator will necessarily be identified. Unfortunately, Bourse *et al.* pointed out that the opposite is not true, except in some very specific settings excluding lattices, as carefully crafted serial numbers and double-spending tags might lead the identification process to output a public key that was not even involved in the fraudulent transactions. Actually, two spendings from different users may even be considered as a double-spending by the system. As a consequence, the security proofs of the e-cash construction of Libert *et al.* [24] and of a subsequent improvement by Yang *et al.* [33] (the only schemes relying on quantum-resistant computational assumptions) are invalid and there is no known simple fix.

Before accusing a user, it is therefore necessary to perform additional verifications on the serial number occurring in a double-spending, in particular to ensure that it was constructed from the same seed and the same identity. This unfortunately seems infeasible given only  $\text{PRF}_k(i)$ , as in [11]. To overcome this problem, the authors of [9] extended the serial number with new elements, each one being protected by a fresh PRF output. To ensure exculpability, it is then necessary to exclude collisions that could result from the PRF, leading to strong and non-standard requirements on the latter in [9]. Indeed, Bourse *et al.* need a notion of collision-resistance where, given the public parameters of a PRF family, the adversary should be unable to output two seeds  $k, k'$  and inputs  $x, x'$  such that  $\text{PRF}_k(x) = \text{PRF}_{k'}(x')$ . This might seem achievable by using a PRF based on symmetric primitives or by applying the techniques of Farshim *et al.* [17] to key-injective PRFs [22]. However, this would result in extremely inefficient e-cash constructions. Indeed, achieving security against cheating spenders requires to have them prove in zero-knowledge that they behaved honestly and correctly evaluated these PRFs, using certified seeds with valid inputs, etc. Such complex relations hardly coexist with the two solutions above. In particular, the Kim-Wu PRF [22] relies on a complex encoding of inputs into matrices which is hardly compatible with zero-knowledge techniques in lattices (recall that the PRF inputs should be part of witnesses). These rather require PRFs with a simpler algebraic structure, in the spirit of [5, 8, 4]. Unfortunately, the latter are not known to achieve collision-resistance. As of today, instantiating the Bourse *et al.* framework [9] from lattices would thus require to translate all statements to be proved into Boolean circuits. This would be much more expensive (by several orders of magnitude) than what we can hope for by leveraging the most efficient zero-knowledge techniques in standard lattices [33].

**Our Contribution.** In this paper, we show that we can dispense with the need for strong collision-resistance requirements by significantly departing from previous frameworks [11, 9]. Our core idea is to perform only one standard PRF evaluation and use the resulting output to mask all the components of the serial number and double-spending tag, thanks to the use of a lossy trapdoor function  $F_{\text{LTF}}$  [29]. Recall that these are function families where injective evaluation keys are computationally indistinguishable from lossy evaluation keys, for which image elements reveal very little information on their preimages. In our construction, during a spending, we reveal  $F_{\text{LTF}}(\text{PRF}_k(i))$  instead of  $\text{PRF}_k(i)$  and then extract randomness from the remaining entropy of  $\text{PRF}_k(i)$  in order to mask the spender's public key. This masked public key constitutes the second part of our serial number. When  $F_{\text{LTF}}$  is set up in its lossy mode in the proof of anonymity, we can show that the resulting serial number is indistinguishable from random and does not leak any sensitive information on the spender. Moreover, as  $F_{\text{LTF}}$  can be generated in injective mode in the real scheme, in case of colliding serial numbers, we are guaranteed that the same value  $\text{PRF}_k(i)$  is used in all the corresponding transactions. Together with the equality of serial numbers, this implies that the involved public keys necessarily coincide.

At this stage, we are ensured that a double-spending alert can only be generated by two transactions involving the same user. Then, it only remains to adapt the same technique to our double-spending tags, which is fairly simple. We can then prove security of our construction based only on the standard security properties of the pseudo-random function and the lossy trapdoor function.

However, as we intend to provide concrete constructions and not just frameworks, we still have to realise the missing component of the coin, namely the non-interactive zero-knowledge (NIZK) proofs that both the serial number and the double-spending tag are well-formed. Indeed, NIZK proofs are notoriously hard to produce in the lattice setting, at least compared to their counterparts in cyclic groups. We start from a very recent result by Yang *et al.* [33] which provides a protocol capturing many interesting lattice-related relations and show that it can be used to prove the statements required by our e-cash system. This is far from trivial as, in particular, spenders need to prove their correct composed evaluation of a pseudo-random function and a lossy trapdoor function using different parameters for the two primitives. We nevertheless manage to propose such NIZK arguments for two different PRF constructions [5,8], leading to two different instantiations. Notably, the GGM-based PRF [21] of Banerjee *et al.* [5] allows for the use of a polynomial modulus.

However, despite this nice asymptotic complexity, one should keep realistic expectations about the concrete performances of our scheme according to the current lattices state-of-the-art. We indeed note that, as of writing, most of our building blocks (zero-knowledge proofs, PRFs, etc) remain complex tools that can hardly compete with their pairing-based counterparts. This is highlighted by the recent paper by Yang et al [33] showing that existing (insecure) lattice e-cash constructions [24,33], which use building blocks similar to ours, generate transactions ranging from 260 MB to 720 TB. Fortunately, any future improvements of these tools could easily be leveraged by our construction. This is particularly true for our zero-knowledge proofs that we manage to represent as a standard instance of the powerful framework from [33].

Eventually, we propose the first concrete e-cash systems based on quantum-resistant hardness assumptions, following the reset of the state-of-the-art resulting from [9]. Unlike [9] that modifies the CHL framework [11] by requiring stronger properties on the original building blocks, we upgrade it by considering alternative building blocks that are instantiable from standard lattice assumptions. Our work does not only lead to concrete constructions, but it also sheds new lights on e-cash by implicitly relying on a new framework which differs from [11,9] and does not require PRFs with non-standard security properties.

## 2 Preliminaries

We use lower-case bold characters (e.g.  $\mathbf{x}$ ) to denote vectors and upper-case bold characters (e.g.  $\mathbf{M}$ ) to denote matrices. The  $(n \times n)$  identity matrix is denoted by  $\mathbf{I}_n$ . A superscript  $\top$  for a vector or matrix denotes its transpose (e.g.  $\mathbf{M}^\top$  is the transpose of  $\mathbf{M}$ ). For any integer  $q > 0$ ,  $\mathbb{Z}_q$  denotes the integers modulo  $q$ .

For integers  $a < b$ ,  $[a, b]$  denotes the set  $\{a, a+1, \dots, b\}$ . Alternatively if  $b > 1$ , we define  $[b] := \{1, \dots, b\}$ . For any real  $x$ , we denote by  $\lfloor x \rfloor$  the greatest integer smaller than or equal to  $x$ . In addition, for positive integers  $n, p, q$  such that  $q > p$ , we define the rounding operation  $\lfloor \cdot \rfloor_p : \mathbb{Z}_q^n \rightarrow \mathbb{Z}_p^n$  as  $\lfloor \mathbf{x} \rfloor_p := \lfloor (p/q) \cdot \mathbf{x} \rfloor$ . For probability distribution  $\mathcal{D}$ , we write  $s \leftarrow \mathcal{D}$  to denote that  $s$  is a sample of the distribution  $\mathcal{D}$ . If  $X$  is a set, then  $s \leftarrow U(X)$  represents the sampling of a uniform element of  $X$ . We also define the min-entropy of a discrete distribution  $\mathcal{D}$  as  $H_\infty(\mathcal{D}) := -\log \max_{x'} \Pr_{x \leftarrow \mathcal{D}}[x = x']$ . The statistical distance between two distributions  $\mathcal{D}_1$  and  $\mathcal{D}_2$  is denoted  $\Delta(\mathcal{D}_1, \mathcal{D}_2)$ . Throughout, we let  $\lambda$  denote a security parameter and use standard asymptotic notation  $\mathcal{O}, \Theta, \Omega, \omega$  etc. We also use the standard notion of a pseudo-random function (PRF) and a zero-knowledge argument of knowledge (ZKAoK).

*Binary decompositions.* We use the same decompositions as those in [24] as explained next. Firstly, for any positive integer  $B$  and  $\delta_B := \lfloor \log(B) \rfloor + 1$ , we define the sequence  $B_1, \dots, B_{\delta_B}$  where  $B_j := \lfloor \frac{B+2^{j-1}}{2^j} \rfloor$  for  $j \in [1, \delta_B]$ . It can be verified that  $\sum_{j \in [1, \delta_B]} B_j = B$ . For any integer  $x \in [0, B]$ , there is an efficiently computable deterministic function  $\text{idec}_B : [0, B] \rightarrow \{0, 1\}^{\delta_B}$  outputting a vector  $\text{idec}_B(x) =: \mathbf{y} \in \{0, 1\}^{\delta_B}$  satisfying  $\sum_{j \in [1, \delta_B]} B_j \cdot y_j = x$ . The function  $\text{idec}_B$  can be extended to handle vector inputs, resulting in  $\text{vdec}_{m,B} : [0, B]^m \rightarrow \{0, 1\}^{m \cdot \delta_B}$ , for any integer  $m > 0$ . Explicitly, for any  $\mathbf{x} \in [0, B]^m$ ,  $\text{vdec}_{m,B}(\mathbf{x}) := (\text{idec}_B(x_1)^\top, \dots, \text{idec}_B(x_m)^\top)^\top$ . In order to invert  $\text{vdec}_{m,B}$ , we define the matrix  $\mathbf{H}_{m,B} := (B_1, \dots, B_{\delta_B}) \otimes \mathbf{I}_m$ . It is easy to see that  $\mathbf{H}_{m,B} \cdot \text{vdec}_{m,B}(\mathbf{y}) = \mathbf{x}$ . In addition, for any  $x \in [0, B]$ , we denote by  $\text{ibin}_B(x)$  the *standard* binary decomposition of  $x$  that fits into  $\lfloor \log(B) \rfloor + 1$  bits. We define the binary representation of a vector to be the concatenation of the binary representations of its entries. Concretely, for any vector  $\mathbf{x} \in [0, B]^m$ , we define its binary representation to be  $\text{bin}_B(\mathbf{x})^\top := (\text{ibin}(x_1), \dots, \text{ibin}(x_m))$ .

## 2.1 Lattice Preliminaries

An  $m$ -dimensional lattice is a discrete subgroup of  $\mathbb{R}^m$ . For any integers  $n$  and  $q$ ,  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$  and  $\mathbf{u} \in \mathbb{Z}_q^n$  we define the full-rank lattice  $\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = 0 \bmod q\}$  and the lattice coset  $\Lambda_q^\mathbf{u}(\mathbf{A}) := \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{u} \bmod q\}$ . Defining  $\rho_\sigma : \mathbb{R}^m \rightarrow \mathbb{R}$  as  $\rho_\sigma(\mathbf{x}) := \exp(-\pi \|\mathbf{x}\|^2 / \sigma^2)$ , the discrete Gaussian distribution over a lattice coset  $L$  with parameter  $\sigma$  (denoted as  $D_{L,\sigma}$ ) is the distribution with support  $L$  and mass function proportional to  $\rho_\sigma$ .

*Hardness Assumptions.* We will be assuming that both the learning with errors (LWE) and short integer solution (SIS) problems (as defined next) are hard for appropriate parameter settings.

**Definition 1.** Let  $m, n, q \in \mathbb{N}$  with  $m > n$  and  $\beta > 0$ . The short integer solution problem  $\text{SIS}_{n,m,q,\beta}$  is, given  $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$ , find a non-zero  $\mathbf{x} \in \Lambda_q^\perp(\mathbf{A})$  with  $0 < \|\mathbf{x}\| \leq \beta$ .

**Definition 2.** Let  $q, \alpha$  be functions of a parameter  $n$ . For a secret  $\mathbf{s} \in \mathbb{Z}_q^n$ , the distribution  $A_{q,\alpha,\mathbf{s}}$  over  $\mathbb{Z}_q^n \times \mathbb{Z}_q$  is obtained by sampling  $\mathbf{a} \leftarrow U(\mathbb{Z}_q^n)$  and a noise  $e \leftarrow D_{\mathbb{Z},\alpha q}$ , and returning  $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e)$ . The learning with errors problem  $\text{LWE}_{n,m,q,\alpha}$  is, for  $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$ , to distinguish between  $m$  independent samples from  $U(\mathbb{Z}_q^n \times \mathbb{Z}_q)$  and the same number of samples from  $A_{q,\alpha,\mathbf{s}}$ .

If  $m$  is omitted in the LWE problem, it is assumed that  $m = \text{poly}(n)$ . If  $q \geq \beta n^\delta$  for any constant  $\delta > 0$  and  $m, \beta = \text{poly}(n)$ , then standard worst-case lattice problems with approximation factors  $\gamma = \max\{1, \beta^2/q\} \cdot \tilde{\mathcal{O}}(\beta\sqrt{n})$  reduce to  $\text{SIS}_{n,m,q,\beta}$  [27]. Alternatively, if  $q \geq \sqrt{n}\beta$  and  $m, \beta = \text{poly}(n)$ , then standard worst-case lattice problems with approximation factors  $\gamma = \mathcal{O}(\beta\sqrt{n})$  reduce to  $\text{SIS}_{m,q,\beta}$  (see, e.g., [19, Sec. 9]). Similarly, if  $\alpha q = \Omega(\sqrt{n})$ , standard worst-case lattice problems with approximation factors  $\gamma = \tilde{\mathcal{O}}(n/\alpha)$  reduce to  $\text{LWE}_{n,q,\alpha}$  [30, 10].

## 2.2 Lossy Trapdoor Functions

We will be using the notion of lossy trapdoor function (LTF) families from [29]. Informally, a lossy trapdoor function family can be used in one of two modes: a lossy mode and an injective mode. In the lossy mode, functions lose information on their inputs and cannot be inverted whereas in the injective mode, a trapdoor enables efficient inversion. In addition, there are generation algorithms that sample functions in either the lossy or injective mode. A crucial requirement is that no efficient adversary can distinguish whether a generation algorithm is outputting lossy functions or injective functions. We now recall the formal syntax and definition of an LTF family.

**Definition 3.** An  $(m, k)$  lossy trapdoor function family with security parameter  $\lambda$  is a 4-tuple of PPT algorithms  $(G_0, G_1, F, F^{-1})$  such that:

- **(Injective Mode)**  $G_0(1^\lambda)$  outputs a function index  $u$  and trapdoor  $\tau$ . For any pair  $(u, \tau)$  output by  $G_0$ ,  $F(u, \cdot)$  computes an injective function  $f_u : \{0, 1\}^m \rightarrow \{0, 1\}^*$  and  $F^{-1}(\tau, F(u, x)) = x$ .
- **(Lossy Mode)**  $G_1(1^\lambda)$  outputs a function index  $u$ . For any  $u$  output by  $G_1$ ,  $F(u, \cdot)$  computes a lossy function  $f_u : \{0, 1\}^m \rightarrow \{0, 1\}^*$ , whose image is of size at most  $2^{m-k}$ .
- **(Indistinguishability)** Let  $(u, \tau) \leftarrow G_0(1^\lambda)$  and  $u' \leftarrow G_1(1^\lambda)$ . Then the distributions of  $u$  and  $u'$  are computationally indistinguishable.

We will use the algorithms of the LTF family given in [29]. This family was reconstructed by Wee [32] where  $n, m, q, \alpha$  are functions of  $\lambda$ ,  $p \leq q/(4n)$  and  $\bar{n} = m/\log p$ . In the following,  $\mathbf{G} \in \mathbb{Z}_q^{m \times \bar{n}}$  is a special public matrix that allows to efficiently solve the bounded error decoding problem [29].

- $G_0(n, m, q, \alpha)$  : Sample  $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$ ,  $\mathbf{S} \leftarrow U(\mathbb{Z}_q^{n \times \bar{n}})$ ,  $\mathbf{E} \leftarrow (\bar{\Psi}_\alpha)^{m \times \bar{n}}$  and output the index  $(\mathbf{A}, \mathbf{B} := \mathbf{S}^\top \mathbf{A} + \mathbf{E}^\top)$  along with trapdoor  $\mathbf{S}$ .
- $G_1(n, m, q, \beta)$  : Sample  $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$ ,  $\mathbf{S} \leftarrow U(\mathbb{Z}_q^{n \times \bar{n}})$ ,  $\mathbf{E} \leftarrow (\bar{\Psi}_\alpha)^{m \times \bar{n}}$  and output the index  $(\mathbf{A}, \mathbf{B} := \mathbf{S}^\top \mathbf{A} + \mathbf{E}^\top)$

- $F$  : On input  $((\mathbf{A}, \mathbf{B}), \mathbf{x})$  where  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{B} \in \mathbb{Z}_q^{\bar{n} \times m}$  and  $\mathbf{x} \in \{0, 1\}^m$ , output  $(\mathbf{Ax}, \mathbf{Bx})$
- $F^{-1}$  : On input  $(\mathbf{S}, (\mathbf{y}_1, \mathbf{y}_2))$  where  $\mathbf{S} \in \mathbb{Z}_q^{n \times \bar{n}}, \mathbf{y}_1 \in \mathbb{Z}_q^n$  and  $\mathbf{y}_2 \in \mathbb{Z}_q^{\bar{n}}$ , compute  $\mathbf{y} := \mathbf{y}_2 - \mathbf{S}^\top \mathbf{y}_1$ . Use the efficient bounded-error decoder with respect to  $\mathbf{G}$  on  $\mathbf{y}$  to recover a vector  $\mathbf{x}^* \in \{0, 1\}^m$  such that  $\mathbf{e}^* + \mathbf{G}^\top \mathbf{x}^* = \mathbf{y}$  for some small  $\mathbf{e}^*$  with  $\|\mathbf{e}^*\|_\infty \leq q/p$ . Output  $\mathbf{x}^*$ .

**Lemma 1** ([32]). *For any constant  $\gamma < 1$  and  $n$ , take  $q = \Theta(n^{1+1/\gamma})$ ,  $p = \Theta(n^{1/\gamma})$  such that  $p \leq q/(4n)$ . Further, take  $m = \mathcal{O}(n \log q)$ ,  $\alpha = \Theta(\sqrt{n}/q)$  and  $\bar{n} = m/\log p$ . Assuming that the  $\text{LWE}_{n,m,q,\alpha}$  problem is hard, the above construction is an  $(m, k)$ -LTF family where  $k = (1 - \gamma)m - n \log q$ .*

The following instantiation of the generalized Leftover Hash Lemma of [15, Lemma 2.4] will be particularly useful:

**Lemma 2.** *Choose  $\gamma, n, q, p, \alpha$  as in Lemma 1, arbitrary integers  $n', q' > 2$  and an arbitrary distribution  $\mathcal{X}$  over  $\{0, 1\}^m$ . Then, for  $\mathbf{A} \leftarrow U(\mathbb{Z}_{q'}^{n' \times m})$ ,  $(\bar{\mathbf{A}}, \bar{\mathbf{B}}) \leftarrow G_1(n, m, q, \alpha)$ ,  $\mathbf{x} \leftarrow U(\mathcal{X})$  and  $\mathbf{u} \leftarrow U(\mathbb{Z}_{q'}^{n'})$ , we have*

$$\begin{aligned} \Delta((\mathbf{Ax}, \mathbf{A}, (\bar{\mathbf{A}}, \bar{\mathbf{B}}, \bar{\mathbf{A}}\mathbf{x}, \bar{\mathbf{B}}\mathbf{x})), (\mathbf{u}, \mathbf{A}, (\bar{\mathbf{A}}, \bar{\mathbf{B}}, \bar{\mathbf{A}}\mathbf{x}, \bar{\mathbf{B}}\mathbf{x}))) \\ \leq \frac{1}{2} \cdot \sqrt{2^{-(H_\infty(\mathcal{X}) - (m\gamma + n \log q + n' \log q'))}}. \end{aligned}$$

### 2.3 Witness Extraction and Forking Lemma

Recall that the transcript of a  $\Sigma$ -protocol consists of three messages starting with a message from a prover to a verifier. The Fiat-Shamir transform [18] provides a well-known method to remove interaction from a  $\Sigma$ -protocol. In particular, the second message (which is a uniformly chosen “challenge” value from the verifier to the prover) is replaced by the evaluation of a random oracle on input given by the first message. When adopting this method, it is important to carefully argue that the resulting non-interactive protocol is still an argument of knowledge. That is, if a prover convinces the verifier to accept with non-negligible probability, then replaying the prover allows for the extraction of a witness to the statement in question. This is usually achieved by applying a “forking lemma”.

We will focus on the argument system of Yang et al.[33] which takes the three-message form of a  $\Sigma$ -protocol. The witness extraction for the interactive ZKAoK of Yang et al. requires any  $\ell = 3$  accepting transcripts, all with the same first prover message but distinct challenge values. We refer to  $\ell$  such accepting transcripts as an  $\ell$ -fork.

When using a random oracle to remove interaction with our chosen argument system, a forking lemma that considers the probability of producing an  $\ell$ -fork for  $\ell = 3$  should be used. The extended/generalised forking lemma of El Kaafarani and Katsumata [16, Lemma 1] provides a forking lemma for any  $\ell \geq 2$ . For simplicity, we state their result in the special case that  $\ell = 3$ .



**Lemma 3 ([16]).** Fix some input  $x \in \{0,1\}^*$  and take some arbitrary set  $\text{accept}$ . Let  $\mathcal{A}$  be an efficient algorithm outputting triples  $(m_1, m_2, m_3)$  on input  $x$  that has oracle access to a random oracle  $H : \{0,1\}^* \rightarrow [h]$  and let  $Q$  be an upper bound on the number of queries that  $\mathcal{A}$  makes to  $H$ . Denote

$$\begin{aligned} \text{acc} &:= \Pr \left[ (m_1, m_2, m_3) \leftarrow \mathcal{A}^{H(\cdot)}(x) : \begin{array}{l} (m_1, m_2, m_3) \in \text{accept} \wedge \\ m_2 \text{ is the result of an } H\text{-query} \end{array} \right] \\ \text{frk}_3 &:= \Pr \left[ ((m_1, m_{2,i}, m_{3,i}))_{i=1}^3 \leftarrow F^{\mathcal{A}}(x) : \begin{array}{l} \forall i \in \{1,2,3\} : (m_1, m_{2,i}, m_{3,i}) \in \text{accept} \\ \wedge (m_{2,i})_{i=1}^3 \text{ are pairwise distinct} \end{array} \right] \end{aligned}$$

for any efficient algorithm  $F^{\mathcal{A}}$  that runs  $\mathcal{A}$  at most 3 times. Then, for a particular choice of  $F^{\mathcal{A}}$ ,

$$\text{frk}_3 \geq \text{acc} \cdot \left( \left( \frac{\text{acc}}{Q} \right)^2 - \frac{3}{h} \right).$$

## 2.4 E-Cash Security Definitions

E-cash systems involve three types of parties: banks denoted  $\mathcal{B}$ , users denoted  $\mathcal{U}$  and merchants denoted  $\mathcal{M}$ . The syntax of an offline compact e-cash system consists of the following algorithms/protocols:

**ParGen**( $1^\lambda, 1^L$ ): On input a security parameter  $\lambda$  and wallet size  $L = \log(\text{poly}(\lambda))$ , outputs public parameters  $\text{par}$  containing  $L$  (amongst other things).

**BKeyGen**( $1^\lambda, \text{par}$ ): On input  $\text{par}$ , outputs a key pair  $(PK_{\mathcal{B}}, SK_{\mathcal{B}})$  for the bank, which allows  $\mathcal{B}$  to issue wallets of size  $2^L$ .

**UKeyGen**( $1^\lambda, \text{par}$ ): On input  $\text{par}$ , generates a key pair  $(PK_{\mathcal{U}}, SK_{\mathcal{U}})$  for the user.

**MKeyGen**( $1^\lambda, \text{par}$ ): On input  $\text{par}$ , generates  $(PK_{\mathcal{M}}, SK_{\mathcal{M}})$  for the merchant.

We henceforth assume that all algorithms implicitly take  $\text{par}$  as input.

**Withdraw**( $\mathcal{U}(PK_{\mathcal{B}}, SK_{\mathcal{U}}), \mathcal{B}(PK_{\mathcal{U}}, SK_{\mathcal{B}})$ ): An interactive protocol that allows  $\mathcal{U}$  to obtain a wallet  $\mathcal{W}$  consisting of  $2^L$  coins or an error message  $\perp$ . The bank  $\mathcal{B}$  obtains tracing information  $\mathsf{T}_{\mathcal{W}}$ .

**Spend**( $\mathcal{U}(\mathcal{W}, PK_{\mathcal{B}}, PK_{\mathcal{M}}), \mathcal{M}(SK_{\mathcal{M}}, PK_{\mathcal{U}}, PK_{\mathcal{B}}, \text{info})$ ): A protocol allowing a user  $\mathcal{U}$  to give a coin from  $\mathcal{W}$  to merchant  $\mathcal{M}$  with respect to transaction metadata  $\text{info}$ . The user outputs an updated wallet  $\mathcal{W}'$  whereas the output of  $\mathcal{M}$  is a coin  $\text{coin}$  consisting of  $\text{info}$ , a serial number, a security tag and a proof of validity or an error symbol  $\perp$ .

**VerifyCoin**( $PK_{\mathcal{B}}, \text{coin}$ ): Outputs 1 if the proof of validity in  $\text{coin}$  verifies correctly with respect to  $PK_{\mathcal{B}}$  and 0 otherwise.

**VerifyDeposit**( $PK_{\mathcal{B}}, PK_{\mathcal{M}}, \text{coin}, \mu$ ): Outputs 1 if the proof of validity in  $\text{coin}$  verifies correctly with respect to  $PK_{\mathcal{B}}$  and if the data  $\mu$  verifies correctly with respect to  $PK_{\mathcal{M}}$ . Else, outputs 0.

**Deposit**( $\mathcal{M}(SK_{\mathcal{M}}, \text{coin}, PK_{\mathcal{B}}), \mathcal{B}(PK_{\mathcal{M}}, SK_{\mathcal{B}}, \text{state}_{\mathcal{B}})$ ): This is a protocol allowing  $\mathcal{M}$  to deposit  $\text{coin}$  (containing some metadata  $\text{info}$ ) in its account with  $\mathcal{B}$ . In the protocol,  $\mathcal{M}$  sends  $\text{coin}$  along with some data  $\mu$ . Then,  $\mathcal{B}$

uses a list  $\text{state}_{\mathcal{B}}$  of previously deposited coins to proceed as follows. If  $\text{VerifyCoin}(PK_{\mathcal{B}}, \text{coin}) = 0$  or  $\text{VerifyDeposit}(PK_{\mathcal{B}}, PK_{\mathcal{M}}, \text{coin}, \mu) = 0$ ,  $\mathcal{B}$  outputs  $\perp$ . If  $\text{info}$  and  $PK_{\mathcal{M}}$  exist in the same entry of  $\text{state}_{\mathcal{B}}$ , then  $\mathcal{B}$  returns this entry  $(\text{coin}, PK_{\mathcal{M}}, \mu')$ . If the serial number  $y_S$  derived from  $\text{coin}$  is not in  $\text{state}_{\mathcal{B}}$ , then  $\mathcal{B}$  adds the tuple  $(\text{coin}, PK_{\mathcal{M}}, \mu, y_S)$  to  $\text{state}_{\mathcal{B}}$ . If there is some tuple  $(\text{coin}', PK'_{\mathcal{M}}, \mu', y_S)$  in  $\text{state}_{\mathcal{B}}$ , then  $\mathcal{B}$  outputs such a tuple.

**Identify**  $(PK_{\mathcal{B}}, \text{coin}_1, \text{coin}_2)$ : An algorithm allowing to identify a double spender  $\mathcal{U}$  whenever  $\text{coin}_1$  and  $\text{coin}_2$  share the same serial number. The output of this algorithm is a public key  $PK_{\mathcal{U}}$  and a proof that this public key corresponds to a double spender  $\Pi_G$ .

E-cash systems should provide the following properties whose formal definitions, adapted from [24,9], are provided below.

- **Anonymity**: no coalition of banks and merchants can identify the wallet that a coin originates from.
- **Traceability**: the bank is always able to identify at least one member of a coalition that has spent more than it has withdrawn. This property introduced by Canard *et al.* [12] simultaneously captures the balance and identification properties considered in [11,7].
- **Strong exculpability**: no coalition of banks and merchants can convincingly accuse an innocent user of double-spending.
- **Clearing**: an honest merchant is always able to deposit the received coins. In particular, no coalition of bank, merchants and users can generate a convincing proof that the latter have already been deposited.

**Definition 4.** *An e-cash system provides **anonymity** if there exists an efficient simulator  $\mathcal{S} = (\text{SimParGen}, \text{SimSpend})$  such that no PPT adversary  $\mathcal{A}$  has non-negligible advantage in the **anonymity** game described below:*

1. *The challenger flips a fair coin  $d \leftarrow U(\{0, 1\})$  and runs  $\text{par} \leftarrow \text{ParGen}(1^\lambda, 1^L)$  if  $d = 1$  and  $(\text{par}, \tau_{\text{sim}}) \leftarrow \text{SimParGen}(1^\lambda, 1^L)$  otherwise. In either case, it gives  $\text{par}$  to  $\mathcal{A}$ .*
2.  *$\mathcal{A}$  outputs some public key  $PK_{\mathcal{B}}$  and adaptively invokes the following oracles:*
  - $\mathcal{Q}_{\text{GetKey}}(i)$ : *this oracle generates  $(SK_{\mathcal{U}_i}, PK_{\mathcal{U}_i}) \leftarrow \text{UKeygen}(\text{par})$  if it does not exist yet and returns  $PK_{\mathcal{U}_i}$ .*
  - $\mathcal{Q}_{\text{Withdraw}}(PK_{\mathcal{B}}, i)$ : *this oracle plays the role of user  $\mathcal{U}_i$  – and creates their key pair if it does not exist yet – in an execution of the withdrawal protocol  $\text{Withdraw}(\mathcal{U}(\text{par}, PK_{\mathcal{B}}, SK_{\mathcal{U}_i}), \mathcal{A}(\text{state}))$ , with the adversary  $\mathcal{A}$  playing the role of the bank. At the  $j$ -th query, we denote by  $\mathcal{W}_j$  the user's output which may be a valid wallet or an error message  $\perp$ .*
  - $\mathcal{Q}_{\text{Spend}}(PK_{\mathcal{B}}, i, j, PK_{\mathcal{M}}, \text{info})$ : *the oracle first checks if the wallet  $\mathcal{W}_j$  has been issued to  $\mathcal{U}_i$  by the bank  $\mathcal{B}$  via an invocation of  $\mathcal{Q}_{\text{Withdraw}}(PK_{\mathcal{B}}, i)$ . If not, the oracle outputs  $\perp$ . Otherwise,  $\mathcal{Q}_{\text{Spend}}$  checks if the internal counter  $J$  of  $\mathcal{W}_j$  satisfies  $J < 2^L - 1$ . If  $J = 2^L - 1$ , it outputs  $\perp$ . Otherwise,  $\mathcal{Q}_{\text{Spend}}$  responds as follows:*

- If  $d = 1$ , it runs  $\text{Spend}(\mathcal{U}_i(\mathcal{W}_j, PK_{\mathcal{B}}, PK_{\mathcal{M}}), \mathcal{A}(\text{state}, \text{info}))$  with the merchant-executing  $\mathcal{A}$  in order to spend a coin from  $\mathcal{W}_j$ .
  - If  $d = 0$ ,  $\mathcal{Q}_{\text{Spend}}$  runs  $\text{SimSpend}(\text{par}, \tau_{\text{sim}}, PK_{\mathcal{B}}, PK_{\mathcal{M}}, \text{info})$ .
3. When  $\mathcal{A}$  halts, it outputs a bit  $d' \in \{0, 1\}$  and wins if  $d' = d$ . The adversary's advantage is the distance  $\text{Adv}_{\mathcal{A}}^{\text{anon}}(\lambda) := |\Pr[d' = d] - 1/2|$ , where the probability is taken over all coin tosses.

**Definition 5.** An e-cash system ensures **traceability** if, for any PPT adversary  $\mathcal{A}$ , the experiment below outputs 1 with negligible probability:

1. The challenger generates public parameters  $\text{par} \leftarrow \text{ParGen}(1^\lambda, 1^L)$  and a public key  $(PK_{\mathcal{B}}, SK_{\mathcal{B}}) \leftarrow \text{BKeyGen}(\text{par})$ . It gives  $\text{par}$  and  $PK_{\mathcal{B}}$  to  $\mathcal{A}$ .
2.  $\mathcal{A}$  is granted access to the oracle  $\mathcal{Q}_{\text{Withdraw}}(PK_{\mathcal{U}})$  that plays the role of the bank  $\mathcal{B}$  in an execution of  $\text{Withdraw}(\mathcal{A}(\text{state}), \mathcal{B}(\text{par}, PK_{\mathcal{U}}, SK_{\mathcal{B}}))$  with  $\mathcal{A}$ , acting as a cheating user. After each query, the challenger stores in a database  $T$  the information  $T_{\mathcal{W}} = PK_{\mathcal{U}}$ , or  $\perp$  if the protocol fails.
3. After  $Q_w$  polynomially many queries,  $\mathcal{A}$  outputs coins  $\{\text{coin}_i\}_{i=1}^N$  which are parsed as  $(\text{info}_i, PK_{\mathcal{M}_i}, S_i, \pi_i)$ . The experiment returns 1, unless (at least) one of the following conditions holds (in which case, it returns 0):
  - $N \leq 2^L \cdot Q_w$ ;
  - $\exists (i, j) \in \{1, \dots, N\}^2$  such that  $(\text{info}_i, PK_{\mathcal{M}_i}) = (\text{info}_j, PK_{\mathcal{M}_j})$ ;
  - $\exists i \in \{1, \dots, N\}$  such that  $\text{VerifyCoin}(PK_{\mathcal{B}}, \text{coin}_i) = 0$ ;
  - $\exists (i, j) \in \{1, \dots, N\}^2$  such that  $\text{Identify}(\text{par}, PK_{\mathcal{B}}, \text{coin}_i, \text{coin}_j)$  returns a public key  $PK_{\mathcal{U}}$  that belongs to the database  $T$ .

**Definition 6.** An e-cash system provides **strong exculpability** if no PPT adversary  $\mathcal{A}$  has noticeable success probability in the game below:

1. The challenger runs  $\text{par} \leftarrow \text{ParGen}(1^\lambda, 1^L)$ , gives  $\text{par}$  to  $\mathcal{A}$  and initializes empty sets of honest users  $\mathcal{HU}$ , wallets  $T_{\text{FW}}$  and double spent coins  $T_{\text{ds}}$ .
2.  $\mathcal{A}$  generates  $PK_{\mathcal{B}}$  on behalf of the bank and interacts with these oracles:
  - $\mathcal{Q}_{\text{GetKey}}(i)$ : this oracle generates  $(SK_{\mathcal{U}_i}, PK_{\mathcal{U}_i}) \leftarrow \text{UKeygen}(\text{par})$  if it does not exist yet and returns  $PK_{\mathcal{U}_i}$ , which is added to  $\mathcal{HU}$ .
  - $\mathcal{Q}_{\text{Withdraw}}(PK_{\mathcal{B}}, i)$ : this oracle plays the role of  $\mathcal{U}_i$  – and creates  $(SK_{\mathcal{U}_i}, PK_{\mathcal{U}_i})$  if it does not exist yet – in an execution of  $\text{Withdraw}(\mathcal{U}(\text{par}, PK_{\mathcal{B}}, SK_{\mathcal{U}_i}), \mathcal{A}(\text{state}))$  where  $\mathcal{A}$  plays the role of the bank. At the  $j$ -th such query, we denote by  $\mathcal{W}_j$  the user's output. If the protocol succeeds ( $\mathcal{W}_j = \perp$ ), then  $(j, \mathcal{W}_j)$  is added to  $T_{\text{FW}}$ .
  - $\mathcal{Q}_{\text{Spend}}(PK_{\mathcal{B}}, i, j, PK_{\mathcal{M}}, \text{info})$ : the oracle first checks if the wallet  $\mathcal{W}_j$  was provided to  $\mathcal{U}_i$  via an invocation of  $\mathcal{Q}_{\text{Withdraw}}(\text{par}, PK_{\mathcal{B}}, i)$  using  $T_{\text{FW}}$ . If not, the oracle outputs  $\perp$ . If the internal counter of  $\mathcal{W}_j$  satisfies  $J = 2^\ell - 1$ , then  $\mathcal{W}_j$  is reset to its original state, where  $J = 0$ . Then,  $\mathcal{Q}_{\text{Spend}}$  spends a coin from  $\mathcal{W}_j$  by running  $\text{Spend}(\mathcal{U}_i(\mathcal{W}_j, PK_{\mathcal{B}}, PK_{\mathcal{M}}), \mathcal{A}(\text{state}, \text{info}))$  with  $\mathcal{A}$ . If the resulting coin has the same serial number  $S$  as a previous query  $\mathcal{Q}_{\text{Spend}}(PK_{\mathcal{B}}, i, j, \cdot, \cdot)$  then add  $(i, j, S)$  to  $T_{\text{ds}}$ .

3. When adversary  $\mathcal{A}$  halts, it outputs two coins  $\text{coin}_1, \text{coin}_2$ . It is declared successful if  $\text{Identify}(\text{par}, PK_{\mathcal{B}}, \text{coin}_1, \text{coin}_2) \in \mathcal{HU}$  and  $\forall (i, j), (i, j, S) \notin T_{\text{ds}}$  where  $S$  is the common serial number shared by  $\text{coin}_1$  and  $\text{coin}_2$

**Definition 7.** An e-cash system ensures **clearing** if for any PPT adversary  $\mathcal{A}$ , the probability of  $\mathcal{A}$  winning the **clearing** game below is negligible:

1. The challenger runs  $\text{par} \leftarrow \text{ParGen}(1^\lambda, 1^L)$ , gives  $\text{par}$  to  $\mathcal{A}$  and initializes a set of honest merchants  $\mathcal{HM}$  which is initially empty.
2.  $\mathcal{A}$  generates  $PK_{\mathcal{B}}$  on behalf of the bank and interacts with these oracles:
  - $\mathcal{Q}_{\text{GetKey}}(i)$ : this oracle generates  $(SK_{\mathcal{M}_i}, PK_{\mathcal{M}_i}) \leftarrow \text{MKeygen}(\text{par})$  if it does not exist yet and returns  $PK_{\mathcal{M}_i}$ , which is added in  $\mathcal{HM}$ .
  - $\mathcal{Q}_{\text{Receive}}(PK_{\mathcal{B}}, i)$ : this oracle plays the role of a merchant – and creates  $(SK_{\mathcal{M}_i}, PK_{\mathcal{M}_i})$  if it does not exist yet – in an execution of  $\text{Spend}(\mathcal{A}(\text{state}), \mathcal{M}_i(SK_{\mathcal{M}}, PK_{\mathcal{U}}, \text{info}))$  where  $\mathcal{A}$  plays the role of the user. At the  $j$ -th query, we denote by  $\text{coin}_j$  the merchant's output.
  - $\mathcal{Q}_{\text{Deposit}}(PK_{\mathcal{B}}, i, j)$ : this oracle plays the role of the merchant in an execution of  $\text{Deposit}(\mathcal{M}_i(SK_{\mathcal{M}_i}, \text{coin}_j, PK_{\mathcal{B}}), \mathcal{A}(\text{state}))$  with  $\mathcal{A}$  playing the role of  $\mathcal{B}$ . It however aborts if  $PK_{\mathcal{M}_i} \notin \mathcal{HM}$ , if  $\text{coin}_j$  has not been received by merchant  $i$  or if it has already been deposited.
3. When  $\mathcal{A}$  halts, it outputs a tuple  $(PK_{\mathcal{M}}, \text{coin}, \mu)$ . The adversary wins if  $PK_{\mathcal{M}} \in \mathcal{HM}$ ,  $\text{VerifyDeposit}(PK_{\mathcal{B}}, PK_{\mathcal{M}}, \text{coin}, \mu) = 1$  and  $\text{coin}$  has not been involved in a previous  $\mathcal{Q}_{\text{Deposit}}$  query.

### 3 Intuition

The core of an e-cash system is the pair constituted by the serial number  $\mathbf{y}_S$  and the double-spending tag  $\mathbf{y}_T$  of a coin. Besides zero-knowledge proofs, they are essentially the only elements made public during a spending and therefore must comply with very strong anonymity requirements while allowing the identification of double-spenders. In addition, it should be possible to (efficiently) prove that they are well-formed, which rules out most simple constructions.

Designing such elements is thus far from trivial which partially explains why most e-cash systems have followed the elegant idea proposed by Camenisch *et al* [11]. It relies on a pseudo-random function PRF as follows. For a wallet of  $N = 2^L$  coins, a first seed  $\mathbf{k}$  is used to generate  $N$  pseudo-random values  $\text{PRF}_{\mathbf{k}}(i)$ , for  $i \in [1, N]$ , acting as the coins' serial numbers. Meanwhile, a second seed  $\mathbf{t}$  allows generating independent values  $\text{PRF}_{\mathbf{t}}(i)$  acting as one-time pads on the spender's identity. The concealed identity constitutes the double-spending tag.

Any user can generate at most  $N$  fresh pairs  $(\text{PRF}_{\mathbf{k}}(i), \text{PRF}_{\mathbf{t}}(i))$  per wallet. In case of double-spending, a pair must have been re-used, meaning that a serial number  $\text{PRF}_{\mathbf{k}}(i)$  will appear twice in the bank database, thus making frauds easy to detect. Moreover, in such a case, the spender's identity will be masked using the same value  $\text{PRF}_{\mathbf{t}}(i)$ . An appropriate combination of the corresponding double-spending tags thus allows to remove  $\text{PRF}_{\mathbf{t}}(i)$  and so to identify the

defrauder. Some adjustments are necessary in the lattice setting [24], but the high-level principle remains the same.

However, Bourse *et al* [9] recently showed that this approach may fail to provide a sound proof of exculpability in many cases. Indeed, the identity returned by the identification algorithm is a complex mix of PRF outputs, public keys and random values, most of them being totally controlled by the adversary. It is therefore impossible to guarantee that the returned identity corresponds to the author of these fraudulent payments nor even to guarantee that both payments have been performed by the same user.

In [9], Bourse *et al.* point out that this problem is partially due to a misidentification of the properties that must be satisfied by the pseudo-random function PRF. They therefore propose to strengthen the requirements on PRF, introducing in particular a notion of collision resistance that essentially states the difficulty of finding  $(s, s', i, i')$  such that  $\text{PRF}_s(i) = \text{PRF}_{s'}(i')$ . Assuming that the PRF satisfies such suitable properties, they prove security of generic constructions that are reminiscent of the seminal scheme proposed by Camenisch *et al.* An interesting aspect of [9] is thus the rehabilitation of the original intuition of compact e-cash [11] that has been common to all following works.

Unfortunately, this is done by relying on unconventional security notions for PRFs that have not been considered by designers of such functions. Bourse *et al.* show that, under suitable assumptions, these notions are actually already satisfied by some PRFs in cyclic groups, but similar results are not known in the lattice setting. Indeed, existing lattice-based PRFs are not known to provide collision-resistance in this strong sense, which prevents instantiation of their frameworks in this setting. Concretely, this means that secure lattice-based e-cash systems are not known to exist for the time being.

In this work, we choose a very different strategy that we believe to be better suited for the lattice setting as it does not rely on collision-resistant PRFs.

Our first step is to modify the construction of the serial numbers to ensure that collisions only occur for spendings performed by the same user. In [9], this is achieved by using the same seed (but different public parameters) to generate all the pseudo-random values used during a spending. Assuming that pseudo-randomness still holds in this context and that collision resistance is achieved by some of the PRFs, they prove that a collision only occurs for spendings involving the same seed. They are then able to prove that the use of the same seed implies the involvement of the same user, and so on until proving exculpability of their construction. Here, we still use a PRF as a source of pseudo-random values but our serial numbers are not simply the outputs of such a function. We indeed want to reuse the same pseudo-random values for different parts of our serial numbers and double-spending tags to rule out the adversarial strategy pointed out in [9]. To achieve this while retaining anonymity, we use the notion of a lossy trapdoor function  $F_{\text{LTF}}$  introduced in [29] and more specifically, the elegant instantiation based on LWE proposed in [32] related to the original construction in [29].

The first element of  $\mathbf{y}_S$  is now  $F_{\text{LTF}}(\text{PRF}_{\mathbf{k}}(i))$ , which still allows to extract random bits from  $\text{PRF}_{\mathbf{k}}(i)$  using a universal hash function  $H_{\text{UH}}$ , as proved in [29].

We can thus incorporate  $PK_{\mathcal{U}} + H_{\text{UH}}(\text{PRF}_{\mathbf{k}}(i))$  in  $\mathbf{y}_S$  while ensuring anonymity of the user  $\mathcal{U}$  that owns the public key  $PK_{\mathcal{U}}$ . In the exculpability proof, we will generate  $F_{\text{LTF}}$  in the injective mode, thus ensuring that a collision  $\mathbf{y}_S = \mathbf{y}_{S'}$  can only occur when the same value  $\text{PRF}_{\mathbf{k}}(i)$  is used for both transactions. Together with  $PK_{\mathcal{U}} + H_{\text{UH}}(\text{PRF}_{\mathbf{k}}(i)) = PK_{\mathcal{U}'} + H_{\text{UH}}(\text{PRF}_{\mathbf{k}}(i))$ , this implies  $PK_{\mathcal{U}} = PK_{\mathcal{U}'}$ .

We then adapt this idea to double-spending tags. We similarly extract random bits from  $\text{PRF}_{\mathbf{k}}(i)$  using a different universal hash function  $H'_{\text{UH}}$  to define  $\mathbf{y}_T = PK_{\mathcal{U}} + \text{FRD}(R) \cdot H'_{\text{UH}}(\text{PRF}_{\mathbf{k}}(i))$ , where  $\text{FRD}(R)$  is some public matrix specific to the transaction. As  $\text{PRF}_{\mathbf{k}}(i)$  and the public key  $PK_{\mathcal{U}}$  are the same for both transactions, the formula  $\mathbf{y}_T - \text{FRD}(R) \cdot [(\text{FRD}(R) - \text{FRD}(R'))^{-1} \cdot (\mathbf{y}_T - \mathbf{y}_{T'})]$  necessarily returns such a public key whose owner is guilty of double-spending.

As far as efficiency goes, we essentially add some matrix-vector products to the construction of [24]. Moreover, since all of these matrices are public, a NIZK proof of correct computations can be produced using the framework provided in [24] or the more efficient techniques in Section 5.

## 4 Construction

We present a new e-cash system that overcomes the aforementioned issues in the proof of exculpability. We use the PRF from [8] that allows for a simpler description of our system. We nevertheless explain in Section 7 how to improve efficiency by using the alternative PRF from [5]. While the **Withdraw** protocol is a simplification of [24], the **Spend** protocol is very different in the way to construct coin serial numbers and security tags. Additional details on the zero-knowledge arguments of knowledge used in our construction are given in Section 5.

**ParGen**( $1^\lambda, 1^L$ ): Given security parameter  $\lambda$  and integer  $L > 0$  such that  $2^L$  is the desired number of coins per wallet issued, perform the following:

1. Choose secure public parameters  $\text{par}_{\text{PRF}} = (m, n, p, q, \mathbf{P}_0, \mathbf{P}_1)$  for the BLMR PRF family [8]. Namely,
  - a. For  $n = \mathcal{O}(\lambda)$ , set  $\alpha = 2^{-\omega(\log^{1+c}(n))}$  for some constant  $c > 0$ ; a prime  $p = 2^{\log^{1+c}(n)}$ ; a prime power  $q = \mathcal{O}(\sqrt{n}/\alpha)$  such that  $p$  divides  $q$ ; and  $m = n \cdot \lceil \log q \rceil$ .
  - b. Sample  $\mathbf{P}_0, \mathbf{P}_1 \leftarrow U(\{0, 1\}^{m \times m})$  over  $\mathbb{Z}_q$ -invertible matrices.
2. Choose parameters  $\text{par}_{\text{sig}} = (q_s, \ell, \sigma, (m_i)_{i=0}^3, m_s, m_f)$  for a signature scheme allowing to sign committed values [23]. Namely,
  - a. Choose a prime power modulus  $q_s = \tilde{\mathcal{O}}(n^3)$  dividing  $q$ , an integer  $\ell = \Theta(\lambda)$  and a Gaussian parameter  $\sigma = \Omega(\sqrt{n \log q_s \log n})$ . Set  $\delta_{q_s-1} = \lceil \log_2(q_s) \rceil$ ,  $\delta_{q-1} = \lceil \log_2(q) \rceil$  and  $\delta_{p-1} = \lceil \log_2(p) \rceil$ . Define the message block lengths  $m_0 = m_s := 2n\delta_{q_s-1}$ , as well as  $m_1 = m$  and  $m_2 = \tilde{m} := m\delta_{q-1}$ .
  - b. Sample  $\mathbf{D}'_0, \mathbf{D}''_0 \leftarrow U(\mathbb{Z}_{q_s}^{n \times m_0})$  and  $\mathbf{D}_i \leftarrow U(\mathbb{Z}_{q_s}^{n \times m_i})$ , for  $i \in \{1, 2\}$ , and define the commitment key to be  $CK := (\mathbf{D}_0 := [\mathbf{D}'_0 | \mathbf{D}''_0], \mathbf{D}_1, \mathbf{D}_2)$ .
  - c. Sample  $\mathbf{F} \leftarrow U(\mathbb{Z}_p^{n \times m})$ .

3. Choose parameters  $\text{par}_{\text{LTF}}$  for the lossy trapdoor function of [29]. In terms of  $n_{\text{LTF}} = \tilde{O}(\lambda)$  for constant  $c > 0$ , these consist of moduli  $q_{\text{LTF}} = \Theta(n_{\text{LTF}}^{1+1/\gamma})$  that divides  $q$  and  $p_{\text{LTF}} = \Theta(n_{\text{LTF}}^{1/\gamma})$  for some constant  $\gamma < 1$ ; matrix dimensions  $n_{\text{LTF}}$  and  $m_{\text{LTF}} = \Theta(n_{\text{LTF}} \log q_{\text{LTF}})$  and  $\bar{n}_{\text{LTF}} = \bar{m}_{\text{LTF}} / \log p_{\text{LTF}}$  such that  $p_{\text{LTF}} < q_{\text{LTF}} / 4n_{\text{LTF}}$ ; and an LWE error rate  $\alpha_{\text{LTF}} = \Theta(\sqrt{n}/q_{\text{LTF}})$ . We additionally require that  $m_{\text{LTF}} = m \cdot \lceil \log p \rceil$ . Then, select an evaluation key  $ek_{\text{LTF}}$  for a lossy trapdoor function in injective mode  $F_{\text{LTF}} : \{0, 1\}^{m_{\text{LTF}}} \rightarrow \mathbb{Z}_{q_{\text{LTF}}}^{n_{\text{LTF}} + \bar{n}_{\text{LTF}}}$ , meaning that  $ek_{\text{LTF}} = (\mathbf{A}_{\text{LTF}}, \mathbf{U}_{\text{LTF}})$  consists of a random  $\mathbf{A}_{\text{LTF}} \leftarrow U(\mathbb{Z}_q^{n_{\text{LTF}} \times m_{\text{LTF}}})$  and a matrix

$$\mathbf{U}_{\text{LTF}} = \mathbf{S}_{\text{LTF}}^\top \cdot \mathbf{A}_{\text{LTF}} + \mathbf{E}_{\text{LTF}}^\top + \mathbf{G}_{\text{LTF}}^\top \in \mathbb{Z}_{q_{\text{LTF}}}^{\bar{n}_{\text{LTF}} \times m_{\text{LTF}}},$$

for some  $\mathbf{S}_{\text{LTF}} \leftarrow U(\mathbb{Z}_{q_{\text{LTF}}}^{n_{\text{LTF}} \times \bar{n}_{\text{LTF}}})$ ,  $\mathbf{E}_{\text{LTF}} \leftarrow D_{\mathbb{Z}^{m_{\text{LTF}} \times \bar{n}_{\text{LTF}}, \alpha_{\text{LTF}} q_{\text{LTF}}}}$  and  $\mathbf{G}_{\text{LTF}}$  referred to in the preliminaries.

4. Choose an integer  $\bar{p} > 0$  such that  $\bar{p} < p/2$  which will define a challenge space  $\{-\bar{p}, \dots, \bar{p}\}$  for the argument system of [33]. Choose a hash function  $H_{\text{FS}} : \{0, 1\}^* \rightarrow \{-\bar{p}, \dots, \bar{p}\}^\kappa$ , for some  $\kappa = \mathcal{O}(\lambda / \log \bar{p})$ , which will be modelled as a random oracle in the security analysis.
5. Choose a full-rank difference function  $\text{FRD} : \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^{n \times n}$  such as the one in [1]; two universal hash functions  $H_{\text{UH}} : \mathbb{Z}_p^{m_{\text{LTF}}} \rightarrow \mathbb{Z}_p^n$ ,  $H'_{\text{UH}} : \mathbb{Z}_p^{m_{\text{LTF}}} \rightarrow \mathbb{Z}_p^n$  keyed by two uniformly random matrices  $\mathbf{U}_{\text{UH}}, \mathbf{U}'_{\text{UH}} \leftarrow U(\mathbb{Z}_p^{n \times m_{\text{LTF}}})$ ; and a collision resistant hash function  $H_0 : \{0, 1\}^* \rightarrow \mathbb{Z}_p^n \setminus \{\mathbf{0}^n\}$ .
6. Select a digital signature algorithm<sup>7</sup>  $\Sigma$  able to sign any bitstring.

The final output is  $\text{par} = (\text{par}_{\text{PRF}}, \text{par}_{\text{sig}}, \text{par}_{\text{LTF}}, \mathbf{F}, \text{FRD}, \mathbf{U}_{\text{UH}}, \mathbf{U}'_{\text{UH}}, H_0, ek_{\text{LTF}}, H_{\text{FS}}, CK, \Sigma)$ .

**BKeyGen**( $1^\lambda, \text{par}$ ): The bank  $\mathcal{B}$  generates a key pair for the signature scheme by conducting the following steps.

1. Sample  $(\mathbf{A}, \mathbf{T}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^{m_s}, q_s)$  (using Lemma 8 in Section A of the Supplementary Material) so that  $\mathbf{T}_{\mathbf{A}}$  is a short basis of  $\Lambda_{q_s}^\perp(\mathbf{A})$  that allows  $\mathcal{B}$  to sample Gaussian vectors in  $\Lambda_{q_s}^\perp(\mathbf{A})$  with parameter  $\sigma$ .
2. Choose uniform  $\mathbf{A}_0, \dots, \mathbf{A}_\ell \leftarrow U(\mathbb{Z}_{q_s}^{n \times m_s})$ .
3. Choose  $\mathbf{D} \leftarrow U(\mathbb{Z}_{q_s}^{n \times m_s/2})$  and  $\mathbf{u} \leftarrow U(\mathbb{Z}_{q_s}^n)$ .

The key pair consists of  $PK_{\mathcal{B}} := (\mathbf{A}, \{\mathbf{A}_j\}_{j=0}^\ell, \mathbf{D}, \mathbf{u})$  and  $SK_{\mathcal{B}} := \mathbf{T}_{\mathbf{A}}$ .

**UKeyGen**( $1^\lambda, \text{par}$ ): Choose a secret key  $SK_{\mathcal{U}} := \mathbf{e}_u \leftarrow U(\{0, 1\}^m)$  and set the public key to be  $PK_{\mathcal{U}} := \mathbf{F} \cdot \mathbf{e}_u \in \mathbb{Z}_p^n$ .

**MKeyGen**( $1^\lambda, \text{par}$ ): Generate and output  $(SK_{\mathcal{M}}, PK_{\mathcal{M}}) \leftarrow \Sigma.\text{Keygen}(1^\lambda)$ .

**Withdraw**( $\mathcal{U}(PK_{\mathcal{B}}, SK_{\mathcal{U}}, 2^L), \mathcal{B}(PK_{\mathcal{U}}, SK_{\mathcal{B}}, 2^L)$ ): A user  $\mathcal{U}$  withdraws a wallet with  $2^L$  coins from a bank  $\mathcal{B}$  by engaging in the following protocol:

1.  $\mathcal{U}$  picks a PRF key  $\mathbf{k} \leftarrow U(\mathbb{Z}_q^m)$  and computes its binary decomposition  $\tilde{\mathbf{k}} = \text{vdec}_{m, q-1}(\mathbf{k}) \in \{0, 1\}^{\bar{m}}$ . Then,  $\mathcal{U}$  commits to the 2-block message  $(\mathbf{e}_u, \tilde{\mathbf{k}}) \in \{0, 1\}^m \times \{0, 1\}^{\bar{m}}$  by sampling  $\mathbf{r}_0 \leftarrow D_{\mathbb{Z}^{m_s}, \sigma}$  and sending

$$\mathbf{c}_{\mathcal{U}} = \mathbf{D}_0' \cdot \mathbf{r}_0 + \mathbf{D}_1 \cdot \mathbf{e}_u + \mathbf{D}_2 \cdot \tilde{\mathbf{k}} \in \mathbb{Z}_{q_s}^n$$

<sup>7</sup> Any EUF-CMA secure scheme  $\Sigma$  can be selected here.

to  $\mathcal{B}$ . In addition,  $\mathcal{U}$  generates an interactive zero-knowledge argument of knowledge of an opening  $(\mathbf{r}_0, \mathbf{e}_u, \tilde{\mathbf{k}})$  such that  $PK_{\mathcal{U}} = \mathbf{F} \cdot \mathbf{e}_u$  with  $\mathcal{B}$ . This argument of knowledge can be instantiated using the methods of [33] by applying the technique of [14] to parallel repetitions.<sup>8</sup>

2. If the argument of  $\mathcal{U}$  verifies, then  $\mathcal{B}$  extends the commitment  $\mathbf{c}_{\mathcal{U}}$  by sampling  $\mathbf{r}_1 \leftarrow D_{\mathbb{Z}^{m_s}, \sigma}$ , and computing  $\mathbf{c}'_{\mathcal{U}} = \mathbf{c}_{\mathcal{U}} + \mathbf{D}_0'' \cdot \mathbf{r}_1$ . Next  $\mathcal{B}$  chooses  $\tau \leftarrow U(\{0, 1\}^\ell)$ , defines  $\mathbf{u}_{\mathcal{U}} = \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{n, q_s-1}(\mathbf{c}'_{\mathcal{U}})$ , sets

$$\mathbf{A}_{\tau} := [\mathbf{A} | \mathbf{A}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{A}_j] \in \mathbb{Z}_{q_s}^{n \times 2m_s}$$

and computes a short basis  $\mathbf{T}_{\tau}$  of  $\Lambda_{q_s}^{\perp}(\mathbf{A}_{\tau})$  using  $\mathbf{T}_A$ . Using  $\mathbf{T}_{\tau}$ , it then samples a short vector  $\mathbf{v} \leftarrow D_{\Lambda_{q_s}^{\mathbf{u}_{\mathcal{U}}(\mathbf{A}_{\tau}), \sigma}}$  and sends  $(\tau, \mathbf{v}, \mathbf{r}_1)$  to  $\mathcal{U}$ .

3.  $\mathcal{U}$  verifies that  $\|\mathbf{v}\| \leq \sigma\sqrt{2m_s}$ ,  $\|\mathbf{r}_1\| \leq \sigma\sqrt{m_s}$  and

$$\mathbf{A}_{\tau} \cdot \mathbf{v} = \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{n, q_s-1}(\mathbf{c}_{\mathcal{U}} + \mathbf{D}_0'' \cdot \mathbf{r}_1) \in \mathbb{Z}_{q_s}^n.$$

If so,  $\mathcal{U}$  sets  $\mathbf{r} = (\mathbf{r}_0^{\top} \parallel \mathbf{r}_1^{\top})^{\top} \in \mathbb{Z}_{q_s}^{2m_s}$  and stores the wallet  $\mathcal{W} := (\mathbf{e}_u, \mathbf{k}, \text{Sig}_{\mathcal{B}} = (\tau, \mathbf{v}, \mathbf{r}), J = 0)$  whereas  $\mathcal{B}$  records a debit of  $2^L$  for the account associated to  $PK_{\mathcal{U}}$ .

**Spend** ( $\mathcal{U}(\mathcal{W}, PK_{\mathcal{B}}, PK_{\mathcal{M}}), \mathcal{M}(SK_{\mathcal{M}}, PK_{\mathcal{B}}, \text{info})$ ): A user  $\mathcal{U}$  in possession of a wallet  $\mathcal{W} = (\mathbf{e}_u, \mathbf{k}, \text{Sig}_{\mathcal{B}} = (\tau, \mathbf{v}, \mathbf{r}), J)$  wants to spend a coin with  $\mathcal{M}$ . If  $J > 2^L - 1$ ,  $\mathcal{U}$  outputs  $\perp$ . Otherwise, they run the following protocol:

1.  $\mathcal{U}$  generates a digital coin by first hashing the transaction information to  $R = H_0(PK_{\mathcal{M}}, \text{info}) \in \mathbb{Z}_p^n$  before conducting the following steps.
  - a. Compute a BLMR PRF evaluation on the standard binary representation of  $J$  in  $\{0, 1\}^L$  using key  $\mathbf{k} \in \mathbb{Z}_q^m$ ; i.e., set

$$\mathbf{y}_{\mathbf{k}} = \left\lfloor \prod_{i=1}^L \mathbf{P}_{J[L+1-i]} \cdot \mathbf{k} \right\rfloor_p$$

and let  $\tilde{\mathbf{y}}_{\mathbf{k}} = \text{bin}_p(\mathbf{y}_{\mathbf{k}}) \in \{0, 1\}^{m_{\text{LTF}}}$  its standard bit-decomposition.

- b. Using  $ek_{\text{LTF}}$ , compute  $\mathbf{y}_1 = F_{\text{LTF}}(\tilde{\mathbf{y}}_{\mathbf{k}})$  and  $\mathbf{y}_2 = PK_{\mathcal{U}} + H_{\text{UH}}(\tilde{\mathbf{y}}_{\mathbf{k}})$  to form the serial number  $\mathbf{y}_S := (\mathbf{y}_1, \mathbf{y}_2) \in \mathbb{Z}_{q_{\text{LTF}}}^{n_{\text{LTF}} + \tilde{n}_{\text{LTF}}} \times \mathbb{Z}_p^n$ .
- c. Compute the security tag  $\mathbf{y}_T = PK_{\mathcal{U}} + \text{FRD}(R) \cdot H'_{\text{UH}}(\tilde{\mathbf{y}}_{\mathbf{k}}) \in \mathbb{Z}_p^n$ .
- d. Generate a non-interactive argument of knowledge  $\pi_K$  to show knowledge of  $(J, \mathbf{k}, \mathbf{e}_u, (\tau, \mathbf{v}, \mathbf{r}))$  such that:
  - The vector  $\mathbf{k}$  and secret key  $\mathbf{e}_u$  associated with  $\mathcal{W}$  and  $PK_{\mathcal{U}}$  have been certified by  $\mathcal{B}$  through the signature  $(\tau, \mathbf{v}, \mathbf{r})$ .
  - $\mathbf{y}_S$  and  $\mathbf{y}_T$  were computed correctly using  $\text{par}$ , the secret key  $\mathbf{e}_u$ , the PRF seed  $\mathbf{k}$  and a valid  $J \in \{0, \dots, 2^{L-1}\}$ .

<sup>8</sup> Technically, we should add a CRS to  $\text{par}$  but we leave this implicit for simplicity



More precisely, letting  $\mathbf{y}_S = (\mathbf{y}_1, \mathbf{y}_2)$ ,  $\pi_K$  argues knowledge of  $(J, \mathbf{k}, \mathbf{e}_u, (\tau, \mathbf{v}, \mathbf{r}))$  where  $J \in \{0, 1\}^L$ ,  $\mathbf{k} \in \mathbb{Z}_q^m$ ,  $\mathbf{e}_u \in \{0, 1\}^m$ ,  $\tau \in \{0, 1\}^\ell$ ,  $\mathbf{v} \in \mathbb{Z}^{2m_s}$  s.t.  $\|\mathbf{v}\|_\infty \leq \sigma\sqrt{2m_s}$  and  $\mathbf{r} \in \mathbb{Z}^{m_s}$  s.t.  $\|\mathbf{r}\|_\infty \leq \sigma\sqrt{2m_s}$ , satisfying the relations

$$\begin{aligned}
& [\mathbf{A} \mid \mathbf{A}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{A}_j] \cdot \mathbf{v} \\
&= \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{n, q_s-1} \left( [\mathbf{D}'_0 \mid \mathbf{D}''_0] \cdot \mathbf{r} + \mathbf{D}_1 \cdot \mathbf{e}_u + \mathbf{D}_2 \cdot \text{vdec}_{m, q-1}(\mathbf{k}) \right) \\
\mathbf{y}_1 &= F_{\text{LTF}} \left( \text{bin}_p \left( \left\lfloor \prod_{i=1}^L \mathbf{P}_{J[L+1-j]} \cdot \mathbf{k} \right\rfloor_p \right) \right) \in \mathbb{Z}_{q_{\text{LTF}}}^{n_{\text{LTF}} + \bar{n}_{\text{LTF}}} \\
\mathbf{y}_2 &= \mathbf{F} \cdot \mathbf{e}_u + H_{\text{UH}} \left( \text{bin}_p \left( \left\lfloor \prod_{i=1}^L \mathbf{P}_{J[L+1-j]} \cdot \mathbf{k} \right\rfloor_p \right) \right) \in \mathbb{Z}_p^n \\
\mathbf{y}_T &= \mathbf{F} \cdot \mathbf{e}_u + \text{FRD}(R) \cdot H'_{\text{UH}} \left( \text{bin}_p \left( \left\lfloor \prod_{i=1}^L \mathbf{P}_{J[L+1-j]} \cdot \mathbf{k} \right\rfloor_p \right) \right) \in \mathbb{Z}_p^n
\end{aligned}$$

The non-interactive argument  $\pi_K$  is produced by running the proof described in Section 5.2  $\kappa = \mathcal{O}(\lambda/\log \bar{p})$  times in parallel and using the Fiat-Shamir heuristic with random oracle  $H_{\text{FS}}$ . We may write

$$\pi_K = \left( \{\text{Comm}_{K,j}\}_{j=1}^{\kappa}, \text{Chall}_K, \{\text{Resp}_{K,j}\}_{j=1}^{\kappa} \right)$$

where  $\text{Chall}_K = H_{\text{FS}}(\text{par}, R, \mathbf{y}_S, \mathbf{y}_T, \{\text{Comm}_{K,j}\}_{j=1}^{\kappa})$ .

$\mathcal{U}$  sends  $\text{coin} = (\text{info}', PK_{\mathcal{M}}, \mathbf{y}_S, \mathbf{y}_T, \pi_K)$  to  $\mathcal{M}$ .

2. If  $\text{info}' = \text{info}$  and  $\text{VerifyCoin}(\text{par}, PK_{\mathcal{B}}, \text{coin})$  outputs 1, then  $\mathcal{M}$  outputs coin. Otherwise,  $\mathcal{M}$  outputs  $\perp$ . In either case,  $\mathcal{U}$  outputs an updated wallet  $\mathcal{W}'$  where  $J$  is increased by 1.

**VerifyCoin** ( $PK_{\mathcal{B}}, \text{coin}$ ): Parse the coin as  $\text{coin} = (\text{info}, PK_{\mathcal{M}}, \mathbf{y}_S, \mathbf{y}_T, \pi_K)$  and output 1 if and only if  $\pi_K$  verifies.

**VerifyDeposit** ( $PK_{\mathcal{B}}, PK_{\mathcal{M}}, \text{coin}, \mu$ ): If  $\text{VerifyCoin}(PK_{\mathcal{B}}, \text{coin}) = 0$ , return 0. Otherwise, return 1 if and only if  $\mu$  is a valid signature on coin with respect to  $PK_{\mathcal{M}}$ : i.e.,  $\Sigma.\text{Verify}(PK_{\mathcal{M}}, \mu, \text{coin}) = 1$ .

**Deposit** ( $\mathcal{M}(SK_{\mathcal{M}}, \text{coin}, PK_{\mathcal{B}}), \mathcal{B}(PK_{\mathcal{M}}, SK_{\mathcal{B}}, \text{state}_{\mathcal{B}})$ ):  $\mathcal{M}$  and  $\mathcal{B}$  interact in the following way.

1.  $\mathcal{M}$  sends  $\text{coin} = (\text{info}, PK_{\mathcal{M}}, \mathbf{y}_S, \mathbf{y}_T, \pi_K)$  to  $\mathcal{B}$  along with a signature  $\mu = \Sigma.\text{Sign}(SK_{\mathcal{M}}, \text{coin})$ .

2. If  $\text{VerifyDeposit}(PK_{\mathcal{B}}, PK_{\mathcal{M}}, \text{coin}, \mu) = 0$  or  $\text{VerifyCoin}(PK_{\mathcal{B}}, \text{coin}) = 0$ , then  $\mathcal{B}$  outputs  $\perp$ . If  $\text{info}$  and  $PK_{\mathcal{M}}$  are found in  $\text{state}_{\mathcal{B}}$ , then  $\mathcal{B}$  outputs the corresponding entry  $(\text{coin}', PK_{\mathcal{M}}, \mu', \mathbf{y}'_S)$ . If the serial number  $\mathbf{y}_S$  contained in  $\text{coin}$  is not found in  $\text{state}_{\mathcal{B}}$ , then  $\mathcal{B}$  accepts the coin, adds the tuple  $(\text{coin}, PK_{\mathcal{M}}, \mu, \mathbf{y}_S)$  to  $\text{state}_{\mathcal{B}}$  and credits  $\mathcal{M}$ 's account. If there exists a tuple  $(\text{coin}, PK'_{\mathcal{M}}, \mu', \mathbf{y}_S)$  in  $\text{state}_{\mathcal{B}}$ , then  $\mathcal{B}$  outputs such a tuple.

**Identify**  $(PK_{\mathcal{B}}, \text{coin}_1, \text{coin}_2)$ : Parse  $\text{coin}_i = (\text{info}_i, PK_{\mathcal{M}_i}, \mathbf{y}_{S,i}, \mathbf{y}_{T,i}, \pi_{K,i})$  for each  $i \in \{1, 2\}$ . If any of the following conditions hold, output  $\perp$ :

- $\mathbf{y}_{S,1} \neq \mathbf{y}_{S,2}$ ,
- $\text{VerifyCoin}(\text{par}, PK_{\mathcal{B}}, \text{coin}_1)$  or  $\text{VerifyCoin}(\text{par}, PK_{\mathcal{B}}, \text{coin}_2) \neq 1$ ,
- $(\text{info}_1, PK_{\mathcal{M}_1}) = (\text{info}_2, PK_{\mathcal{M}_2})$ .

Otherwise, compute  $\mathbf{y}'_T = (\text{FRD}(R_1) - \text{FRD}(R_2))^{-1} \cdot (\mathbf{y}_{T,1} - \mathbf{y}_{T,2}) \in \mathbb{Z}_p^n$  with  $R_i = H_0(PK_{\mathcal{M}_i}, \text{info}_i)$  and set  $PK_{\mathcal{U}} = \mathbf{y}_{T,1} - \text{FRD}(R_1) \cdot \mathbf{y}'_T \in \mathbb{Z}_p^n$ . Note that this calculation is performed using publicly known values, so the proof of guilt of a double spender is simply  $\Pi_G = (\text{coin}_1, \text{coin}_2)$ . The output of this algorithm is then the pair  $(PK_{\mathcal{U}}, \Pi_G)$ .

## 5 Zero-Knowledge Arguments with Soundness Error $1/\text{poly}(\lambda)$ in Standard Lattices

We proceed in two steps to describe the ZKAoK used to spend a coin. We first describe an argument of knowledge of a  $(\text{seed}, \text{input})$  pair generating a given BLMR evaluation. We then extend this to capture the whole statement proved by a user during a spending. For the ZKAoK in the withdrawal protocol, we directly rely on results of [33]. Throughout our construction, we use the argument system of Yang *et al.* [33] which was originally proved computationally honest-verifier ZK (HVZK) with polynomial soundness error. However, we can use known techniques to transform parallel repetitions of this protocol into a 3-round, malicious verifier ZK protocol with negligible soundness error in the CRS model [14]. This is how we instantiate the *interactive* ZKAoK in the withdrawal protocol. In the spend protocol, we use the standard result that the Fiat-Shamir transform [18] applied to parallel repetitions of an HVZK protocol yields a NIZK argument in the ROM. We also note that one may use a statistically hiding configuration of the commitment scheme from [6] instead of the more efficient computationally hiding configuration chosen in [33] to obtain *statistical* ZK arguments.

### 5.1 Zero-Knowledge Arguments for the BLMR PRF

We extend the protocol of Yang *et al.* [33] to build a ZKAoK of a  $(\text{seed}, \text{input})$  pair producing a given BLMR evaluation. A similar result for the GGM-based PRF implied by [5] is provided in Section B of the Supplementary Material, leading to a more efficient instantiation.

In [33], Yang *et al.* provide an argument of knowledge for the “instance-relation” set given by

$$\mathcal{R}^* = \left\{ ((\mathbf{M}', \mathbf{y}', \mathcal{M}), \mathbf{x}') : \forall (h, i, j) \in \mathcal{M}, \mathbf{x}'[h] = \mathbf{x}'[i] \cdot \mathbf{x}'[j] \mod q \wedge \mathbf{M}' \cdot \mathbf{x}' = \mathbf{y}' \mod q \right\}. \quad (1)$$

where  $\mathbf{M}' \in \mathbb{Z}_q^{m' \times n'}$ ,  $\mathbf{y}' \in \mathbb{Z}_q^{m'}$  and  $\mathcal{M} \subseteq [n'] \times [n'] \times [n']$ , for some prime power  $q$ . The tuple  $(\mathbf{M}', \mathbf{y}', \mathcal{M})$  is the instance whereas  $\mathbf{x}' \in \mathbb{Z}_q^{n'}$  is the witness. By carefully crafting each of these elements, we show that a proof of correct evaluation of the BLMR PRF is an instance of this argument of knowledge.

Indeed, recall that, for any seed  $\mathbf{k}$  and input  $x \in \{0, 1\}^L$ , the PRF output is defined as  $\mathbf{y} = \left[ \prod_{i=1}^L \mathbf{P}_{x_{L+1-i}} \cdot \mathbf{k} \right]_p$ , where  $\mathbf{P}_0, \mathbf{P}_1 \in \{0, 1\}^{m \times m}$  are public parameters and  $p$  is a prime power dividing  $q$ . If we write  $\mathbf{y}_j = \prod_{i=L+1-j}^L \mathbf{P}_{x_{L+1-i}} \cdot \mathbf{k}$  for  $j \in [L]$ , we can represent a PRF evaluation using the linear system over  $\mathbb{Z}_q$ :

$$\begin{aligned} \mathbf{y}_1 - \mathbf{P}_0 \cdot (1 - x_1) \mathbf{k} - \mathbf{P}_1 \cdot x_1 \mathbf{k} &= \mathbf{0} \\ \mathbf{y}_2 - \mathbf{P}_0 \cdot (1 - x_2) \mathbf{y}_1 - \mathbf{P}_1 \cdot x_2 \mathbf{y}_1 &= \mathbf{0} \\ &\vdots \\ \mathbf{y}_L - \mathbf{P}_0 \cdot (1 - x_L) \mathbf{y}_{L-1} - \mathbf{P}_1 \cdot x_L \mathbf{y}_{L-1} &= \mathbf{0} \\ \mathbf{y}_L - \mathbf{e} &= \frac{q}{p} \cdot \mathbf{y} \end{aligned}$$

where  $\mathbf{e} \in [0, q/p]^m$ . This system is a linear system in the (quadratic) unknowns  $(1 - x_1) \mathbf{k}, x_1 \mathbf{k}, (1 - x_2) \mathbf{y}_1, x_2 \mathbf{y}_1, \dots, (1 - x_L) \mathbf{y}_{L-1}, x_L \mathbf{y}_{L-1}, \mathbf{y}_L, \mathbf{e}$ . As a first step towards transforming our system into one captured by  $\mathcal{R}^*$ , we can embed the above system in a larger system whose solution is given by

$$(\mathbf{x}')^\top = ((\mathbf{x}'_1)^\top, (\mathbf{x}'_2)^\top, (\mathbf{x}'_3)^\top, \tilde{\mathbf{e}}^\top) \quad (2)$$

where

- $(\mathbf{x}'_1)^\top = ((1 - x_1), x_1, \dots, (1 - x_L), x_L) \in \{0, 1\}^{2L}$ ,
- $(\mathbf{x}'_2)^\top = (\mathbf{y}_0^\top, \mathbf{y}_1^\top, \dots, \mathbf{y}_L^\top) \in \mathbb{Z}_q^{(L+1) \cdot m}$ , with  $\mathbf{y}_0 := \mathbf{k}$ ,
- $\mathbf{x}'_3 \in \mathbb{Z}_q^{2L \cdot m}$  is of the form

$$(\mathbf{x}'_3)^\top = ((1 - x_1) \mathbf{y}_0, x_1 \mathbf{y}_0, (1 - x_2) \mathbf{y}_1, x_2 \mathbf{y}_1, \dots, (1 - x_L) \mathbf{y}_{L-1}, x_L \mathbf{y}_{L-1}),$$

- $\tilde{\mathbf{e}} = \text{vdec}_{m, \frac{q}{p}-1}(\mathbf{e}) \in \{0, 1\}^{m \cdot (\lceil \log(\frac{q}{p}-1) \rceil + 1)}$ , which ensures that  $\|\mathbf{e}\|_\infty < q/p$ .

One aspect of this extended solution is that every consecutive pair of entries of  $\mathbf{x}'_1$  is either  $(0, 1)$  or  $(1, 0)$ . In other words, each consecutive pair of entries of  $\mathbf{x}'_1$  sums to 1 and is binary. The fact that consecutive pairs add to 1 can be captured by a linear constraint that will constitute the first block of our matrix  $\mathbf{M}'$ . Next, the fact that the entries of  $\mathbf{x}'_1$  are binary may be captured by the set of equations  $\mathbf{x}'_1[i] = \mathbf{x}'_1[i] \cdot \mathbf{x}'_1[i]$ . In fact, proving this relation only for even  $i$  is sufficient as  $\mathbf{x}'_1[2i] \in \{0, 1\}$  and  $\mathbf{x}'_1[2i] + \mathbf{x}'_1[2i - 1] = 1$  implies  $\mathbf{x}'_1[2i - 1] \in \{0, 1\}$ .

The next part of a valid solution's structure is that entries of  $\mathbf{x}'_3$  are the result of multiplying entries of  $\mathbf{x}'_1$  and  $\mathbf{x}'_2$ . This can be written as  $\mathbf{x}'_3[h'] = \mathbf{x}'_1[i'] \cdot \mathbf{x}'_2[j']$  for appropriate choices of  $h', i', j'$ . It then only remains to prove that the entries of  $\tilde{\mathbf{e}}$  are binary, which is captured by the equations  $\tilde{\mathbf{e}}[i] = \tilde{\mathbf{e}}[i] \cdot \tilde{\mathbf{e}}[i]$ .

Following the details outlined above, we may represent a BLMR evaluation as the system  $\mathbf{M}' \cdot \mathbf{x}' = \mathbf{y}' \bmod q$  for

$$\begin{aligned}
& - \mathbf{x}' \in \mathbb{Z}_q^{2L+(L+1) \cdot m + 2L \cdot m + (\lfloor \log(q/p-1) \rfloor + 1) \cdot m} \text{ which is subject to the following constraints, when parsed as in Equation 2:} \\
& \quad \bullet \text{ for } i \in [L]: \mathbf{x}'_1[2i] = \mathbf{x}'_1[2i] \cdot \mathbf{x}'_1[2i] \\
& \quad \bullet \text{ for } (i, j) \in [m] \times [L]: \mathbf{x}'_3[2m(j-1) + i] = \mathbf{x}'_1[2j-1] \cdot \mathbf{x}'_2[m(j-1) + i] \text{ and } \\
& \quad \quad \mathbf{x}'_3[2m(j-1) + m + i] = \mathbf{x}'_1[2j] \cdot \mathbf{x}'_2[m(j-1) + i] \\
& \quad \bullet \text{ for } i \in [(\lfloor \log(q/p-1) \rfloor + 1) \cdot m]: \tilde{\mathbf{e}}[i] = \tilde{\mathbf{e}}[i] \cdot \tilde{\mathbf{e}}[i] \\
& - (\mathbf{y}')^\top = (\overbrace{1, \dots, 1}^L, \overbrace{0, \dots, 0}^{m \cdot L}, (q/p)\mathbf{y}^\top) \\
& - \\
& \mathbf{M}' = \left[ \begin{array}{c|c|c|c} \mathbf{I}_L \otimes (1, 1) & & & \\ \hline & 0^{mL \times m} \parallel \mathbf{I}_{m \cdot L} & -\mathbf{I}_L \otimes [\mathbf{P}_0 \parallel \mathbf{P}_1] & \\ \hline & 0^{m \times L \cdot m} \parallel \mathbf{I}_m & & \\ \hline & & & -H_{m, q/p-1} \end{array} \right] \quad (3)
\end{aligned}$$

where all blank blocks consist of 0 entries.

## 5.2 Zero-Knowledge Arguments for the Spend Protocol

The previous protocol enables to prove correct evaluation of the BLMR PRF but is not sufficient to produce the proof  $\pi_K$  expected by the merchant during the Spend protocol. In particular, we also need to prove

- knowledge of (short) solutions to linear systems (e.g., the user's secret key);
- knowledge of solutions to an equation involving a subset sum of known-matrix and secret vector multiplications (i.e. the computation of  $\mathbf{A}_\tau$ );
- correct evaluation of the lossy trapdoor function  $F_{\text{LTF}}$ .

All these statements can be captured by the relation  $\mathcal{R}^*$  from [33], as explained below. Together with our proof of correct PRF evaluation, this means that  $\pi_K$  can be instantiated using only the Yang *et al.* framework. We can then achieve inverse-polynomial soundness error  $1/\bar{p}$  in one ZKAoK protocol run. To achieve a soundness error of  $2^{-\lambda}$ , we only need  $\mathcal{O}(\lambda/\log \bar{p})$  repetitions. This clearly improves upon the Stern-type protocols used in [24], which require  $\mathcal{O}(\lambda)$  repetitions.

*Remark 1.* It should be noted that we have different equations over various moduli in our Spend protocol. However, as long as  $q$  is a prime power and all remaining moduli divide  $q$ , we may lift all equations to use the modulus  $q$ . For example, to lift an equation over  $\mathbb{Z}_{q'}$  to an equation over  $\mathbb{Z}_q$  where  $q'$  divides  $q$ , we simply multiply by  $q/q' \in \mathbb{Z}$ . We will use this trick in what follows.

*The explicit linear system.* Transforming the argument of knowledge produced by a user during the **Spend** protocol into an instance of the Yang *et al.* protocol is far from trivial as there are several details to address. Besides the moduli issue mentioned above, we indeed need to juggle with two different types of binary decomposition in order to ensure security of the whole system.

We use the notation from the **Spend** protocol specification in Section 4. We further parse  $\mathbf{v}$  as  $(\mathbf{v}_1, \mathbf{v}_2)$ , where  $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{Z}^{m_s}$ . Also, we define  $\sigma' := \lfloor \sigma \sqrt{m_s} + 1 \rfloor$  and  $\mathbf{v}_i^+ = \mathbf{v}_i + \sigma' \cdot \mathbf{1}$  for  $i \in \{1, 2\}$ , where  $\mathbf{1}$  denotes the all-one vector. This implies that valid values of  $\mathbf{v}_i$  (i.e., such that  $\|\mathbf{v}_i\|_\infty \leq \sigma'$ ) give rise to  $\mathbf{v}_i^+ \in [0, 2\sigma']^{m_s}$ . We also set  $\mathbf{r}^+ := \mathbf{r} + \sqrt{2}\sigma' \cdot \mathbf{1}$  so that  $\mathbf{r}^+ \in [0, 2\sqrt{2}\sigma']^{2m_s}$  for valid choices of  $\mathbf{r}$  (i.e. values such that  $\|\mathbf{r}\|_\infty \leq \sqrt{2}\sigma'$ ). We can then define  $\tilde{\mathbf{v}}_i := \text{vdec}_{m_s, 2\sigma'}(\mathbf{v}_i^+)$  for  $i \in \{1, 2\}$ ,  $\tilde{\mathbf{r}} := \text{vdec}_{2m_s, 2\sqrt{2}\sigma'}(\mathbf{r}^+)$ ,  $\tilde{\mathbf{k}} := \text{vdec}_{m, q-1}(\mathbf{k})$  and

$$\tilde{\mathbf{w}} := \text{vdec}_{n, q_s-1} \left( [\mathbf{D}'_0 | \mathbf{D}''_0] \cdot \mathbf{r} + \mathbf{D}_1 \cdot \mathbf{e}_u + \mathbf{D}_2 \cdot \tilde{\mathbf{k}} \right).$$

We begin by considering the equation associated to the signature. We can express it as the following linear system over  $\mathbb{Z}_q$

$$\begin{aligned} & \frac{q}{q_s} \left[ \mathbf{A} (\mathbf{H}_{m_s, 2\sigma'} \cdot \tilde{\mathbf{v}}_1 - \sigma' \mathbf{1}) + \right. \\ & \quad \mathbf{A}_0 (\mathbf{H}_{m_s, 2\sigma'} \cdot \tilde{\mathbf{v}}_2 - \sigma' \mathbf{1}) + \\ & \quad \left. \sum_{j=1}^{\ell} \mathbf{A}_j (\mathbf{H}_{m_s, 2\sigma'} \cdot (\tau[j] \cdot \tilde{\mathbf{v}}_2) - \sigma' \tau[j] \cdot \mathbf{1}) - \mathbf{D} \cdot \tilde{\mathbf{w}} \right] = \frac{q}{q_s} \mathbf{u} \\ & \frac{q}{q_s} \left[ \mathbf{H}_{n, q_s-1} \cdot \tilde{\mathbf{w}} - \left( [\mathbf{D}'_0 | \mathbf{D}''_0] \left( \mathbf{H}_{2m_s, 2\sqrt{2}\sigma'} \cdot \tilde{\mathbf{r}} - \sqrt{2}\sigma' \mathbf{1} \right) + \right. \right. \\ & \quad \left. \left. \mathbf{D}_1 \cdot \mathbf{e}_u + \mathbf{D}_2 \cdot \tilde{\mathbf{k}} \right) \right] = \mathbf{0} \\ & \mathbf{H}_{m, q-1} \cdot \tilde{\mathbf{k}} - \mathbf{k} = \mathbf{0}, \end{aligned}$$

whose solution is  $\mathbf{x}_1 := (\tau, \tilde{\mathbf{v}}_1, \tilde{\mathbf{v}}_2, \tau[1] \cdot \tilde{\mathbf{v}}_2, \dots, \tau[\ell] \cdot \tilde{\mathbf{v}}_2, \tilde{\mathbf{w}}, \tilde{\mathbf{r}}, \mathbf{e}_u, \mathbf{k}, \tilde{\mathbf{k}})$ , with some quadratic constraints amongst unknowns.

We next consider the evaluation of  $\mathbf{y}_1$ , as written in the **Spend** protocol. Here a subtlety arises as we need to use two different types of binary decomposition. So far, we have only used the  $\text{vdec}_{m, p-1}$  function because it allows achieving exact soundness with the proofs of Yang *et al.* Unfortunately, the decomposition of an integer according to the sequences  $B_1, \dots, B_{\delta_{p-1}}$  implicitly defined by  $\text{vdec}_{m, p-1}$  (see Section 2) may not be unique, which might lead to undetected frauds in our system. We will then also use the *standard* binary decomposition (that is unique) to ensure that the user is not evaluating  $F_{\text{LTF}}$  on two different decompositions of the same PRF output. It then remains to prove consistency of both decompositions, which is explained below.

Concretely, let  $\tilde{\mathbf{y}}_{\mathbf{k}}$  denote the standard binary decomposition of the PRF output  $\mathbf{y}_{\mathbf{k}} = \left[ \prod_{i=1}^L \mathbf{P}_{J[L+1-j]} \cdot \mathbf{k} \right]_p$ . Importantly, we must ensure that  $\tilde{\mathbf{y}}_{\mathbf{k}}$  does

really correspond to binary decomposition of a vector in  $[0, p-1]^m$  rather than some larger space. Alternatively, we need to ensure that  $\mathbf{y}_k$  (which is unknown) has entries in  $[0, p-1]$ . We achieve this by considering  $\tilde{\mathbf{y}}'_k = \text{vdec}_{m,p-1}(\mathbf{y}_k)$ . By multiplying the evaluation equation of  $\mathbf{y}_1$  by  $q/q_{\text{LTF}}$  and denoting the LTF key  $ek_{\text{LTF}}$  as  $\mathbf{B}_{\text{LTF}} \in \mathbb{Z}_{q_{\text{LTF}}}^{(n_{\text{LTF}} + \tilde{n}_{\text{LTF}}) \times m_{\text{LTF}}}$ , we can derive the following equations over  $\mathbb{Z}_q$ :

$$\boxed{\begin{aligned} \frac{q}{q_{\text{LTF}}} \mathbf{B}_{\text{LTF}} \cdot \tilde{\mathbf{y}}_k &= \frac{q}{q_{\text{LTF}}} \cdot \mathbf{y}_1 \\ \mathbf{y}_k - \mathbf{H}_{m,p-1} \cdot \tilde{\mathbf{y}}'_k &= \mathbf{0} \\ \mathbf{y}_k - \mathbf{I}_m \otimes (1, 2, \dots, 2^{\lceil \log p \rceil}) \cdot \tilde{\mathbf{y}}_k &= \mathbf{0} \end{aligned}}$$

Conveniently, the restriction that the entries of  $\tilde{\mathbf{y}}_k$  and  $\tilde{\mathbf{y}}'_k$  are binary is easily captured using quadratic constraints. Therefore all boxed equations so far constitute a linear system whose solution is  $\mathbf{x}_2 := (\mathbf{x}_1 \| \tilde{\mathbf{y}}_k, \tilde{\mathbf{y}}'_k, \mathbf{y}_k)$ , subject to some quadratic constraints that can easily be handled with the Yang *et al.* framework. However, we still need some equations to ensure that  $\mathbf{y}_k$  is computed correctly as a BLMR PRF output. In order to describe these equations, we will use the observations from Section 5.1 and the matrix  $\mathbf{M}'$  given in Equation (3). In particular, we set the unknown vector

$$\begin{aligned} \mathbf{x}_k = & (1 - J[1], J[1], \dots, 1 - J[L], J[L], \mathbf{y}k_0, \dots, \mathbf{y}k_L, \\ & (1 - J[1])\mathbf{y}k_0, J[1]\mathbf{y}k_0, \dots, (1 - J[L])\mathbf{y}k_{L-1}, J[L]\mathbf{y}k_{L-1}, \mathbf{e}_k) \end{aligned}$$

where  $\mathbf{y}k_i \in \mathbb{Z}_q^m$  for  $i \in [0, L]$  and  $\mathbf{e}_k \in \{0, 1\}^{m \cdot (\lfloor \log(\frac{q}{p}-1) \rfloor + 1)}$ . As noted in Section 5.1 (and shown by the form of  $\mathbf{x}_k$ ), the constraints on these unknown vectors are quadratic as required. To capture the PRF computation, we extend the vector of unknowns by defining  $\mathbf{x}_3 := (\mathbf{x}_2 \| \mathbf{x}_k)$ . We then add the following to the boxed linear equations over  $\mathbb{Z}_q$  above (where  $\mathbf{M}'$  is defined in Equation (3)):

$$\boxed{\begin{aligned} \mathbf{y}k_0 - \mathbf{k} &= \mathbf{0} \\ \mathbf{M}' \cdot \mathbf{x}_k - \left(0^{(m+1) \cdot L}, \frac{q}{p} \mathbf{y}_k^\top\right)^\top &= (1^L, 0^{m \cdot (L+1)})^\top \end{aligned}}$$

Finally, it remains to prove that  $\mathbf{y}_2$  and  $\mathbf{y}_T$  are well-formed. This consists in proving the following relation over  $\mathbb{Z}_q$ :

$$\boxed{\begin{aligned} \frac{q}{p} \mathbf{F} \cdot \mathbf{e}_u + \frac{q}{p} \mathbf{U}_{\text{UH}} \cdot \tilde{\mathbf{y}}_k &= \frac{q}{p} \mathbf{y}_2 \\ \frac{q}{p} \mathbf{F} \cdot \mathbf{e}_u + \frac{q}{p} \text{FRD}(R) \cdot \mathbf{U}'_{\text{UH}} \cdot \tilde{\mathbf{y}}_k &= \frac{q}{p} \mathbf{y}_T, \end{aligned}}$$

where the witnesses are already included in  $\mathbf{x}_3$ .

We have shown that the whole statement proved during the **Spend** protocol can be expressed as the collection of the boxed linear systems with a vector  $\mathbf{x}_3$  of unknowns subject to quadratic constraints supported by the protocol of [33].

## 6 Security Proofs

In this section and Section C of the Supplementary Material, we prove Theorem 1, which states that our construction provides all the required security properties.

**Theorem 1.** *Our construction is a secure e-cash system in the random oracle model assuming that the following conditions hold:*

- The  $SIS_{n,m_s,q_s,\beta'}$  for  $\beta' = \mathcal{O}\left(\sigma^2 m_s^{1/2}(m_s + m \log q)\right)$  and  $SIS_{n,m,p,2\sqrt{m}}$  problems are hard;
- Parameters are chosen so that the interactive AoK  $\Pi_1$  in the withdrawal protocol is zero-knowledge (ZK) and that the non-interactive AoK  $\Pi_2$  in the spend protocol is honest-verifier zero-knowledge (HVZK);
- Parameters  $m, n, q, p$  are chosen so that the BLMR PRF is pseudo-random;
- The  $LWE_{n_{\text{LTF}}, m_{\text{LTF}}, q_{\text{LTF}}, \alpha}$  problem is hard;
- $\Sigma$  is an EUF-CMA secure signature scheme.

**Proof of Exculpability.** Suppose the lossy trapdoor function is sampled in its injective mode. The proof of exculpability relies on the fact that an adversary producing two valid coins with the same serial number must produce at least one fresh proof of knowledge of a secret key underlying an honestly produced public key. In particular, our construction guarantees that this public key is the one that `Identify` points to. The ability to produce fresh arguments of knowledge for an honest public key can be used to solve the SIS problem. We first present a lemma about collision probabilities on PRFs with randomly sampled seeds and polynomial-size domain.

**Lemma 4.** *Let  $\text{PRF} = \{\text{PRF}_{\mathbf{k}} : \{0, 1\}^L \rightarrow \{0, 1\}^M \mid \mathbf{k} \in \mathcal{K}\}$  be a family of pseudo-random functions where  $2^L = \text{poly}(\lambda)$  and  $M = \text{poly}(\lambda)$ . Take any  $N = \text{poly}(\lambda)$  and sample  $\mathbf{k}_1, \dots, \mathbf{k}_N \leftarrow U(\mathcal{K})$ . The probability that  $\exists (i, j, x_1, x_2) \in [N]^2 \times \{0, 1\}^L \times \{0, 1\}^L$  such that  $\text{PRF}_{\mathbf{k}_i}(x_1) = \text{PRF}_{\mathbf{k}_j}(x_2)$  is negligible.*

*Proof.* We first describe a challenger algorithm  $\mathcal{C}$ . In the first stage,  $\mathcal{C}$  samples  $\mathbf{k}_1, \dots, \mathbf{k}_N \leftarrow U(\mathcal{K})$ , samples  $N$  uniform functions  $U_1, \dots, U_N : \{0, 1\}^L \rightarrow \{0, 1\}^M$  and samples a challenge bit  $b \leftarrow U(\{0, 1\})$ . In the second phase,  $\mathcal{C}$  waits for queries  $x \in \{0, 1\}^L$ . If  $b = 1$ , it answers with  $(\text{PRF}_{\mathbf{k}_1}(x), \dots, \text{PRF}_{\mathbf{k}_N}(x))$ . On the other hand, if  $b = 0$ , it responds with  $(U_1(x), \dots, U_N(x))$ . By a standard hybrid argument, no PPT adversary  $\mathcal{A}$  can guess the bit  $b$  with non-negligible advantage under the assumption that PRF is a PRF family and  $N = \text{poly}(\lambda)$ . Consider the following adversary  $\mathcal{A}^*$  that queries  $\mathcal{C}$  on the entire set  $\{0, 1\}^L$ . Denote the response to query  $x$  as  $(y_{1,x}, \dots, y_{N,x})$ . Now,  $\mathcal{A}^*$  outputs  $b^* = 1$  if there exists  $(i, j, x_1, x_2)$  such that  $y_{i,x_1} = y_{j,x_2}$ . Otherwise,  $\mathcal{A}^*$  outputs  $b^* = 0$ . Note that, if  $b = 0$ , the probability that  $\mathcal{A}^*$  outputs  $b^* = 1$  is equal to

$$1 - \prod_{k=1}^{2^L N} \left(1 - \frac{(k-1)}{2^M}\right)$$

which is negligible since  $2^L N = \text{poly}(\lambda)$  and  $2^M = 2^{\text{poly}(\lambda)}$ . Therefore, under the assumption that PRF is a PRF family, the probability that  $\mathcal{A}^*$  outputs  $b^* = 1$  when  $b = 1$  is also negligible.  $\square$

**Lemma 5.** *Our construction provides strong exculpability in the random oracle model assuming that: (i) The  $\text{SIS}_{n,m,p,2\sqrt{m}}$  problem is hard; (ii) Parameters  $(m,n,p,q)$  are chosen so that the BLMR PRF is pseudo-random; (iii)  $\Pi_1$  and  $\Pi_2$  are ZK and HVZK, respectively; (iv) The protocols underlying  $\Pi_1$  and  $\Pi_2$  are arguments of knowledge.*

Recall that a successful adversary returns  $\text{coin}_1$  and  $\text{coin}_2$  such that  $PK_{\mathcal{U}^*} = \text{Identify}(PK_{\mathcal{B}}, \text{coin}_1, \text{coin}_2)$  for honest user  $\mathcal{U}^*$ . This implies two things:

- First, the two coins have been generated using the public key  $PK_{\mathcal{U}^*}$ . Indeed, the fact that the identification procedure succeeds implies that these coins share the same serial number  $\mathbf{y}_S := (\mathbf{y}_1, \mathbf{y}_2)$ . Since the evaluation key of  $F_{\text{LTF}}$  was sampled in injective mode, the serial number  $\mathbf{y}_S$  uniquely determines the value  $PK' = \mathbf{y}_2 - H_{\text{UH}}(F_{\text{LTF}}^{-1}(\mathbf{y}_1))$ , which underlies both  $\text{coin}_1$  and  $\text{coin}_2$ . Then, the soundness of  $\Pi_2$  ensures that

$$\begin{aligned}\mathbf{y}_{T,1} &= PK' + \text{FRD}(R_1) \cdot H'_{\text{UH}}(F_{\text{LTF}}^{-1}(\mathbf{y}_1)), \\ \mathbf{y}_{T,2} &= PK' + \text{FRD}(R_2) \cdot H'_{\text{UH}}(F_{\text{LTF}}^{-1}(\mathbf{y}_1)),\end{aligned}$$

which implies that  $PK'$  is the public key  $PK_{\mathcal{U}^*}$  pointed to by  $\text{Identify}$ .

- Second, there exists  $d \in \{1, 2\}$  such that  $\text{coin}_d = (R_d, \mathbf{y}_{S,d}, \mathbf{y}_{T,d}, \pi_{K,d})$  is *not* the result of a  $\mathcal{Q}_{\text{Spend}}$ -query w.h.p. To see why, consider the case that  $\text{coin}_1$  and  $\text{coin}_2$  are both the result of  $\mathcal{Q}_{\text{Spend}}$ -queries, but do not appear in  $\mathcal{T}_{\text{ds}}$ . This occurs if, when sampling polynomially many seeds, one finds  $\mathbf{k}, \mathbf{k}'$  satisfying  $\text{PRF}_{\mathbf{k}}(J) = \text{PRF}_{\mathbf{k}'}(J')$  for some  $(J, J') \in [0, 2^L - 1]^2$ . By Lemma 4, this occurs with negligible probability  $\text{negl}_1(\lambda)$ .

*Proof.* Using these two observations, we will prove the strong exculpability of our scheme by defining the following sequence of games. Let  $\epsilon$  be the probability that  $\mathcal{A}$  succeeds against the exculpability of our scheme and let  $Q_w$  (resp.  $Q_s$ ) denote the maximal number of  $\mathcal{Q}_{\text{Withdraw}}$  queries (resp.  $\mathcal{Q}_{\text{Spend}}$  queries).

**Game<sub>0</sub>:** This is exactly the strong exculpability experiment, as defined in Section 2. The probability  $\epsilon_0$  that  $\mathcal{A}$  succeeds in this game is then exactly  $\epsilon$ .

**Game<sub>1,0</sub>:** In this game, our reduction  $\mathcal{S}$  (acting as a challenger in the strong exculpability experiment) proceeds as in **Game<sub>0</sub>** except that it defines  $\mathbf{F}$  as  $\bar{\mathbf{A}} \in \mathbb{Z}_p^{n \times m}$ , where  $\bar{\mathbf{A}}$  is a uniform matrix provided in a  $\text{SIS}_{n,m,p,2\sqrt{m}}$  instance. We denote by  $\mathbf{e}_{u^*} \in \{0, 1\}^m$  the secret key generated by  $\mathcal{S}$  for the accused user  $PK_{\mathcal{U}^*} = \mathbf{F} \cdot \mathbf{e}_{u^*}$ . Note that  $\mathcal{A}$  is given black-box access to  $H_{\text{FS}}$  and  $\mathcal{S}$  answers queries to  $H_{\text{FS}}$  by returning uniformly random elements of  $\{-\bar{p}, \dots, \bar{p}\}^\kappa$ . In addition,  $\mathcal{S}$  initialises empty lists of honest users  $\mathcal{HU}$  and double-spent coins  $\mathcal{T}_{\text{ds}}$ . As  $\bar{\mathbf{A}}$  is distributed as  $\mathbf{F}$  in the original setup, the probability that  $\mathcal{A}$  succeeds in this game is  $\epsilon_{1,0} = \epsilon_0$ .



**Game<sub>1,i</sub>:** For  $i \in [1, Q_w]$ , this game is defined as **Game<sub>1,i-1</sub>**, except that  $\mathcal{S}$  now answers the  $i$ -th  $\mathcal{Q}_{\text{Withdraw}}$ -query (if any) by running the simulator of  $\Pi_1$  to simulate the interactive proof generated by the user at this stage. This is done for every user  $PK_{\mathcal{U}}$ , and not just  $PK_{\mathcal{U}^*}$ . Any change of behaviour of  $\mathcal{A}$  can thus be straightforwardly used against the zero-knowledge property of  $\Pi_1$ . We therefore have  $\epsilon_{1,i-1} - \text{Adv}_{ZK}^{\Pi_1}(\mathcal{A}) \leq \epsilon_{1,i}$  for all  $i \in [1, Q_w]$ .

**Game<sub>1,Q\_w+i</sub>:** For  $i \in [1, Q_s]$ , this game is defined as **Game<sub>1,Q\_w+i-1</sub>**, except that  $\mathcal{S}$  now answers the  $i$ -th  $\mathcal{Q}_{\text{Spend}}$ -query (if any) by running the simulator of  $\Pi_2$  to simulate the non-interactive argument generated by the spender at this stage. This can be done (using only the user's public key  $PK_{\mathcal{U}}$ ) by applying the standard technique of programming the random oracle  $H_{\text{FS}}$  on new inputs, which only requires the statistical HVZK property of  $\Pi_2$ . The simulation fails whenever the random oracle  $H_{\text{FS}}$  needs to be programmed at an input that it was previously queried on. However, this happens with negligible probability at most  $\text{Coll}_H := (Q_s + Q_H)/2^\lambda$ , where  $Q_H$  is the total number of queries made by  $\mathcal{A}$  to  $H_{\text{FS}}$  and the denominator  $2^\lambda$  is a lower bound on the domain-size of  $H_{\text{FS}}$ -inputs. Therefore, we can conclude that  $\epsilon_{1,Q_w+i-1} - \text{Adv}_{HVZK}^{\Pi_2}(\mathcal{A}) - \text{Coll}_H \leq \epsilon_{1,Q_w+i}$  for all  $i \in [1, Q_s]$ .

It is important to note that, in **Game<sub>1,Q\_w+Q\_s</sub>**, the reduction  $\mathcal{S}$  only needs  $PK_{\mathcal{U}^*}$  and not  $\mathbf{e}_{u^*}$  to simulate the game. This concretely means that the adversary's view is independent of the preimage  $\mathbf{e}_{u^*}$  of  $PK_{\mathcal{U}^*}$  selected by  $\mathcal{S}$ . Thanks to [25, Lemma 8], we know that this preimage is not unique: i.e., there exists at least one vector  $\mathbf{e} \in \{0, 1\}^m \setminus \{\mathbf{e}_{u^*}\}$  such that  $\bar{\mathbf{A}} \cdot \mathbf{e}_{u^*} = \bar{\mathbf{A}} \cdot \mathbf{e} \pmod p$  with all but negligible probability. This observation will be crucial in what follows.

**Game<sub>2</sub>:** Let  $Q_H$  be a polynomial bounding the number of random oracle queries made by  $\mathcal{A}$  to  $H_{\text{FS}}$ . Up until  $\mathcal{A}$  terminates,  $\mathcal{S}$  answers  $\mathcal{A}$ 's queries as in the previous games, recording the random oracle queries as  $(q_1, q_2, \dots)$  and the corresponding uniformly distributed responses as  $(h_1, h_2, \dots)$ . Our second observation at the beginning of the proof implies that at least one coin  $\text{coin}_d$  returned by  $\mathcal{A}$  is not the result of a  $\mathcal{Q}_{\text{Spend}}$ -query with overwhelming probability (if none of the coins were generated as a response to  $\mathcal{Q}_{\text{Spend}}$ -query, then select a random  $d \in \{1, 2\}$ ). Define

$$\begin{aligned} \pi_{K,d} &:= \left( \{\text{Comm}_{K,d,j}\}_{j=1}^\kappa, \text{Chall}_{K,d}, \{\text{Resp}_{K,d,j}\}_{j=1}^\kappa \right), \\ \text{Chall}_{K,d} &:= H_{\text{FS}} \left( \text{par}, R, \mathbf{y}_{S,d}, \mathbf{y}_{T,d}, \{\text{Comm}_{K,d,j}\}_{j=1}^\kappa \right). \end{aligned}$$

In this game,  $\mathcal{S}$  aborts if the above query was not made to  $H_{\text{FS}}$ . We note that in such a case the proof  $\pi_{K,d}$  would only have been acceptable with probability at most  $(2\bar{p} + 1)^{-\kappa}$ . We then have  $\epsilon_{1,Q_w+Q_s} - (2\bar{p} + 1)^{-\kappa} \leq \epsilon_2$ .

From now on, we know that there exists an index  $i^* \in [Q_H]$  such that the  $i^*$ -th  $H_{\text{FS}}$ -query is used to produce  $\text{Chall}_{K,d}$  (i.e.,  $\text{Chall}_{K,d} = h_{i^*}$ ) and that  $\mathcal{A}$  succeeds in **Game<sub>2</sub>** with probability  $\epsilon_2 \geq \epsilon - Q_w \cdot \text{Adv}_{ZK}^{\Pi_1}(\mathcal{A}) - Q_s (\text{Adv}_{HVZK}^{\Pi_2}(\mathcal{A}) + \text{Coll}_H) - (2\bar{p} + 1)^{-\kappa}$ . We then define our last game **Game<sub>3</sub>** as follows:

1. **Run Game<sub>2</sub> once:**  $\mathcal{S}$  runs  $\mathcal{A}$  by behaving as in Game<sub>2</sub>. If  $\mathcal{A}$  fails to win the game, then  $\mathcal{S}$  aborts. Otherwise, it records  $\text{coin}_d, \pi_{K,d}, \text{Chall}_{K,d}, (q_1, q_2, \dots), (h_1, h_2, \dots), i^*$ , sets a variable  $\text{fork} = 1$  and proceeds to the next step.
2. **(Search for a 3-fork).** This step is repeated twice.  $\mathcal{S}$  runs  $\mathcal{A}$  with the same random tape as in the beginning of the first step. In addition, it sends  $\mathcal{A}$  the same  $\text{par}$  as before, giving  $\mathcal{A}$  oracle access to  $H_{\text{FS}}$ .  $\mathcal{S}$  allows  $\mathcal{A}$  to run until termination, answering queries to  $H_{\text{FS}}$  as follows:
  - Answer queries  $q_1, \dots, q_{i^*-1}$  (which are identical to those of the first run) using the same values  $h_1, \dots, h_{i^*-1}$  as before.
  - At the  $i^*$ -th query  $q_{i^*}$  (which is also the same as the first time  $\mathcal{A}$  was run), pick a fresh uniform response  $h'_{i^*}$ .
  - For the remaining queries made by  $\mathcal{A}$  denoted  $q'_{i^*+1}, \dots, q'_{Q_H}$ , pick fresh uniform random responses  $h'_{i^*+1}, \dots, h'_{Q_H}$ .

If this is the first repetition,  $\mathcal{S}$  sets  $h_{i^*}^{(2)} = h'_{i^*}$ . At the second repetition, it sets  $h_{i^*}^{(3)} = h'_{i^*}$ . If  $\mathcal{A}$  terminates without winning the strong exculpability game, then  $\mathcal{S}$  begins the next repetition of this step. If  $\mathcal{A}$  terminates and wins the game, denote its output as  $(PK'_{\mathcal{B}'}, \text{coin}'_1, \text{coin}'_2)$ . As before, let  $d' \in \{1, 2\}$  denote the index that was not the result of a  $\mathcal{Q}_{\text{Spend}}$ -query (picking  $d' \in \{1, 2\}$  randomly if neither coin was the result of a spend query). Recall that both coins can be the result of  $\mathcal{Q}_{\text{Spend}}$ -queries with at most negligible probability  $\text{negl}(\lambda)_1$ , but if this is the case,  $\mathcal{S}$  skips to the next repetition of this step. Denote  $\text{coin}'_{d'} = (R'_{d'}, \mathbf{y}'_{S,d'}, \mathbf{y}'_{T,d'}, \pi'_{K,d'})$ . Write

$$\pi'_{K,d'} = \left( \{ \text{Comm}'_{K,d',j} \}_{j=1}^{\kappa}, \text{Chall}'_{K,d'}, \{ \text{Resp}'_{K,d',j} \}_{j=1}^{\kappa} \right).$$

$\mathcal{S}$  skips to the next repetition of this step at this point if

$$\left( R_d, \mathbf{y}_{S,d}, \mathbf{y}_{T,d}, \{ \text{Comm}_{K,d,j} \}_{j=1}^{\kappa} \right) \neq \left( R'_{d'}, \mathbf{y}'_{S,d'}, \mathbf{y}'_{T,d'}, \{ \text{Comm}'_{K,d',j} \}_{j=1}^{\kappa} \right)$$

or if  $h_{i^*} = h'_{i^*}$ . Otherwise,  $\mathcal{S}$  sets  $\text{fork} \leftarrow \text{fork} + 1$  and  $\pi_K^{(\text{fork}+1)} = \pi'_{K,d'}$ .

3. **(Derive SIS solution from 3-fork)** If  $\text{fork} < 3$  or,  $\text{fork} = 3$  but there exists no  $j \in [\kappa]$  such that  $(h_{i^*}[j], h_{i^*}^{(2)}[j], h_{i^*}^{(3)}[j])$  take three distinct values, then  $\mathcal{S}$  terminates outputting  $\perp$ . Otherwise,  $\mathcal{S}$  has access to arguments  $\pi_{K,d}, \pi_K^{(2)}, \pi_K^{(3)}$  sharing the same first message which we denote as  $\{ \text{Comm}_j \}_{j=1}^{\kappa}$ . In addition,  $\exists j^* \in [\kappa]$  at where  $h_{i^*}[j^*], h_{i^*}^{(2)}[j^*], h_{i^*}^{(3)}[j^*]$  take three distinct values in  $\{-\bar{p}, \dots, \bar{p}\}$ . Now a witness can be extracted from the transcripts  $\pi_{K,d}, \pi_K^{(2)}, \pi_K^{(3)}$  by considering the  $j^*$ -th parallel repetition and the special-soundness/extractor of the ZKAoK protocol [33]. We denote this witness as  $(\bar{J}, \bar{\mathbf{k}}, \bar{\mathbf{e}}_{u^*})$ . If  $\bar{\mathbf{e}}_{u^*} = \mathbf{e}_{u^*}$ , then  $\mathcal{S}$  aborts. Otherwise,  $\mathcal{S}$  terminates, outputting  $\mathbf{v} := \bar{\mathbf{e}}_{u^*} - \mathbf{e}_{u^*} \in \{-1, 0, 1\}^m$  as a SIS solution.

It then remains to evaluate the probability  $\epsilon_3$  that  $\mathcal{A}$  succeeds in this last game. We begin by noting that the first and second steps corresponds exactly to the forking algorithm denoted as  $F^{\mathcal{A}}$  in Lemma 3. Therefore, a direct application of

this forking lemma implies that the variable `fork` reaches the value `fork = 3` at the beginning of Step 3 with probability at least

$$\text{frk} := \epsilon_2 \cdot \left( \left( \frac{\epsilon_2}{Q_H} \right)^2 - \frac{3}{(2\bar{p} + 1)^\kappa} \right).$$

which is non-negligible if  $\epsilon_2$  is non-negligible as  $1/(2\bar{p} + 1)^\kappa$  is negligible and  $Q_H$  is polynomial. Next, note that  $\mathcal{S}$  extracts a witness  $(\bar{J}, \bar{\mathbf{k}}, \bar{\mathbf{e}}_{u^*})$  if and only if it does not terminate at, or before the beginning of Step 3. In order to analyse the probability that this occurs, we define three events:

- GF (“Good fork”): This is the event that `fork = 3` and there exists an index  $j^* \in [\kappa]$  such that  $(h_{i^*}[j^*], h_{i^*}^{(2)}[j^*], h_{i^*}^{(3)}[j^*])$  is a triple of 3.
- F (“Any fork”): This is the event that `fork = 3` at the beginning of Step 4.
- GH (“Good hashes”): This is the event that there is an index  $j^* \in [\kappa]$  such that  $(h_{i^*}[j^*], h_{i^*}^{(2)}[j^*], h_{i^*}^{(3)}[j^*])$  take 3 distinct values.

It is easy to see that  $\Pr[\overline{\text{GH}}] = ((6\bar{p} + 1)/(2\bar{p} + 1)^2)^\kappa$  is negligible and that  $\Pr[\text{F}] = \text{frk}$ . We also have

$$\Pr[\text{F}] \leq \Pr[\text{F}|\text{GH}] \cdot 1 + 1 \cdot \Pr[\overline{\text{GH}}] = \Pr[\text{F}|\text{GH}] + \text{negl}(\lambda).$$

This implies that  $\mathcal{S}$  does not abort at the beginning of Step 3 or before with non-negligible probability

$$\Pr[\text{GF}] = \Pr[\text{F} \cap \text{GH}] = \Pr[\text{F}|\text{GH}] \cdot \Pr[\text{GH}] \geq (\text{frk} - \text{negl}(\lambda)) \cdot (1 - \text{negl}(\lambda)).$$

The last step is to evaluate the probability that  $\bar{\mathbf{e}}_{u^*} = \mathbf{e}_{u^*}$ , leading  $\mathcal{S}$  to abort. Here we rely on our previous observation, namely that the adversary’s view has been independent of  $\mathbf{e}_{u^*}$  since  $\text{Game}_{1, Q_w + Q_s}$  and that there is, with overwhelming probability, at least another vector  $\bar{\mathbf{e}}_{u^*} \neq \mathbf{e}_{u^*}$  that is a valid secret key for  $PK_{\mathcal{U}^*}$ . We therefore know that the probability of the event  $\bar{\mathbf{e}}_{u^*} \neq \mathbf{e}_{u^*}$  is at least  $\frac{1}{2}$ . In summary, we get the following bound on the probability  $\epsilon_3$  that  $\mathcal{A}$  succeeds in  $\text{Game}_3$ :

$$\epsilon_3 \geq \frac{1}{2} \cdot (\text{frk} - \text{negl}(\lambda)) \cdot (1 - \text{negl}(\lambda))$$

where `frk` is defined above. Any adversary  $\mathcal{A}$  succeeding with non-negligible probability  $\epsilon$  against the exculpability of our scheme can thus be used to solve the SIS problem, distinguish the BLMR PRF from pseudo-random, or break the zero-knowledge property of  $\Pi_1$  or  $\Pi_2$ , which completes the proof.  $\square$

## 7 A More Efficient GGM-based Construction

In Section 4, we use the BLMR PRF because it allows for a simpler description of the argument of knowledge, as it only requires one rounding per evaluation. Unfortunately, this comes at the price of a super-polynomial modulus  $q$ . We

can do better by using a PRF obtained by applying the seminal construction of Goldreich, Goldwasser and Micali [21] to the LWR-based PRG of Banerjee *et al.* [5] for which the LWE-to-LWR reduction of [2] allows the use of a polynomial modulus. This leads to an e-cash construction with  $q = \text{poly}(\lambda)$  which still relies on the hardness of standard worst-case lattice problems. Explicitly, the PRF we have in mind relies on the hardness of the  $\text{LWR}_{m,m,q,p}$  problem (which is at least as hard as  $\text{LWE}_{m',m,q,\alpha'}$  for  $m' \geq \frac{\log q}{\log(2\gamma')}m, q \geq \gamma'm^2\alpha'p$  for any  $\gamma' \geq 1$  [2]). This PRF uses public parameters  $m, p, q, \mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{m \times m}$  where  $\mathbf{A}_0, \mathbf{A}_1 \leftarrow U(\mathbb{Z}_q^{m \times m})$ . The evaluation on seed  $\mathbf{k} \in \mathbb{Z}_q^m$  and input  $x \in \{0, 1\}^L$  is

$$F_{\mathbf{k}}(x) := \left[ \mathbf{A}_{x_L} \cdot \left[ \dots \left[ \mathbf{A}_{x_2} \cdot \left[ \mathbf{A}_{x_1} \cdot \mathbf{k} \right]_p \right]_p \dots \right]_p \right]_p. \quad (4)$$

When replacing the BLMR PRF with the above in our e-cash construction, it is more convenient to keep the parameters  $m$  and  $n$  as described in Section 4. This allows us to reuse our security proofs without any issues. However, in contrast with the BLMR instantiation, we choose polynomially large  $p$  and  $q$  such that  $q^2 > m^{5/2}p$  in the  $\text{ParGen}()$  phase. In addition, the binary public matrices  $\mathbf{P}_0, \mathbf{P}_1$  must be replaced by uniformly sampled  $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{m \times m}$ . In Section B of the Supplementary Material, we show that this alternative PRF is compatible with the ZK relation  $\mathcal{R}^*$  considered in [33], as we did for the BLMR PRF in Section 5.1. Combining this with the reasoning in Section 5.2 allows us to show that the GGM-based PRF is compatible with the ZKAoKs used in *Spend*.

## 7.1 Parameters

We provide in this section some details on the parameters and the complexity of an instantiation of our e-cash system using the GGM-based PRF. Firstly, Theorem 1 states that the security of our construction relies on:

- $\text{LWR}_{m,m,q,p}$  (which is at least as hard as  $\text{LWE}_{m',m,q,\alpha'}$  for  $m' \geq \frac{\log q}{\log(2\gamma')}m, q \geq \gamma'm^2\alpha'p$  for any  $\gamma' \geq 1$  [2])
- $\text{LWE}_{n_{\text{LTF}}, m_{\text{LTF}}, q_{\text{LTF}}, \alpha}$  with  $\alpha = \Theta\left(\frac{\sqrt{n_{\text{LTF}}}}{q_{\text{LTF}}}\right), q_{\text{LTF}} = \Theta(n_{\text{LTF}}^{1+1/\gamma})$  for constant  $\gamma < 1$
- $\text{SIS}_{n,m,p,2\sqrt{m}}$
- $\text{SIS}_{n,m_s,q_s,\beta'}$  for  $\beta' = \mathcal{O}(\sigma^2 m_s^{1/2}(m_s + \bar{m}))$

and also that we use secure ZKAoKs. Since all moduli will be polynomial, we may safely assume that there is a parameter setting such that the argument system of Yang *et al.* is a ZKAoK. Additionally, our proof of the clearing property requires use of a signature scheme. Note that we can use the signature scheme of [19] so that the arising assumption is made redundant by the final item listed above. Recall that for our zero-knowledge proofs, we require that  $q_s, q_{\text{LTF}}$  and  $p$  all divide the prime power  $q$ . In order to achieve this, we now set  $q = q_0^e$  where  $q_0$  is prime and  $e > 1$  is a constant integer. Since all moduli are polynomial, we may take  $n_{\text{LTF}} = \Theta(m) = \Theta(n \log q) = \tilde{\mathcal{O}}(n)$ . Additionally,  $m, \bar{m}, m_s, m_{\text{LTF}}, \bar{n}_{\text{LTF}}$  and  $n'$

are all  $\tilde{O}(n)$ . Note that we will take  $\gamma' = 1$  in the LWE-to-LWR reduction result stated above and  $\gamma = 1/2$ . To comply with hardness results relating standard worst-case lattice problems to SIS [19,27] and LWE [30,10], we require:

$$q^2/p = \tilde{\Omega}(n^{5/2}) \quad q_{\text{LTF}} = \tilde{\Theta}(n^3) \quad p = \tilde{\Omega}(n) \quad q_s = \tilde{\Omega}(\sigma^2 n^2) = \tilde{\Omega}(n^3).$$

Therefore, to base security on worst-case lattice problems, we may take  $n, m, n_{\text{LTF}}, \bar{n}_{\text{LTF}}, m_{\text{LTF}}, m_s$  all  $\tilde{O}(\lambda)$ ,  $p = q_0 = \tilde{O}(\lambda)$  and  $q = q_s = q_{\text{LTF}} = q_0^3 = \tilde{O}(\lambda^3)$ .

**Communication cost of [33].** We assess the communication cost of the protocol of [33] assuming that the linear system in  $\mathcal{R}^*$  (cf. Equation (1)) consists of  $m_{zk}$  equations in  $n_{zk}$  dimensions over  $\mathbb{Z}_{q_{zk}}$  with  $\ell_{zk}$  quadratic relations. The protocol takes the form of a  $\Sigma$ -protocol, where the first message is a cryptographic commitment, the second is a random challenge sampled from  $\{-\bar{p}, \dots, \bar{p}\}$  and the third message is an opening of the original commitment plus some further information depending on the challenge. Since we work in the random oracle model, we may instantiate the commitment using a hash function outputting  $\mathcal{O}(\lambda)$  bits. The second message is simply an integer consisting of  $\log 2\bar{p} = \mathcal{O}(\log q_{zk}) = \tilde{O}(1)$  bits assuming  $q_{zk}$  is polynomial in the security parameter. The final message consists of vectors over  $\mathbb{Z}_{q_{zk}}$  with various dimensions, all of which are  $\mathcal{O}(\max(\lambda, m_{zk}, n_{zk}, \ell_{zk}))$ . Since the protocol of [33] requires  $\lambda/\log \bar{p} = \tilde{O}(\lambda)$  repetitions to achieve soundness error  $2^{-\lambda}$ , the communication complexity of the overall ZKAoK protocol is  $\tilde{O}(\lambda \cdot \max(\lambda, m_{zk}, n_{zk}, \ell_{zk}))$ .

**Communication Cost of the Withdraw Protocol.** We now consider the cost of the Withdraw protocol, which consists of two distinct parts. In the first part, the user sends a commitment  $\mathbf{c}_U \in \mathbb{Z}_{q_s}^n$  and a ZKAoK. The ZKAoK can take the form of a linear system of  $m_{zk} = \mathcal{O}(n + m_s) = \tilde{O}(\lambda)$  equations in dimension  $n_{zk} = \mathcal{O}(m_s \log \sigma + m + 2n \log p) = \tilde{O}(\lambda)$  over  $\mathbb{Z}_{q_s}$  with  $\ell_{zk} = n_{zk} = \tilde{O}(\lambda)$  quadratic constraints. This first component of Withdraw thus incurs a communication cost of  $\tilde{O}(\lambda^2)$ . The second component is the bank's response which consists of  $\mathcal{O}(m_s \log \sigma) = \tilde{O}(\lambda)$  bits. Therefore, the overall communication cost in the Withdraw protocol is  $\tilde{O}(\lambda^2)$  bits.

**Communication Cost of the Spend Protocol.** The Spend protocol consists of a single message from the spender to the merchant, the main components of which are a merchant public key, a serial number, a security tag, and a ZKAoK. Assuming the use of GPV signatures [19], we can assume that the merchant public key consists of  $\tilde{O}(\lambda)$  bits. Next, the serial number and security tag together consists of  $2(n_{\text{LTF}} + \bar{n}_{\text{LTF}}) \log q_{\text{LTF}} + 2n \log p = \tilde{O}(\lambda)$  bits. Finally, if we look at the full ZKAoK described in Section 5, we have a system of  $m_{zk} = \mathcal{O}(m \cdot L) = \tilde{O}(\lambda)$  equations in dimension  $n_{zk} = \mathcal{O}(nL + n \log q) = \tilde{O}(\lambda)$  with  $\ell_{zk} = \mathcal{O}(n_{zk} + nL) = \tilde{O}(\lambda)$  quadratic constraints. This implies that the Spend protocol has a total communication cost of  $\tilde{O}(\lambda^2)$  bits.

## Acknowledgements

This work is supported by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701). Khoa Nguyen is also supported by the Gopalakrishnan – NTU Presidential Postdoctoral Fellowship 2018.

## References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Eurocrypt*, 2010.
2. J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs. Learning with rounding revisited – new reduction, properties and applications. In *Crypto*, 2013.
3. W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296, 1993.
4. A. Banerjee and C. Peikert. New and improved key-homomorphic pseudo-random functions. In *Crypto*, 2014.
5. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *Eurocrypt*, 2012.
6. C. Baum, I. Damgård, V. Lyubashevsky, S. Oechsner, and C. Peikert. More efficient commitments from structured lattice assumptions. In *SCN*, 2018.
7. M. Belenkiy, M. Chase, M. Kohlweiss, and A. Lysyanskaya. Compact e-cash and simulatable VRFs revisited. In *Pairing*, 2009.
8. D. Boneh, K. Lewi, H. Montgomery, and A. Raghunathan. Key-homomorphic prfs and their applications. In *Crypto*, 2013.
9. F. Bourse, D. Pointcheval, and O. Sanders. Divisible e-cash from constrained pseudo-random functions. In *Asiacrypt*, 2019.
10. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. On the classical hardness of learning with errors. In *STOC*, 2013.
11. J. Camenisch, S. Hohenberger, and A. Lysyanskaya. Compact e-cash. In *Eurocrypt*, 2005.
12. S. Canard, D. Pointcheval, O. Sanders, and J. Traoré. Divisible e-cash made practical. In *PKC*, 2015.
13. D. Chaum. Blind signatures for untraceable payments. In *Crypto*, 1982.
14. I. Damgård. Efficient concurrent zero-knowledge in the auxiliary string model. In *Eurocrypt*, 2000.
15. Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *SIAM*, 2008.
16. A. El Kaafarani and S. Katsumata. Attribute-based signatures for unbounded circuits in the rom and efficient instantiations from lattices. In *PKC*, 2018.
17. P. Farshim, C. Orlandi, and R. Rosie. Security of symmetric primitives under incorrect usage of keys. In *ToSC*, volume 2017.
18. A. Fiat and A. Shamir. How to prove yourself – practical solutions to identification and signature problems. In *Crypto*, 1986.
19. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
20. O. Goldreich. *Foundations of Cryptography: Volume 1*. Cambridge University Press, 2006.

21. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. In *FOCS*, 1984.
22. S. Kim and D. Wu. Watermarking cryptographic functionalities from standard lattice assumptions. In *Crypto*, volume 2017.
23. 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*, volume 2016.
24. B. Libert, S. Ling, K. Nguyen, and H. Wang. Zero-knowledge arguments for lattice-based prfs and applications to e-cash. In *Asiacrypt*, 2017.
25. V. Lyubashevsky. Lattice-Based Identification Schemes Secure Under Active Attacks. In *PKC*, 2008.
26. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Eurocrypt*, 2012.
27. D. Micciancio and C. Peikert. Hardness of SIS and LWE with small parameters. In *Crypto*, 2013.
28. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, 2006.
29. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, 2008.
30. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
31. D. Ron and A. Shamir. Quantitative analysis of the full bitcoin transaction graph. In *Financial Cryptography*, 2013.
32. H. Wee. Dual projective hashing and its applications — lossy trapdoor functions and more. In *Eurocrypt*, 2012.
33. R. Yang, M.H. Au, Z. Zhang, Q. Xu, Z. Yu, and W. Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Constructions and applications. In *Crypto*, 2019.

## Supplementary Material

### A Lattice Sampling Algorithms

Our construction uses of the efficient algorithms `SampleLeft` and `SampleRight` of Agrawal *et al.* [1] that satisfy the properties given in the following lemmas. Note that  $\tilde{\mathbf{A}}$  represents the matrix resulting from applying the Gram-Schmidt process to the columns of  $\mathbf{A}$ . Further,  $\|\mathbf{A}\|$  represents the length of the column of  $\mathbf{A}$  with the largest Euclidean norm. We also recall a trapdoor generation result from [26].

**Lemma 6 ([1]).** *Let  $q > 2, n, m > n, m_1$  be positive integers. Also take any  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}; \mathbf{M} \in \mathbb{Z}_q^{n \times m_1}; \mathbf{u} \in \mathbb{Z}_q^n$ ; and full-rank  $\mathbf{T}_\mathbf{A} \in \mathbb{Z}_q^{m \times m}$  such that  $\mathbf{A} \cdot \mathbf{T}_\mathbf{A} = 0 \pmod q$ . For any  $\sigma > \|\tilde{\mathbf{T}}_\mathbf{A}\| \cdot \omega(\sqrt{\log(m + m_1)})$ , `SampleLeft`( $\mathbf{A}, \mathbf{M}, \mathbf{T}_\mathbf{A}, \mathbf{u}, \sigma$ ) outputs a vector distributed statistically close to  $D_{\Lambda_q^\mathbf{u}(\mathbf{A} \parallel \mathbf{M}), \sigma}$ .*

**Lemma 7 ([1]).** *Let  $n, m > n, m_1$  be positive integers and  $q > 2$  be prime. Also take  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}; \mathbf{R} \in \mathbb{Z}_q^{m \times m_1}; \mathbf{B} \in \mathbb{Z}_q^{n \times m_1}; \mathbf{u} \in \mathbb{Z}_q^n$ ; non-zero  $h \in \mathbb{Z}_q$ ; and full rank  $\mathbf{T}_\mathbf{B} \in \mathbb{Z}_q^{m \times m}$  such that  $\mathbf{B} \cdot \mathbf{T}_\mathbf{B} = 0 \pmod q$ . For any  $\sigma > \|\tilde{\mathbf{T}}_\mathbf{B}\| \cdot (s_R + 1)\omega(\sqrt{\log(m)})$  where  $s_R := \max_{\|\mathbf{x}\|=1} \|\mathbf{R} \cdot \mathbf{x}\|_1$ , `SampleRight`( $\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathbf{T}_\mathbf{B}, h, \sigma, \mathbf{u}$ ) outputs a vector distributed statistically close to  $D_{\Lambda_q^\mathbf{u}(\mathbf{A} \parallel \mathbf{A}\mathbf{R} + h\mathbf{B}), \sigma}$ .*

**Lemma 8 ([26]).** *Let  $n = \mathcal{O}(\lambda)$ ,  $q$  be a prime and let  $m \geq 2n \lceil \log q \rceil$ . Then there is an efficient algorithm `TrapGen` such that `TrapGen`( $1^n, 1^m, q$ ) outputs  $(\mathbf{A}, \mathbf{T}_\mathbf{A}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{m \times m}$  where  $\mathbf{A}$  is statistically indistinguishable from uniform, and  $\mathbf{T}_\mathbf{A}$  is a basis of  $\Lambda_q^\perp(\mathbf{A})$  satisfying  $\|\tilde{\mathbf{T}}_\mathbf{A}\| = (2\sqrt{m} + 1) \cdot \sqrt{5}$ .*

### B Zero-Knowledge Arguments Showing the Correct Evaluation of a GGM-based PRF

We now consider reducing the problem of proving knowledge of an input/seed pair of a GGM-based PRF to proving knowledge of a solution to a linear system with quadratic relations over the entries of the solution. In other words, we will show that proving knowledge of an input/seed pair for the PRF in [5] is compatible with  $\mathcal{R}^*$  and the protocol in [33].

Our starting point is to consider a set of equations over  $\mathbb{Q}$  representing a GGM-based PRF evaluation (Equation (4)) leading to output  $\mathbf{y} \in \mathbb{Z}_p^m$ . Using the notation  $(\cdot)_q$  to denote that all operations within the brackets are performed



over  $\mathbb{Z}_q$ , the system of equations (over  $\mathbb{Q}$ ) is given by

$$\begin{aligned}
\mathbf{y}_1 &= \frac{p}{q} \cdot \left( (1 - x_1) \mathbf{A}_0 \cdot \mathbf{k} + x_1 \mathbf{A}_1 \cdot \mathbf{k} \right)_q - \mathbf{e}_1 \\
\mathbf{y}_2 &= \frac{p}{q} \cdot \left( (1 - x_2) \mathbf{A}_0 \cdot \mathbf{y}_1 + x_2 \mathbf{A}_1 \cdot \mathbf{y}_1 \right)_q - \mathbf{e}_2 \\
&\vdots \\
\mathbf{y}_{L-1} &= \frac{p}{q} \cdot \left( (1 - x_{L-1}) \mathbf{A}_0 \cdot \mathbf{y}_{L-2} + x_{L-1} \mathbf{A}_1 \cdot \mathbf{y}_{L-2} \right)_q - \mathbf{e}_{L-1} \\
\mathbf{y} &= \frac{p}{q} \cdot \left( (1 - x_L) \mathbf{A}_0 \cdot \mathbf{y}_{L-1} + x_L \mathbf{A}_1 \cdot \mathbf{y}_{L-1} \right)_q - \mathbf{e}_L
\end{aligned}$$

For  $i \in [L-1]$ ,  $\mathbf{y}_i \in \{0, 1, \dots, p-1\}^m$  and the “error” terms  $\mathbf{e}_i$  belong to the set  $[0, 1]^m \cap \frac{p}{q} \cdot \mathbb{Z}^m$  where  $[0, 1] := \{x : x \in \mathbb{R}, 0 \leq x < 1\}$ . For  $i \in [L]$ , we now define  $\bar{\mathbf{e}}_i := \frac{q}{p} \mathbf{e}_i \in [0, \frac{q}{p} - 1]^m \cap \mathbb{Z}^m$ ,  $x_{i,1} := (1 - x_i)$  and  $x_{i,2} := x_i$ . Multiplying the above system of equations by  $q/p$  and relabelling  $\mathbf{y}_0 := \mathbf{k}$  yields the following system over  $\mathbb{Z}_q$ :

$$\begin{aligned}
\frac{q}{p} \cdot \mathbf{y}_1 &= \mathbf{A}_0 \cdot x_{1,1} \mathbf{y}_0 + \mathbf{A}_1 \cdot x_{1,2} \mathbf{y}_0 - \bar{\mathbf{e}}_1 \bmod q \\
\frac{q}{p} \cdot \mathbf{y}_2 &= \mathbf{A}_0 \cdot x_{2,1} \mathbf{y}_1 + \mathbf{A}_1 \cdot x_{2,2} \mathbf{y}_1 - \bar{\mathbf{e}}_2 \bmod q \\
&\vdots \\
\frac{q}{p} \cdot \mathbf{y}_{L-1} &= \mathbf{A}_0 \cdot x_{L-1,1} \mathbf{y}_{L-2} + \mathbf{A}_1 \cdot x_{L-1,2} \mathbf{y}_{L-2} - \bar{\mathbf{e}}_{L-1} \bmod q \\
\frac{q}{p} \cdot \mathbf{y} &= \mathbf{A}_0 \cdot x_{L,1} \mathbf{y}_{L-1} + \mathbf{A}_1 \cdot x_{L,2} \mathbf{y}_{L-1} - \bar{\mathbf{e}}_L \bmod q
\end{aligned} \tag{5}$$

This system over  $\mathbb{Z}_q$  *almost* captures a GGM PRF evaluation. It remains to ensure that for  $i \in [L-1]$ ,  $\mathbf{y}_i$  is an integer vector reduced modulo  $p$  (i.e. that  $\mathbf{y}_i \in [0, p-1]^n$ ), that for  $i \in [L]$ ,  $\bar{\mathbf{e}}_i \in [0, \frac{q}{p} - 1]^n$ , and also that  $x_{i,1} + x_{i,2} = 1$ , for  $i = 1, \dots, L$ . We will use vector decompositions. In particular, introducing the *binary* vectors  $\tilde{\mathbf{y}}_i := \text{vdec}_{m,p-1}(\mathbf{y}_i)$  and  $\tilde{\mathbf{e}}_i := \text{vdec}_{m,\frac{q}{p}-1}(\bar{\mathbf{e}}_i)$ , we append the

following to Equations (5):

$$\begin{aligned}
x_{1,1} + x_{1,2} &= 1 \bmod q \\
&\vdots \\
x_{L,1} + x_{L,2} &= 1 \bmod q \\
\mathbf{y}_1 - \mathbf{H}_{m,p-1} \cdot \tilde{\mathbf{y}}_1 &= \mathbf{0} \bmod q \\
&\vdots \\
\mathbf{y}_{L-1} - \mathbf{H}_{m,p-1} \cdot \tilde{\mathbf{y}}_{L-1} &= \mathbf{0} \bmod q \\
\bar{\mathbf{e}}_1 - \mathbf{H}_{m,\frac{q}{p}-1} \cdot \tilde{\mathbf{e}}_1 &= \mathbf{0} \bmod q \\
&\vdots \\
\bar{\mathbf{e}}_L - \mathbf{H}_{m,\frac{q}{p}-1} \cdot \tilde{\mathbf{e}}_L &= \mathbf{0} \bmod q
\end{aligned} \tag{6}$$

Note that the system of Equations (5) and (6) is linear in the vector of unknowns

$$(\mathbf{w}')^\top = ((\mathbf{w}'_1)^\top, (\mathbf{w}'_2)^\top, (\mathbf{w}'_3)^\top, (\mathbf{w}'_4)^\top, (\mathbf{w}'_5)^\top, (\mathbf{w}'_6)^\top) \tag{7}$$

where

$$\begin{aligned}
- (\mathbf{w}'_1)^\top &= (x_{1,1}, x_{1,2}, \dots, x_{L,1}, x_{L,2}) \\
- (\mathbf{w}'_2)^\top &= (\mathbf{y}_0^\top, \dots, \mathbf{y}_{L-1}^\top) \\
- (\mathbf{w}'_3)^\top &= (x_{1,1}\mathbf{y}_0^\top, x_{1,2}\mathbf{y}_0^\top, \dots, x_{L,1}\mathbf{y}_{L-1}^\top, x_{L,2}\mathbf{y}_{L-1}^\top) \\
- (\mathbf{w}'_4)^\top &= (\tilde{\mathbf{y}}_0^\top, \dots, \tilde{\mathbf{y}}_{L-1}^\top) \\
- (\mathbf{w}'_5)^\top &= (\bar{\mathbf{e}}_1^\top, \dots, \bar{\mathbf{e}}_L^\top) \\
- (\mathbf{w}'_6)^\top &= (\tilde{\mathbf{e}}_1^\top, \dots, \tilde{\mathbf{e}}_L^\top)
\end{aligned}$$

At this point it can be seen that finding an input/seed pair with GGM evaluation output equal to  $\mathbf{y}$  corresponds *precisely* to finding a solution  $\mathbf{w}'$  to the system  $\mathbf{M}'' \cdot \mathbf{w}' = \mathbf{y}'' \bmod q$  where:

- The matrix defining the linear system is

$$\mathbf{M}'' = \left[ \begin{array}{c|c|c|c|c|c} \mathbf{I}_L \otimes (1, 1) & & & & & \\ \hline & \mathbf{I}_{mL} & \mathbf{0}^{mL \times 2mL} & -\mathbf{I}_L \otimes \mathbf{H}_{m,p-1} & & \\ \hline & & & & \mathbf{I}_{mL} & -\mathbf{I}_L \otimes \mathbf{H}_{m,q/p-1} \\ \hline & \frac{q}{p}\mathbf{I}_{mL} & -\mathbf{I}_{mL} \otimes [\mathbf{A}_0 \parallel \mathbf{A}_1] & & \mathbf{I}_{mL} & \end{array} \right]$$

- $(\mathbf{y})''^\top = (1^L, 0^{2mL}, \frac{q}{p} \cdot \mathbf{y}^\top)$
- when parsed as in Equation (7),  $\mathbf{w}'$  is subject to the following quadratic constraints:
  - $\mathbf{w}'_1[2i] = \mathbf{w}'_1[2i] \cdot \mathbf{w}'_1[2i]$  for  $i \in [L]$
  - $\mathbf{w}'_3[2m(i-1) + j] = \mathbf{w}'_1[2i-1] \cdot \mathbf{w}'_2[m(i-1) + j]$  for  $i \in [L], j \in [m]$
  - $\mathbf{w}'_3[2m(i-1) + m + j] = \mathbf{w}'_1[2i] \cdot \mathbf{w}'_2[m(i-1) + j]$  for  $i \in [L], j \in [m]$
  - $\mathbf{w}'_4[i] = \mathbf{w}'_4[i] \cdot \mathbf{w}'_4[i]$  for  $i \in [m \cdot \lfloor \log(p-1) \rfloor]$
  - $\mathbf{w}'_6[i] = \mathbf{w}'_6[i] \cdot \mathbf{w}'_6[i]$  for  $i \in [m \cdot \lfloor \log(\frac{q}{p}-1) \rfloor]$

Note that one can trivially remove  $\mathbf{w}'_5$  from the system by replacing  $\mathbf{w}'_5 = \mathbf{I}_L \otimes \mathbf{H}_{m,q/p-1} \cdot \mathbf{w}'_6$ , but we choose not to for the sake of simplifying presentation.

## C Deferred Security Proofs

### C.1 Proof of Anonymity

We begin with a high level intuition of the proof structure. We will show that any coin  $(\text{info}, PK_{\mathcal{M}}, \mathbf{y}_S, \mathbf{y}_T, \pi_K)$  that is the result of a real **Spend** protocol run can be replaced by a simulated coin  $(\text{info}, PK_{\mathcal{M}}, \mathbf{y}'_S, \mathbf{y}'_T, \pi'_K)$ . In slightly more detail, recall that  $\mathbf{y}_S$  and  $\mathbf{y}_T$  depend on a PRF evaluations  $\mathbf{y}_{\mathbf{k}}$ , where  $\mathbf{k}$  is a secret value stored in a wallet. The idea of the proof is to first replace a real ZKAoK  $\pi_K$  with a simulated  $\pi'_K$  using the zero-knowledge simulator. Next, we replace  $\mathbf{y}_{\mathbf{k}}$  by a uniform random value and then finally use Lemma 2 to justify replacing the resulting serial number/security tag by a uniformly chosen value.

**Lemma 9.** *Let  $\Pi_1$  be the interactive ZKAoK used during the Withdraw protocol and  $\Pi_2$  be the non-interactive one used during the Spend protocol. Construction 4 provides the anonymity property assuming parameters are chosen so that (i)  $\Pi_1$  and  $\Pi_2$  are both ZK; (ii) the BLMR PRF with respect to  $n, m, p, q$  is pseudo-random; (iii) the  $\text{LWE}_{n_{\text{LTF}}, m_{\text{LTF}}, q_{\text{LTF}}, \alpha}$  problem is hard.*

*Proof.* Let  $\mathcal{A}$  denote a PPT adversary in the anonymity game. Let  $Q_U$  denote the number of users that  $\mathcal{A}$  initialises by making  $\mathcal{Q}_{\text{GetKey}}$  or  $\mathcal{Q}_{\text{Withdraw}}$ -queries and let  $Q_w$  denote the maximal number of  $\mathcal{Q}_{\text{Withdraw}}$ -queries involving any single user. We proceed via a sequence of games as listed next. For notational convenience, we set  $\text{Game}_{2.(0, Q_w)} \equiv \text{Game}_1$ ,  $\text{Game}_{2.(k, Q_w+1)} \equiv \text{Game}_{2.(k+1, 1)}$ ,  $\text{Game}_{4.(0, Q_w)} \equiv \text{Game}_3$  and  $\text{Game}_{4.(k, Q_w+1)} \equiv \text{Game}_{4.(k+1, 1)}$ . In addition, we use  $W_l$  to denote the event that  $\mathcal{A}$  wins  $\text{Game}_l$ .

**Game<sub>0</sub>:** The challenger behaves as described in the real anonymity game with coin toss  $d = 1$ .

**Game<sub>1</sub>:** The challenger behaves as in **Game<sub>0</sub>** except that, when responding to  $\mathcal{Q}_{\text{Spend}}$ -queries, all non-interactive arguments are simulated without using the witnesses. This is achieved by using the honest-verifier zero-knowledge simulator and programming the random oracle  $H_{\text{FS}} : \{0, 1\}^* \rightarrow \{-\bar{p}, \dots, \bar{p}\}^\kappa$ . If any of the zero-knowledge simulations fails, the challenger aborts.

**Game<sub>2.(k, f)</sub>**  $((k, f) \in \{1, \dots, Q_U\} \times \{1, \dots, Q_w\})$ : The challenger proceeds as in **Game<sub>2.(k, f-1)</sub>** except for one modification in the treatment of spending queries of the form  $\mathcal{Q}_{\text{Spend}}(\cdot, k, f, \cdot, \cdot)$ . Using the notations of Section 4, the challenger replaces the real PRF output  $\mathbf{y}_{\mathbf{k}}$  by a uniformly chosen vector  $\mathbf{y}_{\mathbf{k}} \leftarrow U(\mathbb{Z}_p^m)$ . All spending queries of the form  $\mathcal{Q}_{\text{Spend}}(\cdot, i, j, \cdot, \cdot)$  with  $i > k$  or  $(i = k) \wedge (j > f)$  are answered as in **Game<sub>2.(k, f-1)</sub>**.

**Game<sub>3</sub>:** The challenger's behaviour is identical to that in **Game<sub>2.(Q\_U, Q\_w)</sub>** with the sole difference that the evaluation key  $ek_{\text{LTF}}$  of  $F_{\text{LTF}}$  is sampled as a lossy key in the setup phase.

**Game<sub>4.(k, f)</sub>**  $((k, f) \in \{1, \dots, Q_U\} \times \{1, \dots, Q_w\})$ : The challenger behaves as in **Game<sub>4.(k, f-1)</sub>** apart from when responding to spending queries of the form  $\mathcal{Q}_{\text{Spend}}(\cdot, k, f, \cdot, \cdot)$ . For these queries, the challenger computes the first component of a serial number  $\mathbf{y}_S := (\mathbf{y}_1, \mathbf{y}_2) \in \mathbb{Z}_{q_{\text{LTF}}}^{n_{\text{LTF}} + \bar{n}_{\text{LTF}}} \times \mathbb{Z}_p^n$  as before. However,

it now sets the second component to a uniform value  $\mathbf{y}_2 \leftarrow U(\mathbb{Z}_p^n)$ . It also sets  $\mathbf{y}_T \leftarrow U(\mathbb{Z}_p^n)$ . All spending queries of the form  $\mathcal{Q}_{\text{Spend}}(\cdot, i, j, \cdot, \cdot)$  with  $i > k$  or  $(i = k) \wedge (j > f)$  are dealt with as in  $\text{Game}_{5.(k, f-1)}$ .

$\text{Game}_0 \approx_s \text{Game}_1$  : We first argue that no PPT adversary  $\mathcal{A}$  can tell whether it is running playing  $\text{Game}_0$  or  $\text{Game}_1$ . To show this, we bound the probability that the challenger aborts in  $\text{Game}_1$ . Note that an abort occurs when the challenger attempts to program the random oracle on an input where  $H_{\text{FS}}$  was already defined. If  $Q_H$  is the maximal number of random oracle queries made by  $\mathcal{A}$  and  $Q_S$  denotes the number of  $\mathcal{Q}_{\text{Spend}}$ -queries, then we can upper-bound the abort probability when simulating a *single* argument by  $(Q_S + Q_H)/2^\kappa$ . Using a union bound, the probability that an abort occurs when simulating all  $Q_S$   $\mathcal{Q}_{\text{Spend}}$ -queries is upper bounded by  $Q_S \cdot (Q_S + Q_H)/2^\kappa$ , which is also negligible. In addition, the zero-knowledge property implies that  $\mathcal{A}$  cannot distinguish  $\text{Game}_0$  from  $\text{Game}_1$  when no abort occurs. This implies that  $|\Pr[W_0] - \Pr[W_1]| = \text{negl}(\lambda)$ .

$\text{Game}_{2.(k, f-1)} \approx_c \text{Game}_{2.(k, f)}$  : We show that, if  $\mathcal{A}$  can distinguish  $\text{Game}_{2.(k, f-1)}$  from  $\text{Game}_{2.(k, f)}$ , then there is an efficient algorithm  $\mathcal{D}$  distinguishing the BLMR PRF from a truly random function. From now on, we will refer to the BLMR PRF as  $F$ . Let  $\omega$  denote a PRF challenge bit. That is, if  $\omega = 1$ , the PRF challenger answers  $\mathcal{D}$ 's queries pseudo-randomly whereas, if  $\omega = 0$ , it outputs uniformly random vectors over  $\mathbb{Z}_p^m$  at each new query. We now describe  $\mathcal{D}$ , which plays the role of  $\mathcal{A}$ 's challenger as described in both  $\text{Game}_{2.(k, f-1)}$  and  $\text{Game}_{2.(k, f)}$  for all queries that are not of the form  $\mathcal{Q}_{\text{Spend}}(\cdot, k, f, \cdot, \cdot)$ . For queries of the form  $\mathcal{Q}_{\text{Spend}}(\cdot, k, f, \cdot, \cdot)$ , it checks the internal counter  $J$  of  $\mathcal{W}_f$  (aborting if  $J = 2^L - 1$ ) and queries its PRF challenger on the input  $J$  to obtain a response  $\mathbf{y} \in \mathbb{Z}_p^m$ . It uses the latter to set  $\mathbf{y}_k = \mathbf{y}$  and computes the coin as it does in both  $\text{Game}_{2.(k, f-1)}$  and  $\text{Game}_{2.(k, f)}$ . Finally,  $\mathcal{D}$  outputs whatever  $\mathcal{A}$  outputs.

Suppose that  $\omega = 1$ . Denote  $\mathcal{W}_f = (\mathbf{e}_u, \mathbf{k}, \text{Sig}_B, J)$ , where  $\mathbf{k} \leftarrow U(\mathbb{Z}_q^m)$  was randomly chosen by  $\mathcal{D}$  at some  $\mathcal{Q}_{\text{Withdraw}}$ -query. If  $\omega = 1$ ,  $\mathcal{D}$  sets  $\mathbf{y}_k$  to be  $\text{PRF}_{\mathbf{k}^*}(J)$ , where  $\mathbf{k}^* \sim U(\mathbb{Z}_q^m)$  is the secret key chosen by the PRF challenger. Note that  $\mathbf{k}^* \neq \mathbf{k}$  w.h.p. while, in the withdrawal protocol,  $\mathcal{D}$  sent  $\mathcal{A}$  the statistically hiding commitment

$$\mathbf{c}_u = \mathbf{D}'_0 \cdot \mathbf{r}_0 + \mathbf{D}_1 \cdot \mathbf{e}_u + \mathbf{D}_2 \cdot \tilde{\mathbf{k}},$$

where  $\tilde{\mathbf{k}} = \text{vdec}_{m, q-1}(\mathbf{k})$  and  $\mathbf{r}_0 \leftarrow D_{\mathbb{Z}^{m_s}, \sigma}$ . In addition,  $\mathcal{D}$  generated an interactive computational witness indistinguishable (WI)<sup>9</sup> argument that  $\mathbf{c}_u$  is honestly computed for some  $\mathbf{r}_0$  of infinity norm at most  $\beta = \sigma \cdot \sqrt{m_s}$  and binary  $\tilde{\mathbf{k}}$ . Now, define  $\tilde{\mathbf{k}}^* = \text{vdec}_{m, q-1}(\mathbf{k}^*)$  and consider a sub-hybrid where  $\mathcal{D}$  samples uniform  $\mathbf{k}^*$  for itself and  $\mathbf{D}'_0$  with a trapdoor [26] allowing for the efficient sampling of a vector  $\mathbf{r}_0^* \leftarrow D_{\Lambda_{q_s}^{\mathbf{D}'_3(\tilde{\mathbf{k}} - \tilde{\mathbf{k}}^*) + \mathbf{D}'_0 \mathbf{r}_0}(\mathbf{D}'_0), \sigma}$ . Note that  $\|\mathbf{r}_0^*\| \leq \beta$  with all but negligible probability and

$$\mathbf{c}_u = \mathbf{D}'_0 \cdot \mathbf{r}_0^* + \mathbf{D}_1 \cdot \mathbf{e}_u + \mathbf{D}_2 \cdot \tilde{\mathbf{k}}^*,$$

<sup>9</sup> The zero-knowledge property implies witness indistinguishability by, e.g., [20, Proposition 4.6.3]

allowing  $\mathbf{k}$  to be used in a withdrawal protocols and  $\mathbf{k}^*$  to be used in the spend protocol. At this point, we can rely on the witness indistinguishability of the interactive argument generated by  $\mathcal{D}$  at step 1 of the withdrawal protocol to show that  $\mathcal{A}$  cannot distinguish whether it is interacting with the challenger described in  $\text{Game}_{2.(k,f-1)}$ , or with  $\mathcal{D}$  involved in a game with its PRF challenger holding a bit  $\omega = 1$ . If  $\omega = 0$ , then  $\mathcal{D}$  perfectly emulates the challenger of  $\text{Game}_{2.(k,f)}$ . This implies that, if  $\mathcal{A}$  can distinguish  $\text{Game}_{2.(k,f-1)}$  from  $\text{Game}_{2.(k,f)}$ , then  $\mathcal{D}$  can distinguish the PRF from uniform, thus contradicting the assumption that  $n, m, p, q$  are chosen so that the BLMR PRF is pseudo-random. Consequently, under this assumption, it must be that

$$|\Pr[W_{2.(k,f-1)}] - \Pr[W_{2.(k,f)}]| \leq \text{negl}(\lambda).$$

Finally, in  $\text{Game}_{2.(Q_U, Q_w)}$ , the challenger does not need any wallet information to answer  $\mathcal{Q}_{\text{Spend}}$ -queries. However, it still has to know the users' public keys to create the simulated spendings. To remove this dependency, the rest of the proof exploits the lossiness properties of the trapdoor function family  $F_{\text{LTF}}$ .

$\text{Game}_{2.(Q_U, Q_w)} \approx_c \text{Game}_3$  : We now observe that, if the adversary  $\mathcal{A}$  can distinguish between  $\text{Game}_{2.(Q_U, Q_w)}$  and  $\text{Game}_3$ , then it can distinguish between the cases where  $ek_{\text{LTF}}$  is sampled as a lossy mode or injective evaluation key. Since we are using the LWE-based LTF described in [29], such an adversary would violate the  $\text{LWE}_{n_{\text{LTF}}, m_{\text{LTF}}, q_{\text{LTF}}, \alpha}$  assumption by Lemma 1.

$\text{Game}_{4.(k,f-1)} \approx_s \text{Game}_{4.(k,f)}$  : Next, we argue that  $\text{Game}_{4.(k,f)}$  is indistinguishable from  $\text{Game}_{4.(k,f-1)}$  via a statistical argument. Let  $F_{\text{LTF}}$  be the lossy trapdoor function and  $H_{\text{UH}}, H'_{\text{UH}}$  be the universal hash functions (indexed by  $ek_{\text{LTF}}$  and  $\mathbf{U}_{\text{UH}}, \mathbf{U}'_{\text{UH}}$ ) sampled in  $\text{Game}_{4.(k,f-1)}$  and denote the same functions in  $\text{Game}_{4.(k,f)}$  as  $\bar{F}_{\text{LTF}}$  and  $\bar{H}_{\text{UH}}, \bar{H}'_{\text{UH}}$  (indexed by  $\bar{ek}_{\text{LTF}}$  and  $\bar{\mathbf{U}}_{\text{UH}}, \bar{\mathbf{U}}'_{\text{UH}}$ ). To aide presentation, we denote the distribution arising from taking the standard binary decomposition of a uniformly sampled vector in  $\mathbb{Z}_p^m$  as  $\text{bin}_p(U(\mathbb{Z}_p^m))$ . We also define the following matrices:

$$\mathbf{U} := \begin{bmatrix} \mathbf{U}_{\text{UH}} \\ \mathbf{U}'_{\text{UH}} \end{bmatrix}, \quad \bar{\mathbf{U}} := \begin{bmatrix} \bar{\mathbf{U}}_{\text{UH}} \\ \bar{\mathbf{U}}'_{\text{UH}} \end{bmatrix} \in \mathbb{Z}_p^{2n \times m_{\text{LTF}}}.$$

Note that we may apply Lemma 2 to show that for  $\mathbf{y}, \bar{\mathbf{y}} \leftarrow \text{bin}(\mathbb{Z}_p^n)$  and  $\bar{\mathbf{u}}' \leftarrow U(\mathbb{Z}_p^{2n})$ , the distributions of  $(\mathbf{U}\mathbf{y}, \mathbf{U}, (ek_{\text{LTF}}, F_{\text{LTF}}(\mathbf{y})))$  and  $(\bar{\mathbf{u}}', \bar{\mathbf{U}}, (ek_{\text{LTF}}, \bar{F}_{\text{LTF}}(\bar{\mathbf{y}})))$  are at most a statistical distance

$$\delta = \frac{1}{2} \sqrt{2^{-(m \log p - (m_{\text{LTF}} \gamma + n_{\text{LTF}} \log q_{\text{LTF}} + 2n \log p))}}.$$

apart. Recall that  $m = \lceil n \log q \rceil$ ,  $m_{\text{LTF}} = m \cdot \lceil \log p \rceil$ ,  $n_{\text{LTF}} = \tilde{O}(n)$ ,  $\gamma < 1$  is a constant and  $m_{\text{LTF}} = \Theta(n_{\text{LTF}} \log q_{\text{LTF}})$ . From these observations, it can be seen that the exponent in the square root is negative and  $\Omega(n_{\text{LTF}} \log q_{\text{LTF}})$  in absolute value (for appropriate choices of constants), implying that the statistical distance

$\delta$  is negligible. By a hybrid argument, it follows that the statistical distribution between

$$\left( \mathbf{U}, ek_{\text{LTF}}, \{F_{\text{LTF}}(\mathbf{y}_J)\}_{J=1}^{2^L}, \{\mathbf{U}\mathbf{y}_J\}_{J=1}^{2^L} \right) \quad (8)$$

and

$$\left( \bar{\mathbf{U}}, e\bar{k}_{\text{LTF}}, \{\bar{F}_{\text{LTF}}(\bar{\mathbf{y}}_J)\}_{J=1}^{2^L}, \{\bar{\mathbf{u}}'_J\}_{J=1}^{2^L} \right) \quad (9)$$

where  $\mathbf{y}_J, \bar{\mathbf{y}}_J \leftarrow \text{bin}_p(\mathbb{Z}_p^n)$  and  $\bar{\mathbf{u}}_J \leftarrow U(\mathbb{Z}_p^{2n})$  is at most  $2^L \cdot \delta$  which is also negligible. We now argue that an adversary that distinguishes between  $\text{Game}_{4.(k,f-1)}$  and  $\text{Game}_{4.(k,f)}$  with noticeable probability would also distinguish (8) and (9) with noticeable probability. We will use the fact that for any  $R = H_0(\text{info}) \neq \mathbf{0} \in \mathbb{Z}_p^n$ , the matrix  $\text{FRD}(R)$  must be invertible due to the function  $\text{FRD}$  [1] that we use. This follows from the observation that  $\text{FRD}(\mathbf{0}) = \mathbf{0}^{n \times n}$  which implies that  $\text{FRD}(R) = \text{FRD}(R) - \text{FRD}(\mathbf{0})$  must be invertible over  $\mathbb{Z}_p$  by the full-rank differences property. We will show that there is an algorithm  $\mathcal{C}$  that perfectly simulates perfectly  $\text{Game}_{4.(k,f-1)}$  on input distributed as in (8), and simulates  $\text{Game}_{4.(k,f)}$  on input distributed as in (9). In doing so, if  $\mathcal{A}$  manages to distinguish the two cases with noticeable probability, then  $\mathcal{C}$  distinguishes (8) from (9) leading to a contradiction. On parsing its input as

$$\left( \begin{bmatrix} \mathbf{U}_{\text{UH},0}^* \\ \mathbf{U}_{\text{UH},1}^* \end{bmatrix}, ek_{\text{LTF}}^*, \{\mathbf{f}_J^*\}_{J=1}^{2^L}, \left\{ \begin{bmatrix} \mathbf{u}_{J,0}^* \\ \mathbf{u}_{J,1}^* \end{bmatrix} \right\}_{J=1}^{2^L} \right),$$

$\mathcal{C}$  simulates  $\text{ParGen}$  as described in the construction apart from that it sets the LTF key to  $ek_{\text{LTF}}^*$  and uses the universal hash functions associated to the matrices  $\mathbf{U}_{\text{UH},0}^*$  and  $\mathbf{U}_{\text{UH},1}^*$ .  $\mathcal{C}$  also answers all queries apart from those of the form  $\mathcal{Q}_{\text{Spend}}(\cdot, \cdot, k, f, \cdot, \cdot)$  as in  $\text{Game}_{4.(k,f-1)}$  and  $\text{Game}_{4.(k,f)}$ . For a  $\mathcal{Q}_{\text{Spend}}(\cdot, k, f, \cdot, \text{info})$  query where  $\mathcal{W}_f$  has been assigned to the  $k$ -th user, it sets serial number  $\mathbf{y}_S = (f_{J^*}^*, PK_k + \mathbf{u}_{J^*,0}^*)$  and security tag  $\mathbf{y}_T PK_k + \text{FRD}(H_0(\text{info})) \cdot \mathbf{u}_{J^*,1}^*$  where  $J^* < 2^L - 1$  is the internal counter value of  $\mathcal{W}_f$ . Otherwise,  $\mathcal{C}$  behaves as the challenger in  $\text{Game}_{4.(k,f-1)}$  when answering such queries. It should be clear that if  $\mathcal{C}$ 's input is of the form (8), it perfectly simulates  $\text{Game}_{4.(k,f-1)}$ . On the other hand, if  $\mathcal{C}$  is given input of the form (9), then it simulates  $\text{Game}_{4.(k,f)}$ . This can be argued by making two simple observations. Firstly, multiplying a uniform value of  $\mathbb{Z}_p^{n \times n}$  by the invertible matrix  $\text{FRD}(H_0(\text{info}))$  yields a uniform value in  $\mathbb{Z}_p^n$ . Secondly, adding  $PK_k$  to uniform samples of  $\mathbb{Z}_p^n$  is distributed uniformly. Therefore, we can conclude that even if  $\mathcal{A}$  is unbounded, it cannot distinguish between  $\text{Game}_{4.(k,f-1)}$  and  $\text{Game}_{4.(k,f)}$  with noticeable probability.

The anonymity property follows after noting that the parameter generation and answers to  $\mathcal{Q}_{\text{Spend}}$ -queries in  $\text{Game}_{4.(Q_U, Q_w)}$  are independent of both the wallets issued during the game and user public keys. Therefore,  $\text{SimParGen}$  and  $\text{SimSpend}$  are defined by the challenger's behavior in  $\text{Game}_{4.(Q_U, Q_w)}$ .  $\square$

## C.2 Proof of the Traceability Property

**Lemma 10.** *Let  $\Pi_1$  be the interactive ZKAoK used during the Withdraw protocol and  $\Pi_2$  be the non-interactive one used during the Spend protocol. Assuming*

that (i) the  $\text{SIS}_{n,m_s,q_s,\beta'}$  problem is hard for the parameter  $\beta' = \mathcal{O}(\sigma^2 m_s^{1/2}(m_s + \bar{m}))$ , (ii) that  $H_0$  is sampled from a collision-resistant hash function family, and (iii) that  $\Pi_1$  and  $\Pi_2$  are AoKs, our construction possesses the traceability property in the random oracle model.

*Proof.* An adversary  $\mathcal{A}$  breaks traceability of our e-cash scheme if it is able to spend more coins than it has withdrawn without being identified as a double-spender. Concretely, after  $Q_w$  queries to the  $\mathcal{Q}_{\text{Withdraw}}$  oracle, it is able to output  $N > 2^L \cdot Q_w$  coins  $\{\text{coin}_j = (\text{info}_j, PK_{\mathcal{M}_j}, \mathbf{y}_{S,j}, \mathbf{y}_{T,j}, \pi_{K,j})\}_{j=1}^N$  such that

1.  $(PK_{\mathcal{M}_j}, \text{info}_j) \neq (PK_{\mathcal{M}_k}, \text{info}_k)$ , for all pairs  $(j, k) \in \{1, \dots, N\}^2$ ;
2.  $\text{VerifyCoin}(\text{par}, PK_{\mathcal{M}_j}, PK_{\mathcal{B}}, \text{coin}_j) = 1$ , for  $j \in \{1, \dots, N\}$ ;
3.  $\text{Identify}(\text{par}, PK_{\mathcal{B}}, \text{coin}_j, \text{coin}_k)$  does not output a public key in the database  $\mathbb{T}$ , for all pairs  $(j, k) \in \{1, \dots, N\}^2$ .

Each of these coins contains an argument of knowledge of  $\mathbf{v}^{(j)}, \mathbf{r}_0^{(j)}, \mathbf{r}_1^{(j)}, \mathbf{e}_{u,j}$  and  $\tau^{(j)}, \mathbf{k}^{(j)}$  such that

$$\begin{aligned} \mathbf{A}_{\tau^{(j)}} \cdot \mathbf{v}^{(j)} &= \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{n,q_s-1} \left( \mathbf{w}^{(j)} \right) \\ \mathbf{w}^{(j)} &= \mathbf{c}_{\mathcal{U}_j} + \mathbf{D}_0'' \cdot \mathbf{r}_1^{(j)} \\ \mathbf{c}_{\mathcal{U}}^* &= \mathbf{D}_0' \cdot \mathbf{r}_0^{(j)} + \mathbf{D}_1 \cdot \mathbf{e}_{u,j} + \mathbf{D}_2 \cdot \text{vdec}_{m,q-1} \left( \mathbf{k}^{(j)} \right) \end{aligned} \quad (10)$$

At each query to the  $\mathcal{Q}_{\text{Withdraw}}$  oracle, the adversary playing the role of a user  $\mathcal{U}_i$  sends a commitment

$$\mathbf{c}_{\mathcal{U}_i} = \mathbf{D}_0' \cdot \mathbf{r}_0^{(i)} + \mathbf{D}_1 \cdot \mathbf{e}_{u_i} + \mathbf{D}_2 \cdot \text{vdec}_{m,q-1} \left( \mathbf{k}^{(i)} \right) \in \mathbb{Z}_{q_s}^n$$

to the bank and interactively proves knowledge of the corresponding openings. The reduction is thus able to invoke a knowledge extractor  $\mathcal{E}$  to extract such values.

To distinguish elements involved in such queries from those involved in the adversary output, we will always use a superscript  $i$  in the former case and  $j$  in the latter case. For example,  $\mathbf{w}^{(i)} = \mathbf{c}_{\mathcal{U}_i} + \mathbf{D}_0'' \cdot \mathbf{r}_1^{(i)}$  and  $\tau^{(i)}$  are the elements used by the bank during the  $i$ -th  $\mathcal{Q}_{\text{Withdraw}}$ -query.

We distinguish the following cases.

- **Case 1:**  $\exists j^* \in \{1, \dots, N\}$  such that  $\forall i \in \{1, \dots, Q_w\}, \tau^{(j^*)} \neq \tau^{(i)}$ .
- **Case 2:**  $\exists j^* \in \{1, \dots, N\}$  and  $i^* \in \{1, \dots, Q_w\}$  such that  $\tau^{(j^*)} = \tau^{(i^*)}$  but  $\mathbf{w}^{(j^*)} \neq \mathbf{w}^{(i^*)}$ .
- **Case 3:** Cases 1 and 2 do not occur and there is no collision in the set of serial numbers  $\{\mathbf{y}_{S,i}\}_{i=1}^N$ .
- **Case 4:** Cases 1 and 2 do not occur and there is a collision in the set of serial numbers  $\{\mathbf{y}_{S,i}\}_{i=1}^N$ .

In the first two cases, our reduction makes a guess  $j^* \leftarrow U(\{1, \dots, N\})$  for the index  $j^*$ . This guess is correct with probability at least  $1/N$ . For the sake of clarity, the elements associated with this guessed index will only be denoted with a superscript  $*$  (e.g.,  $\mathbf{w}^*$  instead of  $\mathbf{w}^{(j^*)}$ ).

For  $k \in \{1, 2, 3, 4\}$ , we define the algorithm  $\mathcal{B}_k$  designed to solve a SIS instance for matrix  $\mathbf{A} \in \mathbb{Z}_{q_s}^{n \times m_s}$  with non-negligible probability when  $\mathcal{A}$  breaks traceability in **Case k**.

• **Case 1: Description of  $\mathcal{B}_1$**

*Simulation of ParGen and BKeyGen.*  $\mathcal{B}_1$  samples  $\tau^{(1)}, \dots, \tau^{(Q_w)} \leftarrow \{0, 1\}^\ell$  that will be used to answer  $\mathcal{Q}_{\text{Withdraw}}$ -queries at the outset of the game. Next,  $\mathcal{B}_1$  generates  $(\mathbf{C}, \mathbf{T}_{\mathbf{C}}) \leftarrow \text{TrapGen}(1^n, 1^{m_s}, q_s)$ . It then draws random indexes  $i' \leftarrow U([Q_w])$  and  $t' \leftarrow U(\{0, \dots, \ell - 1\})$  and sets  $\tau' = \tau_{[t']^{(i)}}^{(i)}[\tau^{(i)}[t' + 1]]$  where  $\tau_{[t']^{(i)}}$  denotes the vector consisting of the first  $t'$  entries of  $\tau^{(i)}$ . Given the SIS input  $\mathbf{A} \in \mathbb{Z}_{q_s}^{n \times m_s}$ , it samples  $\mathbf{S}_j \leftarrow D_{\mathbb{Z}^{m_s \times m_s}, \sigma}$  for  $j \in \{0, \dots, \ell\}$  and sets

$$\mathbf{A} = \bar{\mathbf{A}}$$

$$\mathbf{A}_0 = \bar{\mathbf{A}} \cdot \mathbf{S}_0 + \sum_{j=1}^{t'} \tau'[j] \cdot \mathbf{C}$$

$$\begin{aligned} \mathbf{A}_j &= \bar{\mathbf{A}} \cdot \mathbf{S}_j + (-1)^{\tau'[j]} \cdot \mathbf{C}, & j &= 1, \dots, t' \\ \mathbf{A}_j &= \bar{\mathbf{A}} \cdot \mathbf{S}_j, & j &= t' + 1, \dots, \ell. \end{aligned}$$

Note that

$$\mathbf{A}_\tau = \left[ \bar{\mathbf{A}} \mid \bar{\mathbf{A}} \cdot \left( \mathbf{S}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{S}_j \right) + h_\tau \cdot \mathbf{C} \right],$$

where  $h_\tau$  is the Hamming distance between  $\tau_{[t'+1]}$  and  $\tau'$ . Furthermore, we have  $h_\tau = 0 \bmod q_s$  if and only if  $\tau'$  is a prefix of  $\tau$  because  $h_\tau \in [0, t' + 1]$  and  $t' + 1 \leq q_s$ . Next,  $\mathcal{B}_1$  samples  $\mathbf{s} \leftarrow D_{\mathbb{Z}^{m_s}, \sigma}$ ,  $\mathbf{R} \leftarrow D_{\mathbb{Z}^{m_s \times m_s/2}, \sigma}$  and sets  $\mathbf{u} = \bar{\mathbf{A}} \cdot \mathbf{s}$ ,  $\mathbf{D} = \bar{\mathbf{A}} \cdot \mathbf{R}$ . It generates the remaining components of the public key of the bank and the public parameters as described in the real construction and feeds  $\mathcal{A}$  with them. Note that the simulated public parameters and  $PK_{\mathcal{B}}$  are sampled from distributions which are statistically close to those used in the real scheme by the standard Leftover Hash Lemma.

*Simulation of  $\mathcal{Q}_{\text{Withdraw}}$ -queries.* At the  $i$ -th withdrawal query,  $\mathcal{B}_1$  uses the value  $\tau^{(i)}$ . If  $\tau'$  is a prefix of  $\tau^{(i)}$ ,  $\mathcal{B}_2$  aborts. Otherwise,  $\mathcal{B}_2$  considers the matrix

$$\mathbf{A}_{\tau^{(i)}} = \left[ \bar{\mathbf{A}} \mid \bar{\mathbf{A}} \cdot \left( \mathbf{S}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{S}_j \right) + h_\tau \cdot \mathbf{C} \right]$$

and uses  $\mathbf{T}_{\mathbf{C}}$  (together with the small-norm matrix  $\mathbf{S}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{S}_j \in \mathbb{Z}^{m_s \times m_s}$  and Lemma 7) to sample Gaussian vectors  $\mathbf{v}^{(i)} \in \mathbb{Z}^{2m_s}$  in the desired coset of



the lattice  $\Lambda_{q_s}^\perp(\mathbf{A}_{\tau^{(i)}})$  with a distribution statistically close to the distribution  $D_{\Lambda_{q_s}^{\text{uU}}(\mathbf{A}_{\tau^{(i)}}), \sigma}$  used in the real withdrawal protocol.

*Output stage.* In case 1, if our guess on  $j^*$  is correct, a successful adversary  $\mathcal{A}$  necessarily outputs an argument of knowledge of short  $\mathbf{v}^*$  and some  $\mathbf{w}^*$  such that  $\mathbf{A}_{\tau^*} \cdot \mathbf{v}^* = \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{m,q-1}(\mathbf{w}^*)$ . Note that, with non-negligible probability at least  $1/(Q_w \cdot \ell)$ , the longest common prefix of  $\tau^*$  amongst prefixes of the  $\tau^{(i)}$ 's is  $\tau_{|t'}^{(i')}$ . We call this event  $\mathbf{E}$  and observe that, if  $\mathbf{E}$  occurs, then  $\tau' = \tau_{|t'}^{(i')} \| \overline{\tau^{(i)}[t' + 1]}$  must be a prefix of  $\tau^*$ , in which case we have

$$\mathbf{A}_{\tau^*} = \bar{\mathbf{A}} \cdot [\mathbf{I} \mid \mathbf{S}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{S}_j].$$

Therefore, if  $\mathcal{A}$  wins the traceability game in Case 1 and event  $\mathbf{E}$  occurs, then  $\mathcal{A}$  proves knowledge of short  $\mathbf{v}^*$  and  $\mathbf{w}^*$  such that

$$\bar{\mathbf{A}} \cdot \left( [\mathbf{I} \mid \mathbf{S}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{S}_j] \cdot \mathbf{v}^* \right) = \bar{\mathbf{A}} \cdot (\mathbf{s} + \mathbf{R} \cdot \text{vdec}_{m,q-1}(\mathbf{w}^*)).$$

Rearranging the above equation shows that  $\mathcal{A}$  argues knowledge of a SIS solution with respect to the instance  $\bar{\mathbf{A}}$  provided that

$$\mathbf{v}_{\text{SIS}} := [\mathbf{I} \mid \mathbf{S}_0 + \sum_{j=1}^{\ell} \tau[j] \cdot \mathbf{S}_j] \cdot \mathbf{v}^* - \mathbf{s} - \mathbf{R} \cdot \text{vdec}_{m,q-1}(\mathbf{w}^*)$$

is non-zero. This is true due to the fact that even when conditioned on  $\mathbf{u}$  (which is the only quantity depending on  $\mathbf{s}$  that  $\mathcal{A}$  is given),  $\mathbf{s}$  retains at least  $m_s$  bits of min-entropy by Lemma 2.11 of [28]. Therefore, the probability that  $\mathcal{A}$  outputs  $\mathbf{v}_{\text{SIS}} = 0$  is negligible. Using the Cauchy-Schwarz inequality and the fact that the Euclidean norm of a sample from  $D_{\mathbb{Z}^{n'}, r}$  is bounded by  $r\sqrt{n'}$  with all but negligible probability [3], we may conclude that  $\|\mathbf{v}_{\text{SIS}}\|_{\infty} \leq \sqrt{1 + \sigma^2 m_s} \cdot \sigma \sqrt{2m_s} + \sigma \sqrt{m_s} + \sigma \sqrt{m_s/2} \cdot \sqrt{m_s/2}$  i.e. that  $\|\mathbf{v}_{\text{SIS}}\| \leq \mathcal{O}(\sigma^2 m_s^{3/2})$  with all but negligible probability. The final step is for  $\mathcal{B}_1$  to extract the SIS solution which is done in the same way as in the proof of exculpability (Section 6). We conclude that, if a PPT  $\mathcal{A}$  wins in Case 1 with non-negligible probability, then  $\mathcal{B}_1$  efficiently computes a non-trivial solution to the SIS instance with non-negligible probability.

## • Case 2: Description of $\mathcal{B}_2$

*Simulation of ParGen and BKeyGen.* Before starting its interaction with  $\mathcal{A}$ ,  $\mathcal{B}_2$  samples random strings  $\tau^{(1)}, \dots, \tau^{(Q_w)} \leftarrow U(\{0, 1\}^{\ell})$  that will be used when answering  $\mathcal{Q}_{\text{Withdraw}}$ -queries. Next, it samples  $i^* \leftarrow U(\{1, \dots, Q_w\})$  along with

uniformly chosen  $h_0, \dots, h_\ell \leftarrow U(\mathbb{Z}_{q_s})$  such that

$$\begin{aligned} h_0 + \sum_{j'=1}^{\ell} \tau^{(i^*)}[j'] \cdot h_{j'} &= 0 \\ h_0 + \sum_{j'=1}^{\ell} \tau^{(i)}[j'] \cdot h_{j'} &\neq 0 \text{ for all } i \neq i^*. \end{aligned} \quad (11)$$

In addition,  $\mathcal{B}_2$  generates  $(\mathbf{C}, \mathbf{T}_{\mathbf{C}}) \leftarrow \text{TrapGen}(1^n, 1^{m_s}, q_s)$ . It also samples vectors  $\mathbf{v}^{(i^*)} \leftarrow D_{\mathbb{Z}^{2m_s}, \sigma}$ ,  $\mathbf{c}'_{\mathcal{U}} \leftarrow U(\mathbb{Z}_{q_s}^m)$  as well as matrices  $\mathbf{R} \leftarrow D_{\mathbb{Z}^{m_s \times m_s/2}, \sigma}$  and  $\mathbf{S}_j \leftarrow D_{\mathbb{Z}^{m_s \times m_s}, \sigma}$  for  $j \in \{0, \dots, \ell\}$ . It sets the various elements of the bank's public key to be:

$$\begin{aligned} \mathbf{D} &= \bar{\mathbf{A}} \cdot \mathbf{R} \\ \mathbf{A} &= \bar{\mathbf{A}} \\ \mathbf{A}_j &= \bar{\mathbf{A}} \cdot \mathbf{S}_j + h_j \cdot \mathbf{C}, & j = 0, \dots, \ell \\ \mathbf{u} &= \bar{\mathbf{A}} \cdot \mathbf{v}^{(i^*)} - \mathbf{D} \cdot \text{vdec}_{n, q_s-1}(\mathbf{c}'_{\mathcal{U}}) \end{aligned}$$

Notice that  $\mathbf{A}_{\tau^{(i^*)}}$  does not depend on  $\mathbf{C}$  by the first constraint (11). Finally,  $\mathcal{B}_2$  samples  $\mathbf{D}_0'', \mathbf{T}_{\mathbf{D}_0''} \leftarrow \text{TrapGen}(1^n, 1^{m_s}, q_s)$  and the remaining public parameters as described in the real scheme. Again, the public parameters and the bank's public key are statistically indistinguishable from those involved in the real construction by the standard Leftover Hash Lemma.

*Simulation of  $\mathcal{Q}_{\text{Withdraw-queries}}$ .* For each  $i \in [Q_w]$ , we have

$$\mathbf{A}_{\tau^{(i)}} = \left[ \bar{\mathbf{A}} \mid \bar{\mathbf{A}} \cdot \left( \mathbf{S}_0 + \sum_{j=1}^{\ell} \tau^{(i)}[j] \cdot \mathbf{S}_j \right) + \left( h_0 + \sum_{j=1}^{\ell} \tau^{(i)}[j] \cdot h_j \right) \cdot \mathbf{C} \right],$$

so that  $\mathcal{B}_2$  can answer the  $i$ -th  $\mathcal{Q}_{\text{Withdraw-query}}$  using trapdoor  $\mathbf{T}_{\mathbf{C}}$  as long as  $i \neq i^*$ . The  $i^*$ -th query is answered as follows. On receipt of  $\mathbf{c}_{\mathcal{U}} \in \mathbb{Z}_{q_s}^n$  from  $\mathcal{A}$ ,  $\mathcal{B}_2$  samples

$$\mathbf{r}_1 \leftarrow D_{\Lambda_{q_s}^{\mathbf{c}'_{\mathcal{U}} - \mathbf{c}_{\mathcal{U}}}(\mathbf{D}_0''), \sigma}$$

using the trapdoor  $\mathbf{T}_{\mathbf{D}_0''}$ . By doing so,  $\mathcal{B}_2$  ensures that  $\mathbf{c}_{\mathcal{U}} + \mathbf{D}_0'' \cdot \mathbf{r}_1^{(i^*)} = \mathbf{c}'_{\mathcal{U}}$ , in which case we have

$$\mathbf{A}_{\tau^{(i^*)}} \cdot \mathbf{v}^{(i^*)} = \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{n, q_s-1}(\mathbf{c}'_{\mathcal{U}}) \pmod{q_s}$$

by construction. Note that the distribution of  $(\tau^{(i^*)}, \mathbf{v}^{(i^*)}, \mathbf{r}^{(i^*)})$  is statistically close to its counterpart in the real construction. To see why, we first note that  $\tau^{(i^*)}$  is chosen uniformly as in the real construction. Next, we use the fact that sampling a discrete Gaussian over a uniformly chosen coset of lattice  $\Lambda$  for parameter  $\sigma$  larger than the smoothing parameter of  $\Lambda$  yields the discrete Gaussian distribution over integer vectors to within negligible statistical distance. This fact

follows from Claim 3.8 in [30]. This implies that the distribution of  $\mathbf{r}_1^{(i^*)}$  is statistically close to a discrete Gaussian as it should be. Finally, conditioned on the values  $\mathbf{r}_1^{(i^*)}, \tau^{(i^*)}$ , the distribution of  $\mathbf{v}^{(i^*)}$  is a Gaussian restricted to a coset depending on  $\mathbf{r}_1^{(i^*)}, \tau^{(i^*)}$  as it is in the real construction.

*Output stage.* If  $\mathcal{A}$  wins in Case 2 and if our guess  $i^*$  is correct, it outputs an argument of knowledge of some  $\mathbf{w}^* \in \mathbb{Z}_{q_s}^n$  and a short  $\mathbf{v}^* \in \mathbb{Z}^{2m_s}$  such that  $\mathbf{A}_{\tau^*} \cdot \mathbf{v}^* = \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{n,q_s-1}(\mathbf{w}^*)$  where  $\tau^* = \tau^{(i^*)}$  and  $\mathbf{w}^* \neq \mathbf{w}^{(i^*)}$ . This happens with probability at least  $1/Q_w$ . If this is the case, then we can extract (as in the proof of exculpability) a short  $\mathbf{v}^* \in \mathbb{Z}^{2m_s}$  and binary  $\text{vdec}_{n,q_s-1}(\mathbf{w}^*)$  such that

$$\begin{aligned} \bar{\mathbf{A}} \cdot [\mathbf{I} \mid (\mathbf{S}_0 + \sum_{j=1}^{\ell} \tau^*[j] \cdot \mathbf{S}_j)] \cdot (\mathbf{v}^* - \mathbf{v}^{(i^*)}) \\ = \bar{\mathbf{A}} \cdot \mathbf{R} (\text{vdec}_{n,q_s-1}(\mathbf{w}^*) - \text{vdec}_{n,q_s-1}(\mathbf{c}'_{\mathcal{U}})) \pmod{q_s}. \end{aligned} \quad (12)$$

Since we are assuming that the adversary  $\mathcal{A}$  wins in Case 2, we have  $\mathbf{w}^* \neq \mathbf{w}^{(i^*)}$ , where  $\mathbf{w}^{(i^*)} := \mathbf{c}'_{\mathcal{U}}$ , which implies  $\text{vdec}_{n,q_s-1}(\mathbf{w}^*) \neq \text{vdec}_{n,q_s-1}(\mathbf{c}'_{\mathcal{U}})$ . From this, Equation (12) implies that

$$\begin{aligned} \mathbf{v}_{\text{SIS}} &:= [\mathbf{I} \mid (\mathbf{S}_0 + \sum_{j=1}^{\ell} \tau^*[j] \cdot \mathbf{S}_j)] \cdot (\mathbf{v}^* - \mathbf{v}^{(i^*)}) \\ &\quad - \mathbf{R} \cdot (\text{vdec}_{n,q_s-1}(\mathbf{w}^*) - \text{vdec}_{n,q_s-1}(\mathbf{c}'_{\mathcal{U}})) \in \mathbb{Z}^{m_s} \end{aligned} \quad (13)$$

is a SIS solution provided it is non-zero. Now, while  $PK_{\mathcal{B}}$  depends on  $\{\mathbf{S}_j\}_{j=0}^{\ell}$  and  $\mathbf{R}$  via the relations  $\mathbf{A}_j = \bar{\mathbf{A}} \cdot \mathbf{S}_j + h_j \cdot \mathbf{C}$  and  $\mathbf{D} = \bar{\mathbf{A}} \cdot \mathbf{R}$ , the columns of  $\mathbf{S}_j$  and  $\mathbf{R}$  retain a lot of entropy conditionally on  $PK_{\mathcal{B}}$ . In particular, by [28, Lemma 2.11], each column of  $\mathbf{R} \in \mathbb{Z}^{m_s \times m_s/2}$  retains at least  $m_s$  bits of entropy conditionally on  $\mathbf{D} = \bar{\mathbf{A}} \cdot \mathbf{R}$ . Moreover, this entropy does not decrease during the game since the matrix  $\mathbf{R}$  is never used to answer  $\mathcal{Q}_{\text{Withdraw}}$ -queries. Since  $(\text{vdec}_{n,q_s-1}(\mathbf{w}^*) - \text{vdec}_{n,q_s-1}(\mathbf{c}'_{\mathcal{U}})) \in \{-1, 0, 1\}^{m_s/2} \setminus \{\mathbf{0}^{m_s/2}\}$ , the second term of (13) can be written as the sum of up to  $m_s/2$  Gaussian vectors, of which at least one is statistically unpredictable. Consequently,  $\mathbf{v}_{\text{SIS}} \in \mathbb{Z}^{m_s}$  is a non-zero vector of  $\Lambda_{q_s}^{\perp}(\bar{\mathbf{A}})$  with all but negligible probability. Similarly to the output phase in **Case 1**, we can argue that  $\|\mathbf{v}_{\text{SIS}}\|_{\infty} \leq \sqrt{1 + \sigma^2 m_s} \cdot \sigma \sqrt{2m_s} + \sigma \sqrt{m_s/2} \cdot \sqrt{m_s/2}$  i.e.  $\|\mathbf{v}_{\text{SIS}}\| \leq \mathcal{O}(\sigma^2 m_s^{3/2})$  with all but negligible probability.

### • Case 3: Description of $\mathcal{B}_3$

In Case 3, all the pairs  $(\tau^{(j)}, \mathbf{w}^{(j)})$  involved in the  $j$ -th coin returned by the adversary have also been used to answer a  $\mathcal{Q}_{\text{Withdraw}}$ -query. However,  $\mathcal{A}$  has returned  $N > 2^L \cdot Q_w$  coins while querying  $\mathcal{Q}_{\text{Withdraw}}$  only  $Q_w$  times and, yet, there is no collision in the  $N$  resulting serial numbers.

Let  $\{(\mathbf{e}_u^{(i)}, \mathbf{k}^{(i)})\}_{i=1}^{Q_w}$  be the elements involved in the  $Q_w$   $\mathcal{Q}_{\text{Withdraw}}$  queries. From them, we can derive at most  $2^L \cdot Q_w$  legitimate serial numbers. Since  $N > 2^L \cdot Q_w$  and since we assume that there is no collision, there is a coin  $\text{coin}^*$  returned by the adversary that contains a serial number  $\mathbf{y}_S^*$  which is not legitimate. However, since we here consider Case 3, we know that there is an index  $i^* \in \{1, \dots, Q_w\}$  such that  $(\tau^*, \mathbf{w}^*) = (\tau^{(i^*)}, \mathbf{w}^{(i^*)})$  (otherwise, this would be one of the first two cases). We construct an algorithm  $\mathcal{B}_3$  that solves its SIS instance  $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m_s}$  as follows.

*Simulation of ParGen and BKeyGen.* Algorithm  $\mathcal{B}_3$  samples matrices with Gaussian entries  $\mathbf{Q}_0'' \leftarrow D_{\mathbb{Z}^{m_s \times m_s}, \sigma}$ ,  $\mathbf{Q}_1 \leftarrow D_{\mathbb{Z}^{m_s \times m}, \sigma}$ ,  $\mathbf{Q}_2, \mathbf{Q}_3 \leftarrow D_{\mathbb{Z}^{m_s \times \tilde{m}}, \sigma}$ . It then sets  $\mathbf{D}_0' = \bar{\mathbf{A}}$ ,  $\mathbf{D}_0'' = \bar{\mathbf{A}} \cdot \mathbf{Q}_0''$  and  $\mathbf{D}_i = \bar{\mathbf{A}} \cdot \mathbf{Q}_i$  for  $i \in \{1, 2\}$ . Note that the resulting matrices are statistically close to uniform by the standard Leftover Hash Lemma. The remaining public parameters and the bank's public key are chosen as described in the real construction.

*Simulation of  $\mathcal{Q}_{\text{Withdraw}}$ -queries.* Since the matrices  $\bar{\mathbf{A}}$  and  $\{\mathbf{A}_i\}_{i=0}^\ell$  were sampled as described in the real construction,  $\mathcal{B}_3$  is able to honestly answer all  $\mathcal{Q}_{\text{Withdraw}}$  queries. However, at each query  $i \in [Q_w]$ ,  $\mathcal{B}_3$  runs the knowledge extractor of the interactive argument by rewinding  $\mathcal{A}$  so as to extract the witnesses  $(\mathbf{r}_0^{(i)}, \mathbf{e}_u^{(i)}, \tilde{\mathbf{k}}^{(i)})$  in the first step of the withdraw protocol. In particular, for the  $i$ -th query to  $\mathcal{Q}_{\text{Withdraw}}$ ,  $\mathcal{B}_3$  can record the resulting wallet

$$\mathcal{W}_i = \left( \mathbf{e}_u^{(i)}, \mathbf{k}^{(i)}, \left( \tau^{(i)}, \mathbf{v}^{(i)}, \mathbf{r}^{(i)} = \begin{pmatrix} \mathbf{r}_0^{(i)} \\ \mathbf{r}_1^{(i)} \end{pmatrix} \right), J = 0 \right)$$

for later use.

*Output stage.* If  $\mathcal{A}$  wins the traceability game in Case 3, it outputs a coin containing a non-trivial serial number

$$\mathbf{y}_S^* = \left( F_{\text{LTF}}(\text{bin}_p(\mathbf{y}_{\mathbf{k}^*})), \mathbf{F} \cdot \mathbf{e}_u^* + H_{\text{UH}}(\text{bin}_p(\mathbf{y}_{\mathbf{k}^*})) \right).$$

Moreover, the coin must contain an argument of knowledge of the underlying  $\mathbf{e}_u^* \in \{0, 1\}^m$ ,  $J^* \in \{0, 1\}^L$ ,  $\mathbf{k}^* \in \mathbb{Z}_q^m$  as well as a binary  $\tau^* \in \{0, 1\}^\ell$  and short vectors  $\mathbf{v}^* \in \mathbb{Z}^{2m_s}$ ,  $\mathbf{r}^* = (\mathbf{r}_0^{*\top} \parallel \mathbf{r}_1^{*\top})^\top \in \mathbb{Z}^{2m_s}$  satisfying the verification equations (10). By the same reasoning as in the proof of exculpability,  $\mathcal{B}_3$  can extract the witness  $(\mathbf{v}^*, \mathbf{r}^*, \mathbf{e}_u^*, \tau^*, \mathbf{k}^*, J^*)$ .

Since we are in Case 3, there exists  $i^* \in [Q_w]$  such that the values

$$\begin{aligned} \mathbf{r}' &:= \mathbf{r}^* - \mathbf{r}^{(i^*)} \\ \mathbf{e}_u' &:= \mathbf{e}_u^* - \mathbf{e}_u^{(i^*)} \\ \mathbf{k}' &:= \text{vdec}_{m, q-1}(\mathbf{k}^*) - \text{vdec}_{m, q-1}(\mathbf{k}^{(i^*)}) \end{aligned}$$

satisfy

$$[\mathbf{D}_0' \parallel \mathbf{D}_0''] \cdot \mathbf{r}' + \mathbf{D}_1 \cdot \mathbf{e}_u' + \mathbf{D}_2 \cdot \mathbf{k}' = \mathbf{0}^n \pmod{q_s}. \quad (14)$$

Moreover, the serial number  $\mathbf{y}_S^*$  does not belong to the set

$$S_{i^*} := \left\{ \left( F_{\text{LTF}}(\text{bin}_p(\text{PRF}_{\mathbf{k}^{(i^*)}}(J))) \right), \right. \\ \left. \mathbf{F} \cdot \mathbf{e}_u^{(i^*)} + H_{\text{UH}}(\text{bin}_p(\text{PRF}_{\mathbf{k}^{(i^*)}}(J))) \right) : J \in \{0, 1\}^L \Big\},$$

since, otherwise, it would be a legitimate serial number. We now rewrite Equation (14) as

$$\bar{\mathbf{A}} \cdot \left( \mathbf{Q} \cdot (\mathbf{r}'^\top \mid \mathbf{e}_u'^\top \mid \mathbf{k}'^\top)^\top \right) = \mathbf{0}^n \pmod{q_s}$$

where  $\mathbf{Q} := [\mathbf{I} \mid \mathbf{Q}_0'' \mid \mathbf{Q}_1 \mid \mathbf{Q}_2]$ . We now argue that

$$\mathbf{v}_{\text{SIS}} := \mathbf{Q} \cdot (\mathbf{r}'^\top \mid \mathbf{e}_u'^\top \mid \mathbf{k}'^\top)^\top$$

is non-zero which implies that it is a valid SIS solution. Since  $\mathbf{y}_S^* \notin S_{i^*}$ , we must have *at least one* of the following cases since  $F_{\text{LTF}}$  was generated in injective mode:

1.  $\text{PRF}_{\mathbf{k}^*}(J^*) \neq \text{PRF}_{\mathbf{k}^{(i^*)}}(J)$  for all  $J \in \{0, 1\}^L$ , which implies that  $\mathbf{k}^* \neq \mathbf{k}^{(i^*)}$ ;
2.  $\mathbf{e}_u^* \neq \mathbf{e}_u^{(i^*)}$ .

In any of the above cases, we have  $(\mathbf{r}'^\top \mid \mathbf{e}_u'^\top \mid \mathbf{k}'^\top)^\top \neq \mathbf{0}$ . To complete the proof, we argue that public matrices  $\mathbf{D}_0'', \mathbf{D}_1, \mathbf{D}_2$  do not reveal significant information on  $\mathbf{Q}_0'', \mathbf{Q}_1$  and  $\mathbf{Q}_2$ . To this end, we rely again on [28, Lemma 2.11] to argue that the columns of  $\mathbf{Q}_0'', \mathbf{Q}_1$  and  $\mathbf{Q}_2$  retain at least  $m_s$  bits of min-entropy. Moreover, this entropy never decreases when  $\mathcal{B}_3$  answers adversarial queries. Since  $(\mathbf{e}_u'^\top \mid \mathbf{k}'^\top)^\top \in \{-1, 0, 1\}^{m+\bar{m}} \setminus \{\mathbf{0}^{m+\bar{m}}\}$ , the sum

$$\mathbf{Q}_1 \cdot \mathbf{e}_u' + \mathbf{Q}_2 \cdot \mathbf{k}' \in \mathbb{Z}^{m_s}$$

can be written as the sum of up to  $m + \bar{m}$  Gaussian vectors, of which at least one is statistically unpredictable. This shows that  $\mathbf{v}_{\text{SIS}} \neq \mathbf{0}^{m_s}$  with all but negligible probability. In a similar manner to the previous cases, we can argue that  $\|\mathbf{v}_{\text{SIS}}\|_\infty \leq \sqrt{1 + \sigma^2 m_s} \cdot \sigma \sqrt{2 \cdot 2m_s} + \sigma \sqrt{m} \cdot \sqrt{m} + 3\sigma \sqrt{m} \cdot \sqrt{m}$  i.e. that  $\|\mathbf{v}_{\text{SIS}}\| = \mathcal{O}(\sigma^2 m_s^{1/2} (m_s + \bar{m}))$  with all but negligible probability.

#### • Case 4: Description of $\mathcal{B}_4$

We will take the behaviour of  $\mathcal{B}_4$  to be identical to that of  $\mathcal{B}_3$ . Therefore, we will simply argue that  $\mathcal{B}_4$  may extract a SIS solution in Case 4 assuming most of the analysis from Case 3. The details follow.

In Case 4, we know that the adversary outputs at least two coins (that we denote by  $\text{coin}_1$  and  $\text{coin}_2$ ) with the same serial number  $\mathbf{y}_S := \mathbf{y}_{S,1} = \mathbf{y}_{S,2}$ .

We distinguish two cases:

- Case 4.a:  $H_0(\text{info}_1, \mathcal{M}_1) = H_0(\text{info}_2, \mathcal{M}_2)$
- Case 4.b:  $H_0(\text{info}_1, \mathcal{M}_1) \neq H_0(\text{info}_2, \mathcal{M}_2)$

The first case is straightforward and implies a collision on  $H_0$  as the success conditions for  $\mathcal{A}$  require  $(\text{info}_1, \mathcal{M}_1) \neq (\text{info}_2, \mathcal{M}_2)$ . We will therefore only consider the second case.

The soundness of arguments  $\pi_{K,1}$  and  $\pi_{K,2}$  (which are included in  $\text{coin}_1$  and  $\text{coin}_2$ , respectively) ensures that, for each  $j \in \{1, 2\}$ , there exists  $\mathbf{y}_{1,j} \in \mathbb{Z}_p^m$  and  $PK_j$  such that

$$\forall j \in \{1, 2\} : \begin{cases} \mathbf{y}_S = (\mathbf{y}_{S,1}, \mathbf{y}_{S,2}) = (F_{\text{LTF}}(\text{bin}_p(\mathbf{y}_{1,j})), PK_j + H_{\text{UH}}(\text{bin}_p(\mathbf{y}_{1,j}))) \\ \mathbf{y}_{T,j} = PK_j + \text{FRD}(H_0(PK_{\mathcal{M}_j}, \text{info}_j)) \cdot H'_{\text{UH}}(\text{bin}_p(\mathbf{y}_{1,j})) \end{cases}$$

Since  $F_{\text{LTF}}$  is injective, the consistency of  $\mathbf{y}_{S,1}$  amongst the two coins implies that  $\mathbf{y}_1 := \mathbf{y}_{1,0} = \mathbf{y}_{1,1}$ . In turn, since  $\text{coin}_1$  and  $\text{coin}_2$  also share the same  $\mathbf{y}_{S,2}$ , this implies that  $PK := PK_0 = PK_1$ . Combining these observations, we find that

$$\begin{aligned} \mathbf{y}_{T,1} &= PK + \text{FRD}(H_0(PK_{\mathcal{M}_1}, \text{info}_1)) \cdot H'_{\text{UH}}(\text{bin}_p(\mathbf{y}_1)) \\ \mathbf{y}_{T,2} &= PK + \text{FRD}(H_0(PK_{\mathcal{M}_2}, \text{info}_2)) \cdot H'_{\text{UH}}(\text{bin}_p(\mathbf{y}_1)) \end{aligned}$$

for  $(\mathcal{M}_1, \text{info}_1) \neq (\mathcal{M}_2, \text{info}_2)$ . In our case, we have  $R_1 := H_0(\mathcal{M}_1, \text{info}_1) \neq H_0(\mathcal{M}_2, \text{info}_2) =: R_2$ , meaning that  $(\text{FRD}(R_1) - \text{FRD}(R_2))$  is invertible over  $\mathbb{Z}_p$ . At this stage, it can be easily verified that  $\text{Identify}(PK_{\mathcal{B}}, \text{coin}_1, \text{coin}_2)$  outputs  $PK$ .

The soundness of the proof implies that the first coin contains a proof of knowledge of  $\mathbf{v}^*, \mathbf{r}_0^*, \mathbf{r}_1^*, \mathbf{e}_u^*$  and  $\tau^*, \mathbf{k}^*$  such that

$$\begin{aligned} \mathbf{A}_{\tau^*} \cdot \mathbf{v}^* &= \mathbf{u} + \mathbf{D} \cdot \text{vdec}_{n, q_s-1}(\mathbf{w}^*) \\ \mathbf{w}^* &= \mathbf{c}_{PK} + \mathbf{D}_0'' \cdot \mathbf{r}_1^* \\ \mathbf{c}_{PK} &= \mathbf{D}_0' \cdot \mathbf{r}_0^* + \mathbf{D}_1 \cdot \mathbf{e}_u^* + \mathbf{D}_2 \cdot \text{vdec}_{m, q-1}(\mathbf{k}^*) \end{aligned} \tag{15}$$

We know that there is some  $i^* \in \{1, \dots, Q_w\}$  such that  $\tau^*, \mathbf{w}^* = (\tau^{(i^*)}, \mathbf{w}^{(i^*)})$  since, otherwise, we would be considering either Case 1 or 2. Additionally,  $\mathcal{A}$  succeeds only if  $PK$  does not belong to  $\mathbb{T}$ . This implies that, for all  $i \in [Q_w]$ , we have  $PK = \mathbf{F} \cdot \mathbf{e}_u^* \neq \mathbf{F} \cdot \mathbf{e}_u^{(i)}$  and thus  $\mathbf{e}_u^* \neq \mathbf{e}_u^{(i^*)}$ . We can then reproduce the proof of Case 3. In particular, re-using the strategy and notation from Case 3, we can conclude that

$$\mathbf{v}_{\text{SIS}} := \mathbf{Q} \cdot (\mathbf{r}'^\top \mid \mathbf{e}_u'^\top \mid \mathbf{k}'^\top)^\top$$

is a solution to the SIS instance  $\bar{\mathbf{A}}$  as long as it is non-zero. Similarly to before, we can argue that  $\mathbf{v}_{\text{SIS}}$  is non-zero due to the fact that  $\mathbf{e}_u' := \mathbf{e}_u^* - \mathbf{e}_u^{(i^*)} \in \{-1, 0, 1\}^m \setminus \mathbf{0}^m$  and that  $\mathbf{Q}_1$  retains high entropy from the perspective of  $\mathcal{A}$ . Therefore, any adversary corresponding to Case 4 can be used to solve SIS with the same norm-bound as in Case 3.  $\square$

### C.3 Proof of the Clearing Property

Recall that the `VerifyDeposit` algorithm run on  $(PK_B, PK_M, \text{coin}, \mu)$  essentially verifies that  $\mu$  is a valid signature under  $PK_M$  on a valid coin `coin`. The success conditions defined by the clearing experiment imply then that the adversary managed to output a valid signature  $\mu$  on a coin `coin` for the public key of an honest merchant. Since `coin` contains a pair  $(\text{info}, PK_M)$  that is necessarily different for each coin deposited through a  $\mathcal{Q}_{\text{Deposit}}$ -query,  $\mu$  is a valid forgery against the signature scheme  $\Sigma$ . The formal security proof is provided below.

*Proof.* Let  $\mathcal{A}$  be an adversary succeeding in the clearing experiment with probability  $\epsilon$ . The challenger  $\mathcal{C}$  makes a guess on the index  $i^*$  of the honest merchant targeted by  $\mathcal{A}$  and runs the EUF-CMA experiment for  $\Sigma$  as an adversary. Let  $PK^*$  be the public key it received from this experiment.  $\mathcal{C}$  thus answers oracle queries as follows:

- $\mathcal{Q}_{\text{GetKey}}(i)$ :  $\mathcal{C}$  proceeds as usual if  $i \neq i^*$  and returns  $PK^*$  otherwise.
- $\mathcal{Q}_{\text{Receive}}(PK_B, i, j)$ : all the computations performed by the merchant during a spending only require public elements, so  $\mathcal{C}$  is perfectly able to handle such oracle queries, even when  $i^* = i$ .
- $\mathcal{Q}_{\text{Deposit}}(PK_B, i, j)$ :  $\mathcal{C}$  proceeds as usual when  $i \neq i^*$  (it knows the corresponding secret key  $SK_{\mathcal{M}_i}$ ) and uses its signing oracle from the EUF-CMA experiment otherwise. In all cases, it can perfectly answer such queries.

At the end of the game,  $\mathcal{A}$  outputs a tuple  $(PK_M, \text{coin}, \mu)$ . The challenger aborts if  $PK_M \neq PK^*$ . Else,  $\mu$  is a valid forgery under  $PK^*$  because `coin` has never been submitted to the signing oracle, as we have explained. Since  $\mathcal{C}$  does not abort with probability at least  $\frac{1}{|\mathcal{HM}|}$ , this means that  $\mathcal{A}$  can be converted into an adversary against the EUF-CMA security of  $\Sigma$  succeeding with probability at least  $\frac{\epsilon}{|\mathcal{HM}|}$ . The probability  $\epsilon$  is therefore necessarily negligible.  $\square$