

Lattice-Based Sanitizable Signature Schemes: Chameleon Hash Functions and More

Sebastian Clermont¹, Samed Düzlü², Christian Janson¹, Laurens Porzenheim³,
and Patrick Struck⁴

¹ Technische Universität Darmstadt, Germany

`sebastian.clermont@tu-darmstadt.de` `christian.janson@cryptoplexity.de`

² Universität Regensburg, Germany

`samed.duzlu@ur.de`

³ Universität Paderborn, Germany

`laurens.porzenheim@uni-paderborn.de`

⁴ Universität Konstanz, Germany

`patrick.struck@uni-konstanz.de`

Abstract. Sanitizable Signature Schemes (SSS) enable a designated party, the sanitizer, to modify predefined parts of a signed message without invalidating the signature, making them useful for applications like pseudonymization and redaction. Since their introduction by Ateniese et al. (ESORICS’05), several classical SSS constructions have been proposed, but none have been instantiated from quantum-resistant assumptions. In this work, we develop the first quantum-secure sanitizable signature schemes based on lattice assumptions. Our primary focus is on SSS constructions that rely on chameleon hash functions (CHF), a key component for enabling the controlled modification of messages. While lattice-based CHF exist, they do not meet the required security guarantees for SSS, becoming insecure under adversarial access to an adapt oracle. To address this, we construct a novel lattice-based CHF that achieves collision resistance even in such settings, called full collision resistance. However, our CHF lacks the *uniqueness* property, a limitation we show to be inherent in lattice-based CHF. As a result, our SSS constructions initially fall short of achieving the critical security property of accountability. To overcome this, we apply a transformation based on verifiable ring signatures (VRS), for which we present the first lattice-based instantiation. Additionally, we provide a comprehensive analysis of existing classical SSS constructions, explore their potential for post-quantum instantiations, and present new attacks on previously assumed secure SSS schemes. Our work closes the gap in constructing quantum-secure SSS and lays the groundwork for further research into advanced cryptographic primitives based on lattice assumptions.

* This work was funded by the German Federal Ministry of Education and Research (BMBF) under reference 16KISQ074, the German Research Foundation (DFG) – SFB 1119 – 236615297, the Ministry of Culture and Science of the State of North Rhine-Westphalia, and the Hector Foundation II.

An extended abstract of this work was first published in PQCrypto 2025, LNCS 15577, pp. 278–311, 2025. https://doi.org/10.1007/978-3-031-86599-2_10

1 Introduction

In modern cryptography, the digital signature scheme is an essential primitive that enjoys ubiquitous usage. However, depending on the use case, one might need more advanced signature schemes which come with additional functionality. These advanced variants usually build upon their base versions, offering strong security guarantees for their specialized functionality and alleviating the need for ad-hoc constructions. Advanced signature-based primitives include sanitizable signatures, ring signatures, group signatures, threshold signatures, and more.

The main focus of this paper are sanitizable signature schemes (SSS). They can be applied, for instance, to pseudonymize sensitive data, like patient names in medical records, while maintaining the authenticity. Sanitizable signature schemes allow the signer of a message to partially delegate signature rights to a semi-trusted third party, the sanitizer. This sanitizer may modify predefined parts of the message without invalidating the signature. In a sanitizable signature scheme, the message is divided into blocks, among which the sanitizer can modify exactly the *admissible blocks*. In contrast, all other blocks, the *immutable blocks*, cannot be modified by the sanitizer without invalidating the signature. The signer defines which blocks are admissible during the signing operation. Another use-case of sanitizable signature schemes are classified documents. SSS constructions allow the signing of such documents (by the signer) such that a censor (the sanitizer) can redact certain parts, depending on which parts should not be made public.

Sanitizable signature schemes were first introduced in Ateniese et al.’s seminal work [3]. Since then, many more works on sanitizable signature schemes have been published, covering new constructions focusing on different, sometimes novel, security notions. Current literature contains only two constructions that exclusively use basic primitives [13, 22], one construction is based on group signatures [12], two are based on ring signatures [15, 30], one is based on signatures with rerandomizable keys [23], while all others, [11, 3, 30, 16, 6], make use of so-called *chameleon hash functions* (CHF) [28].

Chameleon hash functions are families of hash functions with public and secret keys. In contrast to regular hash functions, their input consists of not just the message, but also an auxiliary random value. The public key of a chameleon hash function defines a collision-resistant function mapping pairs (μ, r) to some hash value h , where μ is the message and r is a randomly chosen auxiliary value. With the corresponding secret key, however, one can then *adapt* the message to a new μ' , by finding a suitable auxiliary value r' such that the pair (μ', r') evaluates to the original digest h , i.e., find a collision. Besides weak collision resistance, where the adversary only receives the public key, chameleon hash functions can support stronger forms of collision resistance, where an adapt oracle is available to the adversary. Apart from SSS, chameleon hash functions are a widely used building block in other cryptographic schemes. For example, they can transform a non-adaptively secure signature scheme to an existentially unforgeable one [41].

The gist of constructing a SSS from a CHF is a simple hash-then-adapt approach. In this approach, the admissible blocks are hashed using the CHF

with a public key generated by the sanitizer and then signed by the signer using a regular signature scheme. The sanitizer can thus modify the message without altering the hash value, retaining the original signature’s validity. More advanced constructions of SSS have been developed over the years, some of which attain additional security properties and explore alternative approaches.

Given the general transition to post-quantum cryptography, analyzing which advanced primitives can also be constructed from quantum-hard assumptions is crucial. While there are positive results for various advanced primitives, sanitizable signatures from quantum-hard assumptions have yet to be considered.

1.1 Our Contribution

We close the aforementioned gap by developing post-quantum secure sanitizable signatures from lattices. Our main focus is on constructions based on chameleon hash functions. While lattice-based CHF’s exist [18, 31], we observe that they cannot be used to instantiate existing SSS constructions. The reason is that they do not achieve the necessary security guarantees. Thus, we develop a new chameleon hash function inspired by the one given in [31]. Our construction achieves collision resistance even against adversaries with access to an adapt oracle that provides collision to the adversary—for SSS constructions this is necessary since sanitized signatures are effectively that. This is our first contribution as we believe this strongly secure CHF can be of independent interest.

However, this construction is not sufficient to already instantiate the SSS constructions. The reason is that it lacks another property called *uniqueness*. When a chameleon hash has uniqueness, it is hard to come up with two different auxiliary values r, r' that both map some message μ to the same digest h . Unfortunately, the lack of the uniqueness property seems not to be a flaw of our specific construction but a more general problem of lattice-based CHF’s. The absence of this property implies that the SSS constructions lack crucial security properties like unforgeability and accountability. However, there is a generic transform [15] which turns such a SSS (which they call weakly-secure) into a secure SSS. This transform requires another advanced signature, called verifiable ring signature (VRS), i.e., a variant of a ring signature that allows to prove or deny ownership of a signature at a later point. Thus, as our second contribution, we develop a lattice-based VRS scheme. To do that, we first design a generic construction of a VRS scheme using a so-called linking indistinguishable tag and a non-interactive zero-knowledge proof. Then, we instantiate that generic construction with lattice-based constructions.

With these contributions (illustrated in Fig. 1), we close the aforementioned gap by giving multiple post-quantum sanitizable signatures from lattices.

As additional contributions, we give an attack against an SSS construction [11], breaking the accountability, i.e., an attack that allows the signer to blame the sanitizer for messages of its own choosing. Furthermore, we revisit the SSS constructions that do not rely on chameleon hash functions. We show that for some of them, the necessary building blocks are readily available in the lattice-based literature, meaning they can easily be instantiated as well. Finally, we

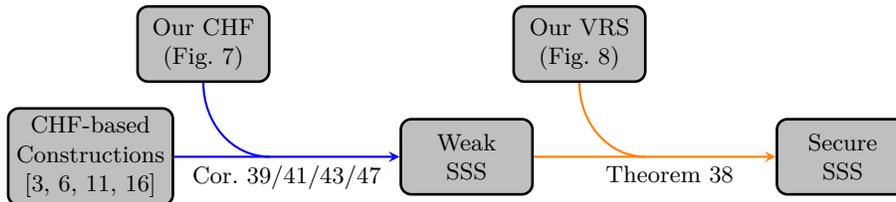


Fig. 1: Overview of our main results. Instantiating the SSS constructions given in [3, 6, 11, 16] with our new lattice-based chameleon hash function results in weak sanitizable signatures. The transformation from [15]—using these weak sanitizable signatures and our new verifiable ring signature scheme—then yields secure sanitizable signatures.

show additional implications between the security notions, which also show some additional security properties for some SSS construction—in particular, security properties that were developed after some of the older SSS constructions.

1.2 Related Work

Current literature provides a variety of constructions of sanitizable signature schemes. Two of them, [13, 22], require standard assumptions only. Some of them, [12, 15, 23] provide constructions based on other advanced signatures. The work [30] by Lai et al. contains two constructions: one based on accountable ring signatures and another one based on chameleon hash functions. Finally, many more constructions based on chameleon hash functions exist, [3, 11, 6, 16, 15]. Besides the different primitives they use, the various works have distinct goals regarding the security notions they achieve. The achieved security properties can be seen in Table 1. Another construction is due to Klonowski and Lauks [27]. Later, however, Canard and Jambert [17] showed it to lack accountability.

Constructions of chameleon hash functions based on lattices are very rare. Cash et al. [18] came forward with the first construction based on a trapdoor in the Ajtai hash function. This chameleon hash function achieves merely a weak form of collision-resistance. Only recently, a new lattice-based construction was brought forward by Li and Liu [31]. They use one-time tags, making the hash function not applicable to known generic constructions of SSS. However, we use and improve their ideas to construct our novel chameleon hash function.

Due to the missing uniqueness of our CHF, we apply the transform introduced in [15] to increase the security of the instantiations from weak SSS to strong SSS. The transform relies on verifiable ring signatures (VRS). In [15], it is stated that other ring signatures with extended functionality, such as linkable ring signatures, imply VRS. However, we were not able to verify that claim.

[39] present (among others) a notion of ring signatures equivalent to VRS, called claimable and repudiable ring signatures. They also show how to construct such a ring signature. However that requires verifiable random functions, which can be difficult to construct (e.g., [21]) or use (non-interactive) proofs in their

Table 1: Overview of our results. The table provides an overview of the SSS constructions that use chameleon hash functions: ACMT05 [3], BCD⁺17 [6], CDK⁺17 [16], and BFF⁺09 [11]. The entries with the color **blue** are the results of the instantiations of the respective constructions with our chameleon hash function defined in Section 3.1. The entries with the color **orange** are the results after using our verifiable ring signature defined in Section 3.2 to apply a transform as described in [15] to ensure unforgeability as well as signer and sanitizer accountability. The various symbols represent: the instantiation achieves the notion (✓), the construction does not achieve the notion (✗), the instantiation does not achieve a notion that the construction otherwise could achieve (◆), the instantiation has not been analyzed with respect to the notion (?), and there is no lattice-based instantiation due to missing building blocks (○). Note that some notions contradict each other, consequently it is impossible to come up with a construction achieving all notions (cf. Appendix B.2).

	Chameleon Hash Function based Sanitizable Signatures				
	ACMT05	LZCS16–1	BCD ⁺ 17	CDK ⁺ 17	BFF ⁺ 09
	[3]	[30]	[6]	[16]	[11]
Unforgeability	✓	○ ✓	◆ ✓	◆ ✓	✓
Signer Accountability	✓	○ ✓	◆ ✓	✓	✗ ✓
Sanitizer Accountability	? ✓	○ ✓	◆ ✓	◆ ✓	✓
NIPA	✗	✗	✗	✗	✗
Immutability	✓	○	✓	✓	✓
Transparency	✓	○	✓	✓	✓
Privacy	✓	○	✓	✓	✓
Unlinkability	✗	✗	✗	✗	✗
Invisibility	?	?	✓	✓	?

construction. Furthermore, the construction of [39] requires a proof over statements including the VRF, i.e., proofs over proofs, which is why we opted for a more direct approach for constructing a VRS.

2 Background

We define some notation. By $x \leftarrow y$ we define deterministic assignment of y to x . If S is some set, we denote by $x \leftarrow \$ S$ sampling a uniform value x from S . Overloading notation, if A is some probabilistic polynomial-time (*ppt*) or quantum polynomial time (*qpt*) algorithm, we denote by $y \leftarrow \$ A(x)$ assigning the random output of $A(x)$ to y . The statement of the kind “for all $y \leftarrow \$ A(x)$ ” means for all y that can be output by the randomized algorithm A on input x . We denote column vectors by lowercase bold-face letters \mathbf{a} , while row-vectors are denoted as \mathbf{a}^t . Matrices \mathbf{A} are denoted in uppercase. By $\|\mathbf{x}\|$ we denote the Euclidian norm of \mathbf{x} . If we want to query the random oracle on some input x , we

denote it by $y \leftarrow \mathcal{RO}(x)$. We will often assume security parameters and public parameters as implicit inputs to algorithms and omit them for readability.

The security analysis under the presence of quantum adversary takes place in the Q1 security model, allowing local quantum computer access to adversaries, but only classical access to oracles including the random oracle. Signature schemes as well as public key encryption schemes follow the commonly established syntax and notation, but for the sake of completeness are defined in Appendix A. We will formalize less ubiquitous primitives in this section.

2.1 Chameleon Hash Function

First, we introduce chameleon hash functions, which we will later build from lattices as one of our core contributions in this paper.

Definition 1 (Chameleon Hash Function). *A chameleon hash function CH consists of four ppt algorithms $\text{CH} = (\text{CKGen}, \text{CHash}, \text{CHashCheck}, \text{CAdapt})$.*

$(\text{pk}_{\text{ch}}, \text{sk}_{\text{ch}}) \leftarrow_{\$} \text{CKGen}(1^\lambda)$: *The key generation algorithm CKGen takes as input the security parameter 1^λ and outputs a private-public key pair.*

$(h, r) \leftarrow_{\$} \text{CHash}(\text{pk}_{\text{ch}}, \mu)$: *The algorithm CHash takes as input a public key pk_{ch} and a message μ . It outputs a hash h and randomness r , under the given public key pk_{ch} .*

$b \leftarrow \text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu, r)$: *Given a public key pk_{ch} , a hash value h , a message μ , and randomness r , the deterministic algorithm CHashCheck outputs a boolean value $b \in \{0, 1\}$, indicating whether the hash h is valid under pk_{ch} or not.*

$r' \leftarrow_{\$} \text{CAdapt}(\text{sk}_{\text{ch}}, h, \mu, r, \mu')$: *The algorithm CAdapt takes as input the secret key sk_{ch} , a hash h , a message μ , randomness r , and a new message μ' . It then outputs new randomness r' , such that $1 \leftarrow \text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu', r')$*

Definition 2 (Correctness of a Chameleon Hash Function). *A chameleon hash function CH is correct, iff*

$$\begin{aligned} & \forall \lambda, \forall (\text{pk}_{\text{ch}}, \text{sk}_{\text{ch}}) \leftarrow_{\$} \text{CKGen}(1^\lambda), \forall \mu, \mu' \in \mathcal{M}, \\ & \forall (h, r) \leftarrow_{\$} \text{CHash}(\text{pk}_{\text{ch}}, \mu), \forall r' \leftarrow_{\$} \text{CAdapt}(\text{sk}_{\text{ch}}, h, \mu, r, \mu') : \\ & \text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu, r) = \text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu', r') = 1 \end{aligned}$$

For security we require that the chameleon hash function is collision-resistant, even in the presence of a collision oracle. Furthermore, we require that it hard to distinguish whether a hash randomness pair (h, r) was output by CH or CAdapt. We model the security games in Figs. 2,3.

Definition 3. *We say that a chameleon hash function CH is fully collision-resistant, if there exists a negligible function $\text{negl}(\lambda)$ such that for all ppt adversaries \mathcal{A} we have that $\Pr[\text{f-CR}_{\Pi, \mathcal{A}}(1^\lambda) = 1] \leq \text{negl}(\lambda)$.*

f-CR _{CH, \mathcal{A}} (1^λ)	Adapt \mathcal{O} (sk _{ch} , \cdot , \cdot , \cdot , \cdot) with (h, μ, r, μ')
$\mathcal{Q} \leftarrow \emptyset$	if CHashCheck(pk _{ch} , h, μ, r) = 0
(pk _{ch} , sk _{ch}) $\leftarrow_{\$}$ CKGen(1^λ)	return \perp
(h, μ, r, μ', r') $\leftarrow_{\$}$ $\mathcal{A}^{\text{Adapt}\mathcal{O}}$ (pk _{ch})	$r' \leftarrow_{\$}$ CAdapt(sk _{ch} , h, μ, r, μ')
return (CHashCheck(pk _{ch} , h, μ, r) = 1	$\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(h, \mu), (h, \mu')\}$
\wedge CHashCheck(pk _{ch} , h, μ', r') = 1	return r'
$\wedge \mu \neq \mu' \wedge (h, \mu) \notin \mathcal{Q}$)	

Fig. 2: f-CR security game.

CHIndist _{II, \mathcal{A}, b} (λ)	HashOrAdapt \mathcal{O}_b (sk _{ch} , \cdot , \cdot) with (μ, μ')
(pk _{ch} , sk _{ch}) $\leftarrow_{\$}$ CKGen(1^λ)	(h, r) $\leftarrow_{\$}$ CHash(pk _{ch} , μ)
$b' \leftarrow_{\$}$ $\mathcal{A}^{\text{HashOrAdapt}\mathcal{O}_b, \text{Adapt}\mathcal{O}}$ (pk _{ch})	(h', r') $\leftarrow_{\$}$ CHash(pk _{ch} , μ')
return b'	$r'' \leftarrow_{\$}$ CAdapt(sk _{ch} , h', μ', r', μ)
	if $b = 0$ return (h, r)
	if $b = 1$ return (h', r'')

Fig. 3: CH Indistinguishability game.

There exist other, weaker notions of collision-resistance of a chameleon hash function, called standard, enhanced, and weak collision-resistance. For definitions and comparisons, see [20]. However, we only concern ourselves with the strongest version, which implies the others.

Definition 4. We say that a chameleon hash function CH is indistinguishable, if there exists a negligible function $\text{negl}(\lambda)$, such that for all qpt adversaries \mathcal{A} we have that $|\Pr[\text{CHIndist}_{\text{CH}, \mathcal{A}, 0}(\lambda) = 1] - \Pr[\text{CHIndist}_{\text{CH}, \mathcal{A}, 1}(\lambda) = 1]| \leq \text{negl}(\lambda)$.

2.2 Lattices

Next, we cover various relevant definitions and results for lattices.

Definition 5. An n -dimensional lattice Λ is a discrete, additive subgroup of \mathbb{R}^n .

Definition 6. In the short integer solution (SIS) problem $\text{SIS}_{n, m, q, \beta}$ a qpt adversary \mathcal{A} is given a uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and is asked to compute a $\mathbf{x} \in \mathbb{Z}^m$ such that $\mathbf{A}\mathbf{x} = \mathbf{0}$ and $0 < \|\mathbf{x}\| \leq \beta$.

We define the advantage of a qpt adversary \mathcal{A} against $\text{SIS}_{n, m, q, \beta}$ as

$$\text{SIS}_{n, m, q, \beta}(\mathcal{A}) := \Pr[\mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}, 0 < \|\mathbf{x}\| \leq \beta : \mathbf{A} \leftarrow_{\$} \mathbb{Z}_q^{n \times m}, \mathbf{x} \leftarrow_{\$} \mathcal{A}(\mathbf{A})].$$

Definition 7. Define the multidimensional Gaussian function with Gaussian parameter s and center \mathbf{c} as $\rho_{s, \mathbf{c}}(\mathbf{x}) = \exp(-\|\mathbf{x} - \mathbf{c}\|/s)$. Then, we define the discrete Gaussian distribution over lattice Λ with Gaussian parameter s and center \mathbf{c} as $D_{\Lambda, s, \mathbf{c}}(\mathbf{x}) = \frac{\rho_{s, \mathbf{c}}(\mathbf{x})}{\rho_{s, \mathbf{c}}(\Lambda)}$.

To construct our chameleon hash, we need a lattice-based trapdoor. We use so-called G-trapdoors.

Definition 8 ([36]). Let $\mathbf{G} \in \mathbb{Z}_q^{n \times w}$ be a so-called gadget matrix, for which SIS is easy. Define a G-trapdoor for a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ to be some matrix $\mathbf{R} \in \mathbb{Z}_q^{m \times w}$ such that $\mathbf{AR} = \mathbf{G} \pmod q$.

We denote the quality of a G-trapdoor by its spectral norm $s_1(\mathbf{R})$. Note that if one has a trapdoor \mathbf{R} for \mathbf{A} , then $[\mathbf{R}^t, \mathbf{0}_{w \times k}]^t \in \mathbb{Z}_q^{m+k \times w}$ is a trapdoor for $[\mathbf{A}|\mathbf{B}]$ for any $\mathbf{B} \in \mathbb{Z}_q^{n \times k}$, where $\mathbf{0}_{a \times b}$ is the all-zero matrix of dimension $a \times b$.

With a G-trapdoor, we can invert matrix-vector multiplication of the trapdoored matrix \mathbf{A} , where the inverted preimage has a certain (conditioned) Gaussian distribution.

Theorem 9 ([36]). There exists a ppt algorithm $\text{TrapGen}(1^n, 1^m, q)$ that, given any integers $n \geq 1$, $q \geq 2$, and sufficiently large $m = \mathcal{O}(n \log q)$, outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a G-trapdoor \mathbf{R} for \mathbf{A} , such that the distribution of \mathbf{A} is negligibly far from uniform. Moreover, there exists a ppt algorithm $\text{PreSample}(\mathbf{A}, \mathbf{R}, \mathbf{u}, s)$ that, given some $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, some G-trapdoor $\mathbf{R} \in \mathbb{Z}_q^{m \times w}$ for \mathbf{A} with $s_1(\mathbf{R}) \in \mathcal{O}(\sqrt{n \log q})$, some $\mathbf{u} \in \mathbb{Z}_q^n$ and large enough $s = \mathcal{O}(\sqrt{n \log q})$, samples from a distribution that is within negligible statistical distance from $D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}), s \cdot \omega(\sqrt{\log n})}$. Furthermore, the distribution of sampling $\mathbf{x} \leftarrow \$ D_{\mathbb{Z}_q^m, s}$, setting $\mathbf{y} = \mathbf{Ax}$, and outputting (\mathbf{x}, \mathbf{y}) is statistically indistinguishable from the distribution of choosing $\mathbf{y}' \leftarrow \$ \mathbb{Z}_q^n$, computing $\mathbf{x}' \leftarrow \$ \text{PreSample}(\mathbf{A}, \mathbf{R}, \mathbf{y}, s)$, and outputting $(\mathbf{x}', \mathbf{y}')$.

If \mathbf{R} is a trapdoor for \mathbf{A} , we sometimes write $\text{PreSample}([\mathbf{A} \ \mathbf{B}], \mathbf{R}, \mathbf{u}, s)$ instead of $\text{PreSample}([\mathbf{A} \ \mathbf{B}], [\mathbf{R}^t \ \mathbf{0}^t]^t, \mathbf{u}, s)$, since \mathbf{R} can be extended to a trapdoor for $[\mathbf{A} \ \mathbf{B}]$ for any \mathbf{B} , as noted above.

2.3 Trapdoor Commitment Scheme

A trapdoor commitment scheme is a commitment scheme that is (computationally) binding and hiding, but if someone is in possession of a trapdoor to the public parameters, then one can break binding. We model that by introducing an additional ppt algorithm Equiv that equivocates a commitment c to any message μ' with the help of the trapdoor.

Definition 10. A trapdoor commitment scheme consists of four ppt algorithms (TdGen , Com , ComCheck , Equiv).

- $(\text{pp}, \text{td}) \leftarrow \$ \text{TdGen}(1^\lambda)$: On input a security parameter 1^λ , the algorithm outputs public parameters pp and a trapdoor td .
- $(c, d) \leftarrow \$ \text{Com}(\text{pp}, \mu)$: On input some pp and a message μ , the algorithm outputs a commitment c and an opening value d .
- $b \leftarrow \text{ComCheck}(\text{pp}, \mu, c, d)$: On input some pp , a message μ , a commitment c and an opening value d , the algorithm outputs a bit b .

$\text{Bind}_{\Pi, \mathcal{A}}(\lambda)$	$\text{DEE}_{\Pi, \mathcal{A}, b}(\lambda)$
$(\text{pp}, \text{td}) \leftarrow_{\$} \text{TdGen}(1^\lambda)$	$(\text{pp}, \text{td}) \leftarrow_{\$} \text{TdGen}(1^\lambda)$
$(c, \mu, d, \mu', d') \leftarrow_{\$} \mathcal{A}(\text{pp})$	$\mu \leftarrow_{\$} \mathcal{A}(\text{pp})$
if $\text{ComCheck}(\text{pp}, \mu, c, d) = 1$	$(c_0, d_0) \leftarrow_{\$} \text{Com}(\text{pp}, \mu)$
$\wedge \text{ComCheck}(\text{pp}, \mu', c, d') = 1$	$c_1 \leftarrow_{\$} \text{Com}(\text{pp}, 0)$
$\wedge \mu \neq \mu'$	$d_1 \leftarrow_{\$} \text{Equiv}(\text{pp}, \text{td}, c_1, \mu)$
return 1	$b' \leftarrow_{\$} \mathcal{A}(c_b, d_b)$
return 0	return b'

Fig. 4: Trapdoor commitment security games. DEE is distributional equivalence of equivocation.

$d' \leftarrow_{\$} \text{Equiv}(\text{pp}, \text{td}, c, \mu')$: On input some pp , a trapdoor td , a commitment c and a message μ' , the algorithm outputs some opening value d' .

A trapdoor commitment scheme is correct, iff

$$\begin{aligned} & \forall \lambda, \forall (\text{pp}, \text{td}) \leftarrow_{\$} \text{TdGen}(1^\lambda), \forall \mu, \mu' \in \mathcal{M}, \\ & \forall (c, d) \leftarrow_{\$} \text{Com}(\text{pp}, \mu), \forall d' \leftarrow_{\$} \text{Equiv}(\text{pp}, \text{td}, c, \mu'), \\ & \text{ComCheck}(\text{pp}, \mu, c, d) = \text{ComCheck}(\text{pp}, \mu', c, d') = 1. \end{aligned}$$

We need two security notions for a trapdoor commitment scheme. One is the typical binding, which states that an adversary cannot open a commitment to two different messages. The other security notion is called *distributional equivalence of equivocation*. This requires that, for fixed pp and μ , the joint distribution of a commitment and an opening value is computationally indistinguishable when either committing to the message or committing to 0 and equivocating to the message. We model these security games in Fig. 4.

Definition 11. We say that a commitment scheme Π is computationally binding, if there exists a negligible function negl such that for all qpt adversaries \mathcal{A} it holds that $\Pr[\text{Bind}_{\Pi, \mathcal{A}}(\lambda)] \leq \text{negl}(\lambda)$.

Definition 12. We say that a commitment scheme Π has distributional equivalence of equivocation (DEE), if there exists a negligible function negl such that for all qpt adversaries \mathcal{A} it holds that

$$|\Pr[\text{DEE}_{\Pi, \mathcal{A}, 0}(\lambda) = 1] - \Pr[\text{DEE}_{\Pi, \mathcal{A}, 1}(\lambda) = 1]| \leq \text{negl}(\lambda).$$

It is possible to adapt known lattice-based commitment schemes such as Ajtai [1] or BDLOP [5] to be trapdoor commitment schemes by replacing their uniform generation of matrices by TrapGen . Then, one uses PreSample to construct Equiv , and due to the properties mentioned in Theorem 9 we can show DEE.

Note that we did not define the typical notion of a hiding commitment scheme as we do not need it explicitly for our construction. However, hiding is implied

by DEE, since in the DEE security game c_1 is generated by committing to 0. If the commitment scheme was not hiding, the adversary could get information about b by c_b alone. Note that binding trapdoor commitment schemes with DEE are strongly related to collision-resistant preimage-sampleable functions [24].

2.4 Linking Indistinguishable Tag

To later construct the VRS we need so-called *linking indistinguishable tags* (LIT) [7, 9]. This is a secret key tagging scheme like a MAC, except that there exists an additional linking algorithm. With this, it is possible to detect if the same person tagged the same message twice. Additionally, there exists a function f that computes a public key from a secret key. While this is not used for the LIT itself, it is useful when using LITs to construct other schemes. For security we require that it is hard to trick the linking algorithm as well as unforgeability similar to a MAC, but also that apart from linking, it is hard to decide which person created a tag. Also it should be hard to compute a secret key from a public key and tags.

Definition 13. *A linking indistinguishable tag scheme consists of a function f and the following ppt algorithms:*

$sk \leftarrow \text{KGen}(1^\lambda)$: On input a security parameter 1^λ , it outputs a secret sk .
 $t \leftarrow \text{Tag}(sk, \mu)$: On input a secret key sk and a message μ , it outputs a tag t .
 $b \leftarrow \text{Vrfy}(sk, \mu, t)$: On input a secret key sk , a message μ and a tag t , it outputs a bit b .
 $b \leftarrow \text{Link}(\mu, t_0, t_1)$: On input a message μ and two tags t_0, t_1 , it outputs a bit b .

We require that a LIT is correct. This is the case if

$$\forall \lambda, \forall sk \leftarrow \text{KGen}(1^\lambda), \forall \mu, \forall t_0, t_1 \leftarrow \text{Tag}(sk, \mu) : \\ \text{Vrfy}(sk, \mu, t_0) = \text{Link}(\mu, t_0, t_1) = 1.$$

For the security model, we use the one from [9], with one change improving the security. We use the same games for tag-indistinguishability, non-invertability, linkability, and unforgeability. However, as can be seen in Fig. 5, when defining the oracles, we no longer return a previously computed tag if a message has been previously queried. Instead, the adversary gets a fresh tag every time it queries the oracle.

Definition 14. *A LIT Π has tag-indistinguishability, if there exists a negligible function negl such that for all ppt adversaries \mathcal{A} it holds that*

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{LITAnon}}(\lambda) := \left| \Pr[\text{Anon}_{\Pi, \mathcal{A}, 0}^{\text{LIT}}(\lambda) = 1] - \Pr[\text{Anon}_{\Pi, \mathcal{A}, 1}^{\text{LIT}}(\lambda) = 1] \right| \leq \text{negl}(\lambda).$$

Definition 15. *A LIT Π has linkability if there exists a negligible function negl such that for all ppt adversaries \mathcal{A} it holds that*

$$\Pr[\text{Linkable}_{\Pi, \mathcal{A}}^{\text{LIT}}(\lambda) = 1] \leq \text{negl}(\lambda).$$

$\text{Anon}_{II,\mathcal{A},b}^{\text{LIT}}(\lambda)$ <hr/> $\mathcal{Q} \leftarrow \emptyset$ $sk_0, sk_1 \leftarrow \mathcal{KGen}(1^\lambda)$ $pk_i \leftarrow f(sk_i), i \in \{0, 1\}$ $\mu^* \leftarrow \mathcal{A}^{\text{Tag}^\mathcal{O}}(pk_0, pk_1)$ $t^* \leftarrow \text{Tag}(sk_b, \mu^*)$ $b' \leftarrow \mathcal{A}^{\text{Tag}^\mathcal{O}}(t^*)$ if $(\mu^*, \cdot) \in \mathcal{Q}$ return 0 return b'	$\text{Tag}^\mathcal{O}(\{sk_j\}_j, \cdot, \cdot) \text{ with } (i, \mu)$ <hr/> $t \leftarrow \text{Tag}(sk_i, \mu)$ $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\mu, t)\}$ return t
$\text{Linkable}_{II,\mathcal{A}}^{\text{LIT}}(\lambda)$ <hr/> $(sk_0, sk_1, \mu, t_0, t_1) \leftarrow \mathcal{A}(1^\lambda)$ if $f(sk_0) \neq f(sk_1)$ return 0 if $\exists i \in \{0, 1\} : \text{Vrfy}(sk_i, \mu, t_i) = 0$ return 0 if $\text{Link}(\mu, t_0, t_1) = 0$ return 1 return 0	$\text{Invert}_{II,\mathcal{A}}^{\text{LIT}}(\lambda)$ <hr/> $sk \leftarrow \mathcal{KGen}(1^\lambda)$ $pk_0 \leftarrow f(sk_0)$ $sk' \leftarrow \mathcal{A}^{\text{Tag}^\mathcal{O}}(pk_0)$ if $pk = f(sk')$ return 1 return 0
	$\text{Forge}_{II,\mathcal{A}}^{\text{LIT}}(\lambda)$ <hr/> $\mathcal{Q} \leftarrow \emptyset$ $sk_0 \leftarrow \mathcal{KGen}(1^\lambda), pk_0 \leftarrow f(sk_0)$ $(sk^*, \mu, t^*) \leftarrow \mathcal{A}^{\text{Tag}^\mathcal{O}}(pk_0)$ if $\text{Vrfy}(sk^*, \mu, t^*) = 0$ return 0. if $\exists (\mu, t) \in \mathcal{Q} : \text{Link}(\mu, t, t^*) = 1$ return 1 return 0

Fig. 5: LIT security games.

Definition 16. A LIT II is *unforgeable*, if there exists a negligible function negl such that for all ppt adversaries \mathcal{A} it holds that

$$\Pr[\text{Forge}_{II,\mathcal{A}}^{\text{LIT}}(\lambda) = 1] \leq \text{negl}(\lambda).$$

Definition 17. A LIT II has *non-invertability*, if there exists a negligible function negl such that for all ppt adversaries \mathcal{A} it holds that

$$\Pr[\text{Invert}_{II,\mathcal{A}}^{\text{LIT}}(\lambda) = 1] \leq \text{negl}(\lambda).$$

Due to the aforementioned change in security model, the construction of a LIT from [9] is not secure in the new model. This is because if this construction tags the same message twice, the adversary receives two *learning with errors* (LWE) samples with the same \mathbf{A} . To remedy this, we propose a construction similar to that of [9], that instead of relying on LWE uses *learning with rounding* (LWR) [4]. Due to this, tagging becomes deterministic and linking becomes an equality

check. Thus, naturally there is only one tag for each message under a secret key. Therefore, with the LWR-based LIT the security model of [9] and our new one look the same for an adversary, which is why we can then do the same reductions as for the LWE-based LIT of [9] to show that the deterministic LIT is secure if LWR is hard. Thus, the following lemma follows directly from [9].

Lemma 18. *There exists a LIT that has tag-indistinguishability, linkability, unforgeability, and invertability if decisional LWR and search LWR are hard.*

For the formal proof, see Section F.1.

2.5 Non-Interactive Zero-Knowledge Proof

We model non-interactive zero-knowledge proof systems in the random oracle model as in [9].

Definition 19 (NIZK). *A non-interactive proof system (NIZK) for a relation \mathfrak{R} in the random oracle model is defined as a triple $\Pi_{\text{NIZK}} = (\text{Setup}, \text{P}, \text{V})$ of ppt algorithms:*

- $\text{crs} \leftarrow \text{Setup}(1^\lambda)$: *On input 1^λ the setup algorithm outputs a common reference string crs .*
- $\pi \leftarrow \text{P}^{\mathcal{RO}}(\text{crs}, x, w, \mu)$: *On input a common reference string crs , an instance x , witness w , and a message μ , and given oracle access to the random oracle \mathcal{RO} , the prover outputs a proof π .*
- $b \leftarrow \text{V}^{\mathcal{RO}}(\text{crs}, x, \mu, \pi)$: *On input a common reference string crs , a statement x , a message μ , and a proof π , and given oracle access to the random oracle \mathcal{RO} , the verifier outputs a bit b .*

To simplify notation, we sometimes omit the random oracle \mathcal{RO} , but assume implicitly that the prover and verifier have access to it. We say that the NIZK is correct, if for all $(x, w) \in \mathfrak{R}$ and $m \in \{0, 1\}^$, we have that*

$$\Pr[\text{V}(\text{crs}, x, m, \text{P}(\text{crs}, x, w, m)) : \text{crs} \leftarrow \text{Setup}(1^\lambda)] = 1.$$

The message in the above definition is not necessary for a NIZK itself, but since the NIZK we use uses the Fiat-Shamir heuristic, we want to be able make the proof dependent on a message. For a relation \mathfrak{R} , we define $L_{\mathfrak{R}} = \{x \mid \exists w : (x, w) \in \mathfrak{R}\}$ as the language of \mathfrak{R} . We define a shorthand notation to quickly show what relation we want to prove.

Definition 20. *We denote the generation of a proof $\pi \leftarrow \text{P}(\text{crs}, x, w, \mu)$ by*

$$\pi \leftarrow \text{NIZK}\{x; w; \mathfrak{R}(x, w)\}(\mu),$$

where P is from a non-interactive proof system Π_{NIZK} for the relation \mathfrak{R} .

For security, we use the standard notions of zero-knowledge and straight-line extractability.

Definition 21 (Zero-Knowledge). A NIZK Π is zero-knowledge if there exists a simulator Sim consisting of three ppt algorithms $\text{Sim} = (\text{Sim.Setup}, \text{Sim.R}\mathcal{O}, \text{Sim.Sim})$ such that for all ppt \mathcal{A} there exists a negligible function negl such that,

$$\left| \begin{array}{l} \Pr[\mathcal{A}^{\text{P}\mathcal{O}, \text{R}\mathcal{O}}(1^\lambda, \text{crs}) = 1 : \text{crs} \leftarrow \text{Setup}(1^\lambda)] \\ - \Pr[\mathcal{A}^{\text{Sim.Sim}, \text{Sim.R}\mathcal{O}}(1^\lambda, \text{crs}) = 1 : \text{crs} \leftarrow \text{Sim.Setup}(1^\lambda)] \end{array} \right| \leq \text{negl}(\lambda)$$

where $\text{R}\mathcal{O}$ denotes a random oracle and $\text{P}\mathcal{O}$, queried on input (x, w, μ) , returns $\text{P}(\text{crs}, x, w, \mu)$. The oracle $\text{Sim}(x, w, \mu)$ checks if $(x, w) \in \mathfrak{R}$ and if so, returns $\text{Sim.Sim}(x, \mu)$. We assume that Sim is stateful, i.e., it implicitly keeps state between invocations of Sim.Setup , $\text{Sim.R}\mathcal{O}$, and Sim.Sim .

Definition 22 (Straight-line extractability). Let $\Pi = (\text{Setup}, \text{P}, \text{V})$ be a NIZK. We say that Π is a straight-line extractable proof of knowledge if there are ppt algorithms $\text{Ext}_0, \text{Ext}_1$ such that for all ppt $\mathcal{A}_0, \mathcal{A}_1$, there exist negligible functions $\text{negl}_0, \text{negl}_1$ such that

$$\left| \begin{array}{l} \Pr[\mathcal{A}_0(1^\lambda, \text{crs}) = 1 : \text{crs} \leftarrow \text{Setup}(1^\lambda)] \\ - \Pr[\mathcal{A}_0(1^\lambda, \text{crs}) = 1 : (\text{crs}, \text{td}) \leftarrow \text{Ext}_0(1^\lambda)] \end{array} \right| \leq \text{negl}_0(\lambda)$$

and

$$\Pr \left[\begin{array}{l} \text{V}^{\text{R}\mathcal{O}}(\text{crs}, x, m, \pi) = 1 \\ \wedge (x, w) \notin \mathfrak{R} \end{array} : \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Ext}_0(1^\lambda), \\ (x, m, \pi) \leftarrow \mathcal{A}_1(1^\lambda, \text{crs}), \\ w \leftarrow \text{Ext}_1(\text{td}, x, m, \pi) \end{array} \right] \leq \text{negl}_1(\lambda)$$

In the random oracle model, Ext_1 gets the list of random oracle queries that \mathcal{A} made as additional input.

2.6 Verifiable Ring Signature

A verifiable ring signature (VRS) is a standard ring signature with an additional functionality. At any point after creating a signature, a signer can output a proof showing that in fact it created the signature. On the other hand, people who did not sign the signature, but are part of the ring, can show that they did *not* create the signature. We formalize this by adding two ppt algorithms Prove and Judge to the standard ring signature model.

Definition 23. A verifiable ring signature consists of five ppt algorithms $(\text{Setup}, \text{KGen}, \text{Sign}, \text{Vrfy}, \text{Link})$.

$\text{pp} \leftarrow \text{Setup}(1^\lambda)$: On input a security parameter 1^λ , the setup algorithm outputs public parameters pp .

$(sk, pk) \leftarrow \text{KGen}(\text{pp})$: On input public parameters pp , the key generation algorithm outputs a secret, public key pair (sk, pk) .

$\sigma \leftarrow \text{Sign}(sk, R, \mu)$: On input a secret key sk , a ring $R = \{pk_i\}_i$ and a message μ , the signing algorithm outputs a signature σ .

$b \leftarrow \text{Vrfy}(R, \mu, \sigma)$: On input a ring R , a message μ and a signature σ , the verifying algorithm outputs a bit b .

$\pi \leftarrow \text{Prove}(sk, R, \mu, \sigma)$: On input a secret key sk , a ring R , a message μ , and a signature σ , the prove algorithm outputs a proof π .

$b \leftarrow \text{Judge}(pk, R, \mu, \sigma, \pi)$: On input a public key pk , a ring R , a message μ , a signature σ , and a proof π , the judging algorithm outputs a bit b .

For security, we require strong unforgeability and anonymity of a ring signature, but now the adversary also has an oracle generating proofs with the `Prove` algorithm. Additionally we require *accountability*, which means that it is hard for an adversary to create a signature, where `Judge` thinks it did not create it. Furthermore we require *non-seizability*, which means that is hard for an adversary to create a signature, where `Judge` attributes the signature to an honest user. To model this, we use the security model of [15]. The definition of the games can be found in Fig. 6.

Definition 24. We say that a VRS Π is strongly unforgeable, if there exists a negligible function negl such that for all qpt adversaries \mathcal{A} and all $\ell \in \text{poly}(\lambda)$ it holds that

$$\Pr[\text{Forge}_{\Pi, \mathcal{A}}^{\ell}(\lambda) = 1] \leq \text{negl}(\lambda).$$

Definition 25. We say that a VRS Π is anonymous, if there exists a negligible function negl such that for all qpt adversaries \mathcal{A} it holds that

$$|\Pr[\text{Anon}_{\Pi, \mathcal{A}}^0(\lambda) = 1] - \Pr[\text{Anon}_{\Pi, \mathcal{A}}^1(\lambda) = 1]| \leq \text{negl}(\lambda).$$

Definition 26. We say that a VRS Π is strongly accountable, if there exists a negligible function negl such that for all qpt adversaries \mathcal{A} and all $\ell \in \text{poly}(\lambda)$ it holds that

$$\Pr[\text{Acc}_{\Pi, \mathcal{A}}^{\ell}(\lambda) = 1] \leq \text{negl}(\lambda).$$

Definition 27. We say that a VRS Π is strongly non-seizable, if there exists a negligible function negl such that for all qpt adversaries \mathcal{A} it holds that

$$\Pr[\text{Seiz}_{\Pi, \mathcal{A}}(\lambda) = 1] \leq \text{negl}(\lambda).$$

3 New Lattice Constructions

3.1 A Fully Collision-Resistant Chameleon Hash Function

We now want to construct our new chameleon hash to later use it when instantiating SSS constructions. Since the security notion that [31] achieves is close to our target security, full collision-resistance, we take their construction as a starting point. The idea of their random oracle model construction is to first generate a trapdoored matrix \mathbf{A} in the setup. To hash a message μ , they query the random oracle on μ and some other values to get another matrix \mathbf{A}_h and choose a small Gaussian value \mathbf{e} . Then, they use the well known Ajtai hash function to compute $[\mathbf{A}|\mathbf{A}_h]\mathbf{e} = \mathbf{h}$, which is their hash value. To adapt a message, they use the trapdoor to compute some short \mathbf{e}' such that $[\mathbf{A}|\mathbf{A}'_h]\mathbf{e}' = \mathbf{h}$ for some different

$\text{Sign}\mathcal{O}(\{sk_i\}_i, \cdot, \cdot, \cdot) \text{ with } (j, R, \mu)$ <hr/> $\sigma \leftarrow \text{Sign}(sk_j, R, \mu)$ $\mathcal{Q}_{\text{Sign}} \leftarrow \mathcal{Q}_{\text{Sign}} \cup \{\mu, \sigma\}$ return σ $\text{Prove}\mathcal{O}(\{sk_i\}_i, \cdot, \cdot, \cdot) \text{ with } (j, R, \mu, \sigma)$ <hr/> if $pk_i \in R \wedge (\mu, \sigma) \notin \mathcal{Q}_{\text{LoR}}$ return $\text{Prove}(sk_j, R, \mu, \sigma)$ else return \perp	$\text{LoRSign}\mathcal{O}_b(sk_0, sk_1, \cdot, \cdot) \text{ with } (R, \mu)$ <hr/> if $\{pk_0, pk_1\} \subseteq R$ $\sigma \leftarrow \text{Sign}(sk_b, R, \mu)$ $\mathcal{Q}_{\text{LoR}} \leftarrow \mathcal{Q}_{\text{LoR}} \cup \{\mu, \sigma\}$ return σ else return \perp
$\text{Forge}_{\Pi, \mathcal{A}}^\ell(\lambda)$ <hr/> $\mathcal{Q}_{\text{Sign}} \leftarrow \emptyset, \text{pp} \leftarrow \text{Setup}(1^\lambda)$ $\forall i \in [\ell], (sk_i, pk_i) \leftarrow \text{KGen}(\text{pp})$ $(R^*, \mu^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}\mathcal{O}, \text{Prove}\mathcal{O}}(\{pk_i\}_{i \in [\ell]})$ $b_0 \leftarrow \text{Vrfy}(R^*, \mu^*, \sigma^*)$ $b_1 \leftarrow R^* \subseteq \{pk_i\}_{i \in [\ell]}$ $b_2 \leftarrow (\mu^*, \sigma^*) \notin \mathcal{Q}_{\text{Sign}}$ return $b_0 \wedge b_1 \wedge b_2$	$\text{Acc}_{\Pi, \mathcal{A}}^\ell(\lambda)$ <hr/> $\mathcal{Q}_{\text{Sign}} \leftarrow \emptyset, \text{pp} \leftarrow \text{Setup}(1^\lambda)$ $\forall i \in [\ell], (sk_i, pk_i) \leftarrow \text{KGen}(\text{pp})$ $(pk^*, R^*, \mu^*, \sigma^*, \pi^*) \leftarrow \mathcal{A}^{\text{Sign}\mathcal{O}, \text{Prove}\mathcal{O}}(\{pk_i\}_{i \in [\ell]})$ $b_0 \leftarrow \text{Vrfy}(R^*, \mu^*, \sigma^*)$ $b_1 \leftarrow (\text{Judge}(pk^*, R^*, \mu^*, \sigma^*, \pi^*) = 0)$ $b_2 \leftarrow R^* \subseteq \{pk_i\}_{i \in [\ell]} \cup \{pk^*\}$ $b_3 \leftarrow (\mu^*, \sigma^*) \notin \mathcal{Q}_{\text{Sign}}$ return $b_0 \wedge b_1 \wedge b_2 \wedge b_3$
$\text{Anon}_{\Pi, \mathcal{A}}^b(\lambda)$ <hr/> $\mathcal{Q}_{\text{LoR}} \leftarrow \emptyset, \text{pp} \leftarrow \text{Setup}(1^\lambda)$ $(sk_0, pk_0) \leftarrow \text{KGen}(\text{pp})$ $(sk_1, pk_1) \leftarrow \text{KGen}(\text{pp})$ $b' \leftarrow \mathcal{A}^{\text{Sign}\mathcal{O}, \text{Prove}\mathcal{O}, \text{LoRSign}\mathcal{O}}(pk_0, pk_1)$ return b'	$\text{Seiz}_{\Pi, \mathcal{A}}(\lambda)$ <hr/> $\mathcal{Q}_{\text{Sign}} \leftarrow \emptyset, \text{pp} \leftarrow \text{Setup}(1^\lambda)$ $(sk_0, pk_0) \leftarrow \text{KGen}(\text{pp})$ $(R^*, \mu^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}\mathcal{O}, \text{Prove}\mathcal{O}}(pk_0)$ $\pi^* \leftarrow \text{Prove}(sk_0, R^*, \mu^*, \sigma^*)$ $b_0 \leftarrow \text{Vrfy}(R^*, \mu^*, \sigma^*)$ $b_1 \leftarrow (\text{Judge}(pk_0, R^*, \mu^*, \sigma^*, \pi^*) \neq 0)$ $b_2 \leftarrow (\mu^*, \sigma^*) \notin \mathcal{Q}_{\text{Sign}}$ return $b_0 \wedge b_1 \wedge b_2$

Fig. 6: VRS security games and oracles.

\mathbf{A}'_h also output by the random oracle. The issue with this idea is that to show security they need to assume having a unique tag τ each time a new \mathbf{h} is created, which they use as additional input to the random oracle. While this is no problem in their application of redactable blockchains, the constructions of SSS that we look at require the chameleon hash to be tag-free. If we simply removed the tag from their construction, their construction would not be fully collision-resistant. An adversary could query the adapt oracle for a collision $[\mathbf{A}|\mathbf{A}_h]\mathbf{e} = [\mathbf{A}|\mathbf{A}'_h]\mathbf{e}'$ and return $(2\mathbf{e}, 2\mathbf{e}')$ as a new collision which wins the security game.

The problem with this approach is that \mathbf{e} can be freely chosen by the adversary without any restrictions. Therefore, the idea for our construction is to use the construction of [31] without tags and to additionally bind \mathbf{e} with a commitment scheme to prohibit this attack. In fact, a commitment scheme is not sufficient, as we then would not be able to implement the CAdapt algorithm correctly. Instead, we use a trapdoor commitment scheme and its DEE property.

We can now construct our chameleon hash function. For this, let $n, q > 1$, $m = \mathcal{O}(n \log q)$, $s = \mathcal{O}(\sqrt{n \log q})$ large enough, and $\beta = s \cdot \sqrt{2m}$. Let $\Pi_{\text{Com}} = (\text{TdGen}, \text{Com}, \text{ComCheck}, \text{Equiv})$ be a trapdoor commitment scheme. Let the random oracle $\mathcal{RO} : \{0, 1\}^* \rightarrow \mathbb{Z}_q^{n \times m}$. We then construct our chameleon hash function as seen in Fig. 7.

Note that we require $\mathbf{e}_2 \neq \mathbf{0}$ for security to hold (which we show in Appendix F.2).

Theorem 28. *If the commitment scheme has DEE, then the chameleon hash function construction given in Fig. 7 has indistinguishability.*

Proof. Let D_{CHash} be the distribution of $(\mathbf{h}, r = (z, \mathbf{e}, c, d) \leftarrow_{\S} \text{CHash}(\text{pk}_{\text{ch}}, \mu))$ and let D_{CAdapt} be the distribution of $(\mathbf{h}', r'' = (z'', \mathbf{e}'', c'', d''))$, where $(\mathbf{h}', r') \leftarrow_{\S} \text{CHash}(\text{pk}_{\text{ch}}, \mu')$, $r'' \leftarrow_{\S} \text{CAdapt}(\text{sk}_{\text{ch}}, \mathbf{h}', \mu', r', \mu)$ for some μ, μ' . Then we can easily see the marginal distributions of \mathbf{h} and \mathbf{h}' are statistically close to uniformly random by Theorem 56. Furthermore, z, z'' are uniform by definition. For \mathbf{e}, \mathbf{e}'' we know that their distribution is statistically indistinguishable due to Theorem 9. Finally, by the DEE property of the commitment scheme we know that $(c, d), (c'', d'')$ are computationally indistinguishable. \square

Lemma 29. *The construction given in Fig. 7 does not have uniqueness (cf. Definition 54).*

This can be shown by the following attack. Since the adversary in the uniqueness game can choose the public key, it can generate a public key honestly together with a trapdoor. Then, it can just hash some message μ to (h, r) and use the CAdapt algorithm with input $(\text{sk}_{\text{ch}}, h, \mu, r, \mu)$ to get some new $r' = (z', \mathbf{e}', c', d')$ for the same message μ . While it is possible to derandomize z', c', d' , such that $z = z', c = c', d = d'$ (for the same μ and \mathbf{e}) with the help of a PRF, by design, $\mathbf{e}' \neq \mathbf{e}$ with overwhelming probability, since the output distribution of PreSample is a (conditioned) discrete Gaussian. We expect it is necessary that in order to construct a chameleon hash function with uniqueness, one needs to use another building block than PreSample.

$\text{CKGen}(1^\lambda)$ <hr style="border: 0.5px solid black;"/> $(\mathbf{A}, \text{td}) \leftarrow \text{TrapGen}(n, m, q, s)$ $(\text{pp}, \text{td}_{\text{Com}}) \leftarrow \text{TdGen}(1^\lambda)$ $\text{pk}_{\text{ch}} \leftarrow (\mathbf{A}, \text{pp})$ $\text{sk}_{\text{ch}} \leftarrow (\text{td}, \text{td}_{\text{Com}})$ $\text{return } (\text{pk}_{\text{ch}}, \text{sk}_{\text{ch}})$ $\text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu, r)$ <hr style="border: 0.5px solid black;"/> $\text{Parse}(z, \mathbf{e} = (\mathbf{e}_1^t, \mathbf{e}_2^t)^t, c, d) \leftarrow r$ <p style="text-align: center; margin-left: 20px;">with $\mathbf{e}_2 \in \mathbb{Z}_q^m$, and $\mathbf{h} \leftarrow h$</p> $\mathbf{B} \leftarrow \mathcal{RO}(\mu, z, c)$ $\text{if } \mathbf{h} = [\mathbf{A} \ \mathbf{B}] \mathbf{e}$ <p style="margin-left: 20px;">$\wedge \text{ComCheck}(\text{pp}, \mathbf{e}, c, d) = 1$</p> <p style="margin-left: 20px;">$\wedge \ \mathbf{e}\ \leq \beta \wedge \mathbf{e}_2 \neq \mathbf{0}$</p> $\text{return } 1$ $\text{return } 0$	$\text{CHash}(\text{pk}_{\text{ch}}, \mu)$ <hr style="border: 0.5px solid black;"/> $z \leftarrow \{0, 1\}^\lambda$ $\mathbf{e} \leftarrow D_{\mathbb{Z}^{2m}, s}$ $(c, d) \leftarrow \text{Com}(\text{pp}, \mathbf{e})$ $\mathbf{B} \leftarrow \mathcal{RO}(\mu, z, c)$ $\mathbf{h} \leftarrow [\mathbf{A} \ \mathbf{B}] \mathbf{e}$ $h \leftarrow \mathbf{h}, r \leftarrow (z, \mathbf{e}, c, d)$ $\text{return } (h, r)$ $\text{CAdapt}(\text{sk}_{\text{ch}}, h, \mu, r, \mu')$ <hr style="border: 0.5px solid black;"/> $\mathbf{h} \leftarrow h$ $\text{if } \text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu, r) = 0$ <p style="margin-left: 20px;">$\text{return } \perp$</p> $z' \leftarrow \{0, 1\}^\lambda$ $c' \leftarrow \text{Com}(\text{pp}, 0)$ $\mathbf{B} \leftarrow \mathcal{RO}(\mu', z', c')$ $\mathbf{e}' \leftarrow \text{PreSample}([\mathbf{A} \ \mathbf{B}], \text{td}, \mathbf{h}, s)$ $d' \leftarrow \text{Equiv}(\text{pp}, \text{td}_{\text{Com}}, c', \mathbf{e}')$ $r' \leftarrow (z', \mathbf{e}', c', d')$ $\text{return } r'$
---	---

Fig. 7: Lattice-based fully-collision-resistant chameleon hash function.

Theorem 30. *If $\text{SIS}_{n,m,q,\beta'}$ is hard, where $\beta' = 2s' \cdot (\sqrt{n} + \sqrt{m} + t + 1) \cdot s\sqrt{2m}$ with $t \geq 0$ and $s' = \omega(\sqrt{\log m})$, and if Π is a computationally binding trapdoor commitment scheme that has DEE, the construction CH given in Fig. 7 is fully collision-resistant in the random oracle model.*

The idea of the proof is as follows. First, we take an adversary \mathcal{A} against the full collision-resistance of the construction. Then, we define an alternative security game for \mathcal{A} to play in, in which we do not generate a trapdoor for \mathbf{A} and forget the trapdoor for the commitment scheme. Instead, when answering a CAdapt query, we use the random oracle to program a trapdoor into the \mathbf{B} we generate during it. However, when we answer a random oracle query, we do not program a trapdoor into \mathbf{B} . Then, we want to construct an adversary \mathcal{B} against SIS that simulates \mathcal{A} in the alternative security game. If \mathcal{A} wins this game by outputting a valid collision (h, μ, r, μ', r') , there are multiple cases that can happen. First, for $(h, \mu, r = (z, \mathbf{e}, c, d))$ we look at $\mathbf{B} \leftarrow \mathcal{RO}(\mu, z, c)$. If \mathbf{B} was generated without a trapdoor, we are fine and can use it to extract an SIS solution. However, we need a second part to do that. In the case that $(h, \mu') \notin \mathcal{Q}$, we hope that $\mathbf{B}' \leftarrow \mathcal{RO}(\mu', z', c')$, where $r' = (z', \mathbf{e}', c', d')$, was

generated without a trapdoor as well. In the other case, i.e., $(h, \mu') \in \mathcal{Q}$, we can argue that there must exist a candidate $(h, \hat{\mu}, \hat{r} = (\hat{z}, \hat{e}, \hat{c}, \hat{d}))$. For this candidate, we hope that $\hat{\mathbf{B}} \leftarrow \mathcal{RO}(\hat{\mu}, \hat{z}, \hat{c})$ was generated without a trapdoor as well. Finally, if \mathbf{B} was generated without trapdoor and either \mathbf{B}' or $\hat{\mathbf{B}}$ was generated without trapdoor (in the respective case), we can extract an SIS solution. We will later formally define this as the event `Free`. If, on the other hand, we have $\neg\text{Free}$, we can instead break the binding of the commitment scheme. This is because with the commitment scheme we bind \mathbf{e} to some \mathbf{B} and therefore to some h . If $\neg\text{Free}$ and thus one of the \mathbf{B} has a trapdoor in it, the adversary used some \mathbf{B} with two different \mathbf{e}, \mathbf{e}' . A full, formal version of this proof can be found in Appendix F.2.

Since we argued before that a lattice-based trapdoor commitment scheme with the required properties exist, we get the following corollary.

Corollary 31. *There exists a construction of a chameleon hash function that is f -CR secure, if SIS is hard.*

The fully-collision-resistant chameleon hash function based on lattice assumptions we constructed here, will be applied to instantiate various SSS constructions in Section 5.

3.2 A Generic Verifiable Ring Signature Construction, Instantiated with Lattices

Let $\text{LIT} = (\text{LIT.KGen}, \text{Tag}, \text{LIT.Vrfy}, \text{Link}, f)$ $\text{NIZK} = (\text{NIZK.Setup}, \text{P}, \text{V})$ respectively be a LIT and NIZK for the relations required in our construction of a VRS, which can be found in Fig. 8.

Theorem 32. *If the LIT has non-invertability and if the NIZK is straight-line extractable and zero-knowledge, then the VRS construction has strong unforgeability.*

Proof. Let $\ell \in \text{poly}(\lambda)$ and let \mathcal{A} be a qpt adversary against the strong unforgeability of the VRS construction. We construct an adversary \mathcal{B} against the non-invertability of the LIT. On input a pk , \mathcal{B} first samples ℓ keys $sk_i \leftarrow_{\$} \text{LIT.KGen}(1^\lambda)$ and sets $pk_i \leftarrow f(sk_i)$. Then, it guesses a $k \leftarrow_{\$} [\ell]$ and replaces the k th public key $pk_k \leftarrow pk$. \mathcal{B} then simulates \mathcal{A} as in the strong unforgeability game, except for the following two changes.

- When \mathcal{A} makes a query (i, R, μ) to its `Sign` oracle with $i = k$, \mathcal{B} generates the tag with $t_{\text{Sign}} \leftarrow_{\$} \text{TagO}((R, \mu, r_{\text{Sign}}))$, where $r_{\text{Sign}} \leftarrow_{\$} \{0, 1\}^\lambda$ as in the construction. The proof π is generated with the simulator of the NIZK.
- When \mathcal{A} makes a query (i, R, μ, σ) to its `ProveO` oracle with $i = k$, and if $pk \in R$ and $\text{Vrfy}(R, \mu, \sigma) = 1$, and where $\sigma = (r_{\text{Sign}}, t_{\text{Sign}}, \pi_{\text{Sign}})$, then \mathcal{B} generates the tag with $t_{\text{Prove}} \leftarrow_{\$} \text{TagO}((R, \mu, r_{\text{Sign}}))$. The proof is generated with the simulator of the NIZK.

<pre> Setup(1^λ) ----- pp \leftarrow NIZK.Setup(1^λ) return pp KGen(pp) ----- sk \leftarrow LIT.KGen(1^λ) pk \leftarrow f(sk) return (sk, pk) Sign(sk, R, μ) ----- if f(sk) \notin R return \perp r_{Sign} \leftarrow {0, 1}^{λ} t_{Sign} \leftarrow Tag(sk, (R, μ, r_{Sign})) $\pi_{\text{Sign}} \leftarrow$ NIZK{(R, μ, r_{Sign}, t_{Sign}); (sk, pk); f(sk) = pk, pk \in R, LIT.Vrfy(sk, (R, μ, r_{Sign}, t_{Sign}) = 1} $\sigma \leftarrow$ (r_{Sign}, t_{Sign}, π_{Sign}) return σ </pre>	<pre> Vrfy(R, μ, σ) ----- return \vee(crs, (R, μ, r_{Sign}, t_{Sign}), π_{Sign}) = 1 Prove(sk, R, μ, σ) ----- if f(sk) \notin R \vee Vrfy(R, μ, σ) = 0 return \perp t_{Prove} \leftarrow Tag(sk, (R, μ, r_{Sign})) $\pi_{\text{Prove}} \leftarrow$ NIZK{(pk, R, μ, r_{Sign}, t_{Prove}); sk; f(sk) = pk, LIT.Vrfy(sk, (R, μ, r_{Sign})) = 1} $\pi \leftarrow$ (t_{Prove}, π_{Prove}) return π Judge(pk, R, μ, σ, π_{Prove}) ----- if \vee(crs, (pk, R, μ, r_{Sign}, t_{Prove}), π_{Prove}) = 0 \vee Vrfy(R, μ, σ) = 0 return \perp if pk \notin R return 0 return Link((R, μ), t_{Sign}, t_{Prove}) </pre>
--	---

Fig. 8: Generic VRS construction.

After \mathcal{A} has output a forgery $(R^*, \mu^*, \sigma^* = (r_{\text{Sign}}^*, t_{\text{Sign}}^*, \pi_{\text{Sign}}^*))$, if $b_0 \wedge b_1 \wedge b_2$ is true, \mathcal{B} uses the straight-line extractor of the NIZK to extract (sk^*, pk^*) from π_{Sign}^* . If $pk = pk^*$, \mathcal{B} outputs sk^* .

By the definition of the non-invertability game and the zero-knowledgeness of the NIZK, \mathcal{A} is perfectly simulated. Since b_0 and b_2 are true, we know that the straight-line extractability is successful and therefore that $f(sk^*) = pk^*$ and that $pk^* \in R^*$. Then, since b_1 is true, there exists a $j \in [\ell]$ such that $pk^* = pk_j$. If $j = k$, we therefore have $f(sk^*) = pk_k = pk^*$, thus sk^* is a valid solution for the invertability game, and we have

$$\Pr[\text{Inv}_{\text{LIT}, \mathcal{B}}(\lambda) = 1] = \frac{1}{\ell} \Pr[\text{Forge}_{\Pi, \mathcal{A}}^\ell(\lambda) = 1]. \quad \square$$

Theorem 33. *If the LIT has tag-indistinguishability and if the NIZK is zero-knowledge, then the VRS construction has anonymity.*

Proof. Let \mathcal{A} be an adversary against the anonymity of the VRS. We first define an alternative security game and then construct an adversary against the tag-indistinguishability of the LIT. The alternative security game Game_1^b

for $b \in \{0, 1\}$ works like $\text{Anon}_{\Pi, \mathcal{A}}^b(\lambda)$, except that during the oracle calls to SignO , ProveO , LoRSignO , the p_{Sign} and π_{Prove} are generated with the zero-knowledge simulator of the NIZK. We know that these games only differ negligibly by the zero-knowledge property of the NIZK. Then, we construct an adversary \mathcal{B} against the tag-indistinguishability of the LIT. On input (\hat{pk}_0, \hat{pk}_1) , \mathcal{B} simulates \mathcal{A} in Game_1 , except for the following changes.

- Instead of generating pk_0, pk_1 , \mathcal{B} instead sets $pk_0 \leftarrow \hat{pk}_0, pk_1 \leftarrow \hat{pk}_1$.
- Whenever \mathcal{B} needs to compute a tag on a message $(R, \mu, r_{\text{Sign}})$ to answer an oracle query of \mathcal{A} , \mathcal{B} instead queries its own tag oracle on input $(R, \mu, r_{\text{Sign}})$.

Then, after \mathcal{A} outputs a bit \hat{b} , \mathcal{B} outputs \hat{b} . We know that by the definition of Game_1 and of the tag-indistinguishability game that if $b = 0$ (or $b = 1$) then \mathcal{A} is simulated as in Game_1^0 (or Game_1^1 , respectively). Thus, we know that

$$\begin{aligned} & \left| \Pr[\text{Anon}_{\Pi, \mathcal{A}}^0(\lambda) = 1] - \Pr[\text{Anon}_{\Pi, \mathcal{A}}^1(\lambda) = 1] \right| \\ &= \left| \Pr[\text{Anon}_{\text{LIT}, \mathcal{B}, 0}^{\text{LIT}}(\lambda) = 1] - \Pr[\text{Anon}_{\text{LIT}, \mathcal{B}, 1}^{\text{LIT}}(\lambda) = 1] \right|, \end{aligned}$$

which concludes the proof. \square

Theorem 34. *If the LIT has non-invertability and linkability and if the NIZK is straight-line extractable and zero-knowledge, then the VRS construction has accountability.*

Proof. Let $\ell \in \text{poly}(\lambda)$ and let \mathcal{A} be a *qpt* adversary against the accountability of the VRS construction. We define a series of games Game_i , in which \mathcal{A} is simulated.

Game_0 is the same as $\text{Acc}_{\Pi, \mathcal{A}}^\ell(\lambda)$. Define the event Win_0 to be the event that $b_0 \wedge b_1 \wedge b_2 \wedge b_3$ is true.

Game_1 is the same as Game_0 , except that after \mathcal{A} outputs a forgery $(pk^*, R^*, \mu^*, \sigma^*, \pi^*)$ with $\sigma^* = (r_{\text{Sign}}^*, t_{\text{Sign}}^*, \pi_{\text{Sign}}^*)$ and $\pi^* = (t_{\text{Prove}}^*, \pi_{\text{Prove}}^*)$ and $b_0 \wedge b_1 \wedge b_2 \wedge b_3$ being true, the game uses the straight-line extractor of the NIZK to extract $(sk_{\text{Sign}}^*, pk_{\text{Sign}}^*)$ from π_{Sign}^* and sk_{Prove}^* from π_{Prove}^* . Define the event Win_1 to be the event that Win_0 is true and that both extractions are successful. Then we know by the straight-line extractability of the NIZK that

$$\Pr[\text{Win}_0] \leq \Pr[\text{Win}_1] + \text{negl}(\lambda).$$

Thus, in Game_1 we always successfully extract if \mathcal{A} wins Game_0 .

Game_2 is the same game as Game_1 , except that we always have $pk^* \in R^*$, i.e., we define Win_2 as $\text{Win}_1 \wedge pk^* \in R^*$. We can show that the probabilities of Win_1 and Win_2 only differ negligibly by using the strong unforgeability of the VRS. Since b_0 and b_3 are true, (R^*, μ^*, σ^*) forms a valid forgery in the strong unforgeability game against the VRS. The input and oracles of \mathcal{A} match in the two games. Thus, since we require non-invertability and linkability, we can construct an adversary \mathcal{C} as in the proof of Theorem 32. Thus, we know that

$$\Pr[\text{Win}_1] \leq \frac{1}{\ell} \Pr[\text{Win}_2].$$

Game_3 is the same as Game_2 , except that the extracted pk_{Sign}^* is always equal to pk^* , i.e., we define Win_3 as $\text{Win}_2 \wedge pk_{\text{Sign}}^* = pk^*$. Again we can use the strong unforgeability of the VRS to show that Game_2 and Game_3 differ negligibly. With the same argument as in Game_2 , we can argue that

$$\Pr[\text{Win}_2] \leq \frac{1}{\ell} \Pr[\text{Win}_3].$$

We can now construct an adversary \mathcal{B} against the linkability of the LIT. \mathcal{B} simulates \mathcal{A} as in Game_3 , i.e., we assume that \mathcal{A} wins, the extraction is successful, $pk^* \in R^*$, and the extracted $pk_{\text{Sign}}^* = pk^*$. Then, \mathcal{B} returns $(sk_{\text{Sign}}^*, sk_{\text{Prove}}^*, (R^*, \mu^*, r_{\text{Sign}}^*), t_{\text{Sign}}^*, t_{\text{Prove}}^*)$ to its challenger. To argue that \mathcal{B} wins the linkability game, we need to show that

1. $f(sk_{\text{Sign}}^*) = f(sk^*)$
2. $\text{LIT.Vrfy}(sk_{\text{Sign}}^*, (R^*, \mu^*, r_{\text{Sign}}^*), t_{\text{Sign}}^*) = \text{LIT.Vrfy}(sk^*, (R^*, \mu^*, r_{\text{Sign}}^*), t_{\text{Prove}}^*) = 1$
3. $\text{Link}((R^*, \mu^*, r_{\text{Sign}}^*), t_{\text{Sign}}^*, t_{\text{Prove}}^*) = 0$.

Condition (2) follows immediately from the successful extraction, while this fact together with the assumption that $pk_{\text{Sign}}^* = pk^*$ implies (1). For condition (3), we know that since b_1 is true, the judge outputs 0. Due to $pk^* \in R^*$, we know that this only happens if condition (3) is true. Thus, we know that if \mathcal{A} wins Game_3 , then \mathcal{B} wins the the linkability game. Therefore, we have that

$$\Pr[\text{Acc}_{\Pi, \mathcal{A}}^\ell(\lambda) = 1] \leq \frac{1}{\ell^2} \Pr[\text{Linkable}_{\text{LIT}, \mathcal{B}}^{\text{LIT}}(\lambda) = 1] + \text{negl}(\lambda),$$

and the VRS is accountable⁵.

Theorem 35. *If the LIT is unforgeable, and if the NIZK is straight-line extractable and zero-knowledge, then the VRS construction has non-seizability.*

Proof. Let \mathcal{A} be a *qpt* adversary against the accountability of the VRS construction. We construct an adversary \mathcal{B} against the unforgeability of the LIT. On input pk , \mathcal{B} simulates \mathcal{A} as in the non-seizability game, except that \mathcal{B} sets $pk_0 \leftarrow pk$, and except for the following two changes.

- When \mathcal{A} makes a query (i, R, μ) to its Sign oracle with $i = k$, \mathcal{B} generates the tag with $t_{\text{Sign}} \leftarrow_{\$} \text{TagO}((R, \mu, r_{\text{Sign}}))$, where $r_{\text{Sign}} \leftarrow_{\$} \{0, 1\}^\lambda$ as in the construction. The proof π is generated with the simulator of the NIZK.
- When \mathcal{A} makes a query (i, R, μ, σ) to its ProveO oracle with $i = k$, and if $pk \in R$ and $\text{Vrfy}(R, \mu, \sigma) = 1$, and where $\sigma = (r_{\text{Sign}}, t_{\text{Sign}}, \pi_{\text{Sign}})$, then \mathcal{B} generates the tag with $t_{\text{Prove}} \leftarrow_{\$} \text{TagO}((R, \mu, r_{\text{Sign}}))$. The proof is generated with the simulator of the NIZK.

⁵ It is possible to reduce the multiplicative loss from $1/\ell^2$ to $1/\ell$ by going from Game_1 to Game_3 directly.

When \mathcal{A} outputs a forgery $(R^*, mu^*, \sigma^* = (r_{\text{Sign}}^*, t_{\text{Sign}}^*, \pi_{\text{Sign}}^*))$ and if $b_0 \wedge b_1 \wedge b_2 \wedge b_3$ is true, \mathcal{B} uses the straight-line extractor of the NIZK to extract (sk^*, pk^*) from π_{Sign}^* . Then, \mathcal{B} queries $t_{\text{Prove}}^* \leftarrow \text{TagO}((R^*, \mu^*, r_{\text{Sign}}^*))$ and uses the simulator to create π_{Prove}^* to define $\pi^* = (t_{\text{Prove}}^*, \pi_{\text{Prove}}^*)$. Afterwards, \mathcal{B} outputs $(sk^*, (R^*, \mu^*, r_{\text{Sign}}^*), t_{\text{Sign}}^*)$.

By definition of the non-invertability game of the LIT and the zero-knowledge of the NIZK, we know that \mathcal{A} is perfectly simulated. Since b_1 is true, $pk_0 \in R^*$. Furthermore, we know that $\text{Judge}(pk_0, R^*, \mu^*, \sigma^*, \pi_{\text{Prove}}^*) = \perp$ if either $\text{NIZK.Vrfy}(\text{crs}, (pk_0, R^*, t_{\text{Prove}}^*), \mu^*, \pi_{\text{Prove}}^*) = 0$ or $\text{Vrfy}(R^*, \mu^*, \sigma^*) = 0$. However, the former case does not happen due to the correctness of the NIZK, while the latter is not true due to b_0 being true. Thus, $\text{Judge}(pk_0, R^*, \mu^*, \sigma^*, \pi_{\text{Prove}}^*) = 1$, and $\text{Link}((R^*, \mu^*), t_{\text{Sign}}^*, t_{\text{Prove}}^*) = 1$. Due to this and b_0 being true, we know that \mathcal{B} wins the unforgeability game and we have

$$\Pr[\text{Forge}_{\text{LIT}, \mathcal{B}}^{\text{LIT}}(\lambda) = 1] = \Pr[\text{Seiz}_{\Pi, \mathcal{A}}(\lambda) = 1]. \quad \square$$

We now want to instantiate the generic construction with lattice-based building blocks. For the LIT, we can use the one described in Lemma 18. The NIZK need not only be lattice-based, but also be able to prove the statements we need. We use the NIZK from [35] made straight-line extractable with Katsumata's transform [26] as seen in [9, 10]. We need to be able to prove that

1. $f(sk) = pk$
2. $pk \in R$
3. $\text{LIT}, \text{Vrfy}(sk, (R, \mu, r_{\text{Sign}}), t_{\text{Sign}}) = 1$.

Since R, μ, r_{Sign} are public, we can use the argument from [9] to show that we can prove (1) and (3) with the NIZK. For (2), in our use case of SSS the ring R will only consist of two public keys. Thus, it is sufficient if we prove (2) by creating an OR-proof over the condition in (1). If one wants to use the VRS with more than two parties, it is advisable to prove (2) by using accumulators (e.g. [32]) or one-out-of-many proofs (e.g. [34]) to be more efficient.

4 Sanitizable Signature Schemes

We now move on to defining sanitizable signature schemes in this section before providing instantiations in the next section.

The definition for sanitizable signature schemes is an extension of regular signature schemes. Due to the addition of the sanitizer, extra operations for generating keys for the sanitizer as well as sanitization itself are required.

Definition 36. *A sanitizable signature scheme SSS is a tuple of seven probabilistic polynomial-time algorithms $\text{SSS} = (\text{KGen}_{\text{Sig}}, \text{KGen}_{\text{San}}, \text{Sign}, \text{Sanit}, \text{Vrfy}, \text{Proof}, \text{Judge})$ defined as follows:*

$(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow \text{KGen}_{\text{Sig}}(1^\lambda)$: The algorithm takes the security parameter as input and outputs a signer key pair $(pk_{\text{Sig}}, sk_{\text{Sig}})$.

- $(pk_{\text{San}}, sk_{\text{San}}) \leftarrow \text{KGen}_{\text{San}}(1^\lambda)$: The algorithm takes the security parameter as input and outputs a sanitizer key pair $(pk_{\text{San}}, sk_{\text{San}})$.
- $\sigma \leftarrow \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$: On input of a message μ in the message space \mathcal{M} , signer secret key sk_{Sig} , a sanitizer public key pk_{San} , as well as the admissible sanitization rights ADM , the algorithm outputs a signature σ or an error message \perp . ADM contains the indices of block that are admissible for modification. Furthermore, we assume that ADM is always valid with regards to the input message. $\text{ADM}_0 \cap \text{ADM}_1$ denotes the intersection of admissible blocks.
- $(\mu', \sigma') \leftarrow \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$: On input of a message μ , a signature σ , a sanitizer secret key sk_{San} , a signer public key pk_{Sig} , as well as the modification instructions MOD . The algorithm outputs the modified message m' along with the corresponding sanitized signature σ' or an error message \perp . We model MOD as a function which takes the old message as input and outputs the modified message $m' \leftarrow \text{MOD}(m)$. We write $\text{MOD}(\text{ADM}) \rightarrow \top/\perp$ to check if the intended modifications are allowed or not.
- $d \leftarrow \text{Vrfy}(\mu, \sigma, pk_{\text{Sig}}, pk_{\text{San}})$: On input of a message μ , a signature σ , a signer public key pk_{Sig} , and a sanitizer public key pk_{San} , the algorithm outputs a decision bit $d \in \{\top, \perp\}$.
- $\pi \leftarrow \text{Proof}(\mu, \sigma, \{(\mu_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{Sig}}, pk_{\text{San}})$: On input of a message μ , a signature σ , a set of additional message-signature pairs $\{(\mu_i, \sigma_i)\}_{i=1}^k$, signer secret key sk_{Sig} and public key pk_{Sig} , and the sanitizer public key pk_{San} , the algorithm outputs a proof $\pi \in \{0, 1\}^*$ or an error message \perp .
- $d \leftarrow \text{Judge}(\mu, \sigma, pk_{\text{Sig}}, pk_{\text{San}}, \pi)$: On input of a message μ , a signature σ , a signer public key pk_{Sig} , a sanitizer public key pk_{San} , and a proof π , the algorithm outputs a bit determining who generated the signature, i.e., $d \in \{\text{Sig}, \text{San}\}$, or returns an error message \perp .

We follow the correctness definition of Brzuska et al. [11, 12] as well as subsequent works that require that genuinely signed or sanitized messages are accepted, and a genuinely generated proof by the signer will lead the judge to determine the accountable party correctly. The formal definition can be found in Appendix B, Definition 72.

The definition for blocks admissible for modification, ADM , and description of desired modifications, MOD , also follows from [11]. MOD is a list of t binary values $\{0, 1\}^t$, where t is the block length of the message. The value at the i -th position of ADM determines whether block i is admissible for modification. $\text{ADM} \subseteq \mathbb{N} \times \mathcal{M}$ describes the desired modification. The first part of the tuple indicates the block to be modified while the second part is the new content.

4.1 Security Notions

Here we briefly describe the intuition for each of the usual security properties for sanitizable signature schemes. Observe that most of these properties were stated in [11, 12, 14, 16]. Furthermore, some later papers like [29] consider “stronger”

versions of the security properties while others like [12, 30, 23] introduce slightly weaker versions. In this work, we will adhere to the basic notions for security.

In the following list we will state all possible security notions for sanitizable signature schemes and give a brief and informal intuition on what they achieve. Formal definitions for all notions can be found in Appendix B.

Unforgeability No efficient adversary should be able to produce valid message-signature pairs that it has not seen before.

Immutability A malicious sanitizer should not be able to modify blocks that the signer has not set to be admissible for modification.

Signer Accountability A signer must not be able to craft a valid message-signature pair and a proof, such that the **Judge** algorithm will output **San**.

Sanitizer Accountability A sanitizer must not be able to craft a valid message-signature pair such that honest proofs generated by the signer will lead to the **Judge** algorithm outputting **Sig**.

Accountability Both sanitizer- and signer accountability hold.

Non-Interactive Public Accountability Given a valid pair of message and signature, a third party can correctly determine who created the signature with no additional information.

Transparency Detecting if a signature has been sanitized must be hard.

Privacy Given a sanitized message, it must be hard to recover any information about the original message’s content before being sanitized.

Unlinkability Two different sanitized messages cannot be identified as belonging to the same initial message.

Invisibility It must be hard to detect which parts of the message may be modified by the sanitizer.

In the literature, there exist several varieties of each of these properties, the use of which we try to limit in this paper for improved readability. To this end, we refer to the notion formally known as “proof-restricted transparency” simply as “transparency”. Definitions for strong/weak variants can be found in Appendix B.

4.2 Signer Accountability

For the main body we limit ourselves to the formal definition of signer accountability, as we later show a new attack on this notion in Theorem 46 for one of the constructions examined as part of this work.

Definition 37 (Signer-Accountability). *A sanitizable signature scheme SSS is (strongly) signer-accountable, if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{SSS, \mathcal{A}}^{(s)\text{sig-acc}}(\lambda)$ as described in Fig. 9, defined as*

$$\text{Adv}_{SSS, \mathcal{A}}^{(s)\text{sig-acc}}(\lambda) := \Pr \left[\text{Exp}_{SSS, \mathcal{A}}^{(s)\text{sig-acc}}(\lambda) = 1 \right]$$

is negligible in the security parameter λ .

$\text{Exp}_{\text{SSS}, \mathcal{A}}^{(\text{s})\text{sig-acc}}(\lambda)$ <hr style="border: 0.5px solid black;"/> $(pk_{\text{San}}, sk_{\text{San}}) \leftarrow_{\$} \text{KGen}_{\text{San}}(1^\lambda)$ $\mathcal{Q}_{\text{San}} \leftarrow \emptyset$ $(pk_{\text{Sig}}^*, \pi^*, \mu^*, \sigma^*) \leftarrow_{\$} \mathcal{A}^{\text{SanitO}}(pk_{\text{San}})$ $\text{if } (*, \mu^*, \sigma^*, pk_{\text{Sig}}^*) \notin \mathcal{Q}_{\text{San}}$ $\quad \wedge \text{Vrfy}(\mu^*, \sigma^*, pk_{\text{Sig}}^*, pk_{\text{San}}) = \top$ $\quad \wedge \text{Judge}(\mu^*, \sigma^*, pk_{\text{Sig}}^*, pk_{\text{San}}, \pi^*) = \text{San}$ $\quad \text{return } 1$ $\text{else return } 0$	$\text{SanitO}(sk_{\text{San}}, \cdot, \cdot, \cdot, \cdot)$ <hr style="border: 0.5px solid black;"/> $\text{with } (\mu, \text{MOD}, \sigma, pk_{\text{Sig}})$ $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$ $\mathcal{Q}_{\text{San}} \leftarrow \mathcal{Q}_{\text{San}} \cup \{(\mu, \mu', \sigma', pk_{\text{Sig}})\}$ $\text{return } (\mu', \sigma')$
---	--

Fig. 9: Game based security definition for signer accountability. Inclusion of the gray box yields strong signer accountability.

5 Lattice-Based Instantiations with CH

This section is devoted to the instantiation of the various SSS constructions with the chameleon hash function. For other types of constructions, see Appendices D and E. We explain the results of utilizing the chameleon hash function we constructed in Section 3.1 first. Due to the missing uniqueness of the chameleon hash function, the resulting sanitizable signature schemes lack important security features like unforgeability and accountability.⁶ Luckily, we can employ the transform by [15]. This transform takes a sanitizable signature scheme, which does not necessarily satisfy unforgeability, and a verifiable ring signature (VRS) to a sanitizable signature scheme that satisfies unforgeability and full accountability. In the present case, we make use of the VRS constructed in Section 3.2, to increase the security guarantees of the lattice-based SSSs.

Before we iterate through the collection of SSS constructions using chameleon hash functions, we explain the transform to ensure unforgeability and accountability from [15].

Theorem 38. *For a sanitizable signature scheme SSS_2 resulting from applying the transformation explained in [15, Fig. 5] to a sanitizable signature scheme SSS_1 and a verifiable ring signature VRS, the following implications hold: If SSS_1 is weakly immutable, then SSS_2 is immutable. If SSS_1 is weakly unlinkable, and VRS is strongly unforgeable, then SSS_2 is unlinkable. If SSS_1 is strongly invisible then SSS_2 is strongly invisible. If VRS is strongly accountable, then SSS_2 is strongly signer accountable. If VRS is strongly non-seizable, then SSS_2 is strongly sanitizer accountable. If VRS is anonymous and SSS_1 is strongly transparent, then SSS_2 is strongly transparent.*

⁶ Only in [3], sanitizer accountability has not been defined and has neither been attacked nor proven.

Finally, we note that standard instantiations of PRFs and PRGs are used in the protocols, without further specification. Post-quantum secure variants are given, for instance, in [25].

5.1 The ACMT05 Construction

The work by Ateniese et al. [3] marked the beginning of sanitizable signatures and gave the first construction of such. This first construction makes use of only two building blocks, a secure digital signature scheme Σ and a chameleon hash function CH which is strongly unforgeable. While the strongly unforgeable notion is not defined explicitly, we inferred that full collision-resistance and indistinguishability is sufficient.

The sanitizable signature scheme in [3] is stated to be *indistinguishable, unforgeable*, and satisfy *identical distribution of sanitized and original signatures*, [3, Section 4.4]. By the general relations between the notions, we have the following results.

Corollary 39. *The construction by Ateniese et al. in [3] instantiated using a DSS from {Dilithium, Falcon} and the chameleon hash function constructed in Fig. 7 satisfies unforgeability, signer accountability, immutability, transparency, and privacy.*

We note that sanitizer accountability and invisibility have been developed after the publication of the work by Ateniese et al. The corresponding construction has not been analyzed regarding this notion. Here, we take the easy road and apply Theorem 38, to additionally achieve sanitizer accountability.

Corollary 40. *Using the transform from Theorem 38 and combining the verifiable ring signature given in Fig. 8 with the construction by Ateniese et al. in [3] as instantiated in Corollary 39, we obtain a sanitizable signature that satisfies unforgeability, signer accountability, sanitizer accountability, immutability, transparency, and privacy.*

5.2 The LZCS16–1 Construction

Recall that [30] gives two constructions. One is based on accountable ring signatures, (cf. Appendix E) while the other uses rerandomizable tagging schemes which are constructed from tag-based trapdoor functions and double-trapdoor chameleon hash functions. While tag-based trapdoor functions are merely weaker versions of chameleon hash functions, and thus, can be instantiated from our lattice-based chameleon hash function given in Fig. 7, it is not clear whether double-trapdoor chameleon hash functions can be instantiated from lattices. Thus, as of now, the construction by Lai et al. using chameleon hash functions, cannot be instantiated from known lattice constructions.

5.3 The BCD⁺17 Construction

The construction by Beck et al. [6] uses a regular digital signature scheme, a *labeled* PKE, and a secure chameleon hash function, by which they mean indistinguishable, unique, and standard collision-resistant. Further, the sanitizable signature construction requires a PRG and a PRF.

As labeled PKEs are non-standard, we present a brief account including how to construct them from regular PKEs. In the construction, the labeled PKE provides strong invisibility of the sanitizable signature scheme, in which the signer public key is used as a label to prevent re-use of ciphertexts from different signer public keys. Further details including full definitions and proofs of the following can be found in Appendix A.5.

Labeled PKEs from regular PKEs. A *labeled PKE* (we refer to Appendix A.5) consists of a triple $(\text{KGen}_\tau, \text{Enc}_\tau, \text{Dec}_\tau)$, closely resembling the properties of regular PKEs. The key generation algorithm KGen_τ creates a key pair (pk, sk) . The probabilistic encryption algorithm Enc_τ additionally requires a label $\tau \in \{0, 1\}^n$, where the length n of the label depends on the security parameter λ . The decryption algorithm Dec_τ additionally requires the label τ . The (perfect) correctness of a labeled PKE is defined as $\text{Enc}_\tau(sk, \text{Dec}_\tau(pk, \mu, \tau), \tau) = \mu$ whenever $(pk, sk) \leftarrow \text{KGen}_\tau$ is honestly generated.

We explain now, how IND-CCA secure PKEs can be transformed to IND-CCA secure labeled PKEs with the proof provided in Appendix A.5. Given a PKE $\Pi = (\text{KGen}, \text{Enc}, \text{Dec})$, we define a labeled PKE $\Pi_\tau = (\text{KGen}_\tau, \text{Enc}_\tau, \text{Dec}_\tau)$ as follows. First, we note that for any label set $\{0, 1\}^n$ we can concatenate a label τ and a message μ as $\tau\|\mu$ so that both, the label and the message can be recovered from $\tau\|\mu$ by either projecting to the first n bits or removing them. Then, $\text{KGen}_\tau = \text{KGen}$ and $\text{Enc}_\tau(pk, \mu, \tau) \leftarrow \text{Enc}(pk, \tau\|\mu)$. Finally, the decryption Dec_τ on input (sk, c, τ) first runs $\text{Dec}(sk, c)$ to get $\tau'\|\mu'$. Then, Dec_τ checks if τ' coincides with the input label τ . If so, Dec_τ returns μ' . Otherwise, it returns \perp .

Instantiation of the SSS construction. The uniqueness, which our construction in Section 3.1 does not achieve, is required for unforgeability, signer accountability, and sanitizer accountability. So that we have the following result that follows from [6, Theorem 1] and the construction given in Section 3.1.

Corollary 41. *The SSS construction by Beck et al. [6] instantiated with a signature scheme among {Dilithium, Falcon}, the PKE Kyber, and the chameleon hash function given in Fig. 7, satisfies immutability, transparency, privacy, and strong invisibility.*

The missing unforgeability and signer and sanitizer accountability are added via the transform Theorem 38.

Corollary 42. *The transform in Theorem 38 applied with the VRS given in Fig. 8 to the construction by Beck in [6] as instantiated in Corollary 41, satisfies unforgeability, signer accountability, sanitizer accountability, immutability, transparency, privacy, and strong invisibility.*

5.4 The CDK⁺17 Construction

In [16], Camenisch et al. introduce a new notion called chameleon hash function with ephemeral trapdoors (CHET). They show how to generically construct CHET from chameleon hash functions such that the required security properties carry over. Thus, it is sufficient to have a chameleon hash function with the required properties. In total, the requirements are a correct, indistinguishable, unique, and standard collision-resistant chameleon hash function, a secure DSS, an IND-CPA secure PKE, PRGs, and PRFs, and an indistinguishable and standard collision-resistant CHET.

In their construction, Camenisch et al. use the uniqueness of the CHF in the proof of unforgeability and sanitizer accountability. Hence, from [16, Theorem 3] and taking into account the missing uniqueness, we get the following.

Corollary 43. *The construction by Camenisch et al. in [16] instantiated using a DSS from {Dilithium, Falcon}, the PKE Kyber, the chameleon hash function given in Fig. 7, which is used twice, for the chameleon hash function and the CHET, satisfies signer accountability, immutability, transparency, privacy, and invisibility.*

To ensure unforgeability and sanitizer accountability, we apply Theorem 38.

Corollary 44. *The transform in Theorem 38 applied with the VRS given in Fig. 8 to the construction by Camenisch et al. in [16] as instantiated in Corollary 43, satisfies unforgeability, signer accountability, sanitizer accountability, immutability, transparency, privacy, and invisibility.*

5.5 The BFF⁺09 Construction

Brzuska et al. [11] give a SSS construction that relies on a chameleon hash function but—unlike the constructions in [6] and [16]—they do not require the chameleon hash function to satisfy the uniqueness property. It therefore is an interesting candidate to be instantiated using the new construction in Section 3.1, without the need for the transform in Theorem 38. However, it turns out that an error in the security proof makes the signer accountability of the construction vulnerable to a generic attack. Therefore, this section is structured differently than the prior ones. First, we show that the security formalization of tagged chameleon hashes is unachievable and needs to be replaced in the construction by the tagged version of full collision resistance. Afterwards, we present the attack on signer accountability, which works irrespectively of the underlying building blocks. Still, the construction by Brzuska et al. can be instantiated with the chameleon hash function we constructed, merely missing the signer accountability. To achieve signer accountability as well, we can apply the transform in Theorem 38.

Idea of the Construction. We begin with a short description of the construction given in [11]. The full construction can be found in Appendix C.2.

First, the message is divided into blocks. Admissible blocks are hashed using a tagged chameleon hash function. The signer computes the tag as the output of a pseudorandom function with a random input value and provides this random input as proof for the accountability. It then signs the output of the tagged chameleon hash function. The sanitizer can use its private key to find collisions on the chameleon hash function, however it chooses the tag randomly. While this ensures sanitizer accountability, we show that signer-accountability for this construction is unachievable, contrary to the claim in [11]. Indeed, there are two distinct errors in the presentation of [11]. First, the notion of *collision-resistance under random-tagging attacks*, which is a strong version of collision-resistance for chameleon hash functions, is too strong and in fact, cannot be achieved by any chameleon hash function. The reason is the oracle being too strong and the winning condition too weak. We present the generic attack here in detail. Second, the claimed signer-accountability does not hold independently of the chameleon hash function's security properties. We present an outline of the generic attack at the end of this section and the detailed proof in Appendix F.3.

Infeasibility of Random Tagging. To show security, Brzuska et al. [11] rely on different properties for the different SSS security notions. For signer accountability, they rely on the *collision-resistance under random-tagging attacks* of the used chameleon hash function. This security notion was developed along with the SSS construction and we describe it in Fig. 10.

```

RndTagACH(λ)
-----
(pk, sk) ←$ CKGen(1λ)
(TAG, μ, r, TAG', μ', r') ←$ AAdaptO(pk)
return 1 if (TAG, μ) ≠ (TAG', μ')
  ∧ CH(pk, TAG, μ, r) = CH(pk, TAG', μ', r')
  ∧ {(TAG, μ), (TAG', μ')} ≠ {(TAGi, μi), (TAG'i, μ'i)}, for all i = 1, ..., q
  ∧ {(TAG, μ), (TAG', μ')} ≠ {(TAG'i, μ'i), (TAG'j, μ'j)}, for all i, j = 1, ..., q

      AdaptO(sk, TAGi, μi, ri, μ'i)
      -----
      TAG'i ←$ {0, 1}2n
      r'i ←$ CAdapt(sk, TAGi, μi, ri, TAG'i, μ'i)
      Store (TAGi, μi), (TAG'i, μ'i)
      return (TAG'i, r'i)

```

Fig. 10: Collision-resistance under random-tagging attacks from [11].

To exclude trivial wins, the game numbers the adversary's queries and stores them as (TAG_i, μ_i) and (TAG'_i, μ'_i) , for $i = 1, \dots, q$. An adversary wins, if it can

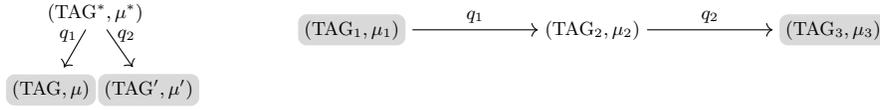


Fig. 11: Left: a trivial collision that is excluded in the security game. Right: a trivial collision that composes a generic attack on the security notion.

output a non-trivial collision $\text{CH}(pk, \text{TAG}, \mu, r) = \text{CH}(pk, \text{TAG}', \mu', r')$ with

$$\begin{aligned} & \{(\text{TAG}, \mu), (\text{TAG}', \mu')\} \neq \{(\text{TAG}_i, \mu_i), (\text{TAG}'_i, \mu'_i)\} \text{ for all } i, \\ & \text{and } \{(\text{TAG}, \mu), (\text{TAG}', \mu')\} \neq \{(\text{TAG}'_i, \mu'_i), (\text{TAG}'_j, \mu'_j)\} \text{ for all } i, j. \end{aligned}$$

The former condition prevents the adversary to output a collision that it obtains from a single query to its oracle $\text{Adapt}\mathcal{O}$. The latter prevents the adversary trivial transitive collisions: By querying the oracle $\text{Adapt}\mathcal{O}$ twice on the same input for different target messages, as depicted on the left of Fig. 11.

It turns out, that the two excluded types of trivial collisions do not suffice. In fact, the adversary can make an arbitrary query to the oracle $\text{Adapt}\mathcal{O}$ and then make a second query on the response of the first query. Then, all three tuples collide, but the first and third as output, are not defined as a trivial win. This is depicted on the right of Fig. 11. In conclusion, *no* chameleon hash function can satisfy this security notion as stated in the theorem below.

Theorem 45. *Collision-resistance under random-tagging attacks is unachievable for any chameleon hash function.*

Proof. The attack idea is depicted in Fig. 11. Given a public key pk we define the attacker against collision-resistance under random-tagging attacks as follows: First, the adversary picks distinct messages μ_1 , μ_2 , and μ_3 . Further it generates TAG_1 and r_1 . Then, the adversary makes two queries to the oracle: first, query $(\text{TAG}_1, \mu_1, r_1, \mu_2)$ to receive (TAG_2, r_2) , second, $(\text{TAG}_2, \mu_2, r_2, \mu_3)$ to receive (TAG_3, r_3) . Thus, we have a collision of the hash values of $(pk, \text{TAG}_1, \mu_1, r_1)$ and $(pk, \text{TAG}_3, \mu_3, r_3)$, the definition of the oracle. The adversary returning $(\text{TAG}_1, \mu_1, r_1)$ and $(\text{TAG}_3, \mu_3, r_3)$ wins the game. Indeed, the collision is valid, as (TAG_1, μ_1) was never an output (hence the collision is not of the second form), and (TAG_1, μ_1) and (TAG_3, μ_3) were not part of the same query. \square

Signer-Accountability. Theorem 45 leaves a gap in the construction of [11] as no instantiation achieves signer accountability. A natural question that arises is whether the construction is secure if one strengthens the requirements towards the underlying chameleon hash function by excluding the transitivity attack that we described above. Here, we answer this question in the negative by providing a generic attack against the signer accountability.

We show that a signer can put the blame on a message to the sanitizer even though the sanitizer did not sanitize to this message. On a high-level, the attack

works as follows. The signer prepares an arbitrary message containing two blocks $\mu_1 \parallel \mu_2$ such that only the second one is admissible. It then lets the sanitizer change the message to $\mu_1 \parallel \mu_2^*$. Finally, the adversary can create a signature on the message $\mu_1^* \parallel \mu_2^*$ by replacing the first message block. When questioned about the message $\mu_1^* \parallel \mu_2^*$ created by the signer, the judge will blame the sanitizer. This is stated formally in the following theorem, the proof is given in Appendix F.3.

Theorem 46. *There is an efficient adversary \mathcal{A} that breaks signer accountability of the BFF⁺09 construction with probability 1, using a single query to $\text{Sanit}\mathcal{O}$.*

Instantiations. Finally, we want to apply our constructions in Section 3 to the construction by Brzuska et al. To account for the unachievable security notion for tagged chameleon hash functions, we refer to Appendix C.1 for the notion of f-CR tagged chameleon hash functions, which can be constructed from a f-CRCHash as in Fig. 7. As uniqueness of the CHF is not required, we get the following result.

Corollary 47. *The construction by Brzuska et al. in [11] instantiated using a DSS from $\{\text{Dilithium}, \text{Falcon}\}$ and the tagged chameleon hash function deduced from the chameleon hash in Fig. 7 as explained in Appendix C.1 satisfies unforgeability, sanitizer accountability, immutability, transparency, and privacy.*

This follows from [11, Theorem 5.3], again noting that signer accountability does not hold. Using Theorem 38, we can increase the security to cover signer accountability.

Corollary 48. *The transform in Theorem 38 applied with the VRS given in Fig. 8 to the construction by Brzuska et al. in [11] as instantiated in Corollary 47, satisfies unforgeability, signer accountability, sanitizer accountability, immutability, transparency, and privacy.*

5.6 General Observations

Unfortunately, the efficiency of the proposed instantiations may be suboptimal for any real world use cases. For example, the signature of the construction of [16] instantiated securely with our building blocks and the transform of [15] contains an encryption of a chameleon hash trapdoor, i.e. an encryption of a G-trapdoor \mathbf{R} for each message block. On the positive side, we do not impose inefficient requirements such as a superpolynomial modulus and the reduction losses are not substantial.

Corollary 49. *There exists an SSS construction that has unforgeability, signer accountability, sanitizer accountability, immutability, transparency, privacy, and invisibility in the random oracle model, if LWE and LWR and SIS with polynomial modulus are hard.*

Furthermore, we expect our building blocks and thus the generic constructions of SSS to carry over to ring or module lattices in a straightforward fashion.

An additional drawback is the necessary inclusion of the transform of [15] for many SSS constructions, which incurs the cost of a NIZK. However, a secure construction without this transformation requires uniqueness of the chameleon hash function, which seems highly non-trivial for lattice-based constructions. Previous non lattice-based constructions used bijections to ensure that only one auxiliary value exists, thus even with a trapdoor an adversary cannot find a second auxiliary value to break uniqueness. Lattice-based trapdoored functions, especially ones used in conjunction with PreSample, have many different preimages for each image by design, which leads us to believe that a chameleon hash achieving uniqueness requires a fundamentally different approach to trapdoor its hash function.

References

- [1] Miklós Ajtai. “Generating Hard Instances of Lattice Problems (Extended Abstract)”. In: *28th ACM STOC*. ACM Press, May 1996, pp. 99–108. DOI: 10.1145/237814.237838.
- [2] Nabil Alkeilani Alkadri, Poulami Das, Andreas Erwig, Sebastian Faust, Juliane Krämer, Siavash Riahi, and Patrick Struck. “Deterministic Wallets in a Quantum World”. In: *ACM CCS 2020*. Ed. by Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna. ACM Press, Nov. 2020, pp. 1017–1031. DOI: 10.1145/3372297.3423361.
- [3] Giuseppe Ateniese, Daniel H. Chou, Breno de Medeiros, and Gene Tsudik. “Sanitizable Signatures”. In: *ESORICS 2005*. Ed. by Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann. Vol. 3679. LNCS. Springer, Berlin, Heidelberg, Sept. 2005, pp. 159–177. DOI: 10.1007/11555827_10.
- [4] Abhishek Banerjee, Chris Peikert, and Alon Rosen. “Pseudorandom Functions and Lattices”. In: *EUROCRYPT 2012*. Ed. by David Pointcheval and Thomas Johansson. Vol. 7237. LNCS. Springer, Berlin, Heidelberg, Apr. 2012, pp. 719–737. DOI: 10.1007/978-3-642-29011-4_42.
- [5] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. “More Efficient Commitments from Structured Lattice Assumptions”. In: *SCN 18*. Ed. by Dario Catalano and Roberto De Prisco. Vol. 11035. LNCS. Springer, Cham, Sept. 2018, pp. 368–385. DOI: 10.1007/978-3-319-98113-0_20.
- [6] Michael Till Beck, Jan Camenisch, David Derler, Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. “Practical Strongly Invisible and Strongly Accountable Sanitizable Signatures”. In: *ACISP 17, Part I*. Ed. by Josef Pieprzyk and Suriadi Suriadi. Vol. 10342. LNCS. Springer, Cham, July 2017, pp. 437–452. DOI: 10.1007/978-3-319-60055-0_23.
- [7] David Bernhard, Georg Fuchsbauer, Essam Ghadafi, Nigel P. Smart, and Bogdan Warinschi. “Anonymous attestation with user-controlled linkability”. In: *Int. J. Inf. Sec.* 12.3 (2013), pp. 219–249. DOI: 10.1007/s10207-013-0191-z. URL: <https://doi.org/10.1007/s10207-013-0191-z>.

- [8] Ward Beullens, Samuel Dobson, Shuichi Katsumata, Yi-Fu Lai, and Federico Pintore. “Group Signatures and More from Isogenies and Lattices: Generic, Simple, and Efficient”. In: *EUROCRYPT 2022, Part II*. Ed. by Orr Dunkelman and Stefan Dziembowski. Vol. 13276. LNCS. Springer, Cham, 2022, pp. 95–126. DOI: 10.1007/978-3-031-07085-3_4.
- [9] Johannes Blömer, Jan Bobolz, and Laurens Porzenheim. “A Generic Construction of an Anonymous Reputation System and Instantiations from Lattices”. In: *Advances in Cryptology - ASIACRYPT 2023 - 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4-8, 2023, Proceedings, Part II*. Ed. by Jian Guo and Ron Steinfeld. Vol. 14439. Lecture Notes in Computer Science. Springer, 2023, pp. 418–452. DOI: 10.1007/978-981-99-8724-5_13. URL: https://doi.org/10.1007/978-981-99-8724-5_13.
- [10] Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Alessandro Sorniotti. “A Framework for Practical Anonymous Credentials from Lattices”. In: *CRYPTO 2023, Part II*. Ed. by Helena Handschuh and Anna Lysyanskaya. Vol. 14082. LNCS. Springer, Cham, Aug. 2023, pp. 384–417. DOI: 10.1007/978-3-031-38545-2_13.
- [11] Christina Brzuska, Marc Fischlin, Tobias Freudenreich, Anja Lehmann, Marcus Page, Jakob Schelbert, Dominique Schröder, and Florian Volk. “Security of Sanitizable Signatures Revisited”. In: *PKC 2009*. Ed. by Stanislaw Jarecki and Gene Tsudik. Vol. 5443. LNCS. Springer, Berlin, Heidelberg, Mar. 2009, pp. 317–336. DOI: 10.1007/978-3-642-00468-1_18.
- [12] Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. “Unlinkability of Sanitizable Signatures”. In: *PKC 2010*. Ed. by Phong Q. Nguyen and David Pointcheval. Vol. 6056. LNCS. Springer, Berlin, Heidelberg, May 2010, pp. 444–461. DOI: 10.1007/978-3-642-13013-7_26.
- [13] Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. “Sanitizable signatures: how to partially delegate control for authenticated data”. In: *BIOSIG 2009*. Bonn: Gesellschaft für Informatik e.V., 2009, pp. 117–128. ISBN: 978-3-88579-249-1.
- [14] Christina Brzuska, Henrich C. Pöhls, and Kai Samelin. “Non-interactive Public Accountability for Sanitizable Signatures”. en. In: *Public Key Infrastructures, Services and Applications*. Ed. by Sabrina De Capitani di Vimercati and Chris Mitchell. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2013, 178–193. ISBN: 978-3-642-40012-4. DOI: 10.1007/978-3-642-40012-4_12.
- [15] Xavier Bultel, Pascal Lafourcade, Russell W. F. Lai, Giulio Malavolta, Dominique Schröder, and Sri Aravinda Krishnan Thyagarajan. “Efficient Invisible and Unlinkable Sanitizable Signatures”. In: *PKC 2019, Part I*. Ed. by Dongdai Lin and Kazue Sako. Vol. 11442. LNCS. Springer, Cham, Apr. 2019, pp. 159–189. DOI: 10.1007/978-3-030-17253-4_6.
- [16] Jan Camenisch, David Derler, Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. “Chameleon-Hashes with Ephemeral Trap-

- doors - And Applications to Invisible Sanitizable Signatures”. In: *PKC 2017, Part II*. Ed. by Serge Fehr. Vol. 10175. LNCS. Springer, Berlin, Heidelberg, Mar. 2017, pp. 152–182. DOI: 10.1007/978-3-662-54388-7_6.
- [17] Sébastien Canard and Amandine Jambert. “On Extended Sanitizable Signature Schemes”. In: *CT-RSA 2010*. Ed. by Josef Pieprzyk. Vol. 5985. LNCS. Springer, Berlin, Heidelberg, Mar. 2010, pp. 179–194. DOI: 10.1007/978-3-642-11925-5_13.
- [18] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. “Bonsai Trees, or How to Delegate a Lattice Basis”. In: *EUROCRYPT 2010*. Ed. by Henri Gilbert. Vol. 6110. LNCS. Springer, Berlin, Heidelberg, 2010, pp. 523–552. DOI: 10.1007/978-3-642-13190-5_27.
- [19] Poulami Das, Andreas Erwig, Michael Meyer, and Patrick Struck. “Efficient Post-Quantum Secure Deterministic Threshold Wallets from Isogenies”. In: *ASIACCS 24*. Ed. by Jianying Zhou, Tony Q. S. Quek, Debin Gao, and Alvaro A. Cárdenas. ACM Press, July 2024. DOI: 10.1145/3634737.3657008.
- [20] David Derler, Kai Samelin, and Daniel Slamanig. “Bringing Order to Chaos: The Case of Collision-Resistant Chameleon-Hashes”. In: *PKC 2020, Part I*. Ed. by Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas. Vol. 12110. LNCS. Springer, Cham, May 2020, pp. 462–492. DOI: 10.1007/978-3-030-45374-9_16.
- [21] Muhammed F. Esgin, Ron Steinfeld, Dongxi Liu, and Sushmita Ruj. “Efficient Hybrid Exact/Relaxed Lattice Proofs and Applications to Rounding and VRFs”. In: *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part V*. Ed. by Helena Handschuh and Anna Lysyanskaya. Vol. 14085. Lecture Notes in Computer Science. Springer, 2023, pp. 484–517. DOI: 10.1007/978-3-031-38554-4_16. URL: https://doi.org/10.1007/978-3-031-38554-4_16.
- [22] Marc Fischlin and Patrick Harasser. “Invisible Sanitizable Signatures and Public-Key Encryption are Equivalent”. In: *ACNS 18 International Conference on Applied Cryptography and Network Security*. Ed. by Bart Preneel and Frederik Vercauteren. Vol. 10892. LNCS. Springer, Cham, July 2018, pp. 202–220. DOI: 10.1007/978-3-319-93387-0_11.
- [23] Nils Fleischhacker, Johannes Krupp, Giulio Malavolta, Jonas Schneider, Dominique Schröder, and Mark Simkin. “Efficient Unlinkable Sanitizable Signatures from Signatures with Re-randomizable Keys”. In: *PKC 2016, Part I*. Ed. by Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang. Vol. 9614. LNCS. Springer, Berlin, Heidelberg, Mar. 2016, pp. 301–330. DOI: 10.1007/978-3-662-49384-7_12.
- [24] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. “Trapdoors for hard lattices and new cryptographic constructions”. In: *40th ACM STOC*. Ed. by Richard E. Ladner and Cynthia Dwork. ACM Press, May 2008, pp. 197–206. DOI: 10.1145/1374376.1374407.

- [25] Christian Janson and Patrick Struck. “Sponge-Based Authenticated Encryption: Security Against Quantum Attackers”. In: *Post-Quantum Cryptography - 13th International Workshop, PQCrypto 2022*. Ed. by Jung Hee Cheon and Thomas Johansson. Springer, Cham, Sept. 2022, pp. 230–259. DOI: 10.1007/978-3-031-17234-2_12.
- [26] Shuichi Katsumata. “A New Simple Technique to Bootstrap Various Lattice Zero-Knowledge Proofs to QROM Secure NIZKs”. In: *CRYPTO 2021, Part II*. Ed. by Tal Malkin and Chris Peikert. Vol. 12826. LNCS. Virtual Event: Springer, Cham, Aug. 2021, pp. 580–610. DOI: 10.1007/978-3-030-84245-1_20.
- [27] Marek Klonowski and Anna Lauks. “Extended Sanitizable Signatures”. In: *ICISC 06*. Ed. by Min Surp Rhee and Byoungcheon Lee. Vol. 4296. LNCS. Springer, Berlin, Heidelberg, 2006, pp. 343–355. DOI: 10.1007/11927587_28.
- [28] Hugo Krawczyk and Tal Rabin. “Chameleon Signatures”. In: *NDSS 2000*. The Internet Society, Feb. 2000.
- [29] Stephan Krenn, Kai Samelin, and Dieter Sommer. “Stronger Security for Sanitizable Signatures”. In: *DPM 2015*. Ed. by Joaquín García-Alfaro, Guillermo Navarro-Arribas, Alessandro Aldini, Fabio Martinelli, and Neeraj Suri. LNCS. Springer, 2015.
- [30] Russell W. F. Lai, Tao Zhang, Sherman S. M. Chow, and Dominique Schröder. “Efficient Sanitizable Signatures Without Random Oracles”. In: *ESORICS 2016, Part I*. Ed. by Ioannis G. Askoxylakis, Sotiris Ioannidis, Sokratis K. Katsikas, and Catherine A. Meadows. Vol. 9878. LNCS. Springer, Cham, Sept. 2016, pp. 363–380. DOI: 10.1007/978-3-319-45744-4_18.
- [31] Yiming Li and Shengli Liu. “Tagged Chameleon Hash from Lattices and Application to Redactable Blockchain”. In: *PKC 2024, Part II*. Ed. by Qiang Tang and Vanessa Teague. Vol. 14603. LNCS. Springer, Cham, Apr. 2024, pp. 288–320. DOI: 10.1007/978-3-031-57725-3_10.
- [32] Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. “Zero-Knowledge Arguments for Lattice-Based Accumulators: Logarithmic-Size Ring Signatures and Group Signatures Without Trapdoors”. In: *EUROCRYPT 2016, Part II*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9666. LNCS. Springer, Berlin, Heidelberg, May 2016, pp. 1–31. DOI: 10.1007/978-3-662-49896-5_1.
- [33] San Ling, Khoa Nguyen, Huaxiong Wang, and Yanhong Xu. “Constant-Size Group Signatures from Lattices”. In: *PKC 2018, Part II*. Ed. by Michel Abdalla and Ricardo Dahab. Vol. 10770. LNCS. Springer, Cham, Mar. 2018, pp. 58–88. DOI: 10.1007/978-3-319-76581-5_3.
- [34] Vadim Lyubashevsky and Ngoc Khanh Nguyen. “BLOOM: Bimodal Lattice One-out-of-Many Proofs and Applications”. In: *ASIACRYPT 2022, Part IV*. Ed. by Shweta Agrawal and Dongdai Lin. Vol. 13794. LNCS. Springer, Cham, Dec. 2022, pp. 95–125. DOI: 10.1007/978-3-031-22972-5_4.

- [35] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. “Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General”. In: *CRYPTO 2022, Part II*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Vol. 13508. LNCS. Springer, Cham, Aug. 2022, pp. 71–101. DOI: 10.1007/978-3-031-15979-4_3.
- [36] Daniele Micciancio and Chris Peikert. “Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller”. In: *EUROCRYPT 2012*. Ed. by David Pointcheval and Thomas Johansson. Vol. 7237. LNCS. Springer, Berlin, Heidelberg, Apr. 2012, pp. 700–718. DOI: 10.1007/978-3-642-29011-4_41.
- [37] Daniele Micciancio and Oded Regev. “Worst-Case to Average-Case Reductions Based on Gaussian Measures”. In: *45th FOCS*. IEEE Computer Society Press, Oct. 2004, pp. 372–381. DOI: 10.1109/FOCS.2004.72.
- [38] *Module-Lattice-Based Digital Signature Standard*. National Institute of Standards and Technology, NIST FIPS PUB 204, U.S. Department of Commerce. Aug. 2024.
- [39] Sunoo Park and Adam Sealfon. “It Wasn’t Me! - Repudiability and Claimability of Ring Signatures”. In: *CRYPTO 2019, Part III*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11694. LNCS. Springer, Cham, Aug. 2019, pp. 159–190. DOI: 10.1007/978-3-030-26954-8_6.
- [40] Kai Samelin and Daniel Slamanig. “Policy-Based Sanitizable Signatures”. In: *CT-RSA 2020*. Ed. by Stanislaw Jarecki. Vol. 12006. LNCS. Springer, Cham, Feb. 2020, pp. 538–563. DOI: 10.1007/978-3-030-40186-3_23.
- [41] Adi Shamir and Yael Tauman. “Improved Online/Offline Signature Schemes”. In: *CRYPTO 2001*. Ed. by Joe Kilian. Vol. 2139. LNCS. Springer, Berlin, Heidelberg, Aug. 2001, pp. 355–367. DOI: 10.1007/3-540-44647-8_21.

A Additional Preliminaries

This section introduces additional details relevant to understanding proofs in detail, but were left out of the main body due to their length.

A.1 Digital Signature Scheme

A signature scheme is a public key based scheme to generate unforgeable signatures over arbitrary messages. Valid signatures can be generated only by a party with access to the secret key, while any party with access to the public key can validate signatures.

Definition 50. A signature scheme Sig consists of three algorithms $\Sigma = (\text{KGen}, \text{Sign}, \text{Vrfy})$.

$(pk, sk) \leftarrow_{\$} \text{KGen}(1^\lambda)$: This probabilistic key generation algorithm generates a key pair for the signature scheme. It takes the security parameter as input and outputs a key pair (pk, sk) .

$\sigma \leftarrow_{\$} \text{Sign}(sk, \mu)$: This potentially probabilistic signing algorithm takes as input a secret key sk and some message m , and outputs a signature σ .

$b \leftarrow \text{Vrfy}(pk, \sigma, \mu)$: This deterministic algorithm takes as input a public key pk , a signature σ and a message m , and outputs a boolean value of either 1 or 0, denoting a valid or invalid signature under the provided public key.

Definition 51. A signature scheme Sig is correct, iff for any security parameter $\lambda \in \mathbb{N}$, all generated key pairs $(pk, sk) \leftarrow_{\$} \text{KGen}(1^\lambda)$, all messages $\mu \in \mathcal{M}$ and all signatures $\sigma \leftarrow_{\$} \text{Sign}(sk, \mu)$, it holds that $\Pr[\text{Vrfy}(pk, \sigma, \mu) = \top] = 1$.

A.2 Public-key Encryption Scheme

A public-key encryption scheme is a public key based scheme which allows parties to securely communicate over an insecure channel. Messages can be encrypted under a public key and decrypted only by the party in possession of the corresponding secret key.

Definition 52. A public key encryption scheme Π consists of three algorithms: $\Pi = (\text{KGen}, \text{Enc}, \text{Dec})$

$(pk, sk) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$: This probabilistic key generation algorithm generates a key pair for the signature scheme. It takes no input and outputs a key pair (pk, sk) .

$c \leftarrow_{\$} \text{Enc}(sk, \mu)$: This potentially probabilistic encryption algorithm takes as input a secret key sk and some message μ , and outputs a ciphertext c .

$\mu \leftarrow \text{Dec}(pk, c)$: This deterministic algorithm takes as input a public key pk , and a ciphertext c . It outputs either the decryption of the provided ciphertext under the given secret key or \perp .

Definition 53. A public-key encryption scheme $\Pi = (\text{KGen}, \text{Enc}, \text{Dec})$ is correct, if $\Pr[\mu' \neq \mu : (pk, sk) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda), c \leftarrow_{\$} \text{Enc}(pk, \mu), \mu' \leftarrow \text{Dec}(sk, c)] = 0$.

A.3 Chameleon Hash Function Uniqueness

While we claimed the construction in Section 3.1 does not have uniqueness, we did not yet define the security game. It should be hard for an adversary to find a message μ and two values r, r' such that μ hashes to the same value with both r and r' . The definition of the security game can be found in Fig. 12.

```

CHUniqueH,A(λ)
-----
(pkch, h, μ, r, r') ←$ A(1λ)
if CHashCheck(pkch, h, μ, r) = 1 ∧ CHashCheck(pkch, h, μ, r') = 1 ∧ r ≠ r'
    return 1
return 0

```

Fig. 12: Uniqueness security game.

Definition 54. *We say that a chameleon hash function CH has uniqueness, if there exists a negligible function $\text{negl}(\lambda)$ such that for all ppt adversaries \mathcal{A} we have that*

$$\Pr[\text{CHUnique}_{\text{CH},\mathcal{A}}(\lambda) = 1] \leq \text{negl}(\lambda).$$

A.4 Additional Results concerning Lattices and Discrete Gaussians

To prove the security of our fully collision-resistant hash function of Section 3.1, we need some additional definitions, theorems and lemmas about lattices and discrete Gaussians.

Theorem 55 ([37]). *For any n -dimensional lattice Λ , vector $\mathbf{c} \in \mathbb{R}^n$, and reals $0 < \epsilon < 1, s > \eta_\epsilon(\Lambda)$, we have*

$$\Pr_{\mathbf{x} \leftarrow D_{\Lambda, s, \mathbf{c}}} [\|\mathbf{x} - \mathbf{c}\| > s\sqrt{n}] \leq \frac{1 + \epsilon}{1 - \epsilon} \cdot 2^{-n}.$$

Here, $\eta_\epsilon(\Lambda)$ describes the smoothing parameter of the lattice Λ [37].

Theorem 56 ([24]). *Let n and q be positive integers with q prime, and let $m \geq 2n \log q$. Then for all but a $2q^{-n}$ fraction of all $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and for any $s \geq \omega(\sqrt{\log m})$, the distribution of the syndrome $\mathbf{u} = \mathbf{A}\mathbf{e} \bmod q$ is statistically close to uniform over \mathbb{Z}_q^n , where $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, s}$.*

Lemma 57 ([36]). *For $\mathbf{R} \leftarrow D_{\mathbb{Z}^n, s}^m$ we know that there is a constant $C > 0$ such that for any $t \geq 0$ we have that $s_1(\mathbf{R}) \leq C \cdot s \cdot (\sqrt{n} + \sqrt{m} + t)$ with overwhelming probability.*

Heuristically, $C < 1$ so we ignore the factor.

A.5 Labeled PKEs

The SSS construction in [6] requires the use of labeled PKEs to ensure strong invisibility. In Section 5.3, we presented a brief outline of labeled PKEs, how these are defined, and how one can generically transform (regular) PKEs to labeled PKEs such that the IND-CCA security of the former translates to the IND-CCA security of the latter.

Here, we explain more details about labeled PKEs, the definition of IND-CCA security of labeled PKEs, and we present the transform and the formal reduction from regular PKEs to labeled PKEs. The definitions are based on [6]. Intuitively, labels are circumstantial data, under which the encryption and decryption algorithms are run. While it is assumed in [6] that the label set can be unbounded, the particular use-case shows that the label will have a specific form, thus, imposing the restriction that the label set is the set of binary strings of a fixed length, which depends on the security level, causes no harm.

Definition 58. *A labeled PKE is a triple $(\text{KGen}_\tau, \text{Enc}_\tau, \text{Dec}_\tau)$ such that*

KGen_τ *Takes a security parameter λ as input and creates key pairs (sk, pk)*

Enc_τ *Takes a public key pk , a message μ , and a label τ , and outputs a ciphertext c*

Dec_τ *Takes a secret key sk , a ciphertext c , and a label τ , and outputs a message μ .*

The labeled PKE is (perfectly) correct, if for any key pair (sk, pk) generated using the algorithm KGen_τ , any message μ and any label τ , the equality

$$\text{Dec}_\tau(sk, \text{Enc}_\tau(pk, \mu, \tau), \tau) = \mu$$

holds. The labeled PKE is ϵ -correct, if

$$\mathbb{P}(\text{Dec}_\tau(sk, \text{Enc}_\tau(pk, \mu, \tau), \tau) = \mu) > \epsilon,$$

where the probability is taken over the probability of the key generation, encryption, choices of the message, and choices of the label.

The security notion that is interesting for the purpose of this work is IND-CCA security, which we define next.

Definition 59. *Let Π_τ be a labeled PKE. Then, Π_τ achieves IND-CCA security, if for all qpt adversaries \mathcal{A} , the advantage in winning the game $\text{IND-CCA}(\lambda)$ as described in Fig. 13, which is given as*

$$\mathcal{A}_{\Pi_\tau}^{\text{IND-CCA}}(\mathcal{A}),$$

is negligible in the security parameter.

IND-CCA $_{\Pi_\tau, \mathcal{A}}(\lambda)$	
(sk, pk) \leftarrow KGen $_\tau(\lambda)$	
b \leftarrow $\{0, 1\}$	
(μ ₀ , μ ₁ , τ*, state) \leftarrow $\mathcal{A}_0^{\text{Dec}\mathcal{O}_\tau}(pk)$	
c* \leftarrow Enc $_\tau(pk, \mu_b, \tau^*)$	
b' \leftarrow $\mathcal{A}_1^{\text{Dec}\mathcal{O}'_\tau}(pk, c^*, \text{state})$	
if b = b'	
return 1	
return 0	
Dec $\mathcal{O}_\tau(sk, \cdot, \cdot)$ with (c, τ)	Dec $\mathcal{O}'_\tau(sk, \cdot, \cdot)$ with (c, τ)
return Dec $_\tau(sk, c, \tau)$	if c = c*
	return \perp
	return Dec$_\tau(sk, c, \tau)$

Fig. 13: IND-CCA security game for a labeled PKE $\Pi_\tau = (\text{KGen}_\tau, \text{Enc}_\tau, \text{Dec}_\tau)$.

Labeled PKEs can be easily constructed from regular PKEs. For this, we assume that the label set is given as $\{0, 1\}^n$ for some n that depends on the security parameter.⁷ Then, for any label τ and any message μ , it is possible to recover both from the concatenation $\tau \parallel \mu$. This is the essential assumption on the behavior of the label used in the decryption in the following transform.

Transform 60 *Given a regular PKE $\Pi = (\text{KGen}, \text{Dec}, \text{Enc})$, the labeled PKE $\Pi_\tau = (\text{KGen}_\tau, \text{Enc}_\tau, \text{Dec}_\tau)$ is defined as follows:*

KGen $_\tau$ *The key generation is exactly KGen*
Enc $_\tau$ *On input (pk, μ, τ) , returns Enc $(pk, \tau \parallel \mu)$*
Dec $_\tau$ *On input (sk, c, τ) , runs Enc with input (sk, c) to receive $\tau' \parallel \mu'$. Then, checks whether $\tau = \tau'$ holds. If not, returns \perp , otherwise, it returns μ' .*

Proposition 61. *Let Π be a PKE and Π_τ the associated labeled PKE under Transform 60. For any adversary \mathcal{A} against IND-CCA of Π_τ , there exists an adversary \mathcal{B} against IND-CCA of Π such that*

$$\mathcal{A}_{\Pi_\tau}^{\text{IND-CCA}}(\mathcal{A}) \leq \mathcal{A}_\Pi^{\text{IND-CCA}}(\mathcal{B}).$$

Proof. The reduction \mathcal{B} needs to simulate the two phases of the IND-CCA game of Π_τ . First, \mathcal{B} receives a public key as input, which it forwards to \mathcal{A} . On decryption queries by \mathcal{A} with input (c, τ) , where c is a ciphertext and τ is a label, \mathcal{B} sends c

⁷ The definition in [15] allows labels to be bit strings of arbitrary length. For the given construction, however, labels will contain the public key which size depends on the security parameter.

to its own decryption oracle. When \mathcal{B} receives $\tau' \parallel \mu'$ as response, it checks whether $\tau = \tau'$ holds. If so, \mathcal{B} returns μ' to \mathcal{A} , otherwise \perp . The first phase finishes with \mathcal{A} returning $(\mu_0, \mu_1, \tau^*, \text{state})$ and \mathcal{B} returns $(\tau^* \parallel \mu_0, \tau^* \parallel \mu_1, \text{state} \parallel \tau^*)$. Note that the reduction explicitly has access to τ^* in the second phase.

The second phase begins with \mathcal{B} receiving c_b , which—by the messages outputted after the first phase—is either a ciphertext for $\tau^* \parallel \mu_0$ or $\tau^* \parallel \mu_1$, and $\text{state} \parallel \tau^*$. The reduction forwards (c_b, state) to \mathcal{A} . Queries to the decryption oracle by \mathcal{A} are of the form (c, τ) . The reduction \mathcal{B} forwards c to its decryption oracle to receive $\tau' \parallel \mu'$ and returns μ' to \mathcal{A} , if $\tau' = \tau$. Note that \mathcal{B} may receive \perp from its decryption oracle. In this case, \mathcal{B} returns \perp to \mathcal{A} . For \mathcal{B} one inadmissible query is c_b itself, while for \mathcal{A} , an inadmissible query is (c_b, τ^*) . A priori, \mathcal{A} can query (c_b, τ) for $\tau \neq \tau^*$. However, as the correct decryption of c_b would yield $\tau^* \parallel \mu_b$, the transformed labeled PKE would return \perp , as the labels do not match. Thus, \mathcal{B} can return \perp whenever a query (c_b, \cdot) is made by \mathcal{A} .

Finally, \mathcal{A} returns a guess b' and wins, if and only if $\text{Dec}_\tau(sk, c_b) = \mu_{b'}$, or, in other words, $\text{Dec}(sk, c_b) = \tau^* \parallel \mu_{b'}$. Hence, \mathcal{B} returning b' wins, if and only if \mathcal{A} wins. \square

B Formal Details about the Security Notions and Correctness of SSS

This section will give formal definitions for all security notions as well as the relations between them.

B.1 Security Notions

Unforgeability Intuitively, unforgeability requires that no adversary should be able to produce a valid message-signature pair that it has not seen before. Note that a valid output can either be signed directly or be the result of sanitizing another unrelated message. It is up to the attacker to break either the signing or the sanitizing part, correspondingly it gets access to both oracles. To exclude trivial wins, the game checks that the output message was not used in a query to either the signing or sanitizing oracle.

Definition 62 (Unforgeability). *A sanitizable signature scheme SSS is (strongly) unforgeable, if for all qpt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS}, \mathcal{A}}^{\text{unf}}(\lambda)$ as described in Fig. 14, defined as*

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

is negligible in the security parameter λ .

$\text{Exp}_{\text{SSS}, \mathcal{A}}^{\text{unf}}(\lambda)$	$\text{SignO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot)$ with $(\mu, \text{ADM}, pk_{\text{San}})$
$(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$	$\sigma \leftarrow \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$
$(pk_{\text{San}}, sk_{\text{San}}) \leftarrow_{\$} \text{KGen}_{\text{San}}(1^\lambda)$	$\mathcal{Q}_{\text{Sig}} \leftarrow \mathcal{Q}_{\text{Sig}} \cup (\mu, \text{ADM}, pk_{\text{San}}, \sigma)$
$\mathcal{Q}_{\text{Sig}} \leftarrow \emptyset, \mathcal{Q}_{\text{San}} \leftarrow \emptyset$	return σ
$(\mu^*, \sigma^*) \leftarrow_{\$} \mathcal{A}^{\text{SignO}, \text{SanitO}, \text{ProofO}}(pk_{\text{Sig}}, pk_{\text{San}})$	$\text{SanitO}(sk_{\text{San}}, \cdot, \cdot, \cdot, \cdot)$
if $\text{Vrfy}(\mu^*, \sigma^*, pk_{\text{Sig}}, pk_{\text{San}}) = \top$	<hr style="border: 0.5px solid black;"/> with $(\mu, \text{MOD}, \sigma, pk_{\text{Sig}})$
$\wedge (\mu^*, *, pk_{\text{San}}, \sigma^*) \notin \mathcal{Q}_{\text{Sig}}$	$(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$
$\wedge (\mu^*, *, \sigma^*, pk_{\text{Sig}}) \notin \mathcal{Q}_{\text{San}}$	$\mathcal{Q}_{\text{San}} \leftarrow \mathcal{Q}_{\text{San}} \cup \{(\mu, \mu', \sigma', pk_{\text{Sig}})\}$
return 1 else return 0	return (μ', σ')
	$\text{ProofO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot, \cdot)$
	<hr style="border: 0.5px solid black;"/> with $(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, pk_{\text{San}})$
	$\pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{San}})$
	return π

Fig. 14: Game based security definition for (strong) unforgeability. Inclusion of the gray box yields strong unforgeability.

(weak) Immutability Intuitively, a malicious sanitizer must not be able to modify blocks that are not admissible for modification. The adversary is only given the public key of a signer and has to output a sanitizer public key (which it can generate honestly or maliciously) and a valid message signature pair. There are two options for the adversary to succeed: Either it has never queried the signing oracle with the sanitizer public key it has output. This would mean that there should not exist any valid signatures which verify under the output public key. The other option is to output a message that none of its queries to the signing oracle can be sanitized to. In other words, it has output an inadmissible modification, which still verifies.

A scheme achieves weak immutability if the adversary is not allowed to query the Sign oracle under a message that is admissible to be modified to its output, but a different sanitizer key.

Definition 63 (Immutability). *A sanitizable signature scheme SSS is (weakly) immutable, if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS}, \mathcal{A}}^{\text{imm}}(\lambda)$ as described in Fig. 15, defined as*

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

is negligible in the security parameter λ .

Privacy Intuitively, no adversary should be able to recover any information about the content of the original message before it was sanitized. The adversary is given access to a sign oracle, a sanitation oracle and a proof oracle (as the proof might leak information about the signer secret key) and a left-or-right oracle. Before

$\text{Exp}_{\text{SSS}, \mathcal{A}}^{(\text{w})\text{imm}}(\lambda)$	$\text{Sign}\mathcal{O}(sk_{\text{Sig}}, \cdot, \cdot, \cdot)$ with $(\mu, \text{ADM}, pk_{\text{San}})$
$(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$	$\sigma \leftarrow \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$
$\mathcal{Q}_{\text{Sig}} \leftarrow \emptyset$	$\mathcal{Q}_{\text{Sig}} \leftarrow \mathcal{Q}_{\text{Sig}} \cup (\mu, \text{ADM}, pk_{\text{San}}, \sigma)$
$(pk_{\text{San}}^*, \mu^*, \sigma^*) \leftarrow_{\$} \mathcal{A}^{\text{Sign}\mathcal{O}, \text{Proof}\mathcal{O}}(pk_{\text{Sig}})$	return σ
if $\text{Vrfy}(\mu^*, \sigma^*, pk_{\text{Sig}}, pk_{\text{San}}^*) = \top$	$\text{Proof}\mathcal{O}(sk_{\text{Sig}}, \cdot, \cdot, \cdot, \cdot)$
$\wedge (\forall (\mu_i, \text{ADM}_i, pk_{\text{San}}, *) \in \mathcal{Q}_{\text{Sig}} :$	$\text{with } (\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, pk_{\text{San}})$
$\mu^* \notin \{\text{MOD}(\mu_i) \mid \text{MOD}(\text{ADM}_i) = \top\}$	$\pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{San}})$
$\vee pk_{\text{San}}^* \neq pk_{\text{San}})$	return π
return 1	
return 0	

Fig. 15: Game based security definition for (weak) immutability. Omission of the gray box yields weak immutability.

making its decision, the adversary may observe an arbitrary amount of signing or sanitizing queries under the challenge public keys. The adversary wins if it can decide which of the two input messages are sanitized by the left-or-right oracle.

Definition 64 (Privacy). *A sanitizable signature scheme SSS is private, if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS}, \mathcal{A}}^{\text{PRV}}(\lambda)$ as described in Fig. 16, defined as*

$$\text{Adv}_{\text{SSS}, \mathcal{A}}^{\text{PRV}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{SSS}, \mathcal{A}}^{\text{PRV}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

is negligible in the security parameter λ .

Unlinkability Intuitively, no adversary given a sanitized message-signature pair should be able to identify the original message-signature pair that has been sanitized. In other words, it must be infeasible to link the given sanitized document to the original signed document. As a consequence, two different sanitized messages cannot be identified as belonging to the same initial message. This is a strictly stronger requirement than privacy and also modeled using a similar left-or-right oracle, but with more control for the adversary. Unlike in the privacy game, the unlinkability adversary can see the signatures before sanitation occurs. The adversary can provide two sets of message-signature pairs with identical admissible blocks (because the scheme might not be invisible) as well as modification instructions that map them to the same resulting message. Depending on the secret bit b , the oracle will sanitize and return either the left or the right message. To succeed, the adversary must correctly determine which of the input messages is being sanitized by the oracle.

In the literature, first introduced in [15], there exists a weaker notion of unlinkability, which restricts any adversary to only sanitizing messages it has

$\text{Exp}_{\text{SSS}, \mathcal{A}}^{\text{prv}}(\lambda)$ <hr style="border: 0.5px solid black;"/> $(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$ $(pk_{\text{San}}, sk_{\text{San}}) \leftarrow_{\$} \text{KGen}_{\text{San}}(1^\lambda)$ $b \leftarrow_{\$} \{0, 1\}, \mathcal{Q}_{\text{Sig}} \leftarrow \emptyset, \mathcal{Q}_{\text{San}} \leftarrow \emptyset$ $b' \leftarrow_{\$} \mathcal{A}^{\text{SignO}, \text{SanitO}, \text{ProofO}, \text{LoRSanit}_{\text{prvO}}}(pk_{\text{Sig}}, pk_{\text{San}})$ <p>return $b = b'$</p> <hr style="border: 0.5px solid black;"/> $\text{SignO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot) \text{ with } (\mu, \text{ADM}, pk_{\text{San}})$ <hr style="border: 0.5px solid black;"/> $\sigma \leftarrow \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$ $\mathcal{Q}_{\text{Sig}} \leftarrow \mathcal{Q}_{\text{Sig}} \cup (\mu, \text{ADM}, pk_{\text{San}}, \sigma)$ <p>return σ</p>	$\text{SanitO}(sk_{\text{San}}, \cdot, \cdot, \cdot, \cdot)$ <hr style="border: 0.5px solid black;"/> <p style="text-align: center;">with $(\mu, \text{MOD}, \sigma, pk_{\text{Sig}})$</p> <hr style="border: 0.5px solid black;"/> $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$ $\mathcal{Q}_{\text{San}} \leftarrow \mathcal{Q}_{\text{San}} \cup \{(\mu, \mu', \sigma', pk_{\text{Sig}})\}$ <p>return (μ', σ')</p> <hr style="border: 0.5px solid black;"/> $\text{ProofO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot, \cdot)$ <hr style="border: 0.5px solid black;"/> <p style="text-align: center;">with $(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, pk_{\text{San}})$</p> <hr style="border: 0.5px solid black;"/> $\pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{San}})$ <p>return π</p> <hr style="border: 0.5px solid black;"/> $\text{LoRSanit}_{\text{prvO}}(sk_{\text{Sig}}, sk_{\text{San}}, b, \cdot, \cdot) \text{ with } ((\mu_0, \text{MOD}_0), (\mu_1, \text{MOD}_1), \text{ADM})$ <hr style="border: 0.5px solid black;"/> $\sigma_b \leftarrow_{\$} \text{Sign}(\mu_b, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$ $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu_b, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$ <p>if $\text{MOD}_0(\mu_0) = \text{MOD}_1(\mu_1)$</p> <p style="padding-left: 20px;">return (μ', σ')</p> <p>return \perp</p>
---	--

Fig. 16: Game based security definition for privacy.

previously honestly generated using the sign oracle. We capture this by the optional check in the $\text{LoRSanit}_{\text{unlO}}$ oracle.

Definition 65 (Unlinkability). *A sanitizable signature scheme SSS is (weakly) unlinkable, if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS}, \mathcal{A}}^{(\text{w})\text{unl}}(\lambda)$ as described in Fig. 17, defined as*

$$\text{Adv}_{\text{SSS}, \mathcal{A}}^{(\text{w})\text{unl}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{SSS}, \mathcal{A}}^{(\text{w})\text{unl}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

is negligible in the security parameter λ .

Transparency Intuitively, no outside adversary should be able to determine the creator of a signature. This property is modeled using a left-or-right oracle, that either sanitizes or signs the result of the modification as a new message. The adversary is successful, if the can determine if the message-signature pairs it is receiving are output from the signing or the sanitizing algorithm. Transparency is mutually exclusive with non-interactive public accountability, as they described opposite requirements for the signatures. This notion of transparency is also frequently referred to as "proof-restricted transparency", to highlight the fact that the proof oracle cannot be misused for trivial wins. This is relevant in cases where the proof oracle works without using the additional message-signature pairs, i.e., it is "history free". If the proof is not history free, the adversary cannot use the proof oracle to trivially win the game, as it does not know learn the original

$\text{Exp}_{\text{SSS}, \mathcal{A}}^{(\text{w})\text{unl}}(\lambda)$ <hr/> $(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$ $(pk_{\text{San}}, sk_{\text{San}}) \leftarrow_{\$} \text{KGen}_{\text{San}}(1^\lambda)$ $b \leftarrow_{\$} \{0, 1\}, \mathcal{Q}_{\text{Sig}} \leftarrow \emptyset$ $b' \leftarrow_{\$} \mathcal{A}^{\text{SignO}, \text{SanitO}, \text{ProofO}, \text{LoRSanit}_{\text{unl}}\text{O}}(pk_{\text{Sig}}, pk_{\text{San}})$ <p>return $b = b'$</p> <hr/> $\text{SignO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot) \text{ with } (\mu, \text{ADM}, pk_{\text{San}})$ <hr/> $\sigma \leftarrow \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$ $\mathcal{Q}_{\text{Sig}} \leftarrow \mathcal{Q}_{\text{Sig}} \cup (\mu, \text{ADM}, pk_{\text{San}}, \sigma)$ <p>return σ</p> <hr/> $\text{ProofO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot, \cdot)$ <hr/> $\text{with } (\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, pk_{\text{San}})$ <hr/> $\pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{San}})$ <p>return π</p>	$\text{SanitO}(sk_{\text{San}}, \cdot, \cdot, \cdot, \cdot)$ <hr/> $\text{with } (\mu, \text{MOD}, \sigma, pk_{\text{Sig}})$ <hr/> $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$ $\mathcal{Q}_{\text{San}} \leftarrow \mathcal{Q}_{\text{San}} \cup \{(\mu, \mu', \sigma', pk_{\text{Sig}})\}$ <p>return (μ', σ')</p> <hr/> $\text{LoRSanit}_{\text{unl}}\text{O}(sk_{\text{Sig}}, sk_{\text{San}}, b, \cdot, \cdot) \text{ with}$ <hr/> $((\mu_0, \text{MOD}_0, \sigma_0), (\mu_1, \text{MOD}_1, \sigma_1))$ <hr/> $\text{if } (\mu_0, \text{MOD}_0, pk_{\text{San}}, \sigma_0) \notin \mathcal{Q}_{\text{Sig}}$ <hr/> $\vee (\mu_1, \text{MOD}_1, pk_{\text{San}}, \sigma_1) \notin \mathcal{Q}_{\text{Sig}}$ <hr/> <p>return \perp</p> $(\text{ADM}_0, \text{ADM}_1) \leftarrow (\sigma_0, \sigma_1)$ $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}_b, \sigma_b, sk_{\text{San}}, pk_{\text{Sig}})$ <p>if $\text{ADM}_0 = \text{ADM}_1$</p> $\wedge \text{MOD}_0(\mu_0) = \text{MOD}_1(\mu_1)$ $\wedge \text{Vrfy}(\mu_0, \sigma_0, pk_{\text{Sig}}, pk_{\text{San}}) = 1$ $\wedge \text{Vrfy}(\mu_1, \sigma_1, pk_{\text{Sig}}, pk_{\text{San}}) = 1$ <p>return (μ', σ')</p> <p>else return \perp</p>
---	---

Fig. 17: Game based security definition for (weak) unlinkability. Inclusion of the gray box yields the weak variant.

signatures from the $\text{Sanit}/\text{SignO}$ oracle which it would have to provide to the proof oracle. Consequently, construction like [13], which are not history free, can achieve the stronger non proof-restricted notion of transparency. This was first observed in [12], For a detailed breakdown of the evolution of the transparency property we refer to [29].

Many works [16, 6, 40] have defaulted to the proof-restricted variant, meaning they do not consider the non proof-restricted variant at all and are calling the proof-restricted variant "transparency". We will follow this approach in this paper and use the term "transparency" to refer to proof-restricted transparency. Note that some works like [30, 15] make this history-freeness explicit by disallowing the input of additional message-signature pairs into the proof oracle, making the proof-restricted variant the only viable definition. There also exists a strong variant of transparency. The term was first introduced by [3], but their description was later formalized by [16] to become the formal notion of invisibility. For strong transparency, the definition by [29] has gained the most traction. For strong transparency, the proof-restriction is slightly weakened, allowing queries where the message, but not the signature, has been output by the $\text{Sanit}/\text{SignO}$ oracle. Interestingly, some works like [16, 15, 6] are using this stronger definition as their regular version of transparency, while others like [12, 23] are using the

weaker variant as their regular version of transparency without discussing the strengthened notion.

Definition 66 (Proof Restricted Transparency). *A sanitizable signature scheme SSS is (strongly)proof restricted transparent (from now on we only say transparency), if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS},\mathcal{A}}^{(s)\text{transp}}(\lambda)$ as described in Fig. 18, defined as*

$$\text{Adv}_{\text{SSS},\mathcal{A}}^{(s)\text{transp}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{SSS},\mathcal{A}}^{(s)\text{transp}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

is negligible in the security parameter λ .

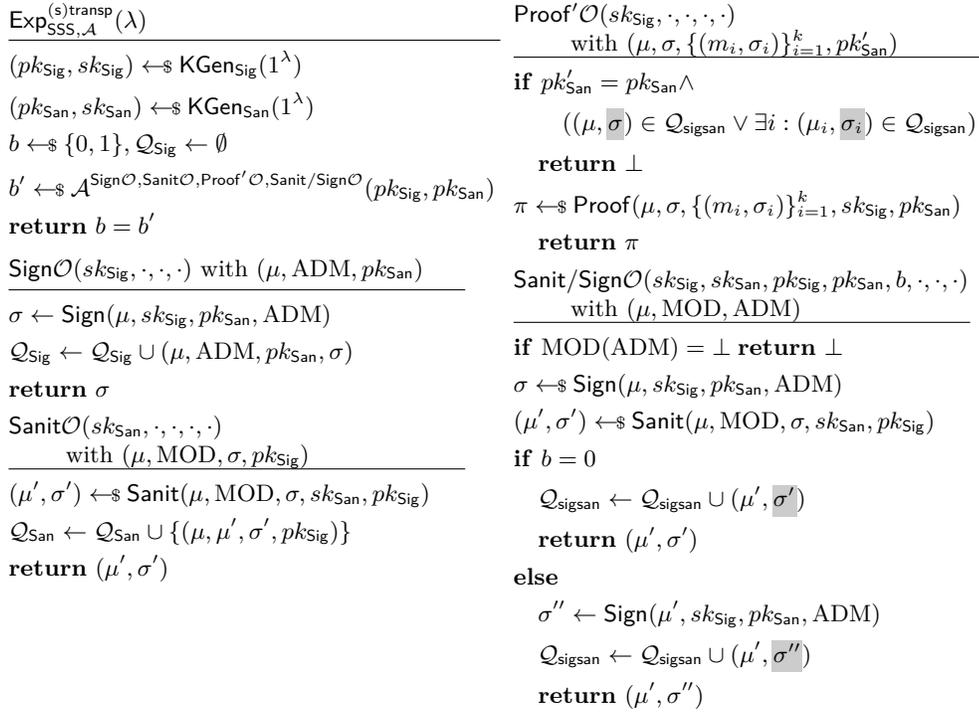


Fig. 18: Game based security definition for (strong) transparency. Inclusion of the gray box yields strong transparency.

Invisibility Intuitively, no efficient adversary should be able to detect whether any specific message block is admissible for modification or not.

To formalize this notion, a left-or-right signing oracle $\text{LoRAdm}\mathcal{O}$ is employed in combination with a restricted sanitizing oracle $\text{Sanit}'\mathcal{O}$. The sanitizing oracle is modified to prevent trivial wins. The adversary is limited to sanitizing

message-signature pairs it previously acquired from the left-or-right oracle and to introducing modifications which are allowed in both ADM_0 and ADM_1 to prevent trivial wins. or to outputs of the modified sanitizer oracle $\text{Sanit}'\mathcal{O}$. The left-or-right oracle behaves like a signing oracle, but takes two sets of admissible blocks ADM_0 and ADM_1 as input. Depending on the secret bit b sampled by the experiment, the message is signed with either the left or the right set of admissible blocks. Additionally, access to $\text{Proof}\mathcal{O}$ is provided to allow the adversary to attempt to learn sk_{Sig} . The adversary wins, if it correctly guesses if the oracle is using ADM_0 or ADM_1 to sign the given messages.

A list \mathcal{Q} is used to track previous queries of the adversary and to evaluate if a given query is allowed or not.

Definition 67 (Invisibility). *A sanitizable signature scheme SSS is invisible, if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{inv}}(\lambda)$ as described in Fig. 19, defined as*

$$\text{Adv}_{\text{SSS},\mathcal{A}}^{\text{inv}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{inv}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

is negligible in the security parameter λ .

$\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{inv}}(\lambda)$ <hr style="border: 0.5px solid black;"/> $(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$ $(pk_{\text{San}}, sk_{\text{San}}) \leftarrow_{\$} \text{KGen}_{\text{San}}(1^\lambda)$ $b \leftarrow_{\$} \{0, 1\}, \mathcal{Q} \leftarrow \emptyset$ $b' \leftarrow_{\$} \mathcal{A}^{\text{Sanit}'\mathcal{O}, \text{Proof}\mathcal{O}, \text{LoRAdm}\mathcal{O}}(pk_{\text{Sig}}, pk_{\text{San}})$ return $b = b'$ $\text{Sanit}'\mathcal{O}(sk_{\text{San}}, pk_{\text{Sig}}, \cdot, \cdot, \cdot)$ with $(\mu, \text{MOD}, \sigma)$ if $\nexists (\mu, \sigma, \text{ADM}) \in \mathcal{Q} : \text{MOD}(\text{ADM}) = \top$ return \perp $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$ $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{\mu', \sigma', \text{ADM}\}$ return (μ', σ)	$\text{Proof}\mathcal{O}(sk_{\text{Sig}}, \cdot, \cdot, \cdot, \cdot)$ with $(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, pk_{\text{San}})$ <hr style="border: 0.5px solid black;"/> $\pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{San}})$ return π $\text{LoRAdm}\mathcal{O}(sk_{\text{Sig}}, b, \cdot, \cdot, \cdot)$ with $(\mu, \text{ADM}_0, \text{ADM}_1)$ <hr style="border: 0.5px solid black;"/> $\sigma \leftarrow_{\$} \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{MOD}_b)$ $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{\mu, \sigma, \text{ADM}_0 \cap \text{ADM}_1\}$ return σ
---	--

Fig. 19: Game based security definition for invisibility.

Definition 68 (Strong Invisibility). *A sanitizable signature scheme SSS is strongly invisible, if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{inv}}(\lambda)$ as described in Fig. 20, defined as*

$$\text{Adv}_{\text{SSS},\mathcal{A}}^{\text{inv}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{inv}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

is negligible in the security parameter λ .

$$\begin{array}{l}
 \text{Exp}_{\text{SSS}, \mathcal{A}}^{\text{s-inv}}(\lambda) \\
 \hline
 (pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda) \\
 (pk_{\text{San}}, sk_{\text{San}}) \leftarrow_{\$} \text{KGen}_{\text{San}}(1^\lambda) \\
 b \leftarrow_{\$} \{0, 1\}, \mathcal{Q} \leftarrow \emptyset \\
 b' \leftarrow_{\$} \mathcal{A}^{\text{Proof}\mathcal{O}, \text{LoRAdm}'\mathcal{O}, \text{Sanit}''\mathcal{O}}(pk_{\text{Sig}}, pk_{\text{San}}) \\
 \mathbf{return} \ b = b' \\
 \hline
 \text{Sanit}''\mathcal{O}(sk_{\text{San}}, \cdot, \cdot, \cdot, \cdot) \text{ with } (pk'_{\text{Sig}}, \mu, \text{MOD}, \sigma) \\
 \hline
 \mathbf{if} \ pk'_{\text{Sig}} = pk_{\text{Sig}} \wedge \nexists(\mu, \sigma, \text{ADM}) \in \mathcal{Q} : \text{MOD}(\text{ADM}) = \top \\
 \quad \mathbf{return} \ \perp \\
 (\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk'_{\text{Sig}}) \\
 \mathbf{if} \ pk'_{\text{Sig}} = pk_{\text{Sig}} \wedge \exists(\mu, \sigma, \text{ADM}') \in \mathcal{Q} : \text{MOD}(\text{ADM}') = \top \\
 \quad \mathcal{Q} \leftarrow \mathcal{Q} \cup \{\mu', \sigma', \text{ADM}'\} \\
 \mathbf{return} \ (\mu', \sigma) \\
 \text{LoRAdm}'\mathcal{O}(sk_{\text{Sig}}, b, \cdot, \cdot, \cdot, \cdot) \text{ with} \\
 \quad (pk'_{\text{San}}, \mu, \text{ADM}_0, \text{ADM}_1) \\
 \hline
 \mathbf{if} \ pk'_{\text{Sig}} = pk_{\text{Sig}} \wedge \text{ADM}_0 = \text{ADM}_1 \\
 \quad \mathbf{return} \ \perp \\
 \sigma \leftarrow_{\$} \text{Sign}(\mu, sk_{\text{Sig}}, pk'_{\text{San}}, \text{MOD}_b) \\
 \mathbf{if} \ pk'_{\text{Sig}} = pk_{\text{Sig}} \wedge \exists(\mu, \sigma, \text{ADM}') \in \mathcal{Q} : \text{MOD}(\text{ADM}') = \top \\
 \quad \mathcal{Q} \leftarrow \mathcal{Q} \cup \{\mu, \sigma, \text{ADM}_0 \cap \text{ADM}_1\} \\
 \mathbf{return} \ \sigma \\
 \text{Proof}\mathcal{O}(sk_{\text{Sig}}, \cdot, \cdot, \cdot, \cdot) \\
 \quad \text{with } (\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, pk_{\text{San}}) \\
 \hline
 \pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{San}}) \\
 \mathbf{return} \ \pi
 \end{array}$$

Fig. 20: Game based security definition for strong invisibility.

Strong invisibility strengthens this notions by allowing an adversary to query its oracles under arbitrary signer public keys, simulating a dishonest signer. In this case the aforementioned restrictions for the left or right and sanitize oracle only apply when queried with the challenge public key. Note that in both the regular and the strong variant of invisibility, a signing oracle can be simulated by querying the corresponding left or right oracle with $\text{ADM}_0 = \text{ADM}_1$.

Accountability Intuitively, accountability formalizes the ability to hold both signer and sanitizer responsible for messages they have signed. The strong accountability

notion extends this to the signature, holding the respective party responsible for the signatures they have created. For technical reasons this notion is split into signer-accountability and sanitizer accountability. Both notions formalize the inability of the respective party to frame the other party for a signature they have not produced; when both notions are achieved one has accountability.

For signer-accountability we refer to Definition 37 in the main body of the paper.

Definition 69 (Sanitizer-Accountability). *A sanitizable signature scheme SSS is (strongly) sanitizer-accountable, if for all qpt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS},\mathcal{A}}^{(\text{s})\text{san-acc}}(\lambda)$ as described in Fig. 21, defined as*

$$\text{Adv}_{\text{SSS},\mathcal{A}}^{(\text{s})\text{sig-acc}}(\lambda) := \Pr \left[\text{Exp}_{\text{SSS},\mathcal{A}}^{(\text{s})\text{sig-acc}}(\lambda) = 1 \right]$$

is negligible in the security parameter λ .

$\text{Exp}_{\text{SSS},\mathcal{A}}^{(\text{s})\text{san-acc}}(\lambda)$	$\text{SignO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot)$ with $(\mu, \text{ADM}, pk_{\text{San}})$
$(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$ $\mathcal{Q}_{\text{Sig}} \leftarrow \emptyset$ $(pk_{\text{San}}^*, \mu^*, \sigma^*) \leftarrow_{\$} \mathcal{A}^{\text{SignO}, \text{ProofO}}(pk_{\text{Sig}})$ $\pi \leftarrow \text{Proof}(sk_{\text{Sig}}, \mu^*, \sigma^*, \{\mu_i, \sigma_i : (\mu_i, *, \sigma_i, *) \in \mathcal{Q}_{\text{Sig}}\}, pk_{\text{San}})$ if $(\mu^*, *, pk_{\text{San}}^*, \sigma^*) \notin \mathcal{Q}_{\text{Sig}}$ $\wedge \text{Vrfy}(\mu^*, \sigma^*, pk_{\text{Sig}}^*, pk_{\text{San}}) = \top$ $\wedge \text{Judge}(\mu^*, \sigma^*, pk_{\text{Sig}}^*, pk_{\text{San}}, \pi) = \text{Sig}$ return 1 else return 0	$\sigma \leftarrow \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$ $\mathcal{Q}_{\text{Sig}} \leftarrow \mathcal{Q}_{\text{Sig}} \cup (\mu, \text{ADM}, pk_{\text{San}}, \sigma)$ return σ $\text{ProofO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot)$ $\text{with } (\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, pk_{\text{San}})$ $\pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(m_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{San}})$ return π

Fig. 21: Game based security definition for (strong) sanitizer accountability. Inclusion of the gray box yields strong sanitizer accountability.

Definition 70 (Accountability). *A sanitizable signature scheme SSS is (strongly) accountable, if it is both (strongly) signer-accountable and (strongly) sanitizer-accountable.*

Non-Interactive Public Accountability Intuitively, non-interactive public accountability requires that for any signature an external observer can reliably decide whether this signature was created by the signer or the sanitizer without any additional information. This is formalized by removing the ability to use the proof algorithm (which takes sk_{Sig} as input). Consequently, the **Judge** algorithm receives the symbol \perp as input instead of the proof π . If the **Judge** algorithm still works as expected the scheme is non-interactively publicly accountable.

Definition 71 (Non-Interactive Public Accountability). A sanitizable signature scheme SSS achieves non-interactive public accountability, if for all ppt adversaries \mathcal{A} , the advantage in winning the game $\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{nipa}}(\lambda)$ as described in Fig. 22, defined as

$$\text{Adv}_{\text{SSS},\mathcal{A}}^{\text{nipa}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{nipa}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

is negligible in the security parameter λ .

$\text{Exp}_{\text{SSS},\mathcal{A}}^{\text{nipa}}(\lambda)$ <hr style="border: 0.5px solid black;"/> $(pk_{\text{Sig}}, sk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$ $(pk_{\text{San}}, sk_{\text{San}}) \leftarrow_{\$} \text{KGen}_{\text{San}}(1^\lambda)$ $\mathcal{Q}_{\text{Sig}} \leftarrow \emptyset, \mathcal{Q}_{\text{San}} \leftarrow \emptyset$ $(pk^*, \mu^*, \sigma^*) \leftarrow_{\$} \mathcal{A}^{\text{SignO}, \text{SanitO}}(pk_{\text{Sig}}, pk_{\text{San}})$ if $(\mu^*, *, pk_{\text{San}}, \sigma) \notin \mathcal{Q}_{\text{Sig}}$ $\wedge \text{Vrfy}(\mu^*, \sigma^*, pk_{\text{Sig}}^*, pk^*) = \top$ $\wedge \text{Judge}(\mu^*, \sigma^*, pk_{\text{Sig}}^*, pk^*, \perp) = \text{Sig}$ return 1 elseif $(*, \mu^*, \sigma^*, pk_{\text{Sig}}^*) \notin \mathcal{Q}_{\text{San}}$ $\wedge \text{Vrfy}(\mu^*, \sigma^*, pk^*, pk_{\text{San}}) = \top$ $\wedge \text{Judge}(\mu^*, \sigma^*, pk^*, pk_{\text{San}}, \perp) = \text{San}$ return 1 else return 0	$\text{SignO}(sk_{\text{Sig}}, \cdot, \cdot, \cdot)$ with $(\mu, \text{ADM}, pk_{\text{San}})$ <hr style="border: 0.5px solid black;"/> $\sigma \leftarrow \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$ $\mathcal{Q}_{\text{Sig}} \leftarrow \mathcal{Q}_{\text{Sig}} \cup (\mu, \text{ADM}, pk_{\text{San}}, \sigma)$ return σ $\text{SanitO}(sk_{\text{San}}, \cdot, \cdot, \cdot)$ with $(\mu, \text{MOD}, \sigma, pk_{\text{Sig}})$ <hr style="border: 0.5px solid black;"/> $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$ $\mathcal{Q}_{\text{San}} \leftarrow \mathcal{Q}_{\text{San}} \cup \{(\mu, \mu', \sigma', pk_{\text{Sig}})\}$ return (μ', σ')
---	---

Fig. 22: Game based security definition for non-interactive public accountability.

Definition 72. A sanitizable signature scheme SSS is correct, if signing correctness, sanitizing correctness, and proof correctness hold.

Signing Correctness. For all security parameter $\lambda \in \mathbb{N}$, any key pair $(sk_{\text{Sig}}, pk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$, any message $\mu \in \mathcal{M}$, any valid ADM , and for every $\sigma \leftarrow_{\$} \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$, it holds that

$$\Pr[\text{Vrfy}(\mu, \sigma, pk_{\text{Sig}}, pk_{\text{San}}) = \top] = 1.$$

Sanitizing Correctness. For all security parameter $\lambda \in \mathbb{N}$, any key pair $(sk_{\text{Sig}}, pk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$, any message $\mu \in \mathcal{M}$, any valid ADM , and for every $\sigma \leftarrow_{\$} \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$, any valid MOD with $\text{MOD}(\text{ADM}) = \top$, and any sanitized message with signature $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$, it holds that

$$\Pr[\text{Vrfy}(\mu', \sigma', pk_{\text{Sig}}, pk_{\text{San}}) = \top] = 1.$$

Proof Correctness. For all security parameters $\lambda \in \mathbb{N}$, any key pair $(sk_{\text{Sig}}, pk_{\text{Sig}}) \leftarrow_{\$} \text{KGen}_{\text{Sig}}(1^\lambda)$, any message $\mu \in \mathcal{M}$, any valid ADM, any signature $\sigma \leftarrow_{\$} \text{Sign}(\mu, sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM})$ with $\text{Vrfy}(\mu, \sigma, pk_{\text{Sig}}, pk_{\text{San}}) = \top$, any valid MOD that does not violate ADM, any sanitized message-signature pair $(\mu', \sigma') \leftarrow_{\$} \text{Sanit}(\mu, \text{MOD}, \sigma, sk_{\text{San}}, pk_{\text{Sig}})$, with $\text{Vrfy}(\mu', \sigma', pk_{\text{Sig}}, pk_{\text{San}}) = \top$, and polynomial many additional signature-message pairs (μ_i, σ_i) for $1 \leq i \leq k$, and for all proofs $\pi \leftarrow_{\$} \text{Proof}(\mu, \sigma, \{(\mu_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{Sig}}, pk_{\text{San}})$ and $\pi' \leftarrow_{\$} \text{Proof}(\mu', \sigma', \{(\mu_i, \sigma_i)\}_{i=1}^k, sk_{\text{Sig}}, pk_{\text{Sig}}, pk_{\text{San}})$ it holds that

$$\begin{aligned} \Pr[\text{Sig} \leftarrow \text{Judge}(\mu, \sigma, pk_{\text{Sig}}, pk_{\text{San}}, \pi)] &= 1 \\ \wedge \Pr[\text{San} \leftarrow \text{Judge}(\mu', \sigma', pk_{\text{Sig}}, pk_{\text{San}}, \pi')] &= 1. \end{aligned}$$

B.2 Relations Between Security Notions

In Fig. 23 we provide an overview of the relations between the different security properties of sanitizable signature schemes and observe an additional implication between non-interactive public accountability and accountability in Theorem 73.

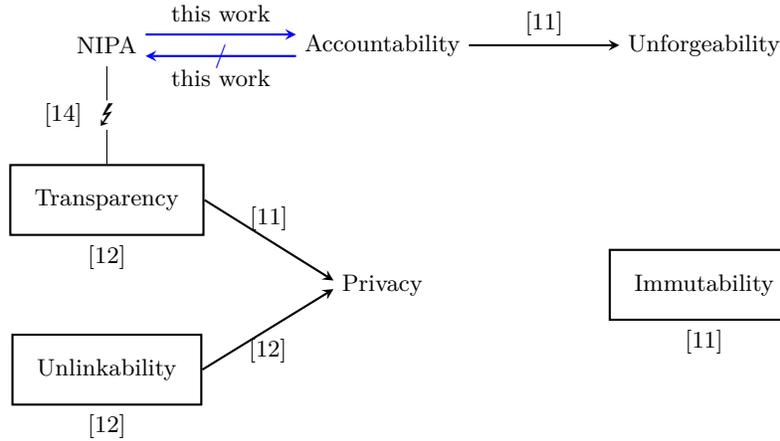


Fig. 23: The relations between different security notions for SSS. Arrows indicate implications, a lightning indicates two notions are conflicting. A box indicates that the notion is independent of all other notions. Citations next to the implications or boxes cite the source where this was first observed.

In [12] it is stated, that accountability is independent of all other security properties. At the time non-interactive public accountability has not yet been established as a security notion for sanitizable signature schemes. Nowadays this claim no longer holds due to the following relation between accountability and non-interactive public accountability.

Theorem 73. *A sanitizable signature scheme SSS that is non-interactively publicly accountable is also accountable.*

Proof. We prove Theorem 73 using reductions. We make use of the fact that, per definition, the proof algorithm of a non-interactively publicly accountable SSS only outputs \perp . Note that accountability is a compound notion consisting of signer and sanitizer accountability. Hence, we need to provide two reductions: one breaking non-interactive public accountability given an adversary $\mathcal{A}_{\text{sig-acc}}$ against signer accountability, and one given an adversary $\mathcal{A}_{\text{san-acc}}$ against sanitizer accountability that both succeed with non-negligible probability.

First, we assume $\mathcal{A}_{\text{san-acc}}$ as described above and construct a reduction \mathcal{B}_{San} that breaks non-interactive public accountability with non-negligible probability. The reduction \mathcal{B}_{San} receives two public keys $(pk_{\text{Sig}}, pk_{\text{San}})$ as part of the experiment and will forward pk_{Sig} to $\mathcal{A}_{\text{san-acc}}$. Oracle queries by $\mathcal{A}_{\text{san-acc}}$ to ProofO can always be answered with \perp per definition, and queries to SignO are forwarded to the reductions own sign oracle. Finally, $\mathcal{A}_{\text{san-acc}}$ will output $(pk_{\text{San}}^*, m^*, \sigma^*)$ where $\text{Vrfy}(m^*, \sigma^*, pk_{\text{Sig}}, pk_{\text{San}}^*) = \top$. Lastly, \mathcal{B}_{San} forwards the adversary's output as its output. The reduction \mathcal{B}_{San} is efficient, perfectly simulates the sanitizer accountability game for \mathcal{A}_{San} , and has the same non-negligible success probability as \mathcal{A}_{San} by definition, which we assume to be non-negligible.

The reduction \mathcal{B}_{Sig} works analogously with two slight modifications: It does not need to simulate the proof oracle and only forwards the output $(pk_{\text{Sig}}^*, m^*, \sigma^*)$, omitting π^* .

We have constructed two reductions, \mathcal{B}_{San} and \mathcal{B}_{Sig} , and have shown that for any SSS it holds that $\text{Adv}_{\text{SSS}, \mathcal{A}_{\text{sig-acc}}}^{\text{nipa}}(\lambda) \leq \text{Adv}_{\text{SSS}, \mathcal{B}_{\text{Sig}}}^{\text{sig-acc}}(\lambda)$ and $\text{Adv}_{\text{SSS}, \mathcal{A}_{\text{san-acc}}}^{\text{nipa}}(\lambda) \leq \text{Adv}_{\text{SSS}, \mathcal{B}_{\text{San}}}^{\text{san-acc}}(\lambda)$, which proves the result. \square

Theorem 74. *A sanitizable signature scheme SSS that is accountable is not necessarily non-interactively publicly accountable.*

Proof (Sketch). Assume any accountable SSS scheme, where the accountability is realized by the judge algorithm evaluating its input π . This scheme is clearly not non-interactively publicly accountable, as in this setting, the judge will receive only the symbol \perp instead of the proof π . \square

C Additional Information on BFF⁺09

In this appendix, we provide additional information on the BFF⁺09 construction by Brzuska et al. [11]. We begin with the notion of tagged chameleon hashes, a corrected security notion for collision-resistance, and a transform to build such fully collision-resistant tagged chameleon hash functions from fully collision-resistant chameleon hash functions (without tags), as constructed in Fig. 7. Afterwards, we present the details on the construction.

C.1 Fully Collision-Resistant Tagged Chameleon Hash Functions

The construction by Brzuska et al. [11] requires a tagged chameleon hash function CH. A tag chameleon hash function takes a tag as additional input to the hashing algorithm. The adapt procedure takes a message, tag, and randomness, and an additional second message as input. It outputs a new tag and a new randomness, such that the hash values of the input triple coincides with the hash of the second input message and the output tag-randomness pair.

As the collision-resistance notion shown in Fig. 10 is unachievable, we present a new formalization which follows closely the f-CR game in Fig. 2. For full collision-resistance, the adversary receives a public key pk and access to an adapt oracle which internally chooses uniformly random tags. On input μ , the oracle that outputs (TAG, r) stores the information (h, TAG, μ) , where h is the hash of (TAG, μ, r) . If the adversary outputs a hash collision by providing two tags (TAG_1, TAG_2) , two messages (μ_1, μ_2) , and two randomnesses (r_1, r_2) with $(TAG_1, \mu_1) \neq (TAG_2, \mu_2)$. Let $h_i = CH(pk, TAG_i, \mu_i, r_i)$. Then, in the full collision-resistance game, the output is accepted, if the one of the triples of (h_1, TAG_1, μ_1) and (h_2, TAG_2, μ_2) has not been part of an adapt query. This description closely resembles the f-CR game in Fig. 2. We note that the security notion for tagged chameleon hash functions here is not related immediately to tagged chameleon hash function security in [31].

Finally, it is easy to see that f-CR tagged chameleon hashes can be constructed from f-CR chameleon hash functions without tags. Indeed, given a chameleon hash function $CH = (CKGen, CHash, CHashCheck, CAdapt)$, we construct the tagged version $(CKGen_{TAG}, CHash_{TAG}, CHashCheck_{TAG}, CAdapt_{TAG})$ as follows:

$CKGen_{TAG} = CKGen$
 $CHash_{TAG}$ takes (TAG, μ, r) and outputs $CHash(TAG||\mu, r)$
 $CHashCheck_{TAG}$ takes (h, TAG, μ, r) and outputs $CHashCheck(h, TAG||\mu, r)$
 $CAdapt_{TAG}$ takes (h, TAG, μ, r, μ') , samples TAG' uniformly randomly, and outputs $(TAG', CAdapt(h, TAG||\mu, r, TAG' || \mu))$.

With standard techniques it can be shown that f-CR of the tagged chameleon hash reduces to f-CR of the underlying chameleon hash function.

C.2 The BFF⁺09 Construction

In this section, we introduce the construction by Brzuska et al. [11], which we have analyzed more closely in Section 5.5.

Construction 75 *Let PRG be a pseudorandom generator, PRF a pseudorandom function, $\Sigma = (KGen, Sign, Vrfy)$ a signature scheme, and CH a fully collision-resistant tagged chameleon hash. A sanitizable signature scheme $SanSig = (KGen_{Sig}, KGen_{San}, Sign, Sanit, Vrfy, Proof, Judge)$ is defined as follows:*

Key Generation. *The signer key (pk_{Sig}, sk_{Sig}) is generated using KGen. The sanitizer key (pk_{San}, sk_{San}) is generated using CKGen.*

Signing. The algorithm **Sign**, on input $\mu \in \{0, 1\}^{tl}$, sk_{Sig} , pk_{San} , and ADM , picks $\text{NONCE} \leftarrow_{\$} \{0, 1\}^n$ at random, computes $x \leftarrow_{\$} \text{PRF}(\kappa, \text{NONCE})$, $\text{TAG} \leftarrow_{\$} \text{PRG}(x)$, and picks $r[j]$ for each $j \in \text{ADM}$ at random. It sets

$$h[j] \leftarrow_{\$} \text{CH}(pk_{\text{San}}, \text{TAG}, (j, \mu_j, pk_{\text{Sig}}); r[j]),$$

if $j \in \text{ADM}$, or μ_j else. Then, it computes $\sigma_0 \leftarrow \text{Sign}(sk_{\text{Sig}}, (h, pk_{\text{San}}, \text{ADM}))$ where $h = (h[1], \dots, h[l])$ and returns

$$\sigma \leftarrow (\sigma_0, \text{TAG}, \text{NONCE}, \text{ADM}, r[j_1], \dots, r[j_k])$$

where $j_i \in \text{ADM}$.

Sanitizing. The algorithm **Sanit** on input a message μ , information MOD , a signature $\sigma = (\sigma_0, \text{TAG}, \text{NONCE}, \text{ADM}, r[j_1], \dots, r[j_k])$, pk_{Sig} , and sk_{San} , first checks that each modification in MOD is admissible according to ADM and that σ_0 is a valid signature for $(h, pk_{\text{San}}, \text{ADM})$. If not, it outputs \perp . Otherwise, for each $j \in \text{ADM}$, it lets $\mu'[j]$ be the modified block of μ_j , possibly $\mu'[j] = \mu_j$, picks NONCE' and TAG' randomly and replaces $r[j]$ for each admissible j by $r'[j] = \text{CAdapt}(sk_{\text{San}}, \text{TAG}, (j, \mu_j, pk_{\text{Sig}}), r[j], \text{TAG}', (j, \mu'[j], pk_{\text{Sig}}))$. It outputs μ' and $\sigma' = (\sigma_0, \text{TAG}', \text{NONCE}', \text{ADM}, r'[j_1], \dots, r'[j_k])$.

Verification. The algorithm **Vrfy** on input a message $\mu \in \{0, 1\}^{tl}$ and a signature

$$\sigma \leftarrow (\sigma_0, \text{TAG}, \text{NONCE}, \text{ADM}, r[i_1], \dots, r[i_k]),$$

pk_{Sig} , and pk_{San} , computes the hashes $h[j]$ as above, sets $h = (h[1], \dots, h[l])$ and runs $\Sigma.\text{Vrfy}(pk_{\text{Sig}}, (h, pk_{\text{San}}, \text{ADM}), \sigma_0)$.

Proof. The algorithm **Proof** on input sk_{Sig} , μ , σ , a sequence (μ_i, σ_i) , and pk_{San} , searches for $(\text{TAG}_i, (j, \mu_i[j], pk_{\text{Sig}}), r_i[j])$ such that

$$\text{CH}(pk_{\text{San}}, \text{TAG}_i, (j, \mu_i[j], pk_{\text{Sig}}), r_i[j]) = \text{CH}(pk_{\text{San}}, \text{TAG}, (j, \mu_j, pk_{\text{Sig}}), r[j]),$$

for some distinct $(\text{TAG}, (j, \mu_j, pk_{\text{Sig}}), r[j])$ in μ , σ , where $\text{TAG}_i = \text{PRG}(x_i)$ for $x_i = \text{PRF}(\kappa, \text{NONCE}_i)$ for the value NONCE_i in σ_i . If such data is found, it returns

$$\pi \leftarrow (\text{TAG}_i, (j, \mu_i[j], pk_{\text{Sig}}), r_i[j], x_i),$$

else it outputs \perp .

Judge. The algorithm **Judge**, on input μ , σ , pk_{Sig} , pk_{San} , and

$$\pi = (\text{TAG}_\pi, (j, \mu_\pi[j], pk_{\text{Sig}}), r_\pi[j], x_\pi)$$

checks if

$$\begin{aligned} & pk_{\text{Sig}} = pk_{\text{Sig}, \pi}, \\ & \wedge \text{CH}(pk_{\text{San}}, \text{TAG}, (j, \mu_j, pk_{\text{Sig}}), r[j]) \\ & \quad = \text{CH}(pk_{\text{San}}, \text{TAG}_\pi, (j, \mu_\pi[j], pk_{\text{Sig}}), r_\pi[j]), \\ & \wedge \text{block } j \text{ is admissible,} \\ & \wedge \text{TAG}_\pi = \text{PRG}(x_\pi). \end{aligned}$$

If all conditions are satisfied, **Judge** outputs **San**, otherwise, it outputs **Sig**.

D Lattice-Based Instantiations from Standard Primitives

Following the different classes of constructions, we start with a discussion of those instantiations that rely solely on standard primitives. This includes the FH18 construction by Fischlin and Harasser [22], followed by the BFLS09 construction by Brzuska et al. [13].

The FH18 Construction. Fischlin and Harasser [22] show in their paper that public-key encryption is equivalent to invisible sanitizable signature schemes. They give a construction that relies solely on standard primitives: public-key encryption schemes and digital signature schemes.

We only give a the high-level idea of their construction and refer to [22] for the details.

The construction works as follows. The signer key pair consists of a key pair for the underlying signature scheme Σ . The sanitizer key pair also consists of a key pair of the underlying signature scheme Σ and additionally contains a key pair of the public-key encryption scheme Π . To sign a message μ , it is chopped into blocks μ_1, \dots, μ_l which are then signed individually using freshly generated key pairs of Σ , yielding signatures $\sigma_1, \dots, \sigma_l$. The message is chopped into blocks and each block is signed with a freshly generated key pair. For each block, the signer encrypts either the corresponding secret key—if the block is sanitizable—or the zero-string—if the block is fixed. This enables the sanitizer to obtain the secret key for each to, which allows to change the message and sign it with that secret key. Finally, the signer (or the sanitizer, if it sanitized a message) uses its secret key to sign the message alongside the individual signatures and ciphertexts.

Theorem 76. *If the signature scheme Σ is instantiated with one of the candidates of the set $\{\text{Dilithium}, \text{Falcon}\}$, and the public-key encryption scheme instantiated with *Kyber.CPA*, then the FH18 construction [22] is an unforgeable, immutable, private, publicly accountable, and invisible sanitizable signature scheme against *qpt* adversaries.*

Proof. The theorem follows straightforwardly from [22]. Their proof directly transfers to *qpt* adversaries, as none of the oracles are available in superposition and there is no random oracle. Furthermore, the adversaries in their reductions are black box, therefore there is no rewinding or cloning taking place, which would also contradict *qpt* security. \square

The above theorem shows that sanitizable signatures are possible from lattices and, in fact, any assumptions that allow the construction of both public-key encryption and signature schemes.⁸ We note, however, that [22] is a more theoretical work. Their main result shows that sanitizable signatures and public-key encryption are equivalent. The construction is fairly inefficient—signatures are very large since they contain a public key for each message block.

⁸ This precludes hash-based cryptography which only allows for signatures but not public-key encryption

The BFLS09 Construction. A more efficient construction is the BFLS09 construction by Brzuska et al. [13]. The construction works as follows. The key pairs for signer and sanitizer consists of a key pair for the underlying signature scheme Σ . To sign a message μ , it is chopped into blocks μ_1, \dots, μ_l and the signer creates two signatures. First, σ_{fix} , which essentially signs all immutable message blocks. Second, σ_{full} , which signs the entire message—immutable and admissible blocks. In order to sanitize a signature, the sanitizer changes the message μ to μ^* and replaces σ_{full} by signing the μ^* . Verification works by verifying that σ_{fix} is a valid signature under the public key of the signer and that σ_{full} is a valid signature under the public key of either signer or sanitizer.

The signature σ_{fix} on the immutable blocks prevents the sanitizer to change any immutable blocks, as these are always verified against the public key of the signer. The simplicity of the construction comes at the expense of not achieving transparency: whether σ_{full} verifies against the public key of the signer or the sanitizer immediately reveals if a message was sanitized or not. On the other hand, the construction is minimal in the sense that it relies solely on signature schemes—allowing easy instantiation from, say, Dilithium.

Below we state that the BFLS09 construction can be instantiated from lattices, which follows straightforward from [13].

Theorem 77. *If the signature scheme Σ is instantiated with one of the candidates of the set $\{\text{Dilithium}, \text{Falcon}\}$, then the BFLS09 construction [13] is an unforgeable, immutable, private, and accountable sanitizable signature scheme against qpt adversaries.*

E Lattice-Based Instantiations from Advanced Signatures

In this section, we cover the different SSS constructions that rely on various types of advanced signature schemes (and not on chameleon hash functions). These include the following constructions: the FKM⁺16 construction [23], the BFLS10 construction [12], the LZCS16–2 construction [30], and the BLL⁺19 construction [15].

The FKM⁺16 Construction. The construction by Fleischhacker et al. [23] relies on signature schemes with rerandomizable keys. On a high-level, these are signature schemes that allow to randomize a key pair (sk, pk) , using some randomness r , into a new key pair (sk', pk') such that (1) (sk', pk') is a matching key pair, i.e., signatures generated using sk' verify under pk' , and (2) re-randomized keys are indistinguishable from keys generated via KGen.

Alkeilani Alkadri et al. [2] showed that rerandomizable signature schemes can be constructed from lattices. However, they consider a weaker notion, where only the re-randomized public key should be hard to distinguish from a public key obtained via KGen.⁹ In fact, [2] identifies obstacles due to usage of Gaussian distributions for lattice-based constructions which typically result in different

⁹ This modification is motivated via their goal of showing public keys to be unlinkable.

distributions for keys generated via KGen and re-randomized ones. Thus, the construction from [2] cannot be used to instantiate the construction by Fleischhacker et al. [23]. This leaves the open problem of constructing rerandomizable signature schemes according to [23] from lattice¹⁰ assumptions.

The BFLS10 Construction. The construction by Brzuska et al. [12] uses group signatures and is historically the first sanitizable signature scheme that does make use of advanced signature notions: In this work, the construction is based on group signatures, which are required to satisfy *anonymity*, *traceability*, and *non-frameability*.

The construction can be instantiated using the group signature that has been developed by Ling et al. [33].

Theorem 78. *The BFLS10 construction [12] instantiated with a signature from {Dilithium, Falcon} and the group signature from [33] is a lattice-based sanitizable signature scheme that satisfies unforgeability, immutability, signer- and sanitizer accountability, transparency, privacy, and unlinkability.*

The LZCS16–2 Construction. In [30], Lai et al. give two distinct constructions. One deploys chameleon hash functions, which we discuss later. Here, we focus on the other construction, which uses accountable ring signatures as its primary component in combination with deterministic signature schemes. For both, we have lattice based primitives available, which achieve the required security properties and can be used for a generic instantiation. For deterministic signatures ML-DSA [38] can be used in the deterministic mode, as standardized by NIST. For accountable ring signatures, Beullens et al. [8] provide a lattice based as well as an isogeny based instantiation, both of which are suitable for instantiating this construction in practice.

The BLL⁺19 Construction. Finally, we review the construction by Bultel et al. [15], which already played a crucial role in the instantiations of SSS constructions using chameleon hash functions. There, we used the transform described in [15] to increase the security to include unforgeability and accountability, if their security proofs in the respective constructions required the chameleon hash to satisfy uniqueness—the property that our lattice-based CHF fails to achieve.

In [15], so-called equivalence class signatures (EQS) have been used to create a weak SSS. Then, requiring (strongly) unforgeable VRS, it is shown how to add unforgeability and accountability features to the weak SSS resulting in a secure SSS. The equivalence class signatures in [15] rely on pairing problems which are not post-quantum secure. Further investigations regarding the weak SSS construction may provide useful insights on generating other weak sanitizable signatures. However, the pairing problem has to be replaced to ensure post-quantum security. On the other hand, we make extensive use of the transform in Theorem 38.

¹⁰ Recently, Das et al. [19] showed a construction from isogenies.

F Full Proofs for Theorem Statements

F.1 Learning with Rounding-Based LITs

In Lemma 18 we stated that there exists an LWR-based LIT. We now prove this lemma by constructing and proving secure such a LIT. First, we need to define learning with rounding and its rounding function. We define the rounding function $\lfloor \cdot \rfloor_p : \mathbb{Z}_q \rightarrow \mathbb{Z}_p, x \mapsto \lfloor (p/q) \cdot x \rfloor$. If we use that function on vectors, we use it component-wise.

Definition 79 (Learning with Rounding). *In the learning with rounding (LWR) problem $\text{LWR}_{n,q,p,\chi}$, the challenger picks a secret $\mathbf{s} \in \mathbb{Z}_q^n$ from a distribution χ and defines the LWR distribution that samples a uniform $\mathbf{a} \leftarrow \mathbb{Z}_q^n$ and outputs $(\mathbf{a}, \lfloor \mathbf{s}^t \mathbf{a} \rfloor_p)$. Then, a qpt adversary \mathcal{A} can ask for samples from either the LWR distribution or from the uniform distribution over \mathbb{Z}_q^{n+1} and has to decide from what distribution the samples are.*

In the search version of LWR $\text{sLWR}_{n,q,p,\chi}$ the qpt adversary gets samples from the LWR distribution and has to find \mathbf{s} .

Furthermore, we need a lemma telling us when an LWR secret is unique, similar to [9, Lemma 1] for module LWE. For that, define $B_a^b = \{\mathbf{x} \in \mathbb{Z}_q^b : \|\mathbf{x}\|_\infty \leq a\}$.

Lemma 80. *Let $n, q, p, \beta > 0$ with $q > p$ and $0 < \beta < q$. Let $q = \prod_{i=1}^d q_i^{e_i}$, where q_i are prime and let j be such that $q_j = \min\{q_i\}$. Choose $m > 0$ with $m \in \text{poly}(n)$ such that $\left(\frac{2\lfloor q/p \rfloor}{q_j}\right)^{m/n} \cdot 2\beta$ is smaller than some constant $c < 1$. Then, it holds that*

$$\Pr[\exists \mathbf{s}, \hat{\mathbf{s}} \in B_\beta^n : \mathbf{s} \neq \hat{\mathbf{s}}, \lfloor \mathbf{s}^t \mathbf{A} \rfloor_p = \lfloor \hat{\mathbf{s}}^t \mathbf{A} \rfloor_p] \leq \left(\frac{q_j^{e_j-1}}{q}\right)^m \cdot (2\beta)^n \cdot (\lfloor q/p \rfloor)^m,$$

which is negligible in n .

Proof. Assume that we have $\mathbf{s}, \hat{\mathbf{s}} \in \mathbb{Z}_q^n$ with $\mathbf{s} \neq \hat{\mathbf{s}}$ and $\lfloor \mathbf{s}^t \mathbf{A} \rfloor_p = \lfloor \hat{\mathbf{s}}^t \mathbf{A} \rfloor_p$. We can interpret this LWR equation as an LWE equation with fixed errors. Since the rounding function divides \mathbb{Z}_q into p intervals of size $\lfloor q/p \rfloor$ and maps an element from \mathbb{Z}_q to the index of the interval, we know that by scaling this index by $\lfloor q/p \rfloor$ we get the rounded representation of x in \mathbb{Z}_q . Furthermore, since elements in the interval $[k \cdot \lfloor q/p \rfloor, (k+1) \cdot \lfloor q/p \rfloor - 1]$ are mapped to k and thus scaled to $k \lfloor q/p \rfloor$, we know that the difference between x and its rounded representation in \mathbb{Z}_q differ by at most $\lfloor q/p \rfloor$. Thus, there exist $\mathbf{e}, \hat{\mathbf{e}} \in \mathbb{Z}_q^m$ with $\|\mathbf{e}\|_\infty, \|\hat{\mathbf{e}}\|_\infty \leq \lfloor q/p \rfloor$ such that $\mathbf{s}^t \mathbf{A} + \mathbf{e} = \lfloor \mathbf{s}^t \mathbf{A} \rfloor_p \cdot \lfloor q/p \rfloor \pmod q$ and $\hat{\mathbf{s}}^t \mathbf{A} + \hat{\mathbf{e}} = \lfloor \hat{\mathbf{s}}^t \mathbf{A} \rfloor_p \cdot \lfloor q/p \rfloor \pmod q$. Then, we know that

$$\begin{aligned} (\mathbf{s}^t - \hat{\mathbf{s}}^t) \mathbf{A} + \mathbf{e} - \hat{\mathbf{e}} &= \mathbf{0} \\ \Leftrightarrow (\mathbf{s}^t - \hat{\mathbf{s}}^t) \mathbf{A} &= -\mathbf{e} + \hat{\mathbf{e}} \\ \Leftrightarrow \bar{\mathbf{s}}^t \mathbf{A} &= \bar{\mathbf{e}} \end{aligned}$$

for some $\bar{\mathbf{s}} \in B_{2\beta}^n$ and some $\bar{\mathbf{e}} \in B_{2\lfloor q/p \rfloor}^m$. Thus, we can apply [9, Theorem 6] to say that such $\bar{\mathbf{s}}, \bar{\mathbf{e}}$ exist only with probability

$$\begin{aligned} \epsilon(n) &:= \Pr[\exists(\bar{\mathbf{s}}, \bar{\mathbf{e}}) \in B_{2\beta}^n \times B_{2\lfloor q/p \rfloor}^m : \bar{\mathbf{s}}^t \mathbf{A} = \bar{\mathbf{e}}^t; \mathbf{A} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^{n \times m}] \\ &\leq \left(\frac{1 + z_q^{\max}}{q} \right)^m \cdot (2\beta)^n \cdot (2\lfloor q/p \rfloor)^m, \end{aligned}$$

where $z_w^{\max} := \max\{|\mathcal{Z}_{w,s}| : s \in \mathbb{Z}_w \setminus \{0\}\}$ with $\mathcal{Z}_{w,s} := \{a \in \mathbb{Z}_w \setminus \{0\} : a \cdot s = 0\}$. Since $q = \prod_{i=1}^d q_i^{e_i}$, we know that $\mathbb{Z}_q \cong \mathbb{Z}_{q_1}^{e_1} \times \dots \times \mathbb{Z}_{q_d}^{e_d}$. Furthermore, we know that $z_{q_i}^{\max} = q_i^{e_i-1} - 1$ for all i , since there are $q_i^{e_i-1} - 1$ non-zero multiples of q_i that are smaller than $q_i^{e_i}$ and $q_i^{e_i-1}$ times any other multiple is zero. Therefore, an element $(0, \dots, 0, q_j^{e_j-1}, 0, \dots, 0)$, where the index j is such that $q_j^{e_j} = \min_i \{q_i^{e_i}\}$, maximises z_q^{\max} . Thus, we have

$$z_{\max} = -1 + q_j^{e_j-1} \prod_{i \neq j} q_i = \frac{q}{q_j} - 1$$

Using this and the definition of m on the bound for $\epsilon(n)$, we get

$$\begin{aligned} \epsilon(n) &\leq \frac{1}{q_j^m} \cdot (2\beta)^n \cdot (2\lfloor q/p \rfloor)^m \\ &\leq \left(\left(\frac{2\lfloor q/p \rfloor}{q_j} \right)^{m/n} \cdot 2\beta \right)^n \\ &\leq c^n, \end{aligned}$$

which is negligible in n since $c < 1$. □

Corollary 81. *For q prime, $n, p, \beta > 0$ with $p < q$ and $0 < \beta < q$. Choose m such that $4\beta < (2p)^{-m/n}$. Let χ be a distribution over \mathbb{Z}_q such that for all x in the support of χ we have $x \leq \beta$. Then, the secrets of $\text{LWR}_{n,q,p,\chi}$ are unique with overwhelming probability.*

The construction of the LWR-based LIT is, as previously mentioned, very similar to the LWE-based construction of [9] with errors replaced by rounding.

Construction 82 *For q prime, $n, p, \beta > 0$ with $p < q$ and $0 < \beta < q$. Choose m such that $4\beta < (2p)^{-m/n}$. Let χ be a distribution over \mathbb{Z}_q such that for all x in the support of χ we have $x \leq \beta$. For a uniform $\mathbf{A} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^{n \times m}$, let $f_{\mathbf{A}} : \mathbb{Z}_q^n \rightarrow \mathbb{Z}_q^m, \mathbf{x} \mapsto \lfloor \mathbf{x}^t \mathbf{A} \rfloor_p$. Then, the LIT Π_{LIT} is defined as in Fig. 24.*

We now need to prove that our construction has the four required security properties tag-indistinguishability (Def. 14), non-invertability (Def. 17), linkability (Def. 15), and unforgeability (Def. 16). Since the proofs are very similar to the proofs for the module LWE-based LIT of [9], we only give an outline for each one.

$\frac{\text{KGen}(1^\lambda)}{\text{s} \leftarrow \mathcal{X}^k}$ $sk \leftarrow \text{s}$ $\frac{\text{Tag}(sk, \mu)}{\text{A}_\mu \leftarrow \mathcal{RO}(\mu) \in \mathbb{Z}_q^{k \times m}}$ $\mathbf{b}^t \leftarrow \lfloor \text{s}^t \mathbf{A} \rfloor_p$ $t \leftarrow \mathbf{b}$ $\text{return } t$	$\frac{\text{Vrfy}(sk, \mu, t)}{\text{A}_\mu \leftarrow \mathcal{RO}(\mu) \in \mathbb{Z}_q^{k \times m}}$ $\hat{\mathbf{b}}^t \leftarrow \lfloor \text{s}^t \mathbf{A} \rfloor_p$ $\text{if } t = \hat{\mathbf{b}}$ $\quad \text{return } 1$ $\text{return } 0$ $\frac{\text{Link}(\mu, t_0, t_1)}{\text{if } t_0 = t_1}$ $\quad \text{return } 1$ $\text{return } 0$
---	--

Fig. 24: Learning with Errors-based LIT.

Lemma 83. *Construction 82 has tag-indistinguishability in the random oracle model if $\text{LWR}_{n,q,p,\chi}$ is hard.*

To prove this, one uses the LWR assumption twice to replace the public keys and tags generated with sk_0 and sk_1 by uniformly random values, once in $\text{Anon}_{\text{LIT}, \mathcal{A}, 0}^{\text{LIT}}(\lambda)$ and once in $\text{Anon}_{\text{LIT}, \mathcal{A}, 1}^{\text{LIT}}(\lambda)$. Then, the modified games are equal and thus the adversary has advantage 0 to distinguish between them.

Lemma 84. *Construction 82 has non-invertability if $\text{sLWR}_{n,q,p,\chi}$ is hard.*

To prove this, one uses LWR samples to generate the public key and the tag oracle answers. When the adversary returns an sk' , we know by Corollary 81 that the LWR secrets are unique, meaning we have $sk' = sk$, therefore sk' is a solution for the LWR challenge.

Lemma 85. *Construction 82 is linkable in the random oracle model.*

In order for the adversary to be able to win, he needs to output sk_0, sk_1 such that $f(sk_0) = f(sk_1)$. By Corollary 81 we know that LWR secrets are unique, thus we have $sk_0 = sk_1$. Now the tags t_0, t_1 output by the adversary must both verify under pk_0 for the message μ also output by it. Since tagging is deterministic, this can only be if $t_0 = t_1$, thus we have that t_0, t_1 link and the adversary always loses.

Lemma 86. *Construction 82 is unforgeable in the random oracle model if $\text{sLWR}_{n,q,p,\chi}$ is hard.*

To prove this, one uses LWR samples to generate the public key and the tag oracle answers. When the adversary outputs some sk^*, μ, t^* , we know $t = t^*$ since $\text{Linkable}^{\text{LIT}}(\mu, t, t^*) = 1$. Furthermore, by Corollary 81 we know that LWR secrets are unique, thus we know that $sk^* = sk_0$ and sk^* is a solution to the LWR challenge.

F.2 Proof of Theorem 30

Proof. Let \mathcal{A} be an adversary against the full collision-resistance of the construction. Let $s' = \omega(\sqrt{\log m})$. We define an alternative security game for full collision-resistance called Game_1 in Fig. 25, where we neither need the trapdoor of the commitment scheme nor the trapdoor td of \mathbf{A} . This is achieved by leveraging the power of the random oracle, i.e., by being able to program the random oracle. In this game, we also remember some additional values, which will be used to define the event **Free** mentioned above and to extract either an SIS solution or to break the binding of the commitment scheme. In the original security game, we can implement the random oracle by choosing uniform outputs and storing what we answered so far in a set $\mathcal{Q}_{\mathcal{RO}}$. In Game_1 , if we answer a random oracle query, we generate $\mathbf{B} = \mathbf{A}\mathbf{U}$, i.e., we do not put a trapdoor into \mathbf{B} . We indicate that by also storing \mathbf{U} and an indicator td_{no} . If, on the other hand, we are answering an $\text{Adapt}\mathcal{O}$ oracle call, we generate $\mathbf{B} = \mathbf{A}\mathbf{R} + \mathbf{G}$ with a trapdoor. Again, we indicate this, this time by storing \mathbf{R} and an indicator td_{yes} .

In the original game, \mathcal{Q} only stored for which (h, μ) the adversary has seen a collision. In Game_1 , \mathcal{Q} also remembers if the (h, μ) is "fresh", denoted by \mathfrak{f} : if the adversary wants $\text{Adapt}\mathcal{O}$ to not return \perp , it needs to call the oracle with a valid (h, μ) pair as the first half of the argument. If the adversary has not seen a collision on this (h, μ) , it is marked as fresh. If that is the case, \mathcal{Q} also remembers the r that was used to generate the h from μ . Pairs (h, μ') for which the adversary has seen a collision before (this includes the μ' of a $\text{Adapt}\mathcal{O}$ query) are marked as "old", denoted by \mathfrak{o} . Note that by this definition, for every old (h, μ') , there must exist a fresh (h, μ) , because in order for the adversary to see the first collision on h , it has to query $\text{Adapt}\mathcal{O}$ with a valid hash, which is then considered fresh.

We can argue that the advantages of \mathcal{A} in the original and the alternative game are negligibly close. This is due to Theorem 56, Theorem 9, and due to the DEE property of the commitment scheme, as well as z being chosen from a superpoly large space (such that no z is repeated with overwhelming probability). For a similar proof, see [31].

We now define the event **Free**, by the following four subevents:

FirstFree is true if there exist \mathbf{B}, \mathbf{U} s.t. $((\mu, z, c), \mathbf{B}, \mathbf{U}, \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$, where $(\cdot, \mu, r = (z, \mathbf{e}, c, d), \cdot, \cdot)$ is output by \mathcal{A} .

SecondQueried is true if $(h, \mu') \in \mathcal{Q}$, where $(h, \cdot, \cdot, \mu', \cdot)$ is output by \mathcal{A} .

SecondFree is true if there exist \mathbf{B}', \mathbf{U}' s.t. $((\mu', z', c'), \mathbf{B}', \mathbf{U}', \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$, where $(\cdot, \cdot, \cdot, \mu', r' = (z', \mathbf{e}', c', d'))$ is output by \mathcal{A} .

FreshFree is true if there exist $\hat{\mu}, \hat{z}, \hat{c}, \hat{\mathbf{B}}, \hat{\mathbf{U}}, \hat{\mathbf{e}}$ s.t. $((\hat{\mu}, \hat{z}, \hat{c}), \hat{\mathbf{B}}, \hat{\mathbf{U}}, \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$ and $(h, \hat{\mu}, \mathfrak{f}, \hat{r} = (\hat{z}, \hat{\mathbf{e}}, \hat{c}, \hat{d})) \in \mathcal{Q}$, where $(h, \cdot, \cdot, \cdot, \cdot)$ is output by \mathcal{A} .

Then, we can define

$$\text{Free} = \text{FirstFree} \wedge (\neg \text{SecondQueried} \wedge \text{SecondFree} \vee \text{SecondQueried} \wedge \text{FreshFree}).$$

We will argue that if **Free** holds and the adversary wins, we can reduce to SIS. The idea is if **Free** holds, we can find a collision $((\mu, r), (\mu', r'))$ such that \mathbf{B}, \mathbf{B}' ,

$\text{Game}_1(1^\lambda)$	$\text{Adapt}\mathcal{O}'(\text{sk}_{\text{ch}}, \cdot, \cdot, \cdot, \cdot)$ with (h, μ, r, μ')
$\mathcal{Q}, \mathcal{Q}_{\mathcal{RO}} \leftarrow \emptyset$ $\mathbf{A} \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$ $(\text{pp}, \cdot) \leftarrow_{\$} \text{TdGen}(1^\lambda)$ $\text{pk}_{\text{ch}} \leftarrow (\mathbf{A}, \text{pp})$ $(h, \mu, r, \mu', r') \leftarrow_{\$} \mathcal{A}^{\mathcal{RO}', \text{Adapt}\mathcal{O}'}(\text{pk}_{\text{ch}})$ return $(\text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu, r) = 1$ $\wedge \text{CHashCheck}(\text{pk}_{\text{ch}}, h, \mu', r') = 1$ $\wedge \mu \neq \mu'$ $\wedge (h, \mu, \cdot, \cdot) \notin \mathcal{Q})$	if $\text{CHashCheck}(pk, h, \mu, r) = 0$ return \perp $z' \leftarrow_{\$} \{0, 1\}^\lambda$ $\mathbf{R} \leftarrow_{\$} D_{\mathbb{Z}^{2m}, s'}$ $\mathbf{B} \leftarrow_{\$} \mathbf{A}\mathbf{R} + \mathbf{G}$ $\mathbf{e}' \leftarrow_{\$} \text{PreSample}([\mathbf{A} \ \mathbf{B}], \mathbf{R}, h, s)$ $(c', d') \leftarrow_{\$} \text{Com}(\text{pp}, \mathbf{e}')$ $x \leftarrow (\mu', z', c')$ $r' \leftarrow (z', \mathbf{e}', c', d')$ $\mathcal{Q}_{\mathcal{RO}} \leftarrow \mathcal{Q}_{\mathcal{RO}} \cup \{(x, \mathbf{B}, \mathbf{R}, \text{td}_{\text{yes}}, r')\}$ if $\nexists (h, \mu, \cdot, \cdot) \in \mathcal{Q}$ $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(h, \mu, \cdot, r)\}$ $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(h, \mu', \cdot, \perp)\}$ return r'
$\mathcal{RO}'(\cdot)$ with $x \in \{0, 1\}^*$	
if $\exists (x, \mathbf{B}, \cdot, \cdot) \in \mathcal{Q}_{\mathcal{RO}}$ return \mathbf{B} $\mathbf{U} \leftarrow_{\$} D_{\mathbb{Z}^{2m}, s'}$ $\mathbf{B} \leftarrow_{\$} \mathbf{A}\mathbf{U}$ $\mathcal{Q}_{\mathcal{RO}} \leftarrow \mathcal{Q}_{\mathcal{RO}} \cup \{(x, \mathbf{B}, \mathbf{U}, \text{td}_{\text{no}}, \perp)\}$ return \mathbf{B}	

Fig. 25: The alternative security game Game_1 .

which are defined by the first and second pair of the collision, respectively, do not contain a gadget \mathbf{G} . This collision always includes the (μ, r) output by the adversary, but the second element is either the (μ', r') output by the adversary or the fresh pair $(\hat{\mu}, \hat{r})$ if (μ', r') is old. Then, we can combine the collision into a solution for SIS. On the other hand, if **Free** does not hold, we are able to break the binding of the commitment scheme. To show this, we construct two adversaries, \mathcal{B} against SIS in Fig. 26 and \mathcal{C} against the binding of the commitment scheme in Fig. 27 that both simulate \mathcal{A} .

Let (h, μ, r, μ', r') be the forgery output by \mathcal{A} with $r = (z, \mathbf{e}, c, d)$ and $r' = (z', \mathbf{e}', c', d')$. If **Free** happens, we can easily see that \mathcal{B} is able to find $\mathbf{B}, \mathbf{e}, \hat{\mathbf{B}}, \hat{\mathbf{e}}$, where the matrices are free of \mathbf{G} and such that

$$[\mathbf{A} \ \mathbf{B}] \mathbf{e} = [\mathbf{A} \ \hat{\mathbf{B}}] \hat{\mathbf{e}}.$$

Therefore, we can easily see that $\mathbf{A}\mathbf{v} = \mathbf{0}$ and that

$$\begin{aligned} \|\mathbf{v}\| &= \|[\mathbf{I} \ \mathbf{U}] \mathbf{e} - [\mathbf{I} \ \hat{\mathbf{U}}] \hat{\mathbf{e}}\| \\ &\leq (s_1(\mathbf{U}) + 1) \cdot \beta + (s_1(\hat{\mathbf{U}}) + 1) \cdot \beta \\ &\leq 2s' \cdot (\sqrt{n} + \sqrt{m} + t + 1) \cdot s\sqrt{2m} \end{aligned}$$

```

 $\mathcal{B}(\mathbf{A})$ 


---


 $\mathcal{Q}, \mathcal{Q}_{\mathcal{RO}} \leftarrow \emptyset$ 
 $(\text{pp}, \cdot) \leftarrow_{\S} \text{TdGen}(1^\lambda)$ 
 $\text{pk}_{\text{ch}} \leftarrow (\mathbf{A}, \text{pp})$ 
 $(h, \mu, r, \mu', r') \leftarrow_{\S} \mathcal{A}^{\mathcal{RO}', \text{Adapt}\mathcal{O}'}(\text{pk}_{\text{ch}})$ 
  with  $r = (z, \mathbf{e}, c, d), r' = (z', \mathbf{e}', c', d')$ 
if FirstFree
  Let  $\mathbf{B}, \mathbf{U}$  be s.t.  $((\mu, z, c), \mathbf{B}, \mathbf{U}, \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$ 
if SecondQueried  $\wedge$  SecondFree
  Let  $\mathbf{B}', \mathbf{U}'$  be s.t.  $((\mu', z', c'), \mathbf{B}', \mathbf{U}', \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$ 
  return  $\mathbf{v} \leftarrow [\mathbf{I} \mathbf{U}] \mathbf{e} - [\mathbf{I} \mathbf{U}'] \mathbf{e}'$ 
if  $\neg$ SecondQueried  $\wedge$  FreshFree
  Let  $\hat{\mu}, \hat{z}, \hat{c}, \hat{\mathbf{B}}, \hat{\mathbf{U}}, \hat{\mathbf{e}}$  be s.t.  $(h, \hat{\mu}, \mathbf{f}, (\hat{z}, \hat{\mathbf{e}}, \hat{c}, \cdot)) \in \mathcal{Q}$ 
  and  $((\hat{\mu}, \hat{z}, \hat{c}), \hat{\mathbf{B}}, \hat{\mathbf{U}}, \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$ 
  return  $\mathbf{v} \leftarrow [\mathbf{I} \mathbf{U}] \mathbf{e} - [\mathbf{I} \hat{\mathbf{U}}] \hat{\mathbf{e}}$ 
return  $\perp$ 

```

Fig. 26: Adversary against SIS.

by definition of β and Theorem 57. Furthermore, by definition of `CHashCheck` we know that $\mathbf{0} \neq \mathbf{e}_2, \hat{\mathbf{e}}_2 \in \mathbb{Z}_q^m$, where $(\mathbf{e}_1^t, \mathbf{e}_2^t)^t = \mathbf{e}^t, (\hat{\mathbf{e}}_1^t, \hat{\mathbf{e}}_2^t)^t = \hat{\mathbf{e}}^t$. Due to this fact, together with the distributions of $\mathbf{U}, \hat{\mathbf{U}}$ having high min-entropy and a standard entropy argument (see [31] for an example), we can argue that $\mathbf{v} \neq \mathbf{0}$. Thus, \mathbf{v} is an SIS solution and we know that

$$\Pr[\text{f-CR}_{\text{CH}, \mathcal{A}}(1^\lambda) = 1 \wedge \text{Free}] = \text{SIS}_{n, m, q, \beta'}(\mathcal{B}).$$

If \neg Free happens, we make a case distinction which part of the event is not true.

Case 1 (\neg FirstFree): In this case, we know that $\nexists((\mu, z, c), \mathbf{B}, \mathbf{R}, \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$. We assume wlog. that the adversary queries the random oracle (or by extension, `AdaptO`) to get any \mathbf{B} , because else it has to predict the random oracle. Since $\nexists((\mu, z, c), \mathbf{B}, \mathbf{R}, \text{td}_{no}, \perp) \in \mathcal{Q}_{\mathcal{RO}}$ but \mathcal{A} has output a valid hash on μ , we know that \mathbf{B} must have been generated in an `AdaptO` call. Therefore $\exists((\mu, z, c), \mathbf{B}, \mathbf{R}, \text{td}_{yes}, r^*) \in \mathcal{Q}_{\mathcal{RO}}$ which \mathcal{B} can find. Furthermore, since that tuple was generated during an `AdaptO` call, we know that $(\mathbf{B}, \mathbf{e}^*, \mu, \cdot, \cdot) \in \mathcal{Q}$. However, by definition of the security game we know that $(h, \mu, \cdot, \cdot) \notin \mathcal{Q}$. Therefore $\mathbf{e} \neq \mathbf{e}^*$ and \mathcal{C} has broken the binding of the commitment.

Case 2 (\neg SecondQueried \wedge \neg SecondFree): In this case, since $(h, \mu', \cdot, \cdot) \notin \mathcal{Q}$ we argue as in Case 1.

Case 3 (SecondQueried \wedge \neg FreshFree): In the case that $(h, \mu', \cdot, \cdot) \in \mathcal{Q}$ we first argue that there must exist some $(h, \hat{\mu}, \mathbf{f}, \hat{r}) \in \mathcal{Q}$. Either $(h, \mu', \mathbf{f}, \hat{r}) \in \mathcal{Q}$ and

```

C(pp)


---


 $\mathcal{Q}, \mathcal{Q}_{\mathcal{RO}} \leftarrow \emptyset$ 
 $\mathbf{A} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^{n \times m}$ 
 $\text{pk}_{\text{ch}} \leftarrow (\mathbf{A}, \text{pp})$ 
 $(h, \mu, r, \mu', r') \leftarrow_{\mathcal{S}} \mathcal{A}^{\mathcal{RO}', \text{Adapt}\mathcal{O}'}(\text{pk}_{\text{ch}})$ 
  with  $r = (z, \mathbf{e}, c, d), r' = (z', \mathbf{e}', c', d')$ 
if  $\neg \text{FirstFree}$ 
  Let  $\mathbf{e}^*, d^*$  be s.t.  $((\mu, z, c), \cdot, \cdot, \text{td}_{\text{yes}}, (z, \mathbf{e}^*, c, d^*)) \in \mathcal{Q}_{\mathcal{RO}}$ 
  return  $(c, \mathbf{e}, d, \mathbf{e}^*, d^*)$ 
if  $\text{SecondQueried} \wedge \neg \text{SecondFree}$ 
  Let  $\mathbf{e}^*, d^*$  be s.t.  $((\mu', z', c'), \cdot, \cdot, \text{td}_{\text{yes}}, z', \mathbf{e}^*, c', d^*) \in \mathcal{Q}_{\mathcal{RO}}$ 
  return  $(c', \mathbf{e}', d', \mathbf{e}^*, d^*)$ 
if  $\neg \text{SecondQueried} \wedge \neg \text{FreshFree}$ 
  Let  $\hat{\mu}, \hat{z}, \hat{c}, \hat{d}, \bar{h}, \bar{\mathbf{e}}, \bar{d}$  be s.t.  $(h, \hat{\mu}, \mathbf{f}, \cdot) \in \mathcal{Q}$ 
  such that  $((\hat{\mu}, \hat{z}, \hat{c}), \cdot, \cdot, \text{td}_{\text{yes}}, (\hat{z}, \hat{\mathbf{e}}, \hat{c}, \hat{d})) \in \mathcal{Q}_{\mathcal{RO}}$ 
  and  $(\bar{h}, \hat{\mu}, \mathbf{f}, (\hat{z}, \bar{\mathbf{e}}, \hat{c}, \bar{d})) \in \mathcal{Q}$ 
  return  $(\hat{c}, \hat{\mathbf{e}}, \hat{d}, \bar{\mathbf{e}}, \bar{d})$ 
return  $\perp$ 

```

Fig. 27: Adversary against the binding of the commitment scheme.

we are done. Else we have added $(h, \mu', \mathbf{o}, \perp)$ to \mathcal{Q} since $(h, \mu', \cdot, \cdot) \in \mathcal{Q}$, which happened during an $\text{Adapt}\mathcal{O}(h, \mu^*, r^*, \mu')$ query. This means that μ' was the message for which $\text{Adapt}\mathcal{O}$ was supposed to find a collision, i.e., it was the second message during the oracle call. During that same oracle call either $(h, \mu^*, \mathbf{f}, r^*)$ was added to \mathcal{Q} , or it was already the case that $(h, \mu^*, \cdot, \cdot) \in \mathcal{Q}$. In the latter case, we can recursively apply the above argument until we find $(h, \hat{\mu}, \mathbf{f}, \hat{r})$.

Once we have $(h, \hat{\mu}, \mathbf{f}, \hat{r})$, since $\neg \text{FreshFree}$ holds and by our assumption that \mathcal{A} does not try to predict the random oracle, we know that $\mathcal{RO}(\hat{\mu}, \hat{z}, \hat{c}) = \hat{\mathbf{B}} = \mathbf{A}\hat{\mathbf{R}} + \mathbf{G}$. But since $\hat{\mathbf{B}}$ has this form, it must have been created during an $\text{Adapt}\mathcal{O}$ call. Because we argued that $(\hat{h}, \hat{\mu}, \mathbf{f}, \cdot)$ is fresh, this must mean when $\hat{\mathbf{B}}$ was created, it was created for some different hash \bar{h} . Thus, there exists some $((\bar{\mu}, \bar{z}, \bar{c}), \hat{\mathbf{B}}, \cdot, \text{td}_{\text{yes}}, \bar{r}) \in \mathcal{Q}_{\mathcal{RO}}$ with $\hat{r} = (\hat{z}, \hat{\mathbf{e}}, \hat{c}, \hat{d}) \neq \bar{r} = (\hat{z}, \bar{\mathbf{e}}, \bar{c}, \bar{d})$. However, due to $\mathbf{A}\hat{\mathbf{R}}$ being statistically close to uniform, we have with overwhelming probability that the inputs to the oracle that correspond to \hat{h}, \bar{h} were equal, i.e., $\hat{\mu} = \bar{\mu}, \hat{z} = \bar{z}, \hat{c} = \bar{c}$. Now since $h = \hat{\mathbf{B}}\hat{\mathbf{e}} \neq \hat{\mathbf{B}}\bar{\mathbf{e}} = \bar{h}$, we know that $\hat{\mathbf{e}} \neq \bar{\mathbf{e}}$, which are the messages committed in \hat{c}, \bar{c} respectively. Thus, \mathcal{C} breaks the binding of the commitment scheme. Therefore, we know that

$$\Pr[\text{f-CR}_{\text{CH}, \mathcal{A}}(1^\lambda) = 1 \wedge \neg \text{Free}] = \Pr[\text{Bind}_{\Pi_{\text{com}}, \mathcal{C}}(\lambda) = 1]$$

and thus

$$\Pr[\text{f-CR}_{\text{CH},\mathcal{A}}(1^\lambda) = 1] = \text{SIS}_{n,m,q,\beta'}(\mathcal{B}) + \Pr[\text{Bind}_{\Pi_{\text{com}},\mathcal{C}}(\lambda) = 1]. \quad \square$$

F.3 Proof of Theorem 46

Proof. For simplicity, we only consider the case of two blocks. The adversary against signer accountability is given a public key pk_{San} of a sanitizer and has access to a sanitize oracle.

The attacker honestly generates a key pair $(pk_{\text{Sig}}, sk_{\text{Sig}})$ of the signature scheme Σ . The attacker picks randomly a message $\mu = (\mu_1, \mu_2)$, another message $\mu'[2]$ such that $\mu_2 \neq \mu'[2]$ and sets $\mu' = (\mu_1, \mu'[2])$, a random NONCE, and an auxiliary value $r[2]$. Then, TAG is set to $\text{PRG}(x)$ where $x := \text{PRF}(\kappa, \text{NONCE})$, ADM is defined to contain the information that only block 2 is admissible for change, MOD is set such that $\text{MOD}(\mu) = \mu'$. Finally, it sets σ_0 , the underlying signature, to $\text{Sign}(sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM}, \text{TAG}, \mu, r[2])$.

The attacker queries the sanitize oracle with message μ , modifications MOD, and signature $(\sigma_0, \text{TAG}, \text{NONCE}, \text{ADM}, r[2])$, and receives NONCE' , TAG' , and $r'[2]$, which constitute the new signature. In particular, it holds that

$$\begin{aligned} \text{CH}(pk_{\text{San}}, \text{TAG}, (2, \mu_2, pk_{\text{Sig}}), r[2]) &= \\ \text{CH}(pk_{\text{San}}, \text{TAG}', (2, \mu'[2], pk_{\text{Sig}}), r'[2]). & \end{aligned} \quad (1)$$

The attacker returns $pk_{\text{Sig}}^* = pk_{\text{Sig}}$, $\mu^* = (\mu'[1], \mu'[2])$, for some $\mu'[1] \neq \mu_1$, $\text{TAG}^* = \text{TAG}'$, $r^*[2] = r'[2]$, $\text{NONCE}^* = \text{NONCE}'$, σ^* the signature $\text{Sign}(sk_{\text{Sig}}, pk_{\text{San}}, \text{ADM}, \text{TAG}^*, \mu^*, r^*[2])$, and the proof π^* given by $pk_{\text{Sig},\pi} = pk_{\text{Sig}}$, $\text{TAG}_\pi = \text{TAG}$, $(2, \mu_\pi[2], pk_{\text{Sig}})$, where $\mu_\pi[2] = \mu_2$, $r_\pi[2] = r[2]$ and $x_\pi = x$.

The attacker wins, as the following conditions are satisfied:

- (i) $\mu^* \neq \mu$ and $\mu^* \neq \mu'$,
- (ii) $\text{Vrfy}(pk_{\text{Sig}}^*, pk_{\text{San}}, \mu^*, \sigma^*) = 1$ as the μ^* was honestly signed,
- (iii) $\text{Judge}(\mu^*, \sigma^*, pk_{\text{Sig}}^*, pk_{\text{San}}, \pi^*) = \text{Sanit}$,

where the final step holds as

- (a) $pk_{\text{Sig}}^* = pk_{\text{Sig}} = pk_{\text{Sig},\pi}$,
- (b) $\text{CH}(pk_{\text{San}}, \text{TAG}^*, (2, \mu^*[2], pk_{\text{Sig}}^*), r^*[2])$
 $= \text{CH}(pk_{\text{San}}, \text{TAG}_\pi, (2, \mu_\pi, pk_{\text{Sig},\pi}), r_\pi[2])$,
- (c) Block 2 is admissible,
- (d) $\text{TAG}_\pi = \text{PRG}(x_\pi)$.

Indeed, (a), (c), and (d) hold by definition, (b) is Equation (1). \square

Remark 87. The attacker does not need to construct a new hash collision, but makes use of a collision produced by the sanitizer. The reason the present construction is vulnerable is that it suffices to modify the message in a single block to ensure condition (i), although the judge only tests a single block for collisions. Thus, the attacker can modify the message in one block to pass the check in condition (i), and use a collision generated by the sanitizer to convince the judge in condition (iii).