





Lattice-based Proof-Friendly Signatures from Vanishing Short Integer Solutions

Adrien Dubois¹, Michael Kloob², Russell W. F. Lai³, and Ivy K. Y. Woo³

¹ ENS de Lyon, France

² Karlsruhe Institute of Technology, Germany

³ Aalto University, Finland

Abstract. Efficient anonymous credentials are typically constructed by combining proof-friendly signature schemes with compatible zero-knowledge proof systems. Inspired by pairing-based proof-friendly signatures such as Boneh- Boyen (BB) and Boneh-Boyen-Shacham (BBS), we propose a wide family of lattice-based proof-friendly signatures based on variants of the vanishing short integer solution (vSIS) assumption [Cini-Lai-Malavolta, Crypto’23]. In particular, we obtain natural lattice-based adaptations of BB and BBS which, similar to their pairing-based counterparts, admit nice algebraic properties.

[Bootle-Lyubashevsky-Nguyen-Sorniotti, Crypto’23] (BLNS) recently proposed a framework for constructing lattice-based proof-friendly signatures and anonymous credentials, based on another new lattice assumption called ISIS_f parametrised by a fixed function f , with focus on f being the binary decomposition. We introduce a generalised ISIS_f framework, called GenISIS_f , with a keyed and probabilistic function f . For example, picking $f_b(\mu) = 1/(b - \mu)$ with key b for short ring element μ leads to algebraic and thus proof-friendly signatures. To better gauge the robustness and proof-friendliness of $(\text{Gen})\text{ISIS}_f$, we consider what happens when the inputs to f are chosen selectively (or even adaptively) by the adversary, and the behaviour under relaxed norm checks. While bit decomposition quickly becomes insecure, our proposed function families seem robust.

Keywords: lattice cryptography · proof-friendly signatures · BBS signature · vanishing SIS · ISIS_f assumption

1 Introduction

Constructing secure and concretely efficient lattice-based signature schemes is by now a well solved problem. Indeed, schemes following both of the main construction paradigms, Hash-and-Sign [GPV08] and Fiat-Shamir-with-abort (FSwA) [Lyu12], have been standardised [PFH+22, LDK+22] with signature size in single-digit kilobytes and with security connected to the hardness of worst-case lattice problems⁴ in the random oracle model. However, the verification relations of most efficient schemes in either paradigm, including standardised ones, inherently require evaluating a hash function on the signed message and additionally on part of the signature in the case of FSwA. Since this hash function needs to be modelled as a random oracle for security proofs to go through, it is typically instantiated with a *non-algebraic* (hence not proof-friendly) hash function. While signatures based on non-algebraic hashes suffice for standalone uses, they may not be well suited as building blocks for efficient constructions of privacy-preserving authentication primitives, such as anonymous credentials, blind signatures and group signatures.

Constructing efficient privacy-preserving authentication primitives is of high practical interest, e.g. in the context of the European Digital Identity framework. Concretely efficient constructions typically require a user to prove knowledge of a message-signature pair which satisfies the verification relation of a signature scheme, in zero-knowledge. For example, the message could be the secret attributes of the user, and the signature could be issued by an authority who asserts that the attributes are genuine. For such a proof to be computed efficiently, a common approach is to instantiate the constructions with “proof-friendly” signatures and zero-knowledge proofs (ZKP), such that the verification relations of the signatures are “natively” supported by the proof system.⁵

⁴ We say connected because parameters chosen for the schemes differ from those which admit worst-case to average-case reductions.

⁵ Such a combination of signatures and ZKPs is sometimes called “signatures with efficient protocols” [CL03]. In the pairing-based setting, a typical choice is to combine the BBS signatures [BBS04] with the Groth-Sahai proof system [GS08].

Lattice-based Proof-Friendly Signatures. In the context of lattice-based signatures, we regard a signature scheme as *proof-friendly* if 1) it natively supports signing committed messages (via a hiding and binding commitment), and 2) its verification relation can be expressed as the bounded-norm satisfiability of a system of low-degree polynomial equations. Combined with efficient lattice-based ZKPs for proving well-formedness of commitments and bounded-norm relations (e.g. [LNP22b]), a signature scheme with the above properties can be efficiently turned into constructions of privacy-preserving authentication primitives, as demonstrated in [BLNS23b]. In this area, two competing approaches represent the state of the art:

1. Jeudy, Roux-Langlois and Sanders [JRS23], building upon [LLM⁺16], considered signature schemes of the following form: A signature of a short message vector \mathbf{m} is a tuple $(x, \mathbf{s}, \mathbf{r})$, where x is an invertible element, and \mathbf{s}, \mathbf{r} are short vectors satisfying

$$[\mathbf{A}|\mathbf{B} + x\mathbf{G}] \cdot \mathbf{s} = \mathbf{v} + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \bmod q,$$

where $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are public random matrices, \mathbf{v} is a public random vector, and \mathbf{G} is the so-called gadget matrix [MP12]. This type of signatures relies on the gadget lattice trapdoor machinery [MP12], which tends to be concretely less efficient than GPV trapdoors [GPV08]. Indeed, [JRS23] reported signature sizes in the hundreds of kilobytes.

2. Bootle, Lyubashevsky, Nguyen and Sorniotti [BLNS23b] considered signature schemes of the following form: A signature of a short message vector \mathbf{m} is a tuple $(x, \mathbf{s}, \mathbf{r})$, where x is chosen uniformly at random from an appropriate domain, and \mathbf{s}, \mathbf{r} are short vectors satisfying

$$\mathbf{A} \cdot \mathbf{s} = f(x) + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \bmod q,$$

where \mathbf{A}, \mathbf{C} are public random matrices and f is a function. This signature scheme can be instantiated efficiently, with [BLNS23b] reporting signature sizes as low as dozens of kilobytes. However, the security of such scheme is based on a new lattice assumption, called ISIS_f , introduced in the same work, whose hardness crucially depends on the choice of f . Indeed, it is very easy to come up with (linear) functions f for which the assumption and the scheme are completely broken. The authors advocated picking f to be the binary decomposition function, but were light on the evidence supporting the hardness of ISIS_f for this f . In Section 5.4, we illustrate that this choice of f is not very robust. Security breaks down under the (relatively benign) relaxation to one selective f -query and norm relaxation by a factor of $\sqrt{2}$.

Translating BB(S) Signatures. In view of the scarcity of lattice-based proof-friendly signatures, a natural strategy is to translate proof-friendly pairing-based signatures to the lattice setting, for example, using the general translation strategy proposed in [ACL⁺22]. Signature schemes which utilise only generic pairing group operations are abundant. Of particular importance are the related signature schemes of Boneh and Boyen (BB) [BB08], whose signature consists of a single group element, and Boneh, Boyen and Shacham (BBS) [BBS04, ASM06, TZ23], which allows to sign messages committed via Pedersen’s commitment. (More discussion in Section 1.2.) Below, we outline a translation attempt of the simpler BB signatures and highlight the difficulty behind.

To recall, using implicit notation for group elements, a public key in the BB signature scheme is a tuple of group elements $([1], [b]) \in \mathbb{G}^2$, the secret key is $b \in \mathbb{Z}_q$, and a signature of $\mu \in \mathbb{Z}_q \setminus \{b\}$ is $[u] = [1/(b - \mu)]$. Signature verification simply checks if $([b] - [1] \cdot \mu) \cdot [u] \stackrel{?}{=} [1]$, where \cdot denotes the pairing operation.

Adopting the translation strategy of [ACL⁺22], a natural lattice-analogue of BB signatures would be as follows: The public key consists of a random matrix \mathbf{A} and a random vector \mathbf{b} , the secret key is a trapdoor $\text{td}_{\mathbf{A}}$, and a signature of μ is a short vector \mathbf{s} satisfying $\mathbf{A}\mathbf{s} = \mathbf{1}_n \oslash (\mathbf{b} - \mathbf{1}_n \cdot \mu) \bmod q$, where \oslash denotes component-wise division. Equivalently, the verification equation is $(\mathbf{A}\mathbf{s}) \odot (\mathbf{b} - \mathbf{1}_n \cdot \mu) \stackrel{?}{=} \mathbf{1}_n \bmod q$, where \odot denotes the component-wise product, which apparently shares structural similarities with that of the BB signatures.

Despite the above natural translation of the BB signature scheme, its original security proof, based on the Q -strong Diffie-Hellman assumption (Q -SDH), unfortunately fails to translate to the lattice setting.

In brief, a core argument in the security proof of BB signatures relies on constructing a polynomial

$$f(\tilde{b}) := \prod_{i=1}^Q (\tilde{b} - \mu_i)$$

where μ_1, \dots, μ_Q are selective signing oracle queries, and using properties of the quotients $f(\tilde{b})/(\tilde{b} - \mu)$ for $\mu \in \{\mu_1, \dots, \mu_Q, \mu^*\}$, where μ^* is the target message of a forgery, to answer signing oracle queries and extract a Q -SDH solution. A major difficulty in carrying this argument over to the lattice setting, among others, lies in the inability to control the norm of the coefficients of both the quotient and remainder of $f(\tilde{b})/(\tilde{b} - \mu)$. This suggests that, rather than proving security of the lattice-BB signatures based on a lattice-analogue of Q -SDH, an alternative strategy is needed.

1.1 Our Contributions

In this work, we present a wide family of lattice-based proof-friendly signatures, including those obtained by translating the pairing-based BB and BBS signatures. We prove security of these signature schemes under new but natural extensions of existing lattice-based assumption, specifically, the (strong) hinted variants of the vanishing short integer solution (vSIS) assumption [CLM23] family, which can also be seen as variants of the kRISIS assumption [ACL⁺22] family with slightly more flexible adversaries. Our results are summarised in Figure 1.

(Strong) Hinted vSIS Assumptions, Plausibility Criteria, Reduction. We propose the hinted vSIS assumption and its strong variant in Section 3. The original vSIS assumption, introduced by [CLM23] and parametrised by a set of rational functions \mathcal{F} , asserts hardness of the following task:

Given a random matrix \mathbf{A} , find a short linear combination of $(f(\mathbf{A}))_{f \in \mathcal{F}}$ vanishing to $\mathbf{0}$ modulo q .

The hinted vSIS assumption, further parametrised by two (possibly intersecting) sets of rational functions \mathcal{G}, \mathcal{H} , asserts hardness of the following task:

Pick a Q -subset $\mathcal{Q} = (h_1, \dots, h_Q)$ of \mathcal{H} and some g^* in $\mathcal{G} \setminus \mathcal{Q}$, receive a random matrix \mathbf{A} and short linear combinations of $(f(\mathbf{A}))_{f \in \mathcal{F}}$ which evaluate to $h_i(\mathbf{A})$ modulo q for each $i \in [Q]$, and find a short linear combination of $(f(\mathbf{A}))_{f \in \mathcal{F}}$ which evaluates to $g^*(\mathbf{A})$ modulo q .

The strong variant, which is strong in the same sense as in Q -SDH, asks to perform the above task with the flexibility that g^* can be picked after seeing \mathbf{A} .

We suggest general criteria for the hinted vSIS assumptions⁶ to be plausible. Further, under the Evasive SIS assumption envisioned by [Wee22] (but which was not formalised nor used), we show that the (non-strong) hinted vSIS assumption is implied by the (plain) vSIS assumption for certain parameter choices (Thm. 2). Similar to the gaps between strong and non-strong assumptions in the group setting, e.g. (strong) Diffie-Hellman and (strong) RSA, formal reductions from non-strong to strong hinted vSIS are out of reach except in trivial cases.

Lattice-based Adaption of the BB(S) Signatures. We construct a family of lattice-based signatures in Section 4, capturing the lattice-BB signatures sketched above as a special case. In brief, suppose $\mathcal{H} = \{h_{\mu, \chi}\}_{\mu, \chi}$ is a set of rational functions indexed by messages μ and signing randomness χ . For a public key (\mathbf{A}, \mathbf{b}) , a signature is simply a tuple (χ, \mathbf{s}) , where \mathbf{s} is a short vector satisfying $\mathbf{A}\mathbf{s} = h_{\mu, \chi}(\mathbf{b}^\top) \bmod q$. Assuming strong hinted vSIS holds for \mathcal{G} , then the signature scheme has strong selective-query security (Thm. 5). By instantiating \mathcal{G} appropriately, we obtain natural lattice-analogues of the BB and the BBS signature schemes, elaborated in Section 4.2.

Generalised ISIS_f. We generalise the ISIS_f assumption of [BLNS23b] to allow the function f inputting additional randomness, which we call the GenISIS_f assumption, presented in Section 5. Analogous to ISIS_f of [BLNS23b], the GenISIS_f assumption can be generically lifted to an interactive GenISIS_f assumption without additional overhead (Thm. 6). Applying the transformation to our strongly selective-query-secure signature scheme yields a fully strongly secure one. Moreover, we show that the GenISIS_f assumption implies a weakened version of the strong hinted vSIS assumption, where the set of hints \mathcal{Q} is sampled uniformly randomly.

⁶ Apply also to the (plain) vSIS and kRISIS assumptions upon appropriate adaption.




Fig. 1: Overview of results. An arrow from A to B means “Assumption/Security A implies Assumption/Security B ”. Dashed arrows denote trivial reductions.

1.2 Related Work

The BBS signature scheme was implicit in their group signature construction [BBS04] and can be seen as an extension of another signature scheme by Boneh and Boyen (BB) [BB08]. The BBS scheme was later explicitly cast as a standalone signature scheme [CL04]. The BBS+ signature scheme [ASM06] is a slightly modified and provably secure version of BBS, under the q -strong Diffie Hellman (q -SDH) assumption. For almost two decades since its introduction, the BBS+ signature scheme is a de facto standard building block for pairing-based anonymous credentials. Only until recently [TZ23] it is shown that the original BBS signature scheme is also provably secure under the same assumption.

The combination of (lattice-based) proof-friendly signatures with a tailored zero-knowledge proof system is a general template for privacy-preserving authentication primitives (see e.g. [CGT23]). Anonymous credentials tend to be the hardest to construct, as typically both the signature and parts of the signed message should remain hidden when the credential is shown. Thus, efficient proof-friendly signatures [ABB10, MP12, DM14, BLNS23b] along with suitable proof systems [BLS19, YAZ⁺19, LNP22b] had to be devised first, and still much optimisation was (and is) required [LLM⁺16, BLNS23b, JRS23, BBP23].

For group and blind signatures, using a random oracle to hash the message provides some leverage. Indeed, in the lattice setting, we have seen earlier and steady development with group signatures [GKV10, LLS13, LLNW16, BCN18, dPLS18, BDK⁺22, LNPS21]. The situation with blind signatures is less fortunate, where all prior works based on the blind Schnorr-type template [Rüc10, AEB20a, AEB20b, AHJ21] have been found gaps in their proofs [HKLN20], and later broken by the so-called ROS attacks [BLL⁺21, BLL⁺22, KLR24]. Schemes that remain standing [LNP22a, AKSY22, dPK22, BLNS23a] follow Fischlin’s two-move template [Fis06]. From the efficiency perspective, the most competitive blind signatures to-date are based on either the ISIS_f assumption [BLNS23b] which our work extends, the one-more-ISIS assumption [AKSY22], or heuristic assumptions [BLNS23a], namely succinct arguments proving about random oracles, which is concretely expensive but achieves small signature size.

2 Preliminaries

Let $\lambda \in \mathbb{N}$ denote the security parameter. For two (ensembles of) distributions $\mathcal{D}_0, \mathcal{D}_1$, we write $\mathcal{D}_0 \approx_c \mathcal{D}_1$ if they are computationally indistinguishable. We write $\text{poly}(\lambda)$ and $\text{negl}(\lambda)$ for the sets of functions polynomial and negligible in λ , respectively. We write matrices and vectors in bold upper and lower case letters, e.g. \mathbf{A} and \mathbf{x} , respectively. For matrices and vectors of compatible dimensions, we write \odot and \oslash for the Hadamard (i.e. component-wise) product and division, respectively. We write $\mathbf{1}_n$ for the all-1 vector of dimension n over whichever ring within context. For real vectors $\mathbf{x} \in \mathbb{R}^n$, we write $\|\mathbf{x}\| := \|\mathbf{x}\|_2$ for its Euclidean norm. If S is a finite set, we write $\mathcal{U}(S)$ for the uniform distribution over S and $x \leftarrow S$ for the sampling of a uniformly random element x from S .

For a sequence of k formal variables $\tilde{\mathbf{x}}$ and a ring \mathcal{X} , we write $\mathcal{X}[\tilde{\mathbf{x}}^T]$ and $\mathcal{X}(\tilde{\mathbf{x}}^T) = \{f/g : f, g \in \mathcal{X}[\tilde{\mathbf{x}}^T]\}$ for the set of k -variate polynomial and rational functions over \mathcal{X} respectively.⁷ We use $\tilde{\cdot}$ to denote formal variables using the same letter as the intended input. For example, we write $f(\tilde{x})$ for a function f with variable \tilde{x} , which is intended to be evaluated at a point x . We will use the following shorthand for vectors

⁷ The transposes in $\mathcal{X}[\tilde{\mathbf{x}}^T]$ and $\mathcal{X}(\tilde{\mathbf{x}}^T)$ matter due to the notation of evaluating functions at matrices defined below.

consisting of evaluations of one or multiple functions at multiple points. For $f : \mathcal{X}^k \rightarrow \mathcal{X}$ a k -variate function, $\mathcal{F} = (f_j : \mathcal{X}^k \rightarrow \mathcal{X})_{j=1}^m$ a sequence of k -variate functions, and $\mathbf{A} \in \mathcal{X}^{n \times k}$ a \mathcal{X} -matrix with the i -th row given by $\mathbf{a}_i^\top \in \mathcal{X}^k$, we write

$$f(\mathbf{A}) := (f(\mathbf{a}_1^\top) \dots f(\mathbf{a}_n^\top))^\top \in \mathcal{X}^n,$$

$$\mathcal{F}(\mathbf{A}) := (f_1(\mathbf{A}) \dots f_m(\mathbf{A})) = \begin{pmatrix} f_1(\mathbf{a}_1^\top) & \dots & f_m(\mathbf{a}_1^\top) \\ \vdots & \ddots & \vdots \\ f_1(\mathbf{a}_n^\top) & \dots & f_m(\mathbf{a}_n^\top) \end{pmatrix} \in \mathcal{X}^{n \times m}.$$

2.1 Algebraic Number Theory

We state our results over the cyclotomic field $\mathbb{Q}(\zeta)$, where $\zeta = \zeta_{\mathfrak{f}}$, with conductor \mathfrak{f} and degree $\varphi := \varphi(\mathfrak{f})$, and its ring of integers $\mathcal{R} = \mathbb{Z}[\zeta]$. All results can be specialised to the integer setting, i.e. $\mathcal{R} = \mathbb{Z}$. For $q \in \mathbb{N}$, we write $\mathcal{R}_q := \mathcal{R}/q\mathcal{R}$. Let $\sigma = (\sigma_i)_{i \in \mathbb{Z}_{\mathfrak{f}}^\times} : \mathbb{Q}(\zeta) \rightarrow \mathbb{C}^\varphi$ denote the canonical embedding of $\mathbb{Q}(\zeta)$, with its definition naturally extended to $\mathbb{Q}(\zeta)$ -vectors by concatenation. We norm a $\mathbb{Q}(\zeta)$ -vector \mathbf{x} geometrically by the ℓ_p -norm of its canonical embedding, i.e. $\|\mathbf{x}\|_p := \|\sigma(\mathbf{x})\|_p$. For any $a, b \in \mathbb{Q}(\zeta)$, it holds that $\|a \cdot b\|_p \leq \|a\|_p \cdot \|b\|_\infty$. We omit the subscript p when $p = 2$.

Any \mathcal{R} -module $\mathcal{M} \subseteq \mathcal{R}^m$ can be viewed as a lattice via $\sigma(\mathcal{M})$. In particular, for $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ and $\mathbf{v} \in \mathcal{R}_q^n$, we consider the following lattice (cosets):

$$A_q^\perp(\mathbf{A}) := \{\mathbf{x} \in \mathcal{R}^m : \mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}\}, \quad A_q^\mathbf{v}(\mathbf{A}) := \{\mathbf{x} \in \mathcal{R}^m : \mathbf{A}\mathbf{x} = \mathbf{v} \pmod{q}\}.$$

2.2 Discrete Gaussians, Lattice Trapdoors

The Gaussian function with parameter $s > 0$ is $\rho_s(\mathbf{x}) := \exp(-\pi\|\mathbf{x}\|^2/s^2)$ for all $\mathbf{x} \in \mathbb{R}^n$. For a discrete set $A \subseteq \mathbb{R}^n$, the discrete Gaussian distribution with parameter s is $\mathcal{D}_{A,s}(\mathbf{x}) := \rho_s(\mathbf{x})/\rho_s(A)$ for any $\mathbf{x} \in A$, where $\rho_s(A) := \sum_{\mathbf{x} \in A} \rho_s(\mathbf{x})$.

Lemma 1 ([Ban93, Lemma 1.5]). *For any lattice $\Lambda \subseteq \mathbb{R}^n$ and $s > 0$, it holds $\Pr[\|\mathcal{D}_{\Lambda,s}\| > s\sqrt{n}] \leq 2^{-n}$.*

We summarise the properties of lattice trapdoors as a ‘‘lattice trapdoor scheme’’.

Definition 1 (Lattice Trapdoors [GPV08]). *Let \mathcal{R} be parametrised by λ . A lattice trapdoor scheme over \mathcal{R} consists of PPT algorithms (TrapGen, SampPre):*

$(\mathbf{A}, \text{td}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, 1^n, 1^m, q)$: *Sample a matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ together with a trapdoor $\text{td}_{\mathbf{A}}$.*
 $\mathbf{u} \leftarrow \text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{v}, s)$: *Given the trapdoor $\text{td}_{\mathbf{A}}$, a target image vector $\mathbf{v} \in \mathcal{R}_q^n$ and a Gaussian parameter s , sample a preimage vector $\mathbf{u} \in \mathcal{R}^m$.*

A tuple of parameters $\text{params}_{\text{td}} = (\mathcal{R}, n, m, q, s)$ is said to be admissible if they satisfy the following properties:

1. It holds that $\{\mathbf{A} : (\mathbf{A}, \text{td}_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, 1^n, 1^m, q)\} \approx_c \mathcal{U}(\mathcal{R}_q^{n \times m})$.
2. For any $s' \geq s$ and for all but a $\text{negl}(\lambda)$ -fraction of $(\mathbf{A}, \text{td}_{\mathbf{A}})$ in the support of $\text{TrapGen}(1^\lambda, 1^n, 1^m, q)$, the following hold:
 - For any $\mathbf{v} \in \mathcal{R}_q^n$, it holds that $\text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{v}, s') \approx_c \mathcal{D}_{A_q^\mathbf{v}(\mathbf{A}), s'}$.
 - It holds that $\left\{ (\mathbf{u}, \mathbf{v}) \left| \begin{array}{l} \mathbf{v} \leftarrow \mathcal{R}_q^n \\ \mathbf{u} \leftarrow \mathcal{D}_{A_q^\mathbf{v}(\mathbf{A}), s'} \end{array} \right. \right\} \approx_c \left\{ (\mathbf{u}, \mathbf{v}) \left| \begin{array}{l} \mathbf{u} \leftarrow \mathcal{D}_{\mathcal{R}^m, s'} \\ \mathbf{v} := \mathbf{A}\mathbf{u} \pmod{q} \end{array} \right. \right\}$.

We refer, for example, to [GPV08, MP12] for how to instantiate a lattice trapdoor scheme with admissible parameters.

$\text{Pre}_{\mathcal{A}}(1^\lambda)$	$\text{Post}_{\mathcal{B}}(1^\lambda)$
$(\tilde{\mathbf{P}}, \tilde{\mathbf{A}}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$	$(\tilde{\mathbf{P}}, \tilde{\mathbf{A}}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$
$\mathbf{B} \leftarrow_{\mathcal{R}} \mathcal{R}_q^{n \times m}$	$\mathbf{B} \leftarrow_{\mathcal{R}} \mathcal{R}_q^{n \times m}$
	$\mathbf{U} \leftarrow_{\mathcal{R}, s} \mathcal{D}_{\mathcal{R}, s}^{m \times m_P}$ conditioned on $\mathbf{B}\mathbf{U} = \tilde{\mathbf{P}} \bmod q$
$\mathbf{u}^* \leftarrow \mathcal{A}(\mathbf{B}, \tilde{\mathbf{P}}, \tilde{\mathbf{A}}, \text{aux})$	$\mathbf{u}^* \leftarrow \mathcal{B}(\mathbf{B}, \tilde{\mathbf{P}}, \tilde{\mathbf{A}}, \mathbf{U}, \text{aux})$
$b_0 := ((\mathbf{B} \tilde{\mathbf{P}} \tilde{\mathbf{A}})\mathbf{u}^* = \mathbf{0} \bmod q)$	$b_0 := ((\mathbf{B} \tilde{\mathbf{A}})\mathbf{u}^* = \mathbf{0} \bmod q)$
$b_1 := (0 < \ \mathbf{u}^*\ \leq \beta_1)$	$b_1 := (0 < \ \mathbf{u}^*\ \leq \beta_0)$
return $b_0 \wedge b_1$	return $b_0 \wedge b_1$

Fig. 2: Experiments Pre and Post for evasive SIS assumption.

2.3 Lattice Assumptions

The vanishing SIS (vSIS) assumption [CLM23] is parametrised by, among others, a set of rational functions \mathcal{F} . It states that, given a random matrix \mathbf{A} , it is hard to find a short linear combination of $\{f(\mathbf{A})\}_{f \in \mathcal{F}}$ which vanishes modulo q .

Definition 2 (Vanishing-SIS ([CLM23])). *Let $\text{params} = (\mathcal{R}, n, k, q, \beta, \mathcal{F})$ be parametrised by λ , where n, k, q are positive integers, $\beta \in \mathbb{R}^+$ and \mathcal{F} is a set of k -variate functions over \mathcal{R} . The $\text{vSIS}_{\text{params}}$ assumption states that, for any PPT adversary \mathcal{A} , it holds that*

$$\Pr \left[\begin{array}{l} \mathcal{F}(\mathbf{A}) \cdot \mathbf{u}^* = \mathbf{0} \bmod q \\ \wedge 0 < \|\mathbf{u}^*\| \leq \beta \end{array} \middle| \begin{array}{l} \mathbf{A} \leftarrow_{\mathcal{R}} \mathcal{R}_q^{n \times k} \\ \mathbf{u}^* \leftarrow \mathcal{A}(\mathbf{A}) \end{array} \right] \leq \text{negl}(\lambda).$$

In this work, we consider also settings where $|\mathcal{F}|$ could be super-polynomial in λ , as long as \mathcal{F} and the answer \mathbf{u}^* by the adversary admits a succinct description.

The evasive SIS assumption was informally introduced by Wee [Wee22] (in conjunction with the public-coin evasive LWE assumption) and envisioned as a tool for analysing the plausibility of SIS-based hinted lattice assumptions.

Definition 3 (EvasiveSIS). *Let $\text{params} = (\mathcal{R}, q, n, m, m_P, m_A, s, \beta_0, \beta_1)$ be parametrised by λ , where \mathcal{R} is a ring admitting an embedding as a lattice in \mathbb{R}^φ for some $\varphi \in \mathbb{N}$, and $s, \beta_0, \beta_1 > 0$. Let Samp be a PPT algorithm which, on input 1^λ , outputs*

$$(\tilde{\mathbf{P}} \in \mathcal{R}_q^{n \times m_P}, \tilde{\mathbf{A}} \in \mathcal{R}_q^{n \times m_A}, \text{aux} \in \{0, 1\}^*)$$

where aux contains all coin tosses used by Samp . The $\text{EvasiveSIS}_{\text{params}}$ assumption states that, for any PPT Samp and \mathcal{B} there exists a PPT \mathcal{A} such that

$$\Pr[\text{Pre}_{\mathcal{A}}(1^\lambda) = 1] \geq \Pr[\text{Post}_{\mathcal{B}}(1^\lambda) = 1] / \text{poly}(\lambda) - \text{negl}(\lambda),$$

where the experiments Pre and Post are defined in Figure 2.

Analogous to the evasive LWE assumption [Wee22], the evasive SIS assumption says that “if SIS is hard for the matrix $(\mathbf{B}, \tilde{\mathbf{P}}, \tilde{\mathbf{A}})$, then SIS is also hard for $(\mathbf{B}, \tilde{\mathbf{A}})$ even when given short preimages \mathbf{U} of $\tilde{\mathbf{P}}$ w.r.t. \mathbf{B} ”. This stems from the intuition that, there seems no alternative meaningful use of \mathbf{U} , other than multiplying which to \mathbf{B} to obtain $\tilde{\mathbf{P}}$ and solve (the potentially easier) SIS problem for $(\mathbf{B}, \tilde{\mathbf{P}}, \tilde{\mathbf{A}})$ jointly. Following [Wee22], Definition 3 is “public-coin” in that we insist Samp to output all its random coins, which avoids obfuscation-based counterexamples.

We note that Definition 3 is heuristically no stronger than the evasive LWE assumption of [Wee22], in the following sense: Suppose there exists a PPT solver \mathcal{B} for Post in Figure 2, then immediately \mathcal{B} is also a successful distinguisher for the analogous LWE problem – distinguish $\mathbf{s}^\top(\mathbf{B}, \tilde{\mathbf{A}}) + \mathbf{e}^\top \bmod q$ from random given \mathbf{U} . Assuming evasive LWE, there exists a PPT \mathcal{A} distinguishing $\mathbf{s}^\top(\mathbf{B}, \tilde{\mathbf{P}}, \tilde{\mathbf{A}}) + \mathbf{e}^\top \bmod q$ from random. At this point, under the common heuristic that solving decision-LWE is no easier than solving SIS (which is quantumly true at least for uniformly random matrices [SSTX09]), we arrive at a solver for Pre in Figure 2.

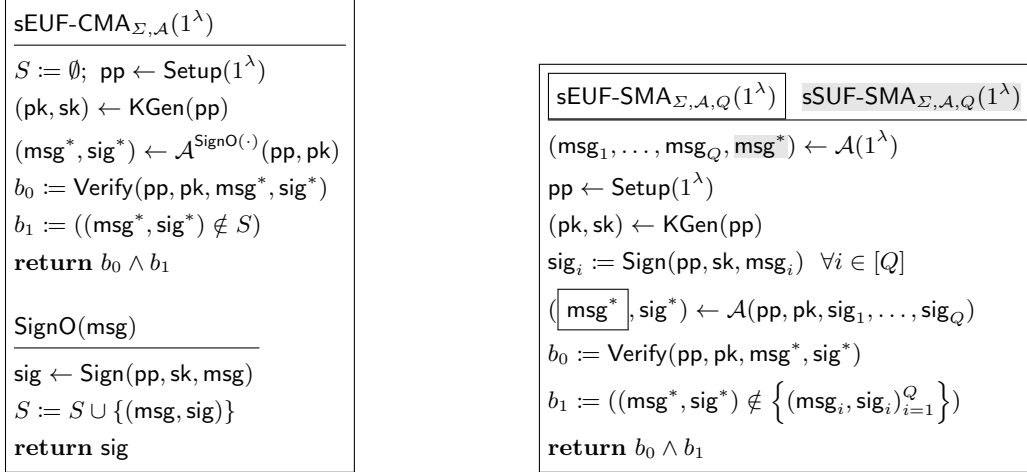


Fig. 3: Security experiments for sEUF-CMA, sEUF-SMA and sSUF-SMA. For sEUF-RMA, modify the sEUF-SMA experiment to have $\text{msg}_1, \dots, \text{msg}_Q \leftarrow \$ \mathcal{M}$.

2.4 Signatures

Definition 4 (Signature Scheme). A signature scheme for a message space \mathcal{M} is a tuple of PPT algorithms $\Sigma = (\text{Setup}, \text{KGen}, \text{Sign}, \text{Verify})$:

- $\text{pp} \leftarrow \text{Setup}(1^\lambda)$: Generate the public parameters pp.
- $(\text{pk}, \text{sk}) \leftarrow \text{KGen}(\text{pp})$: Generate a public key pk and a secret key sk.
- $\text{sig} \leftarrow \text{Sign}(\text{sk}, \text{msg})$: Sign a message $\text{msg} \in \mathcal{M}$ with a signature sig.
- $b \leftarrow \text{Verify}(\text{pk}, \text{msg}, \text{sig})$: Decide if sig is a valid signature of msg under pk.

A signature scheme Σ is said to be correct if, for any message $\text{msg} \in \mathcal{M}$,

$$\Pr \left[\text{Verify}(\text{pp}, \text{pk}, \text{msg}, \text{sig}) = 1 \mid \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda) \\ (\text{pk}, \text{sk}) \leftarrow \text{KGen}(\text{pp}) \\ \text{sig} \leftarrow \text{Sign}(\text{sk}, \text{msg}) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

A signature scheme Σ is said to have strong existential unforgeability under chosen message attack (sEUF-CMA) if, for any PPT \mathcal{A} , it holds

$$\text{Adv}_{\Sigma, \mathcal{A}}^{\text{sEUF-cma}}(\lambda) := \Pr[\text{sEUF-CMA}_{\Sigma, \mathcal{A}}(1^\lambda) = 1] \leq \text{negl}(\lambda),$$

where $\text{sEUF-CMA}_{\Sigma, \mathcal{A}}$ is defined in Figure 3. It is said to have strong existential unforgeability under selective message attack (sEUF-SMA) and strong selective unforgeability under selective message attack (sSUF-SMA) respectively, if for any $Q \leq \text{poly}(\lambda)$, the above holds for $\text{sEUF-SMA}_{\Sigma, \mathcal{A}, Q}$ and $\text{sSUF-SMA}_{\Sigma, \mathcal{A}, Q}$ defined in Figure 3 respectively. Strong existential unforgeability under random message attack (sEUF-RMA) is similarly defined in Figure 3.

Remark 1 (Key-dependent message space). We also consider a relaxed definition of signature schemes where the public parameters specify a subspace of the message space. In this case, we say that the signature scheme is correct if for all but a $\text{negl}(\lambda)$ -fraction of $\text{pp} \in \text{Setup}(1^\lambda)$ and for any message msg in the message subspace defined by pp, it holds that $\text{Verify}(\text{pp}, \text{pk}, \text{msg}, \text{Sign}(\text{pp}, \text{sk}, \text{msg})) = 1$ except with $\text{negl}(\lambda)$ probability. Corresponding, we modify the security experiments so that the Sign algorithm aborts whenever it is called on messages outside the message subspace defined by pp.

3 (Strong) Hinted vSIS Assumptions

Aiming to construct algebraic lattice signatures akin to pairing-based ones based on assumptions such as strong Diffie Hellman (SDH), we introduce hinted variants of the vanishing short integer solution (vSIS) assumption.

We define in Section 3.1 general families of two vSIS assumption variants – hinted, and strong hinted – which extend the existing vSIS assumption. The strong hinted variant is strong in the same sense as in the SDH (and others such as strong RSA) assumption – the adversary is allowed to choose its “target” freely. More discussions on the relations to existing assumptions follow.

Since the hinted vSIS assumptions have numerous parameters, we discuss in Section 3.2 what we believe to be plausible choices of them. To install confidence in the new assumptions, we show in Section 3.3 that some of them admit reductions from the plain vSIS assumption of [CLM23].

3.1 Assumptions Statements

We define extended variants of the vSIS assumption. These variants aim to capture natural lattice-analogues of group-based assumptions such as the strong Diffie-Hellman (SDH) assumption.

Definition 5 ((Strong) Hinted vSIS Assumptions). *Let*

$$\text{params} = (\mathcal{R}, n, k, q, Q, \beta, s, \mathcal{F}, \mathcal{G}, \mathcal{H}, \mathcal{P})$$

be parametrised by λ , where n, k, q are positive integers, Q is a non-negative integer, $\beta, s \in \mathbb{R}^+$, $\mathcal{F}, \mathcal{G}, \mathcal{H}$ are k -variate rational functions over \mathcal{R}_q such that $Q \leq |\mathcal{H}|$, and \mathcal{P} is a predicate over sets of k -variate rational functions. We define the following hinted variants of the vSIS assumption.

Hinted. *The Hint-vSIS_{params} assumption states that, for any PPT stateful adversary \mathcal{A} , it holds that*

$$\Pr \left[\begin{array}{l} \mathcal{F}(\mathbf{A}) \cdot \mathbf{u}^* = g^*(\mathbf{A}) \bmod q \\ \wedge 0 < \|\mathbf{u}^*\| \leq \beta \\ \wedge Q \subseteq_Q \mathcal{H} \\ \wedge g^* \in \mathcal{G} \setminus Q \\ \wedge \mathcal{P}(\mathcal{F} \cup Q \cup (\{g^*\} \setminus \{0\})) = 1 \end{array} \middle| \begin{array}{l} (Q, g^*) \leftarrow \mathcal{A}(1^\lambda) \\ \mathbf{A} \leftarrow \mathcal{R}_q^{n \times k} \\ \mathbf{V} := Q(\mathbf{A}) \bmod q \\ \mathbf{u}_i \leftarrow \mathcal{D}_{\mathcal{A}_q^{v_i}(\mathcal{F}(\mathbf{A})), s} \quad \forall i \in [Q] \\ \mathbf{u}^* \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{u}_1, \dots, \mathbf{u}_Q) \end{array} \right] \leq \text{negl}(\lambda).$$

Strong Hinted. *The s-Hint-vSIS_{params} assumption states that, for any PPT stateful adversary \mathcal{A} ,*

$$\Pr \left[\begin{array}{l} \mathcal{F}(\mathbf{A}) \cdot \mathbf{u}^* = g^*(\mathbf{A}) \bmod q \\ \wedge 0 < \|\mathbf{u}^*\| \leq \beta \\ \wedge Q \subseteq_Q \mathcal{H} \\ \wedge g^* \in \mathcal{G} \setminus Q \\ \wedge \mathcal{P}(\mathcal{F} \cup Q \cup (\{g^*\} \setminus \{0\})) = 1 \end{array} \middle| \begin{array}{l} Q \leftarrow \mathcal{A}(1^\lambda) \\ \mathbf{A} \leftarrow \mathcal{R}_q^{n \times k} \\ \mathbf{V} := Q(\mathbf{A}) \bmod q \\ \mathbf{u}_i \leftarrow \mathcal{D}_{\mathcal{A}_q^{v_i}(\mathcal{F}(\mathbf{A})), s} \quad \forall i \in [Q] \\ (g^*, \mathbf{u}^*) \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{u}_1, \dots, \mathbf{u}_Q) \end{array} \right] \leq \text{negl}(\lambda).$$

Strong Random-Hinted. *The s- \mathcal{H} -Hint-vSIS_{params} assumption is almost identical to the s-Hint-vSIS_{params} assumption, except that Q is sampled as a uniformly random Q -subset of \mathcal{H} , not chosen by \mathcal{A} .*

Compared to the plain vSIS assumption (Definition 5, [CLM23]), the (strong) hinted vSIS assumption is further parametrised by an integer Q , two sets of rational functions \mathcal{G}, \mathcal{H} , and a predicate \mathcal{P} . The sets \mathcal{G} and \mathcal{H} are where the adversary could choose the target g^* and the Q queries Q respectively. The adversary is considered successful if $g^* \notin Q$ i.e. not queried and $\mathcal{F} \cup Q \cup \{g^*\}$ (or $\mathcal{F} \cup Q$ if g^* is the all zero function) satisfies the predicate \mathcal{P} . The non-strong and strong variants differ by whether g^* is chosen before or after seeing \mathbf{A} .

Analogous to the discussion below Definition 2, we allow the sets $\mathcal{F}, \mathcal{G}, \mathcal{H}$ of rational functions to have cardinalities super-polynomial in λ , as long as they and the answer \mathbf{u}^* output by \mathcal{A} admit succinct representations. We note, however, that the number Q of queries Q output by \mathcal{A} must be polynomial in λ , for otherwise \mathcal{A} could not input all preimages $\mathbf{u}_1, \dots, \mathbf{u}_Q$ while still being PPT.

Remark 2 (Hint-vSIS \Rightarrow s-Hint-vSIS). If $|\mathcal{G}| \leq \text{poly}(\lambda)$, then the Hint-vSIS_{params} assumption implies the s-Hint-vSIS_{params} assumption (for the same params), where a trivial reduction simply guesses g^* upfront.

Beyond obvious connections to the vSIS assumption, we discuss further connections of Definition 5 to other existing assumptions.

Relation to SDH. The vSIS assumption variants are defined in a way intended to translate certain group-based assumptions such as SDH. Recall that the $(Q-1)$ -SDH problem asks to find, given $[1], [b], \dots, [b^{Q-1}]$, a tuple $(\mu, [1/(b + \mu)])$. To translate this to a strong hinted vSIS problem, let $\tilde{\mathbf{x}}^T = (\tilde{\mathbf{a}}^T, \tilde{b})$ denote a sequence of formal variables, $\gamma > 0$ and $B_\gamma = \{\mu \in \mathcal{R} : \|\mu\| \leq \gamma\}$. For suitably selected parameters, in particular

$$\mathcal{F}(\tilde{\mathbf{x}}^T) = \tilde{\mathbf{a}}^T, \quad \mathcal{H} = \{1, \tilde{b}, \dots, \tilde{b}^{Q-1}\}, \quad \mathcal{G} = \{1/(\tilde{b} + \mu) : \mu \in B_\gamma\},$$

the s-Hint-vSIS assumption can be seen as a natural lattice-analogue of the $(Q-1)$ -SDH assumption.

Relation to kRISIS. The kRISIS assumption family is introduced in [ACL⁺22], stronger than the (plain) vSIS assumption family, but weaker than the BASIS assumption family introduced by [WW23].⁸ We observe that certain members of the (strong) hinted vSIS family are in between (plain) vSIS (obviously) and kRISIS. More precisely, for suitably selected parameters, in particular $\mathcal{H} \cap \mathcal{G} = \emptyset$, $|\mathcal{H}| = Q$ (so $\mathcal{Q} = \mathcal{H}$) and $|\mathcal{G}| = 1$ (so there is only one choice for g^*), the s-Hint-vSIS assumption is essentially a kRISIS assumption.⁹ Note also that, since $|\mathcal{G}| = 1$, the strong and non-strong variants are equivalent.

3.2 Criteria for Plausible vSIS Assumptions

In Definition 5, the vSIS assumption variants are parametrised by a predicate \mathcal{P} which dictates which combinations of functions are admissible. We discuss criteria for \mathcal{P} for the vSIS assumption variants to plausibly hold.

One uninteresting way to violate a vSIS assumption is to consider ill-formed sets of functions \mathcal{F}, \mathcal{G} . For example, if \mathcal{F} contains the all-zero function, then it is trivial to find short \mathbf{u}^* satisfying $\mathcal{F}(\mathbf{A}) \cdot \mathbf{u}^* = \mathbf{0} \pmod q$. Similarly, if $\mathcal{F} \cap \mathcal{G}$ is not empty and contains some g^* , then it is trivial to find short \mathbf{u}^* satisfying $\mathcal{F}(\mathbf{A}) \cdot \mathbf{u}^* = g^*(\mathbf{A}) \pmod q$. A more sophisticated example is where $\mathcal{F} \supseteq \{1, \tilde{x}^q\}$ and \mathcal{R}_q fully splits into a product of fields. In this case, we have $1 - x^q = 0 \pmod q$ for any $x \in \mathcal{R}_q$. To rule out this type of counterexamples which exploits linear dependency between the chosen functions, we suggest to restrict \mathcal{P} so that it only accepts tuples satisfying a “strong linear independence” property defined below.

Definition 6 (Strong Linear Independence). Let $\mathcal{F} \subseteq \mathcal{R}_q(\tilde{\mathbf{x}}^T)$ be a set of k -variate rational functions where $1 \in \mathcal{R}_q\text{-span}(\mathcal{F})$. For $\epsilon > 0$, we say that \mathcal{F} is ϵ -strong linear independent if, for any not-all-zero coefficients $(c_f)_{f \in \mathcal{F}}$ over \mathcal{R}_q ,

$$\Pr \left[\sum_{f \in \mathcal{F}} c_f \cdot f(\mathbf{a}^T) = 0 \pmod q \mid \mathbf{a} \leftarrow_{\$} \mathcal{R}_q^k \right] \leq \epsilon.$$

If $1 \notin \mathcal{R}_q\text{-span}(\mathcal{F})$, then we say \mathcal{F} is ϵ -strong linear independent if $\mathcal{F} \cup \{1\}$ is.

Remark 3. In case the denominator of f vanishes at \mathbf{a} , we let $f(\mathbf{a}) = \text{undef}$. We take “ $a \times \text{undef} = a + \text{undef} = \text{undef}$ ” for any value a , meaning that if one of the terms is undefined then the whole sum is (which does not equal zero).

⁸ By an assumption family A being stronger than another family B, we mean that for any member in family B there exists a member in family A which implies the former.

⁹ In the kRISIS assumption definition stated in [ACL⁺22], the images for which preimages are given to an adversary take the form $\mathbf{t} \cdot h_i(\mathbf{v}^T)$. We believe this is an oversight, since it would mean that solving kRISIS by solving the vSIS instance $(h_i(\mathbf{v}^T))_{i=1}^Q$ is significantly easier (due to lower lattice dimension) than solving SIS w.r.t. \mathbf{A} .

Remark 4 (Implicit guaranteed min-entropy). In Definition 6, we distinguish between $1 \in \mathcal{R}_q\text{-span}(\mathcal{F})$ or $1 \notin \mathcal{R}_q\text{-span}(\mathcal{F})$ for technical reasons: We actually want that for any fixed $d \in \mathcal{R}_q$ and $(c_f)_{f \in \mathcal{F}}$, we have $\Pr[\sum_{f \in \mathcal{F}} c_f \cdot f(\mathbf{a}^\top) = d \bmod q] \leq \epsilon$. That is, we want that the min-entropy (over \mathbf{a}) is at least $\log_2(\epsilon)$. However, if $1 \in \mathcal{R}_q\text{-span}(\mathcal{F})$, this is trivially false (for $\epsilon < 1$), e.g. if $1 \in \mathcal{F}$ then simply set $c_1 = d$ and all other $c_f = 0$. In fact, if $1 \in \mathcal{R}_q\text{-span}(\mathcal{F})$, then the min-entropy guarantee is already implied (by an analogous reasoning). Hence, if $1 \notin \mathcal{R}_q\text{-span}(\mathcal{F})$, we consider $\mathcal{F} \cup \{1\}$ (instead of introducing d explicitly).

The strong linear independence property could be unwieldy to work with. We show that for certain sets of rational functions, strong linear independence is equivalent to linear independence.

Theorem 1. *Let $q \in \mathbb{N}$ prime and $\mathcal{R}_q \cong \mathbb{F}^e$ split into e fields.¹⁰ Let $\mathcal{F} \subseteq \mathcal{R}_q(\tilde{\mathbf{x}}^\top)$ be a set of k -variate rational functions. Suppose $m \in \mathcal{R}_q[\tilde{\mathbf{x}}]$ is such that, for each rational function in \mathcal{F} represented as $f/g \in \mathcal{F}$ where $f, g \in \mathcal{R}_q[\tilde{\mathbf{x}}]$, it holds that $m \in \langle g \rangle_{\mathcal{R}_q[\tilde{\mathbf{x}}]}$, i.e. m is a common multiple of $(g)_{f/g \in \mathcal{F}}$. Let $d := \deg(m) + \max_{f/g \in \mathcal{F}} \deg(f)$ and $\epsilon := d/|\mathbb{F}|$. If \mathcal{F} is linearly independent, then it is ϵ -strong linearly independent.*

The proof is deferred to Appendix A.1.

Next, we show that the property of being strongly linearly independent is closed under “proper” set union.

Lemma 2. *Suppose $\mathcal{F}, \mathcal{G} \subseteq \mathcal{R}_q(\tilde{\mathbf{x}}^\top)$ are $\epsilon_{\mathcal{F}}$ - and $\epsilon_{\mathcal{G}}$ -strong linearly independent respectively, and $\text{span}_{\mathcal{R}_q}(\mathcal{F}) \cap \text{span}_{\mathcal{R}_q}(\mathcal{G}) = \{0\}$. Then $\mathcal{F} \cup \mathcal{G}$ is $\max(\epsilon_{\mathcal{F}}, \epsilon_{\mathcal{G}})$ -strong linearly independent.*

The proof is deferred to Appendix A.2.

Finally, we highlight a counterexample against the strong-hinted-vSIS assumption, which exploits the denominators of rational functions and the adaptivity to specify g^* after seeing $a \leftarrow \mathcal{R}_q$. Specifically, let

$$1 \in \mathcal{F} \quad \text{and} \quad \mathcal{G} = \mathcal{H} = \{1/(\tilde{a} - b) : b \in \mathcal{R}_q\}.$$

Consider this choice of $\mathcal{F}, \mathcal{G}, \mathcal{H}$, and an adversary \mathcal{A} whose set of queries \mathcal{Q} contains $1/(\tilde{a} - b)$ for some $b \in \mathcal{R}_q$. Upon receiving $a \leftarrow \mathcal{R}_q$, \mathcal{A} specifies $g^*(\tilde{a}) := 1/(\tilde{a} - a + 1)$. Note that $1 - g^*(a) = 0 \bmod q$, where the coefficients 1 and -1 for the functions 1 and g^* respectively are both short. One way to avoid this kind of counterexamples is to let \mathcal{P} accept only rational functions represented as f/g where the denominator g has short coefficients.

To summarise, for a norm bound γ and probability ϵ , we propose the following “natural” predicate $\mathcal{P}_{\gamma, \epsilon}$.

Definition 7 (Natural vSIS Predicates). *For a ring \mathcal{R} , a modulus $q \in \mathbb{N}$, a norm bound $\gamma > 0$, and a probability $\epsilon \in [0, 1]$, the predicate $\mathcal{P}_{\gamma, \epsilon}$ inputs a set \mathcal{F} of rational functions over \mathcal{R}_q and outputs 1 if and only if the following hold:*

- \mathcal{F} is ϵ -strongly linearly independent.
- For any rational function represented as $f/g \in \mathcal{F}$, where f, g are polynomials over \mathcal{R}_q written in expanded form, each coefficient of g is of norm at most γ .

Heuristically, we think of γ to be as small as the norm bound for the plain SIS assumption to hold and $\epsilon \leq \text{negl}(\lambda)$.

We remark that although the plain vSIS assumption (Definition 2, [CLM23]) is not parametrised by a predicate, it is advisable to only rely on a vSIS assumption for \mathcal{F} satisfying $\mathcal{P}_{\gamma, \epsilon}(\mathcal{F}) = 1$. Similar could be recommended for the kRISIS assumption [ACL⁺22].

3.3 Reductions from vSIS

To gain confidence in the hinted vSIS assumptions, we give two hardness reductions. Theorem 2 says that under certain parameters, the (non-strong) hinted vSIS assumption is implied by the vSIS and evasive SIS assumptions together.

¹⁰ This happens when q has multiplicative order φ/e modulo f .

Theorem 2 (EvasiveSIS + vSIS \Rightarrow Hint-vSIS). Let $k_f, k_g \in \mathbb{N}$ and $k = k_f + k_g$. Let $\mathcal{F}, \mathcal{G}, \mathcal{H}$ be sequences of k -variate functions, such that for any $\mathbf{A} = [\mathbf{A}_f, \mathbf{A}_g] \in \mathcal{R}_q^{n \times (k_f + k_g)}$, it holds $\mathcal{F}(\mathbf{A}) = \mathbf{A}_f$, $\mathcal{G}(\mathbf{A}) = \hat{\mathcal{G}}(\mathbf{A}_g)$ and $\mathcal{H}(\mathbf{A}) = \hat{\mathcal{H}}(\mathbf{A}_g)$ for some $\hat{\mathcal{G}}, \hat{\mathcal{H}}$ independent of \mathbf{A}_f .

Let $\beta, \beta_1 > 0$, $\beta_0 = \sqrt{\beta^2 + 1}$. Let $\text{params}_0 = (\mathcal{R}, q, n, k_f, Q, 1, s, \beta_0, \beta_1)$, $\text{params}_1 = (\mathcal{R}, n, k, q, \beta_1, \mathcal{F} \cup \hat{\mathcal{G}} \cup \hat{\mathcal{H}})$, and $\text{params}_2 = (\mathcal{R}, n, k, q, Q, \beta, s, \mathcal{F}, \mathcal{G}, \mathcal{H}, \mathcal{P})$. If the EvasiveSIS $_{\text{params}_0}$ and vSIS $_{\text{params}_1}$ assumptions hold, then the Hint-vSIS $_{\text{params}_2}$ assumption holds.

Proof. Suppose there exists a PPT solver against Hint-vSIS $_{\text{params}_2}$. Below we show that under EvasiveSIS $_{\text{params}_0}$, there exists a PPT solver against vSIS $_{\text{params}_1}$, hence a contradiction and the theorem follows.

To begin, we observe that a successful PPT solver against Hint-vSIS $_{\text{params}_2}$ implies a PPT $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that

$$\Pr \left[\begin{array}{l} (\mathbf{A}_f | \hat{g}^*(\mathbf{A}_g)) \mathbf{u}^* = \mathbf{0} \bmod q \\ \wedge 0 < \|\mathbf{u}^*\| \leq \beta_0 \\ \wedge \mathcal{Q} \cup \{\hat{g}^*\} \subseteq_{Q+1} \hat{\mathcal{G}} \cup \hat{\mathcal{H}} \\ \wedge \mathcal{P}(\mathcal{F} \cup \mathcal{Q} \cup (\{g^*\} \setminus \{0\})) = 1 \end{array} \middle| \begin{array}{l} (\mathcal{Q}, \hat{g}^*, \text{st}) \leftarrow \mathcal{A}_1(1^\lambda) \\ [\mathbf{A}_f, \mathbf{A}_g] \leftarrow \mathcal{R}_q^{n \times k_f} \times \mathcal{R}_q^{n \times k_g} \\ \mathbf{V} := \mathcal{Q}(\mathbf{A}_g) \bmod q \\ \mathbf{u}_i \leftarrow \mathcal{D}_{\Lambda_q^{v_i}(\mathbf{A}_f), s} \quad \forall i \in [Q] \\ \mathbf{u}^* \leftarrow \mathcal{A}_2(\mathbf{A}_f, \mathbf{A}_g, \mathbf{u}_1, \dots, \mathbf{u}_Q, \text{st}) \end{array} \right] > \text{negl}(\lambda),$$

where for all $i \in [Q]$, $\hat{h}_i \in \mathcal{Q} \subset \hat{\mathcal{H}}$ is independent of \mathbf{A}_f , similarly for \hat{g}^* , and st is arbitrary internal state of \mathcal{A} which we assume w.l.o.g. to contain (\mathcal{Q}, \hat{g}^*) . Indeed, given a valid solution \mathbf{u}' to Hint-vSIS $_{\text{params}_2}$, we have

$$\mathcal{F}(\mathbf{A}) \cdot \mathbf{u}' = g^*(\mathbf{A}) \bmod q \iff (\mathbf{A}_f | \hat{g}^*(\mathbf{A}_g)) \cdot \begin{pmatrix} \mathbf{u}' \\ -1 \end{pmatrix} = \mathbf{0} \bmod q,$$

where we make use of that $\mathcal{F}(\mathbf{A}) = \mathbf{A}_f$ and $g^*(\mathbf{A}) = \hat{g}^*(\mathbf{A}_g)$. Similarly, the distribution of each preimage \mathbf{u}_i in the above is identical to that in a Hint-vSIS $_{\text{params}_2}$ instance, as $h_i(\mathbf{A}) = \hat{h}_i(\mathbf{A}_g)$. Therefore \mathcal{A} succeeds by outputting $\mathbf{u}^* = \begin{pmatrix} \mathbf{u}' \\ -1 \end{pmatrix}$, whose norm satisfies $\|\mathbf{u}^*\|^2 \leq \|\mathbf{u}'\|^2 + 1 \leq \beta^2 + 1 = \beta_0$.

Next, consider a PPT Samp which runs $(\mathcal{Q}, \hat{g}^*, \text{st}) \leftarrow \mathcal{A}_1(1^\lambda)$ and outputs

$$\tilde{\mathbf{P}} := \mathcal{Q}(\mathbf{A}_g) \bmod q, \quad \tilde{\mathbf{A}} := \hat{g}^*(\mathbf{A}_g) \bmod q$$

where $\mathbf{A}_g \leftarrow \mathcal{R}_q^{n \times k_g}$ is sampled uniformly randomly, together with aux containing $\mathbf{A}_g, \mathcal{Q}, \hat{g}^*, \text{st}$, and all random coins used. From above, we have that \mathcal{A}_2 is a successful solver for the Post experiment in Figure 2 w.r.t. Samp . Invoking the EvasiveSIS $_{\text{params}_0}$ assumption, there exists a PPT solver \mathcal{B}_2 such that

$$\Pr \left[\begin{array}{l} (\mathbf{A}_f | \mathcal{Q}(\mathbf{A}_g) | \hat{g}^*(\mathbf{A}_g)) \mathbf{u}^* = \mathbf{0} \bmod q \\ \wedge 0 < \|\mathbf{u}^*\| \leq \beta_1 \\ \wedge \mathcal{Q} \cup \{\hat{g}^*\} \subseteq_{Q+1} \hat{\mathcal{G}} \cup \hat{\mathcal{H}} \\ \wedge \mathcal{P}(\mathcal{F} \cup \mathcal{Q} \cup (\{g^*\} \setminus \{0\})) = 1 \end{array} \middle| \begin{array}{l} (\mathcal{Q}(\mathbf{A}_g), \hat{g}^*(\mathbf{A}_g), \text{aux}) \leftarrow \text{Samp}(1^\lambda) \\ \mathbf{A}_f \leftarrow \mathcal{R}_q^{n \times k_f} \\ \mathbf{u}^* \leftarrow \mathcal{B}_2(\mathbf{A}_f, \mathcal{Q}(\mathbf{A}_g), \hat{g}^*(\mathbf{A}_g), \text{aux}) \end{array} \right]$$

is $> \text{negl}(\lambda)$. We note that to invoke EvasiveSIS $_{\text{params}_0}$, we rely on that \mathbf{A}_f is uniformly random over $\mathcal{R}_q^{n \times k_f}$, and that $\tilde{\mathbf{P}}, \tilde{\mathbf{A}}$ in above are independent of \mathbf{A}_f .

Expressing the code of Samp inline and rewriting, the above is equivalent to

$$\Pr \left[\begin{array}{l} (\mathbf{A}_f | \mathcal{Q}'(\mathbf{A}_g)) \mathbf{u}^* = \mathbf{0} \bmod q \\ \wedge 0 < \|\mathbf{u}^*\| \leq \beta_1 \\ \wedge \mathcal{Q}' \subseteq_{Q+1} \hat{\mathcal{G}} \cup \hat{\mathcal{H}} \\ \wedge \mathcal{P}(\mathcal{F} \cup \mathcal{Q} \cup (\{g^*\} \setminus \{0\})) = 1 \end{array} \middle| \begin{array}{l} (\mathcal{Q}' = \mathcal{Q} \cup \{g^*\}, \text{st}) \leftarrow \mathcal{A}_1(1^\lambda) \\ [\mathbf{A}_f, \mathbf{A}_g] \leftarrow \mathcal{R}_q^{n \times k_f} \times \mathcal{R}_q^{n \times k_g} \\ \mathbf{u}^* \leftarrow \mathcal{B}_2(\mathbf{A}_f, \mathbf{A}_g, \text{st}) \end{array} \right] \geq \text{negl}(\lambda).$$

Finally, given the above $(\mathcal{A}_1, \mathcal{B}_2)$, we construct a PPT solver \mathcal{B}^* against vSIS $_{\text{params}_1}$:

- Obtain $\mathbf{A} \in \mathcal{R}_q^{n \times k}$ from the vSIS $_{\text{params}_1}$ challenger.

- Run \mathcal{A}_1 to obtain $(\mathcal{Q}', \text{st})$, pass $[\mathbf{A}_f, \mathbf{A}_g] = \mathbf{A}$ and st to \mathcal{B}_2 to obtain \mathbf{u}^* . Write $\mathbf{u}^* = \begin{pmatrix} \mathbf{u}_{\mathcal{F}} \\ \mathbf{u}_{\mathcal{Q}'} \end{pmatrix}$, where $\mathbf{u}_{\mathcal{F}} \in \mathcal{R}^{|\mathcal{F}|}$ and $\mathbf{u}_{\mathcal{Q}'} \in \mathcal{R}^{Q+1}$.
- Let π be the permutation mapping $(\mathcal{Q}', (\hat{\mathcal{G}} \cup \hat{\mathcal{H}}) \setminus \mathcal{Q}') \mapsto \hat{\mathcal{G}} \cup \hat{\mathcal{H}}$, where the order of $(\hat{\mathcal{G}} \cup \hat{\mathcal{H}}) \setminus \mathcal{Q}'$ is arbitrary. Return $\tilde{\mathbf{u}}^* = (\mathbf{u}_{\mathcal{F}}^T, \pi(\mathbf{u}_{\mathcal{Q}'}^T, \mathbf{0}^T))^T$.¹¹

Whenever $(\mathcal{A}_1, \mathcal{B}_2)$ succeeds, we have $0 < \|\mathbf{u}^*\| = \|\tilde{\mathbf{u}}^*\| \leq \beta_1$ and

$$(\mathcal{F} \cup \hat{\mathcal{G}} \cup \hat{\mathcal{H}})(\mathbf{A}) \cdot \tilde{\mathbf{u}}^* = (\mathcal{F} \cup \mathcal{Q}')(\mathbf{A}) \cdot \begin{pmatrix} \mathbf{u}_{\mathcal{F}} \\ \mathbf{u}_{\mathcal{Q}'} \end{pmatrix} + ((\hat{\mathcal{G}} \cup \hat{\mathcal{H}}) \setminus \mathcal{Q}')(\mathbf{A}) \cdot \mathbf{0} = \mathbf{0} \pmod{q}$$

so that \mathcal{B}^* has the same advantage against $\text{vSIS}_{\text{params}_1}$, non-negligible. \square

The second reduction, summarised by Theorem 3 below, says that the (plain) vSIS assumption family over sets of rational functions is implied by its much smaller subclass which restricts to monomials, up to an exponential blow up in norm bound in the worst case.

Theorem 3 (Monomial-vSIS \Rightarrow vSIS). *Let $d, \beta, \beta_f > 0$. Let $\mathcal{F} \subseteq \mathcal{R}_q(\tilde{\mathbf{x}}^T)$ be a set of k -variate rational functions, such that for each rational function in \mathcal{F} represented as $f/g \in \mathcal{F}$ where $f, g \in \mathcal{R}_q[\tilde{\mathbf{x}}]$, it holds $\deg(f), \deg(g) \leq d$ and $\|f\| \leq \beta_f$. Define the following:*

- $\mathcal{F}_{\text{monomial}} := \left\{ \tilde{\mathbf{b}}^{\mathbf{i}} : \mathbf{i} \in \mathbb{N}_0^k, 0 \leq \|\mathbf{i}\|_1 \leq 2d \right\}$, the set of all k -variate monomials of degree at most $2d$ (independent of \mathcal{F}),
- h an arbitrary common multiple of the denominators $\{g : f/g \in \mathcal{F}\}$ of \mathcal{F} ,
- $\beta' = \sqrt{n} \cdot \|h\| \cdot \beta_f \cdot \beta \cdot \min(|\mathcal{F}|, \text{poly}(\lambda))$.

Let $\text{params}_0 = (\mathcal{R}, n, k, q, \beta', \mathcal{F}_{\text{monomial}})$ and $\text{params}_1 = (\mathcal{R}, n, k, q, \beta, \mathcal{F})$. If the $\text{vSIS}_{\text{params}_0}$ assumption holds, then the $\text{vSIS}_{\text{params}_1}$ assumption holds.

Proof. The idea is to clear the denominators of \mathcal{F} by multiplying with their common multiple. Concretely, assume that \mathcal{A} is a PPT solver against $\text{vSIS}_{\text{params}_1}$, we construct a PPT solver against $\text{vSIS}_{\text{params}_0}$ as follows:

- Receive \mathbf{A} from the $\text{vSIS}_{\text{params}_1}$ challenger, pass which to \mathcal{A} , and receive $\mathbf{u}^* = (u_1^*, \dots, u_{|\mathcal{F}|}^*)^T$.
- Let $I \subseteq \mathcal{F}$ be the index set of the non-zero entries of \mathbf{u}^* , i.e. $u_i^* \neq 0$ if and only if $i \in I$. Note that $|I| = \min(|\mathcal{F}|, \text{poly}(\lambda))$ since \mathcal{A} is PPT.
- For each $f \in I$, write $h \cdot f$ as $\mathcal{F}_{\text{monomial}} \cdot \mathbf{c}_f$ for some $\mathbf{c}_f \in \mathcal{R}^{|\mathcal{F}_{\text{monomial}}|}$, i.e. a linear combination of monomials in $\mathcal{F}_{\text{monomial}}$ with coefficient vector \mathbf{c}_f .
- For each $f \in \mathcal{F} \setminus I$, let $\mathbf{c}_f = \mathbf{0}$ be all-zero. Let $\mathbf{C} := (\mathbf{c}_f)_{f \in \mathcal{F}} \in \mathcal{R}^{|\mathcal{F}_{\text{monomial}}| \times |\mathcal{F}|}$ (whose description size $\leq \text{poly}(\lambda)$). Return $(\mathbf{1}_n \otimes \mathbf{C}) \cdot \mathbf{u}^*$.

First we note that Step 3 in the above is possible, since for any $f \in I \subseteq \mathcal{F}$, it holds $h \cdot f$ is a k -variate polynomial of degree at most $2d$, so that $h \cdot f$ can always be written as a linear combination of the monomials in $\mathcal{F}_{\text{monomial}}$. Further, $\|\mathbf{c}_f\| = \|h \cdot f\| \leq \|h\| \cdot \beta_f$. Suppose \mathcal{A} is successful, then

$$\begin{aligned} \mathcal{F}(\mathbf{A}) \cdot \mathbf{u}^* &= \mathbf{0} \pmod{q}, \\ (h \cdot \mathcal{F})(\mathbf{A}) \cdot \mathbf{u}^* &= (\mathcal{F}_{\text{monomial}})(\mathbf{A}) \cdot (\mathbf{1}_n \otimes \mathbf{C}) \cdot \mathbf{u}^* = \mathbf{0} \pmod{q}, \end{aligned}$$

where $\mathbf{1}_n$ is the all-one vector. Also, we have $\|(\mathbf{1}_n \otimes \mathbf{C})\mathbf{u}^*\| \leq \sqrt{n}\|\mathbf{C}\|\|\mathbf{u}^*\| \leq \sqrt{n} \cdot \|h\| \cdot \beta_f \cdot \beta \cdot \min(|\mathcal{F}|, \text{poly}(\lambda)) = \beta'$. \square

Remark 5 (Norm of common multiple.). For Theorem 3, suppose for any $f/g \in \mathcal{F}$ it holds $\|g\| \leq \beta_g$ and there are at most $\ell \leq |\mathcal{F}|$ distinct denominators, then $\|h\| \leq \beta_g^\ell$. For specially chosen \mathcal{F} a tighter bound is possible.

¹¹ $\tilde{\mathbf{u}}^*$ has $\leq \min(|\mathcal{F}|, \text{poly}(\lambda)) + Q + 1$ non-zero entries, although of dimension $|\mathcal{F} \cup \hat{\mathcal{G}} \cup \hat{\mathcal{H}}|$. Both π and $\tilde{\mathbf{u}}^*$ admit representation of size $\leq \text{poly}(\lambda)$ since $\mathcal{F}, \hat{\mathcal{G}}, \hat{\mathcal{H}}$ do.

Setup (1^λ)	KGen (pp)
$\mathbf{B} \leftarrow \mathcal{R}_q^{n \times \ell}$	$(\mathbf{A}, \text{td}_{\mathbf{A}}) \leftarrow \text{TrapGen}(\mathcal{R}, 1^n, 1^m, q)$
pp := \mathbf{B}	return (pk, sk) := $(\mathbf{A}, \text{td}_{\mathbf{A}})$
Sign (pp, sk, msg)	Verify (pp, pk, msg, sig)
$\chi \leftarrow \mathcal{X}$	$\mathbf{t} := h_{\mu, \chi}(\mathbf{B}) \bmod q$
$\mathbf{t} := h_{\mu, \chi}(\mathbf{B}) \bmod q$	$b_0 := (\mu \in \mathcal{M}_{\mathcal{X}, \mathcal{G}, \mathbf{B}}) \wedge (\chi \in \mathcal{X})$
$\mathbf{s} \leftarrow \text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{t}, s)$	$b_1 := (\mathbf{A}\mathbf{s} = \mathbf{t} \bmod q) \wedge (0 < \ \mathbf{s}\ \leq \beta)$
return sig := (χ, \mathbf{s})	return $b_0 \wedge b_1$

Fig. 4: vSIS-based Signatures Σ_{vSIS} .

4 Proof-Friendly Signatures from Strong Hinted vSIS

We present a general family of algebraic lattice-based signature schemes inspired by the ISIS_f framework and pairing-based signatures such as that of Boneh and Boyen (BB) [BB08] and of Boneh, Boyen and Shacham (BBS) [BBS04, ASM06, TZ23]. The sEUF-SMA security of the construction is tightly connected to the s-Hint-vSIS assumption, which we defined and analysed in Section 3. We then suggest concrete instantiations of the general construction, which can be seen as translations of BB and BBS to the lattice setting.

4.1 General Construction

Our general construction is parametrised by a ring \mathcal{R} , dimensions $n, m, \ell \in \mathbb{N}$, a modulus $q \in \mathbb{N}$, a Gaussian parameter $s > 0$, a norm bound $\beta > 0$, a failure probability $\delta \geq 0$, a message space \mathcal{M} , a randomness space¹² \mathcal{X} , and a set $\mathcal{H} = \{h_{\mu, \chi} : \mu \in \mathcal{M}, \chi \in \mathcal{X}\} \subseteq \mathcal{R}(\tilde{\mathbf{b}}^T)$ of ℓ -variate rational functions over \mathcal{R} indexed by the set $\mathcal{M} \times \mathcal{X}$. The public parameter space is $\mathcal{R}_q^{n \times \ell}$. For any public parameter $\mathbf{B} \in \mathcal{R}_q^{n \times \ell}$, define the message subspace

$$\mathcal{M}_{\mathcal{X}, \mathcal{H}, \mathbf{B}} := \{\mu \in \mathcal{M} : \Pr[h_{\mu, \chi}(\mathbf{B}) = \perp \mid \chi \leftarrow \mathcal{X}] \leq \delta\}.$$

That is, a valid message $\mu \in \mathcal{M}$ w.r.t. the sets \mathcal{X} and \mathcal{H} and public parameter \mathbf{B} is such that, over the randomness of $\chi \leftarrow \mathcal{X}$, the probability of $h_{\mu, \chi}(\mathbf{B})$ being undefined is at most δ . The full construction Σ_{vSIS} is presented in Figure 4.

To explain, the public parameters consists of a random matrix $\mathbf{B} \leftarrow \mathcal{R}_q^{n \times \ell}$. A public key is a trapdoored matrix $\mathbf{A} \in \mathcal{R}_q^{n \times m}$ and the corresponding secret key is the trapdoor $\text{td}_{\mathbf{A}}$. To sign a message μ , sample randomness χ , and evaluate the function $h_{\mu, \chi}$ at \mathbf{B} to obtain an image \mathbf{t} . The signature of μ then consists of the randomness χ and a short preimage \mathbf{s} of \mathbf{t} with respect to \mathbf{A} , sampled using $\text{td}_{\mathbf{A}}$. To verify a signature, check that μ belongs to the message subspace $\mathcal{M}_{\mathcal{X}, \mathcal{H}, \mathbf{B}}$ and χ belongs to the randomness space \mathcal{X} . Also check that \mathbf{s} is of norm bounded by β and satisfies $\mathbf{A}\mathbf{s} = h_{\mu, \chi}(\mathbf{B}) \bmod q$.

Theorem 4 (Correctness). *If $(\mathcal{R}, n, m, q, s)$ are admissible parameters for $(\text{TrapGen}, \text{SampPre})$, $\beta \geq s\sqrt{\varphi m}$ and $\delta \leq \text{negl}(\lambda)$, then the signature scheme in Figure 4 is correct in the sense of Remark 1.*

Proof. By Definition 1, for all but a negligible fraction of $(\mathbf{A}, \text{td}_{\mathbf{A}})$, it holds that $\text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{t}, s) \approx_c \mathcal{D}_{\Lambda_q^+(\mathbf{A}), s}$. Combining with Lemma 1, for any $\mathbf{t} \in \mathcal{R}_q^n$, $(\mathbf{A}, \text{td}_{\mathbf{A}}) \leftarrow \text{TrapGen}(\mathcal{R}, 1^n, 1^m, q)$ and $\mathbf{s} \leftarrow \text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{t}, s)$, it holds that $\mathbf{A}\mathbf{s} = \mathbf{t} \bmod q$ and $\|\text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{t}, s)\| \leq s\sqrt{\varphi m} \leq \beta$ with overwhelming probability in λ . \square

Theorem 5 (Security). *Let $\text{params} = (\mathcal{R}, n, k, q, \beta, \mathcal{F}, \mathcal{G}, \mathcal{H}, \mathcal{P}, Q)$, $m, \ell, s, \delta, \delta', \mathcal{M}, \mathcal{X}$ be parametrised by λ and satisfy the following constraints:*

¹² We assume for simplicity that the randomness space is a finite set and randomness are drawn from the uniform distribution over this set. In general, \mathcal{X} could be a distribution over a possibly infinite set.

- $(\mathcal{R}, n, m, q, s)$ are admissible parameters for $(\text{TrapGen}, \text{SampPre})$, $\beta > 0$, $0 \leq \delta, \delta' \leq \text{negl}(\lambda)$, and $k = m + \ell$.
- For formal variables $(\tilde{\mathbf{a}}^\top, \tilde{\mathbf{b}}^\top)$ of dimension $m + \ell$, $\mathcal{F}(\tilde{\mathbf{a}}^\top, \tilde{\mathbf{b}}^\top) = \tilde{\mathbf{a}}^\top$.
- $\mathcal{H} = \{h_{\mu, \chi} : \mu \in \mathcal{M}, \chi \in \mathcal{X}\} \subseteq \mathcal{R}(\tilde{\mathbf{b}}^\top)$ is a set of ℓ -variate rational functions over \mathcal{R} , as defined above.
- $\mathcal{G} = \mathcal{H} \cup \{0\}$ extends the set \mathcal{H} with the all zero function.
- For random $\chi_1, \dots, \chi_Q \leftarrow_{\$} \mathcal{X}$ and arbitrarily chosen $\mu_1, \dots, \mu_Q, \mu' \in \mathcal{M}$ where $\mu' \notin \{\mu_1, \dots, \mu_Q\}$, and $\chi' \in \mathcal{X}$ possibly dependent on (χ_1, \dots, χ_Q) , it holds except with probability δ' that

$$\mathcal{P}(\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q \cup \{h_{\mu', \chi'}\}) = 1.$$

1. If the \mathbf{s} -Hint-vSIS_{params} (resp. \mathbf{s} - \mathcal{H} Hint-vSIS_{params}) assumption holds for every polynomial $Q(\lambda)$, then Σ_{vSIS} is sEUF-SMA (resp. sEUF-RMA) secure.
2. If \mathcal{X} is a singleton set (so that an image \mathbf{t} is deterministic in μ), and the Hint-vSIS_{params} assumption holds for every polynomial $Q(\lambda)$, then Σ_{vSIS} satisfies a relaxed sSUF-SMA security with the restriction that all signing queries are distinct.¹³

Proof. We prove the implication from \mathbf{s} -Hint-vSIS_{params} to sEUF-SMA security, and then highlight differences of the proofs for the other implications. Define the following hybrid security experiments:

Hyb₀: Identical to $\text{sEUF-SMA}_{\Sigma_{\text{vSIS}}, \mathcal{A}, Q}$, the sEUF-SMA security experiment.

Hyb₁: Recall that, in **Hyb₀**, the step $\text{Sign}(\text{sk}, \text{msg}_i)$ for $\text{msg}_i = \mu_i$ is computed as follows: Sample $\chi_i \leftarrow_{\$} \mathcal{X}$, compute $\mathbf{t}_i := h_{\mu_i, \chi_i}(\mathbf{B})$, compute $\mathbf{s}_i \leftarrow \text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{t}_i, s)$ and return (χ_i, \mathbf{s}_i) . In **Hyb₁**, we change to sample $\mathbf{s}_i \leftarrow_{\$} \mathcal{D}_{\Lambda_q^{\mathbf{t}_i}(\mathbf{A}), s}$ directly from the Gaussian distribution (inefficiently).

Hyb₂: We change how the public key \mathbf{A} is sampled. In **Hyb₂**, we sample $\mathbf{A} \leftarrow_{\$} \mathcal{R}_q^{n \times m}$ uniformly at random.

By the properties of lattice trapdoors (Definition 1), the above hybrids are clearly computationally indistinguishable.

Next, we show that if \mathcal{A} is such that **Hyb₂** outputs 1 with a non-negligible probability $\epsilon > \text{negl}(\lambda)$, then there exists a PPT algorithm \mathcal{B} for \mathbf{s} -Hint-vSIS_{params}. Our reduction \mathcal{B} simulates **Hyb₂** for \mathcal{A} as follows:

- Run \mathcal{A} to obtain Q messages μ_1, \dots, μ_Q .
- Sample $\chi_i \leftarrow_{\$} \mathcal{X}$ for $i \in [Q]$.
- Return the subset $\mathcal{Q} := \{h_{\mu_i, \chi_i} : i \in [Q]\} \subseteq \mathcal{G}$. Except with probability at most $Q\delta$, for a random $\mathbf{B} \leftarrow_{\$} \mathcal{R}_q^{n \times \ell}$, it holds that $h_{\mu_i, \chi_i}(\mathbf{B}) \neq \perp$ for all $i \in [Q]$. In the following, we assume that this is the case.
- Receive in return $([\mathbf{A}, \mathbf{B}], \mathbf{s}_1, \dots, \mathbf{s}_Q)$ where $[\mathbf{A}, \mathbf{B}] \leftarrow_{\$} \mathcal{R}_q^{n \times m} \times \mathcal{R}_q^{n \times \ell}$, $\mathbf{s}_i \leftarrow_{\$} \mathcal{D}_{\Lambda_q^{\mathbf{t}_i}(\mathbf{A}), s}$ and $\mathbf{t}_i = h_{\mu_i, \chi_i}(\mathbf{B})$ for all $i \in [Q]$.
- Set $\text{pp} = \mathbf{B}$, $\text{pk} = \mathbf{A}$ and $\text{sig}_i = (\chi_i, \mathbf{s}_i)$ for all $i \in [Q]$, and run $(\mu^*, \text{sig}^*) \leftarrow \mathcal{A}(\text{pp}, \text{pk}, \text{sig}_1, \dots, \text{sig}_Q)$ where $\text{sig}^* = (\chi^*, \mathbf{s}^*)$. Assuming that \mathcal{A} is successful, we have $\mu^* \in \mathcal{M}_{\mathcal{X}, \mathcal{G}, \mathbf{B}}$, $\chi^* \in \mathcal{X}$, $(\mu^*, \chi^*, \mathbf{s}^*) \notin \{(\mu_1, \chi_1, \mathbf{s}_1), \dots, (\mu_Q, \chi_Q, \mathbf{s}_Q)\}$, $\mathbf{A}\mathbf{s}^* = h_{\mu^*, \chi^*}(\mathbf{B}) \bmod q$, and $0 < \|\mathbf{s}^*\| \leq \beta$.
- If $(\mu^*, \chi^*) \notin \{(\mu_1, \chi_1), \dots, (\mu_Q, \chi_Q)\}$, set $g^* := h_{\mu^*, \chi^*}$ and return (g^*, \mathbf{s}^*) .
- Otherwise, let i^* be such that $(\mu^*, \chi^*) = (\mu_{i^*}, \chi_{i^*})$. Set g^* to be the zero function and return $(g^*, \mathbf{s}^* - \mathbf{s}_{i^*})$.

Clearly, the reduction \mathcal{B} runs in PPT and the failure event of $h_{\mu_i, \chi_i}(\mathbf{B}) = \perp$ for some $i \in [Q]$ happens with probability at most $Q\delta$. Moreover, by the constraint on \mathcal{P} , we have $\mathcal{P}(\mathcal{F} \cup \mathcal{Q} \cup (\{g^*\} \setminus \{0\})) = 1$ except with probability δ' . If $(\mu^*, \chi^*) \notin \{(\mu_1, \chi_1), \dots, (\mu_Q, \chi_Q)\}$, then clearly $g^* \in \mathcal{G} \setminus \mathcal{Q}$. Otherwise, we have $(\mu^*, \chi^*) = (\mu_{i^*}, \chi_{i^*})$ and thus

$$\begin{aligned} \mathbf{A} \cdot \mathbf{s} &= h_{\mu^*, \chi^*}(\mathbf{B}) = h_{\mu_{i^*}, \chi_{i^*}}(\mathbf{B}) = \mathbf{A} \cdot \mathbf{s}_{i^*} \bmod q, \\ \mathbf{A} \cdot (\mathbf{s} - \mathbf{s}_{i^*}) &= \mathbf{0} = g^*(\mathbf{B}) \bmod q, \end{aligned}$$

where $\mathbf{s}^* - \mathbf{s}_{i^*} \neq \mathbf{0}$ since \mathbf{s}^* is a valid forgery. In either case, \mathcal{B} solves the \mathbf{s} -Hint-vSIS_{params} instance with probability at least $\epsilon - Q\delta - \delta' > \text{negl}(\lambda)$.

¹³ The restriction can be lifted by derandomising the signing algorithm with a pseudorandom function.

	Message μ	Randomness χ	Function $h_{\mu,\chi}(\tilde{\mathbf{b}}^\top)$
BB-lite	u	-	$1/(\tilde{b} - u)$
BB-full	\mathbf{u}	\mathbf{x}	$1/((1, \mathbf{u}^\top, \mathbf{x}^\top) \cdot \tilde{\mathbf{b}})$
BB-tran	\mathbf{u}	(\mathbf{x}_0, x_1)	$(\mathbf{u}^\top, \mathbf{x}_0^\top) \cdot \tilde{\mathbf{b}}_0 + 1/(\tilde{b}_1 - x_1)$
BBS	\mathbf{u}	(\mathbf{x}_0, x_1)	$((1, \mathbf{u}^\top, \mathbf{x}_0^\top) \cdot \tilde{\mathbf{b}}_0)/(\tilde{b}_1 - x_1)$

Table 1: Instantiations to obtain lattice-analogues of BB and BBS signatures.

For the implication from $\mathfrak{s}\text{-Hint-vSIS}_{\text{params}}$ to sSUF-RMA security, the argument is identical except that \mathcal{B} samples \mathcal{Q} as a uniformly random Q -subset of \mathcal{H} . For the implication from $\text{Hint-vSIS}_{\text{params}}$ to sSUF-SMA security, we make the following modifications. First, we define Hyb_0 to be identical to $\text{sSUF-SMA}_{\Sigma, \text{vSIS}, \mathcal{A}, Q}$, and propagate this change to Hyb_1 and Hyb_2 . Then, the reduction \mathcal{B} changes as follows. At the beginning, \mathcal{B} receives from \mathcal{A} in addition to μ_1, \dots, μ_Q the target message μ^* . Since $\mathcal{X} = \emptyset$, \mathcal{B} no longer needs to sample χ_1, \dots, χ_Q . If $\mu^* \notin \{\mu_1, \dots, \mu_Q\}$, it sets $g^* := h_{\mu^*}$. Otherwise, say $\mu^* = \mu_{i^*}$, it sets g^* to be the all zero function. It outputs the set $\mathcal{Q} = \{h_{\mu_i} : i \in [Q]\}$ and g^* . Towards the end, upon receiving \mathbf{s}^* from \mathcal{A} , \mathcal{B} simply returns \mathbf{s}^* if $\mu^* \notin \{\mu_1, \dots, \mu_Q\}$, and $\mathbf{s} - \mathbf{s}_{i^*}$ if $\mu^* = \mu_{i^*}$. Clearly, the reduction \mathcal{B} runs in PPT and solves the $\text{Hint-vSIS}_{\text{params}}$ instance with probability at least $\epsilon - Q\delta - \delta' > \text{negl}(\lambda)$. \square

Remark 6 (Constraint on \mathcal{P}). In Theorem 5, we require that $\mathcal{P}(\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q \cup \{h_{\mu', \chi'}\}) = 1$ with high probability. We remark that this indeed captures both cases – whether the adversary chooses to forge a signature on μ^* belonging to $\{\mu_1, \dots, \mu_Q\}$ or not. If $\mu^* \notin \{\mu_1, \dots, \mu_Q\}$, by setting $(\mu', \chi') = (\mu^*, \chi^*)$ we directly recover the constraint $\mathcal{P}(\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q \cup (\{g^*\} \setminus \{0\})) = 1$ where $g^* = h_{\mu^*, \chi^*}$. If $\mu^* \in \{\mu_1, \dots, \mu_Q\}$, then for any $\mu' \notin \{\mu_1, \dots, \mu_Q\}$, the condition $\mathcal{P}(\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q \cup \{h_{\mu', \chi'}\}) = 1$ implies that $\mathcal{P}(\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q) = 1$, equivalent to $\mathcal{P}(\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q \cup (\{g^*\} \setminus \{0\})) = 1$ with $g^* = 0$ the all-zero function.

4.2 Candidate Instantiations

In Table 1, we suggest a few natural candidate instantiations of the parameters, as inspired by the BB [BB08] and BBS [BBS04, ASM06, TZ23] signatures. The first two rows are attempts to translate the selective-query secure (BB-lite) and fully secure versions of the BB signatures. The third row is obtained by interpreting the BB-lite instantiation as a (generalised) ISIS_f instance, and then applying the transformation from ISIS_f to interactive ISIS_f presented in [BLNS23b] and generalised in Section 5.2. The last row is an attempt to translate the BBS signatures which, after adapting notation, take the following form:

$$[(1, \mathbf{u}^\top) \cdot \mathbf{b}_0]/(b_1 - x)$$

where $([1], [\mathbf{b}_0], [b_1])$ is the public key, \mathbf{u} a message, and x the signing randomness.

We remark that, for all instantiated signature schemes suggested in Table 1, we are not aware of any efficient attacks against the sEUF-CMA (i.e. adaptive query) security of the schemes, although Theorem 5 only guarantees their sEUF-SMA (i.e. selective-query) security. Furthermore, in Section 5.2, we show that the BB-tran scheme indeed provably achieves sEUF-CMA security under the same strong hinted vSIS assumption along with other mild parameter constraints.

Satisfiability of Natural vSIS Predicates. Recall that, for Theorem 5 to apply, we require the following constraint:

For random $\chi_1, \dots, \chi_Q \leftarrow \mathcal{X}$ and arbitrarily chosen $\mu_1, \dots, \mu_Q, \mu' \in \mathcal{M}$, with $\mu' \notin \{\mu_1, \dots, \mu_Q\}$, and $\chi' \in \mathcal{X}$ possibly dependent on (χ_1, \dots, χ_Q) , it holds except with negligible probability that

$$\mathcal{P}(\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q \cup \{h_{\mu', \chi'}\}) = 1.$$

As discussed in Remark 6, such a constraint on \mathcal{P} captures both cases – where the adversary chooses to forge a signature of a previously queried message or a new message. Let $\mathcal{P} = \mathcal{P}_{\gamma, \epsilon}$ for some $0 < \gamma \ll q$ and $\epsilon \leq \text{negl}(\lambda)$ as defined in Definition 7, and let \mathcal{X} contain vectors of norm at most γ . We sketch below an argument for why the predicate would be satisfied for all candidates listed in Table 1. We observe the following:

- The denominators of all choices of $h_{\mu, \chi}(\tilde{\mathbf{b}}^T)$ mentioned in Table 1 clearly satisfy the norm bound constraint of $\mathcal{P}_{\gamma, \epsilon}$.
- For χ with sufficient entropy and for $\mathcal{R}_q \cong \mathbb{F}^e$ splitting into large enough copies of \mathbb{F} , it is clear that $\{h_{\mu_i, \chi_i}\}_{i=1}^Q$ for all choices in Table 1 is linearly independent with overwhelming probability, and thus ϵ -strongly linearly independent by Theorem 1.
- The intersection between the span of \mathcal{F} and that of $\{h_{\mu_i, \chi_i}\}_{i=1}^Q$ (or $\{h_{\mu', \chi'}\}$) is clearly zero, since \mathcal{F} depends only on $\tilde{\mathbf{a}}$ but not on $\tilde{\mathbf{b}}$.

To show that $\mathcal{F} \cup \{h_{\mu_i, \chi_i}\}_{i=1}^Q \cup \{h_{\mu', \chi'}\}$ is ϵ -strongly linearly independent with overwhelming probability, by Lemma 2, it suffices to show that the intersection of the spans of $\{h_{\mu_i, \chi_i}\}_{i=1}^Q$ and $h_{\mu', \chi'}$ is trivial with overwhelming probability.

By assumption, we have that $\mu' \notin \{\mu_1, \dots, \mu_Q\}$, but (μ', χ') may arbitrarily depend on $(\mu_i, \chi_i)_{i=1}^Q$. From the constraint $\mu' \notin \{\mu_1, \dots, \mu_Q\}$, the (strong) linear independence between $\{h_{\mu_i, \chi_i}\}_{i=1}^Q$ and $h_{\mu', \chi'}$ is clear for the BB-lite and BB-full cases, since they consist of distinct denominator polynomials. For the BB-tran and BBS cases, write $\chi = (\chi_0, \chi_1) = (\mathbf{x}_0, x_1)$. If $x'_1 \notin \{x_{1,i}\}_{i=1}^Q$, then again the denominator polynomials are distinct and therefore the claim holds. Now, suppose $x'_1 = x_{1,i'}$ for some $i' \in [Q]$. Then clearly the sets $\{h_{\mu_i, \chi_i}\}_{i \neq i'}$ and $\{h_{\mu_{i'}, \chi_{i'}}, h_{\mu', \chi'}\}$ are (strong) linearly independent because denominator polynomials are distinct.

To finish the claim, we must show that $\{h_{\mu_{i'}, \chi_{i'}}, h_{\mu', \chi'}\}$ is (strongly) linearly independent. We first handle the BB-tran case which is clearer. We have

$$\begin{aligned} h_{\mu_{i'}, \chi_{i'}} &= (\mathbf{u}_{i'}^T, \mathbf{x}_{0,i'}^T) \cdot \tilde{\mathbf{b}}_0 + 1/(\tilde{b}_1 - x'_1) \quad \text{and} \\ h_{\mu', \chi'} &= ((\mathbf{u}')^T, (\mathbf{x}'_0)^T) \cdot \tilde{\mathbf{b}}_0 + 1/(\tilde{b}_1 - x'_1), \end{aligned}$$

to cancel out the second term, the only way is to take the difference. However, since $\mu_{i'} = \mathbf{u}_{i'}$ and $\mu' = \mathbf{u}'$ are distinct, taking difference does not make the first term vanish, thus proving the claim. For the BBS case, any linear combination of

$$((1, \mathbf{u}_{i'}^T, \mathbf{x}_{0,i'}^T) \cdot \tilde{\mathbf{b}}_0) / (\tilde{b}_1 - x'_1) \quad \text{and} \quad ((1, (\mathbf{u}')^T, (\mathbf{x}'_0)^T) \cdot \tilde{\mathbf{b}}_0) / (\tilde{b}_1 - x'_1)$$

to zero would also be a linear combination of

$$(1, \mathbf{u}_{i'}^T, \mathbf{x}_{0,i'}^T) \cdot \tilde{\mathbf{b}}_0 \quad \text{and} \quad (1, (\mathbf{u}')^T, (\mathbf{x}'_0)^T) \cdot \tilde{\mathbf{b}}_0$$

to zero. However, since $\mu_{i'} = \mathbf{u}_{i'}$ and $\mu' = \mathbf{u}'$ are distinct, the two coefficients of such a linear combination must either be zero or have different magnitudes. In the latter case, the coefficient of the first variable in $\tilde{\mathbf{b}}_0$ in the linear combination would not be zero. We thus conclude that the only possible linear combination to make zero is the one with zero coefficients.

Proof-Friendliness of Suggested Candidates. We briefly comment on the proof-friendliness on the BB-tran and BBS instantiations suggested in Table 1, since we believe they are of the most practical relevance. Essentially, the verification relation of both schemes are simple bounded-norm satisfiability relations of quadratic equations, which can be handled concretely efficiently by state-of-the-art lattice-based proof systems such as [LNP22b].

For simplicity, for signature verification we collect all supposedly bounded-norm components as a vector and check the norm of such a vector.¹⁴ The verification relation of the BB-tran instantiation then takes the form:

$$\mathbf{A} \cdot \mathbf{s} = \mathbf{B}_0 \cdot \begin{pmatrix} \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix} + \mathbf{1}_n \odot (\mathbf{b}_1 - \mathbf{1}_n \cdot x_1) \pmod{q} \quad \text{and} \quad \|(\mathbf{s}^T, \mathbf{u}^T, \mathbf{x}_0^T, x_1)\| \leq \beta$$

¹⁴ We note that, depending on the application, the norm bounds for the signature vector \mathbf{s} and the message \mathbf{u} , for example, could differ.

which can be rearranged to the form

$$(\mathbf{b}_1 \square (\mathbf{A} - \mathbf{B})) \cdot \begin{pmatrix} \mathbf{s} \\ \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix} - (\mathbf{A} - \mathbf{B}) \cdot \begin{pmatrix} \mathbf{s} \\ \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix} \cdot x_1 = \mathbf{1}_n \bmod q,$$

$$\|(\mathbf{s}^\top, \mathbf{u}^\top, \mathbf{x}_0^\top, x_1)\| \leq \beta$$

where $\mathbf{b}_1 \square (\mathbf{A} - \mathbf{B})$ denotes the matrix obtained by taking the Hadamard products between \mathbf{b}_1 and every column of $(\mathbf{A} - \mathbf{B})$. To produce a zero-knowledge proof of such a relation, one strategy is to commit to the expanded witnesses

$$\mathbf{w}_1 = \begin{pmatrix} \mathbf{s} \\ \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix}, \quad w_2 = x_1, \quad \mathbf{w}_3 = \begin{pmatrix} \mathbf{s} \\ \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix} \cdot x_1,$$

prove that \mathbf{w}_1 and w_2 are of bounded norm and satisfy the linear relation, and prove that $\mathbf{w}_3 = \mathbf{w}_1 \cdot w_2$. Similarly, the verification relation of the BBS instantiation takes the form:

$$\mathbf{A} \cdot \mathbf{s} = \left(\mathbf{B}_0 \begin{pmatrix} 1 \\ \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix} \right) \odot (\mathbf{b}_1 - \mathbf{1}_n \cdot x_1) \bmod q \quad \text{and} \quad \|(\mathbf{s}^\top, \mathbf{u}^\top, \mathbf{x}_0^\top, x_1)\| \leq \beta$$

which can be rearranged to the proof-friendly form

$$(\mathbf{b}_1 \square \mathbf{A}) \cdot \mathbf{s} - \mathbf{A} \cdot \mathbf{s} \cdot x_1 - \mathbf{B}_0 \cdot \begin{pmatrix} 1 \\ \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix} = \mathbf{0}_n \bmod q, \quad \|(\mathbf{s}^\top, \mathbf{u}^\top, \mathbf{x}_0^\top, x_1)\| \leq \beta.$$

We expect proving knowledge of a signature for the BBS instantiation is slightly more efficient, since the quadratic part $\mathbf{s} \cdot x_1$ is of lower dimension (independent of the message \mathbf{u}) than $\begin{pmatrix} \mathbf{s} \\ \mathbf{u} \\ \mathbf{x}_0 \end{pmatrix} \cdot x_1$ in the BB-tran instantiation.

We further note that proving knowledge of a signature for either instantiation suggested above only requires proving that the witness is norm-bounded. In contrast, the concrete scheme suggested in [BLNS23b] requires proving that the witness has a binary component, on top of the witness being norm-bounded.

Concrete Parameters. To estimate performance, we heuristically assume that breaking the sEUF-CMA security of the BB-tran and BBS instantiations are both as hard as solving SIS with dimension φn . As in [BLNS23b], we consider instantiations based on NTRU trapdoors [DLP14] for power-of-2 cyclotomic rings, meaning that $n = 1$, $m = 2$, and φ is a power of 2. NTRU trapdoors [DLP14] allow sampling preimages statistically close to Gaussian with parameter $s \geq 1.17 \cdot \sqrt{q} \cdot \eta_\epsilon(\mathbb{Z}^{2\varphi})$ where $\eta_\epsilon(\mathbb{Z}^{2\varphi}) \leq \sqrt{\log(4\varphi(1 + 2^\lambda))}/\pi$. Recall that a message-signature tuple is $(\mathbf{s}^\top, \mathbf{u}^\top, \mathbf{x}_0^\top, x_1) \in \mathcal{R}^{2+\ell_m+\ell_r+1}$. We hence set the norm bound β to be such that $s\sqrt{(3 + \ell_m + \ell_r)} \cdot \varphi \leq \beta < q$. We adopt the standard optimisation of omitting one ring element from the signature (which can be derived from the verification equation), and thus a signature is of size $(2 + \ell_r) \cdot \varphi \cdot \log \beta$ bits. Since the signature schemes' security anyway rely on the hardness of vSIS, we also adopt the optimisation of deriving the entire public key from a single \mathcal{R}_q element (given by the NTRU trapdoor algorithm), hence the public key size is $\varphi \cdot \log q$ bits. Following [BLNS23b], we set $\ell_r = 2$ and $\ell_m = 1$, and additionally consider $\ell_m = 128$. We run the Lattice Estimator [APS15]¹⁵ with these parameter constraints and obtain the parameters and sizes presented in Table 2. The script is [attached](#).

5 Generalised ISIS_f

We define a generalised version of the ISIS_f assumption [BLNS23b] and relate it to the security of the signature scheme Σ_{vSIS} in Section 4, therefore also the hinted vSIS assumptions. We provide cryptanalytic discussions around these assumptions.

¹⁵ Commit 162c5053 of <https://github.com/malb/lattice-estimator/>.

Security Level	φ	q	β	s	ℓ_m	ℓ_r	$ \text{pk} $	$ \text{sig} $
193	1024	2^{20}	$2^{18.99}$	2^{13}	1	2	2.5	9.5
150	1024	2^{25}	$2^{23.73}$	2^{16}	128	2	3.1	11.9
399	2048	2^{22}	$2^{20.77}$	2^{14}	1	2	5.5	20.8
312	2048	2^{27}	$2^{25.50}$	2^{17}	128	2	6.8	25.5

Table 2: Estimated parameters for BB-tran and BBS. Sizes are in KB.

5.1 ISIS_f and Interactive ISIS_f

We first recall the ISIS_f assumption and its interactive variant IntISIS_f defined in [BLNS23b], both parametrised by some fixed public function $f: [N] \rightarrow \mathcal{R}_q^n$.

In the ISIS_f experiment, the adversary receives a random matrix \mathbf{A} over \mathcal{R}_q and polynomially many tuples (\mathbf{s}_i, μ_i) , where $\mu_i \leftarrow_{\$} [N]$ and \mathbf{s}_i is a short preimage $\mathbf{A} \cdot \mathbf{s}_i = f(\mu_i) \bmod q$. Its goal is to find a new tuple (\mathbf{s}^*, μ^*) which satisfies $\mathbf{A}\mathbf{s}^* = f(\mu^*) \bmod q$, where \mathbf{s}^* is short and $\mu^* \in [N]$.

Note that an ISIS_f instance immediately yields a signature scheme which is sEUF-RMA secure. However, instead of building cryptographic primitives directly from the ISIS_f assumption, [BLNS23b] introduced an intermediate assumption which they call *interactive* ISIS_f.

In the interactive ISIS_f experiment, the adversary receives random matrices \mathbf{A}, \mathbf{C} over \mathcal{R}_q and is given access to an oracle which, on the i -th (adaptive) query short vectors $(\mathbf{m}_i, \mathbf{r}_i)$, returns a short preimage \mathbf{s}_i and a value μ_i satisfying $\mathbf{A} \cdot \mathbf{s}_i = f(\mu_i) + \mathbf{C} \begin{bmatrix} \mathbf{m}_i \\ \mathbf{r}_i \end{bmatrix} \bmod q$, where $\mu_i \leftarrow_{\$} [N]$. Its goal is to find new tuple $(\mathbf{s}^*, \mu^*, \mathbf{m}^*, \mathbf{r}^*)$ which satisfies $\mathbf{A}\mathbf{s}^* = f(\mu^*) + \mathbf{C} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \bmod q$, where $\mathbf{s}^*, \mathbf{m}^*, \mathbf{r}^*$ are short, $\mu^* \in [N]$, and $\mathbf{m}^* \notin \{\mathbf{m}_1, \dots, \mathbf{m}_Q\}$.

An important step in [BLNS23b] is to show that the interactive ISIS_f assumption is implied by the (non-interactive) ISIS_f assumption, hence providing a convenient interface for yielding simple constructions of ordinary, group and blind signatures. Furthermore, [BLNS23b] showed how it can be turned into an efficient anonymous credential system when combined with compatible zero-knowledge proof systems.

5.2 Generalised Assumptions

Recall that the ISIS_f assumption is parametrised by a fixed (deterministic) function f . Below, we generalise the ISIS_f assumption and its interactive variant by extending the input space of f as follows: 1) It additionally inputs a function key $\kappa \in \mathcal{K}$ which is sampled from a distribution at the beginning of the security experiment. 2) It additionally inputs a randomness χ chosen from \mathcal{X} . We denote the non-interactive and interactive variants of the generalised assumptions GenISIS_f and IntGenISIS_f respectively. Unlike [BLNS23b], we consider a “strong unforgeability” flavour for IntGenISIS (i.e, the set S contains full “signatures”, not just \mathbf{m}). Further, in contrast to [BLNS23b] where the input μ to f in the ISIS_f experiment is always random, we additionally consider an adaptive variant GenISIS_f⁺, where the adversary can freely choose μ .

Definition 8 (Generalised ISIS_f Assumptions (GenISIS_f)). Let $\mathcal{R}, n, m, \ell_{\text{msg}}, \ell_{\text{tag}}, q, Q, \beta, \gamma, s, f$ be parametrised by λ , where $n, m, \ell_{\text{msg}}, \ell_{\text{tag}}, q$ are positive integers, Q is a non-negative integer, $\beta, \gamma, s \in \mathbb{R}^+$, and $f: \mathcal{K} \times \mathcal{M} \times \mathcal{X} \rightarrow \mathcal{R}_q^n$ is a function where the domain is efficiently sampleable¹⁶. Let $\text{params} = (\mathcal{R}, n, m, q, Q, \beta, s, f)$. The GenISIS_{params} (resp. GenISIS_{params}⁺) assumption states that, for any PPT adversary \mathcal{A} ,

$$\begin{aligned} \text{AdvGenISIS}_{\mathcal{A}'}^{\text{params}'}(\lambda) &:= \Pr[\text{ExpGenISIS}_{\text{params}, \mathcal{A}}(\lambda) = 1] \leq \text{negl}(\lambda), \\ (\text{resp. } \text{AdvGenISIS}_{\mathcal{A}'}^{+, \text{params}'}(\lambda) &:= \Pr[\text{ExpGenISIS}_{\text{params}, \mathcal{A}}^+(\lambda) = 1] \leq \text{negl}(\lambda)) \end{aligned}$$

where the experiments are defined in Figure 5. Similarly, for $\text{params} = (\mathcal{R}, n, m, \ell_{\text{msg}}, \ell_{\text{tag}}, q, Q, \beta, \gamma, s, f)$, the IntGenISIS_{params} assumption states that, for any PPT adversary \mathcal{A} ,

$$\text{AdvIntGenISIS}_{\mathcal{A}'}^{\text{params}'}(\lambda) := \Pr[\text{ExpIntGenISIS}_{\text{params}, \mathcal{A}}(\lambda) = 1] \leq \text{negl}(\lambda).$$

¹⁶ I.e. there are efficient algorithms for uniformly sampling from $\mathcal{K}, \mathcal{M}, \mathcal{X}$, respectively.

ExpGenISIS _{params, A} ⁺ (1 ^λ)	ExplntGenISIS _{params, A} (1 ^λ)
$S := \emptyset$	$S := \emptyset$
$\kappa \leftarrow \mathcal{K}; \mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$	$\kappa \leftarrow \mathcal{K}; \mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}; \mathbf{C} \leftarrow \mathcal{R}^{n \times (\ell_{\text{msg}} + \ell_{\text{tag}})}$
$(\mathbf{s}^*, \mu^*, \chi^*) \leftarrow \mathcal{A}^{\text{OPre}}(\kappa, \mathbf{A})$	$(\mathbf{s}^*, \mu^*, \chi^*, \mathbf{m}^*, \mathbf{r}^*) \leftarrow \mathcal{A}^{\text{OSign}}(\kappa, \mathbf{A}, \mathbf{C})$
$\mathbf{t}^* := f(\kappa, \mu^*, \chi^*)$	$\mathbf{t}^* = f(\kappa, \mu^*, \chi^*) + \mathbf{C} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix}$
$b_0 := (\mathbf{A} \cdot \mathbf{s}^* = \mathbf{t}^* \text{ mod } q)$	$b_0 := (\mathbf{A} \cdot \mathbf{s}^* = \mathbf{t}^* \text{ mod } q)$
$b_1 := (0 < \ \mathbf{s}^*\ \leq \beta)$	$b_1 := (0 < \ \mathbf{s}^*\ \leq \beta)$
$b_2 := ((\mathbf{s}^*, \mu^*, \chi^*) \notin S)$	$b_2 := ((\mathbf{s}^*, \mu^*, \chi^*, \mathbf{m}^*, \mathbf{r}^*) \notin S)$
return $b_0 \wedge b_1 \wedge b_2$	$b_3 := (\ (\mathbf{m}^*, \mathbf{r}^*)\ \leq \gamma)$ return $b_0 \wedge b_1 \wedge b_2 \wedge b_3$
OPre(μ')	OSign(\mathbf{m}, \mathbf{r})
assert $ S < Q$	assert $(\ (\mathbf{m}, \mathbf{r})\ \leq \gamma) \wedge (S < Q)$
$\mu \leftarrow \mathcal{M}; \chi \leftarrow \mathcal{X}$	$\mu \leftarrow \mathcal{M}; \chi \leftarrow \mathcal{X}$
$\mu := \mu' \quad // \text{ Overwrite with input}$	$\mathbf{t} = f(\kappa, \mu, \chi) + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix}$
$\mathbf{t} = f(\kappa, \mu, \chi)$	$\mathbf{s} \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{t}}(\mathbf{A}), s}$
$\mathbf{s} \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{t}}(\mathbf{A}), s}$	$S := S \cup \{(\mathbf{s}, \mu, \chi, \mathbf{m}, \mathbf{r})\}$
$S := S \cup \{(\mathbf{s}, \mu, \chi)\}$	return (\mathbf{s}, μ, χ)
return (\mathbf{s}, μ, χ)	

Fig. 5: The GenISIS, GenISIS⁺ and IntGenISIS experiments.

When the parameters are clear from the context, we drop most of them and simply write GenISIS_f, GenISIS_f⁺ and IntGenISIS_f.

The GenISIS_f and IntGenISIS_f experiments (and hence assumptions) essentially coincide with the ISIS_f and IntISIS_f assumptions from [BLNS23b] respectively if f is a trivial family of deterministic functions, i.e. if $\mathcal{K} = \mathcal{X} = \{0\}$ are sets with a single element.¹⁷ A minor change is the “strong unforgeability” flavour in our definition of IntGenISIS_f where, compared to [BLNS23b], S contains $(\mathbf{s}, \mu, \chi, \mathbf{m}, \mathbf{r})$ instead of just \mathbf{m} .

Theorem 6 (Adapted from [BLNS23b, Theorem 3.3]). *Let the parameters $\text{params} = (\mathcal{R}, q, n, m, \ell_{\text{msg}}, \ell_{\text{tag}}, s, \beta, \gamma, f)$ be such that \mathcal{R} is a power-of-two cyclotomic ring, $q/2 > \gamma \geq 1$, $m = n \log q + \omega(\log \lambda)$. Let $\epsilon \leq \text{negl}(\lambda)$. Suppose that $s \geq \max\left(\eta_{\min}(\epsilon), \sqrt{\lambda} \gamma \varphi \sqrt{(\ell_{\text{msg}} + \ell_{\text{tag}})m}\right)$, where $\eta_{\min}(\epsilon)$ is such that $\eta_{\epsilon}(\Lambda_q^{\perp}(\mathbf{A})) \geq \eta_{\min}(\epsilon)$ with probability at most $2^{-\varphi}$ over the randomness of $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$, and that $\varphi \in \omega(\log \lambda)$. Then for every PPT adversary \mathcal{A} in ExplntGenISIS_{params, A}, there is a PPT adversary \mathcal{A}' in ExpGenISIS_{params', A'} such that*

$$\text{AdvGenISIS}_{\mathcal{A}'}^{\text{params}'}(\lambda) \geq \text{AdvIntGenISIS}_{\mathcal{A}}^{\text{params}}(\lambda) / \text{poly}(\lambda) - \text{negl}(\lambda)$$

where $\text{params}' = (\mathcal{R}, q, n, m, \text{poly}(\lambda) \cdot Q, s, \beta' = \beta + \gamma, \varphi \sqrt{(\ell_{\text{msg}} + \ell_{\text{tag}})m})$.

The proof of Theorem 6 is almost identical to that of [BLNS23b, Theorem 3.3] (despite the “strong unforgeability” of IntGenISIS_f). We make the following translations:

- (From probabilistic to deterministic f .) The function f in [BLNS23b] is *deterministic*, i.e. there is no randomness space \mathcal{X} . However, since we sample both $\mu \in \mathcal{M}$ and $\chi \in \mathcal{X}$ uniformly (in the non-adaptive setting), we could as well consider $f': \mathcal{K} \times \mathcal{M}' \rightarrow \mathcal{R}_q^n$ where $\mathcal{M}' = \mathcal{M} \times \mathcal{X}$. Thus, for the sake of proving Theorem 6, we can w.l.o.g. assume f is deterministic.

¹⁷ The mostly syntactical difference is that we give \mathcal{A} access to a sampling oracle, whereas [BLNS23b] chose to produce Q samples up front and give them to \mathcal{A} .

- (Function families.) The (deterministic) function f in [BLNS23b, Theorem 3.3] is used in a completely black-box way. Thus, as long as f is specified and fixed at the beginning of the experiment (so that it can be used in the proof), the proof applies verbatim to any choice (of distributions) of f , in particular any function sampled from a function family over \mathcal{K} .

For the “strong unforgeability” of IntGenISIS_f , first suppose (as above) w.l.o.g. that $\mathcal{X} = \{0\}$, so that GenISIS_f coincides with ISIS_f (up to being a function family). Given a forgery $(\mathbf{s}^*, \mu^*, 0, \mathbf{m}^*, \mathbf{r}^*)$, the proof of [BLNS23b, Theorem 3.3] proceeds through indistinguishable game hops to prepare for the reduction to ISIS_f that is explained after [BLNS23b, Lemma 3.12]. There, it is asserted that (\mathbf{s}^*, μ^*) must be fresh for the adversary to win, which means the reduction wins its ISIS_f instance. This assertion holds even if $(\mathbf{m}^*, \mathbf{r}^*)$ were reused.

Remark 7. We expect that Theorem 6 generalises to any ring \mathcal{R} for which there is a suitable regularity lemma.

5.3 Generalised ISIS and Strong Hinted vSIS

We establish connection between GenISIS_f and the sEUF-RMA security of the vSIS-based signature scheme Σ_{vSIS} (Figure 4). In essence, we show that these two are equivalent experiments by simple renaming of variables. Chaining with Theorem 5 yields a reduction from s- $\text{\$Hint-vSIS}$ to GenISIS_f . The reverse reduction, provided by Theorem 8, establish the equivalence of s- $\text{\$Hint-vSIS}$ and GenISIS_f for given parameters. Note that it is the parameters of GenISIS_f that are restricted, as it corresponds to a more general assumption. In particular, the equivalence holds only when the set of functions f consists of ℓ -variate rational functions.

Theorem 7 ($\text{GenISIS}_f \Leftrightarrow \text{sEUF-RMA}_{\Sigma_{\text{vSIS}}}$). *Let $\mathcal{R}, n, m, \ell, q, s, \beta, \delta, \mathcal{M}, \mathcal{X}, \mathcal{H}$ be parameters of the signature scheme Σ_{vSIS} in Figure 4, in particular $\mathcal{H} := \{h_{\mu, \chi} \mid \mu \in \mathcal{M}, \chi \in \mathcal{X}\}$ is a set of ℓ -variate rational functions. Let $\text{params}_1 := (\mathcal{R}, n, m, q, Q, \beta, s, f)$ with $f : \mathcal{R}_q^{n \times \ell} \times \mathcal{M} \times \mathcal{X} \rightarrow \mathcal{R}_q^n$ such that $f(\mathbf{B}, \mu, \chi) = h_{\mu, \chi}(\mathbf{B})$ for any $\mathbf{B} \in \mathcal{R}_q^{n \times \ell}$, and $Q \leq \text{poly}(\lambda)$. Suppose $(\mathcal{R}, n, m, q, s)$ are admissible parameters for $(\text{TrapGen}, \text{SampPre})$. There is a PPT adversary \mathcal{B} winning the $\text{ExpGenISIS}_{\text{params}_1, \mathcal{B}}$ experiment with non-negligible probability if and only if there is a PPT adversary \mathcal{A} winning the $\text{sEUF-RMA}_{\Sigma_{\text{vSIS}}, \mathcal{A}, Q}$ experiment with non-negligible probability.*

Theorem 8 ($\text{GenISIS}_f \Rightarrow \text{s-}\text{\$Hint-vSIS}$). *Let $\text{params}_1 := (\mathcal{R}, n, m, q, Q, \beta, s, f)$, $\mathcal{H} := \{h_{\mu, \chi} \mid \mu \in \mathcal{M}, \chi \in \mathcal{X}\}$ a set of ℓ -variate rational function, $\mathcal{G} := \mathcal{H} \cup \{0\}$, and \mathcal{P} a predicate on sets of rational functions with $f : \mathcal{R}_q^{n \times \ell} \times \mathcal{M} \times \mathcal{X} \rightarrow \mathcal{R}_q^n$ such that $f(\mathbf{B}, \mu, \chi) = h_{\mu, \chi}(\mathbf{B})$. Let $\text{params}_0 := (\mathcal{R}, n, k, q, Q, \beta, s, \mathcal{F}, \mathcal{G}, \mathcal{H}, \mathcal{P})$, where $k = m + \ell$ and \mathcal{F} is such that for formal variables $(\tilde{\mathbf{a}}^T, \tilde{\mathbf{b}}^T)$ of dimension k , $\mathcal{F}(\tilde{\mathbf{a}}^T, \tilde{\mathbf{b}}^T) = \tilde{\mathbf{a}}^T$. Suppose also that for all $\mathcal{Q} \subseteq \mathcal{H}$, all $g^* \in \mathcal{G} \setminus \mathcal{Q}$, it holds $\mathcal{P}(\mathcal{F} \cup \mathcal{Q} \cup (\{g^*\} \setminus \{0\})) = 1$. If the $\text{GenISIS}_{\text{params}_1}$ assumption holds, then the s- $\text{\$Hint-vSIS}_{\text{params}_0}$ assumption holds.*

The proofs of Theorems 7 and 8 are deferred to Appendices A.3 and A.4. Applying Theorems 5 and 7, we have the following immediate corollary.

Corollary 1 (s- $\text{\$Hint-vSIS} \Rightarrow \text{GenISIS}_f$). *Let $\text{params}_0 = (\mathcal{R}, n, k, q, \beta, \mathcal{F}, \mathcal{G}, \mathcal{H}, \mathcal{P}, Q)$ and $m, \ell, s, \delta, \delta', \mathcal{M}, \mathcal{X}$ be such that the constraints in Theorem 5 are satisfied. Let $\text{params}_1 := (\mathcal{R}, n, m, Q, q, \beta, s, f)$ with $f : \mathcal{R}_q^{n \times \ell} \times \mathcal{M} \times \mathcal{X} \rightarrow \mathcal{R}_q^n$ such that $f(\mathbf{B}, \mu, \chi) = h_{\mu, \chi}(\mathbf{B})$ for any $\mathbf{B} \in \mathcal{R}_q^{n \times \ell}$. If the s- $\text{\$Hint-vSIS}_{\text{params}_0}$ assumption holds for every polynomial $Q(\lambda)$, then so does the $\text{GenISIS}_{\text{params}_1}$ assumption.*

5.4 Robustness and Linear Combination Attacks against (Gen)ISISf

We study the robustness of (generalised) ISIS_f assumptions, in particular in relation to a natural attack that was also discussed in [BLNS23b] for the special case where the function f was binary decomposition.¹⁸ We extend this discussion to our setting, and connect it to the (strong) hinted vSIS assumption (Definition 5). For simplicity, we omit the function key κ for f and the randomness \mathcal{X} in the first discussion, and consider a fixed $f : \mathcal{M} \rightarrow \mathcal{R}_q^n$ as in the original ISIS_f assumption.

¹⁸ More precisely, in [BLNS23b], f is a function which first binary decomposes the input, linear-maps it to the domain of the coefficient embedding, then outputs the corresponding ring vector.

Preimage resistance. A necessary requirement on $f: \mathcal{M} \rightarrow \mathcal{R}_q^n$ is that it is hard to invert on random elements in the codomain. Otherwise, the hash-then-sign paradigm is clearly broken, and $\mu^* = f^{-1}(\mathbf{As}^*)$ for a random short \mathbf{s}^* breaks the assumption. Note here that “hard-to-invert” can be statistically satisfied, if $|\mathcal{M}|/|\mathcal{R}_q^n|$ is negligible. For example, this is the case in [BLNS23b] and our BB-lite instantiation. Examples for computational hardness include the non-BB-lite instantiations in Table 1, and f being a collision-resistant hash or random oracle.

Linear combination attacks (LCA). The attack idea of an LCA is quite simple: Given tuples (\mathbf{s}_i, μ_i) with $\mathbf{As}_i = f(\mu_i)$, somehow compute short coefficients c_i together with μ^* such that

$$\sum_i c_i f(\mu_i) = f(\mu^*) \quad (1)$$

holds. Then, for $\mathbf{s}^* = \sum_i c_i \mathbf{s}_i$, we get

$$\mathbf{As}^* = \sum_i c_i \mathbf{As}_i = \sum_i c_i f(\mu_i) = f(\mu^*). \quad (2)$$

Intuitively, the attack exploits (approximate) linearity of f , or rather, within the image $\text{im}(f)$. For example, in [BLNS23b], the binary decomposition function $f: [N] \rightarrow \mathcal{R}_q^n$ appears to be highly non-linear at first glance. However, the required linearity for the attack is present in the image $\text{im}(f) = \mathcal{R}_2^n \subseteq \mathcal{R}_q^n$. Indeed, since f is far from one-way over its image, for the purpose of analysing LCA against f , we can equivalently replace f by $g: \mathcal{R}_2^n \rightarrow \mathcal{R}_q^n$, where g is simply the identity embedding.

LCAs against binary decomposition. We first discuss how and why linear combination attacks do (not) break the ISIS_f assumption, when instantiated with $f(x) = \text{bindecomp}(x) \in \mathcal{R}_2^n$ as in [BLNS23b]. As noted earlier, we consider $g: \mathcal{R}_2^n \rightarrow \mathcal{R}_q^n$ where $g(\mathbf{m}) = \mathbf{m}$ and $\mathcal{M} = \mathcal{R}_2^n$, instead of the function f .

The presumed security of $g: \mathcal{R}_2^n \rightarrow \mathcal{R}_q^n$ against linear combination attacks crucially relies on two properties. On the one hand, the domain $\mathcal{R}_2^n \subseteq \mathcal{R}_q^n$ should be high dimensional. This makes it unlikely that for two random $\mathbf{m}_1, \mathbf{m}_2 \in \mathcal{R}_2^n$, their sum or difference $\pm \mathbf{m}_1 \pm \mathbf{m}_2$ is again in \mathcal{R}_2^n , i.e. again a bitstring; indeed, the probability is precisely $2 \cdot (3/4)^{\varphi n} - 2^{-\varphi n}$ as is easily checked.¹⁹ Hence, we can hope that finding a *sufficiently short* linear combination of $\mathbf{m}_i = f(\mu_i)$ with

$$\sum_i c_i \mathbf{m}_i = \mathbf{m}^* \in \mathcal{R}_2^n$$

for any choice of \mathbf{m}^* remains difficult, given sufficiently high dimensions (at the very least $\varphi n \geq \lambda / \log_2(4/3) \approx 2.41\lambda$). Moreover, one should account for attacks which use integer linear programming (ILP) to tackle similar problems [BDE⁺18, HM17], see [BLNS23b, Section 3.1.2]. On the other hand, the dimension φm of preimages should be high and the norm bound β very strict, so that even if we have two ISIS_f pairs $(\mathbf{s}_1, \mathbf{m}_1), (\mathbf{s}_2, \mathbf{m}_2)$ where $\mathbf{m}_1 \pm \mathbf{m}_2$ is again binary, we get $\|\mathbf{s}_1 \pm \mathbf{s}_2\| > \beta$ almost certainly. To thwart the attacks, [BLNS23b] exploits the high dimension φm and sets the bound β close to $\sqrt{\varphi m s}$.

Limited robustness of binary decomposition. To investigate the security-efficiency tradeoff, we strengthen the GenISIS_f assumption slightly: (1) we consider *selective* GenISIS_f^+ where the adversary chooses (μ'_1, \dots, μ'_Q) , receives $\text{OPre}(\mu'_i)$ for $i \in [1, \dots, Q]$, and has no further access to OPre ; (2) we allow a relaxed norm bound, namely $\sqrt{2}\beta$ instead of β , for forgeries. While this seems like a mild strengthening, security completely breaks down for $f = \text{bindecomp}$, even with a single selective query: By selectively querying a standard unit vector \mathbf{e} (in coefficient embedding), say $\mathbf{e}_1 = (1, 0, \dots) \in \{0, 1\}^{\varphi n}$ (or even $\mathbf{0}$), and one (random) query \mathbf{m}_1 , the naive linear combination $\mathbf{e}_1 + \mathbf{m}_1$ breaks with probability 1/2. Of course, this attack generalises: Any \mathbf{e} of low Hamming weight is beneficial for linear combination attacks, and success can be amplified by trying many \mathbf{m}_i 's. The above total break with just a single selective query suggests, that one should consider other (presumably) more robust functions f , which perhaps allows a better security-efficiency tradeoff and gives higher confidence in the hardness of the problem.

¹⁹ We just need to count strings which are pointwise \leq or \geq . Perhaps surprisingly, for $\varphi n = 256$, this happens with probability roughly $2^{-105.25} \gg 2^{-128}$.

LCA against Candidate Rational Functions. We elaborate on why we believe our candidate instantiations of GenISIS_f (cf. Section 4.2) offer better security than the binary decomposition function. First, by Theorem 8, to break GenISIS_f it suffices to break the corresponding s - $\$$ Hint- v SIS assumption. For concreteness, we focus on the function $h_{\mu,\chi}(\tilde{b}_0, \tilde{b}_1) = \frac{\tilde{b}_0 - \mu}{\tilde{b}_1 - \chi}$ which captures all important structures of the examples in Table 1. For simplicity, we consider module rank $n = 1$.

To begin, we observe that, for $\beta_\mu \geq \sqrt{q}$, the distribution $1/\mu \pmod{q}$ of inverses of short elements $\mu \leftarrow \mathcal{M} = \{|\mu| \leq \beta_\mu \mid \mu \in \mathbb{Z}\}$ is “close to uniformly” [Shp12] for some non-cryptographic measure of closeness. Nevertheless, let us adopt the heuristic that $(b_0 - \mu)/(b_1 - \chi) \pmod{q}$ for short (μ, χ) is “uniformly random enough” for the purpose of attaining SIS hardness. Given many short samples $(\mathbf{s}_i, \mu_i, \chi_i)$ satisfying $\langle \mathbf{a}, \mathbf{s}_i \rangle = (b_0 - \mu_i)/(b_1 - \chi_i) \pmod{q}$, there are two natural strategies to find a new short solution $(\mathbf{s}^*, \mu^*, \chi^*)$.

The first is to ignore the hints and attempt to solve the problem directly. Towards this, an idea is to sample a random short \mathbf{s}^* , compute $t^* := \langle \mathbf{a}, \mathbf{s}^* \rangle \pmod{q}$, and try to find (μ^*, χ^*) such that $h_{\mu^*, \chi^*}(b_0, b_1) = t^* \pmod{q}$. In our example, this means solving $t^* \chi^* - \mu^* = t^* b_1 - b_0 \pmod{q}$ for a short solution (μ^*, χ^*) . However, if we sampled \mathbf{s}^* from a sufficiently entropic distribution, then t^* would be close to uniformly random, and thus we essentially need to solve a random ISIS instance.

Another strategy is to somehow make use of the hints. The idea is essentially that of the LCA – to find a short linear combination of $h_{\mu_i, \chi_i}(b_0, b_1)$ which yields $h_{\mu^*, \chi^*}(b_0, b_1)$ for some (μ^*, χ^*) . As before, if we pick the linear combination blindly and then attempt to solve the equation for (μ^*, χ^*) , then we will face an ISIS instance. Since $h_{\mu_1, \chi_1}, \dots, h_{\mu_Q, \chi_Q}, h_{\mu^*, \chi^*}$ are strongly linearly independent, choosing the linear combination without taking (b_0, b_1) into account is unlikely to succeed. We are thus left with the option of picking the short linear combination by somehow exploiting our knowledge on (b_0, b_1) , which appears difficult.

We remark that the difficulty of the above attacks crucially rely on the restriction that μ, χ are short. Indeed, if the shortness condition is dropped, then there are simple attacks by simple arithmetic. Also note that all attack strategies discussed are non-adaptive. At present, we are unaware of meaningful ways to exploit adaptivity. This holds true even when additionally considering a relaxed norm bound $\beta' \gg \beta$ on forgeries. As long as the related problems (such as SIS, inverting f) remain hard, we do not know how to exploit a norm check for $\beta' \gg \beta$, even if β' is much larger than the expected norm $\sqrt{\varphi m s}$ of preimages \mathbf{s} which are obtained by an (adaptive) OPre . Intuitively, the non-linearity of f obstructs the LCA attack, which is the only way we know how to take advantage of the short preimages \mathbf{s} to obtain a slightly longer preimage (e.g. as for $f = \text{bindecomp}$).

To summarise: The $\text{GenISIS}_f^{(+)}$ assumption(s) for our families f seems quite robust, but better cryptanalysis is crucial to gauge its hardness and robustness.

GenISIS_f vs. one-more ISIS. The one-more ISIS (OM-ISIS) assumption was introduced in [AKSY22] for building lattice-based blind signatures. There, an adversary is given a *preimage oracle* OPre which produces preimages of images *freely chosen* by the adversary. It also has oracle access to (an arbitrary number of) random challenge vectors (that is, ISIS syndromes). To win the game, the adversary must produce preimages of $\ell + 1$ challenge vectors, if it has made ℓ queries to OPre .

OM-ISIS does not seem directly comparable to any variant (generalised, adaptive, interactive, etc.) of $(\text{Gen})\text{ISIS}_f$. We focus our discussion on GenISIS_f . For OM-ISIS, the images for which OPre provides preimages are chosen freely, but the challenge images are truly random. For GenISIS_f , the preimages are for random images $y = f(\kappa, \mu, \chi)$ and there is no specified challenge, i.e. the adversary chooses freely (μ^*, χ^*) . Interestingly, the selective and adaptive variants of GenISIS_f^+ are closer to OM-ISIS because OPre now works on adversarial inputs, but the choice is still not free as the inputs are “hashed” by f .

As a qualitative difference, we note that in OM-ISIS it is possible to learn a trapdoor of \mathbf{A} , e.g. by requesting sufficiently many preimages of $\mathbf{0}$. Under a relaxed norm check $\beta' > \beta$, the above implies that β' cannot be too large, e.g. OM-ISIS with $\beta' > \sqrt{m} \cdot \sqrt{\varphi m s}$ is broken, cf. [AKSY22]. Thus, similar to ISIS_f , there is a close tie between admissible $\beta' > \beta$ and the expected norm $\sqrt{\varphi m s}$ of the sampled preimages. In contrast, for (adaptive) GenISIS_f , we are not aware of how to obtain a short basis, or even take advantage of relaxed norm checks.

Overall, we can hope that our GenISIS assumptions are as secure (or more so) than OM-ISIS, while leading to smaller round-optimal blind signatures than [AKSY22].

Acknowledgments. The research of Russell W. F. Lai and Ivy K. Y. Woo are supported by Research Council of Finland grants 358951 and 358950 respectively.

References

- ABB10. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 98–115. Springer, Berlin, Heidelberg, August 2010. 4
- ACL⁺22. Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri Aravinda Krishnan Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 102–132. Springer, Cham, August 2022. 2, 3, 9, 10
- AEB20a. Nabil Alkeilani Alkadri, Rachid El Bansarkhani, and Johannes Buchmann. BLAZE: Practical lattice-based blind signatures for privacy-preserving applications. In Joseph Bonneau and Nadia Heninger, editors, *FC 2020*, volume 12059 of *LNCS*, pages 484–502. Springer, Cham, February 2020. 4
- AEB20b. Nabil Alkeilani Alkadri, Rachid El Bansarkhani, and Johannes Buchmann. On lattice-based interactive protocols: An approach with less or no aborts. In Joseph K. Liu and Hui Cui, editors, *ACISP 20*, volume 12248 of *LNCS*, pages 41–61. Springer, Cham, November / December 2020. 4
- AHJ21. Nabil Alkeilani Alkadri, Patrick Harasser, and Christian Janson. BlindOR: an efficient lattice-based blind signature scheme from OR-proofs. In Mauro Conti, Marc Stevens, and Stephan Krenn, editors, *CANS 21*, volume 13099 of *LNCS*, pages 95–115. Springer, Cham, December 2021. 4
- AKSY22. Shweta Agrawal, Elena Kirshanova, Damien Stehlé, and Anshu Yadav. Practical, round-optimal lattice-based blind signatures. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 39–53. ACM Press, November 2022. 4, 22
- APS15. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3):169–203, October 2015. 17
- ASM06. Man Ho Au, Willy Susilo, and Yi Mu. Constant-size dynamic k-TAA. In Roberto De Prisco and Moti Yung, editors, *SCN 06*, volume 4116 of *LNCS*, pages 111–125. Springer, Berlin, Heidelberg, September 2006. 2, 4, 13, 15
- Ban93. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993. 5
- BB08. Dan Boneh and Xavier Boyen. Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology*, 21(2):149–177, April 2008. 2, 4, 13, 15
- BBP23. Johannes Blömer, Jan Bobolz, and Laurens Porzenheim. A generic construction of an anonymous reputation system and instantiations from lattices. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part II*, volume 14439 of *LNCS*, pages 418–452. Springer, Singapore, December 2023. 4
- BBS04. Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 41–55. Springer, Berlin, Heidelberg, August 2004. 1, 2, 4, 13, 15
- BCN18. Cecilia Boschini, Jan Camenisch, and Gregory Neven. Floppy-sized group signatures from lattices. In Bart Preneel and Frederik Vercauteren, editors, *ACNS 18International Conference on Applied Cryptography and Network Security*, volume 10892 of *LNCS*, pages 163–182. Springer, Cham, July 2018. 4
- BDE⁺18. Jonathan Bootle, Claire Delaplace, Thomas Espitau, Pierre-Alain Fouque, and Mehdi Tibouchi. LWE without modular reduction and improved side-channel attacks against BLISS. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 494–524. Springer, Cham, December 2018. 21
- BDK⁺22. 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 Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 95–126. Springer, Cham, May / June 2022. 4
- BLL⁺21. Fabrice Benhamouda, Tancrede Lepoint, Julian Loss, Michele Orrù, and Mariana Raykova. On the (in)security of ROS. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 33–53. Springer, Cham, October 2021. 4
- BLL⁺22. Fabrice Benhamouda, Tancrede Lepoint, Julian Loss, Michele Orrù, and Mariana Raykova. On the (in)security of ROS. *Journal of Cryptology*, 35(4):25, October 2022. 4
- BLNS23a. Ward Beullens, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Lattice-based blind signatures: Short, efficient, and round-optimal. In Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda, editors, *ACM CCS 2023*, pages 16–29. ACM Press, November 2023. 4

- BLNS23b. Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Alessandro Sorniotti. A framework for practical anonymous credentials from lattices. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 384–417. Springer, Cham, August 2023. 2, 3, 4, 15, 17, 18, 19, 20, 21
- BLS19. Jonathan Bootle, Vadim Lyubashevsky, and Gregor Seiler. Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 176–202. Springer, Cham, August 2019. 4
- CGT23. Alishah Chator, Matthew Green, and Pratyush Ranjan Tiwari. SoK: Privacy-preserving signatures. Cryptology ePrint Archive, Report 2023/1039, 2023. 4
- CL03. Jan Camenisch and Anna Lysyanskaya. A signature scheme with efficient protocols. In Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano, editors, *SCN 02*, volume 2576 of *LNCS*, pages 268–289. Springer, Berlin, Heidelberg, September 2003. 1
- CL04. Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 56–72. Springer, Berlin, Heidelberg, August 2004. 4
- CLM23. Valerio Cini, Russell W. F. Lai, and Giulio Malavolta. Lattice-based succinct arguments from vanishing polynomials - (extended abstract). In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 72–105. Springer, Cham, August 2023. 3, 6, 8, 10
- DLP14. Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Berlin, Heidelberg, December 2014. 17
- DM14. Léo Ducas and Daniele Micciancio. Improved short lattice signatures in the standard model. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 335–352. Springer, Berlin, Heidelberg, August 2014. 4
- dPK22. Rafaël del Pino and Shuichi Katsumata. A new framework for more efficient round-optimal lattice-based (partially) blind signature via trapdoor sampling. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 306–336. Springer, Cham, August 2022. 4
- dPLS18. Rafaël del Pino, Vadim Lyubashevsky, and Gregor Seiler. Lattice-based group signatures and zero-knowledge proofs of automorphism stability. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 574–591. ACM Press, October 2018. 4
- Fis06. Marc Fischlin. Round-optimal composable blind signatures in the common reference string model. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 60–77. Springer, Berlin, Heidelberg, August 2006. 4
- GKV10. S. Dov Gordon, Jonathan Katz, and Vinod Vaikuntanathan. A group signature scheme from lattice assumptions. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 395–412. Springer, Berlin, Heidelberg, December 2010. 4
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 1, 2, 5
- GS08. Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 415–432. Springer, Berlin, Heidelberg, April 2008. 1
- HKLN20. Eduard Hauck, Eike Kiltz, Julian Loss, and Ngoc Khanh Nguyen. Lattice-based blind signatures, revisited. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 500–529. Springer, Cham, August 2020. 4
- HM17. Gottfried Herold and Alexander May. LP solutions of vectorial integer subset sums — cryptanalysis of Galbraith’s binary matrix LWE. In Serge Fehr, editor, *PKC 2017, Part I*, volume 10174 of *LNCS*, pages 3–15. Springer, Berlin, Heidelberg, March 2017. 21
- JRS23. Corentin Jeudy, Adeline Roux-Langlois, and Olivier Sanders. Lattice signature with efficient protocols, application to anonymous credentials. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 351–383. Springer, Cham, August 2023. 2, 4
- KLR24. Shuichi Katsumata, Yi-Fu Lai, and Michael Reichle. Breaking parallel ROS: Implication for isogeny and lattice-based blind signatures. In Qiang Tang and Vanessa Teague, editors, *PKC 2024, Part I*, volume 14601 of *LNCS*, pages 319–351. Springer, Cham, April 2024. 4
- LDK⁺22. Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. 1
- LLS13. Fabien Laguillaumie, Adeline Langlois, Benoît Libert, and Damien Stehlé. Lattice-based group signatures with logarithmic signature size. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 41–61. Springer, Berlin, Heidelberg, December 2013. 4

- LLM⁺16. Benoît Libert, San Ling, Fabrice Mouhartem, Khoa Nguyen, and Huaxiong Wang. Signature schemes with efficient protocols and dynamic group signatures from lattice assumptions. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 373–403. Springer, Berlin, Heidelberg, December 2016. [2](#), [4](#)
- LLNW16. 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 Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 1–31. Springer, Berlin, Heidelberg, May 2016. [4](#)
- LNP22a. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Efficient lattice-based blind signatures via gaussian one-time signatures. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part II*, volume 13178 of *LNCS*, pages 498–527. Springer, Cham, March 2022. [4](#)
- LNP22b. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 71–101. Springer, Cham, August 2022. [2](#), [4](#), [16](#)
- LNPS21. Vadim Lyubashevsky, Ngoc Khanh Nguyen, Maxime Plançon, and Gregor Seiler. Shorter lattice-based group signatures via “almost free” encryption and other optimizations. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 218–248. Springer, Cham, December 2021. [4](#)
- Lyu12. Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Berlin, Heidelberg, April 2012. [1](#)
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. [2](#), [4](#), [5](#)
- PFH⁺22. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. [1](#)
- Rüc10. Markus Rückert. Lattice-based blind signatures. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 413–430. Springer, Berlin, Heidelberg, December 2010. [4](#)
- Shp12. Igor E Shparlinski. Modular hyperbolas. *Japanese Journal of Mathematics*, 7(2):235–294, 2012. [22](#)
- SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 617–635. Springer, Berlin, Heidelberg, December 2009. [6](#)
- TZ23. Stefano Tessaro and Chenzhi Zhu. Revisiting BBS signatures. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 691–721. Springer, Cham, April 2023. [2](#), [4](#), [13](#), [15](#)
- Wee22. Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Cham, May / June 2022. [3](#), [6](#)
- WW23. Hoeteck Wee and David J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 385–416. Springer, Cham, April 2023. [9](#)
- YAZ⁺19. Rupeng Yang, Man Ho Au, Zhenfei Zhang, Qiuliang Xu, Zuoxia Yu, and William Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 147–175. Springer, Cham, August 2019. [4](#)

A Missing Proofs

A.1 Proof of Theorem 1

Proof. Case $\mathcal{F} \subseteq \mathcal{R}_q[\tilde{\mathbf{x}}]$ is a set of polynomials. In this case, $g = m = 1$ for any $f/g \in \mathcal{F}$. We aim to verify that, for any not-all-zero coefficients $(c_f)_{f \in \mathcal{F}}$ over \mathcal{R}_q ,

$$\Pr \left[\sum_{f \in \mathcal{F}} c_f \cdot f(\mathbf{a}^T) = 0 \pmod q \mid \mathbf{a} \leftarrow_{\$} \mathcal{R}_q^k \right] \leq \epsilon.$$

We show this by using the Schwartz-Zippel lemma over \mathbb{F} , which states that for any non-zero k -variate degree- d polynomial f over \mathbb{F} ,

$$\Pr [f(\mathbf{a}^T) = 0 \mid \mathbf{a} \leftarrow_{\$} \mathbb{F}^k] \leq d/|\mathbb{F}| = \epsilon.$$

Consider a non-zero k -variate degree- d polynomial f over \mathcal{R}_q , which splits via the Chinese Remainder Theorem (for \mathcal{R}_q) into a tuple of polynomials (f_1, \dots, f_e) over \mathbb{F} . Since f is non-zero, one of (f_1, \dots, f_e) , say f_{i^*} , is non-zero. We thus have

$$\Pr[f(\mathbf{a}^\top) = 0 \mid \mathbf{a} \leftarrow \mathcal{R}_q^k] \leq \Pr[f_{i^*}(\mathbf{a}_{i^*}^\top) = 0 \mid \mathbf{a}_i \leftarrow \mathbb{F}^k] \leq d/|\mathbb{F}| = \epsilon.$$

Now fix any not-all-zero coefficients $(c_f)_{f \in \mathcal{F}}$ over \mathcal{R}_q . Since all $f \in \mathcal{F}$ are linearly independent as polynomials over \mathcal{R}_q , for each $i \in [e]$, the i -th CRT components $(f_i)_{f \in \mathcal{F}}$ are also linearly independent as polynomials over \mathbb{F} . Thus, there exists $i^* \in [e]$ such that $\sum_{f \in \mathcal{F}} c_{f,i^*} \cdot f_{i^*}$, where c_{f,i^*} is the i^* -th CRT component of c_f , is a non-zero polynomial over \mathbb{F} . Since the i^* -th CRT component is non-zero, the polynomial $\sum_{f \in \mathcal{F}} c_f \cdot f$ over \mathcal{R}_q is also non-zero.

Case $\mathcal{F} \subseteq \mathcal{R}_q(\bar{\mathbf{x}}^\top)$ is a set of rational function. The idea is to clear out the denominator using the common multiple m . Consider the polynomial

$$h := m \cdot \sum_{f/g \in \mathcal{F}} c_{f/g} \cdot f/g.$$

By the same argument as in the polynomial case, $\sum_{f/g \in \mathcal{F}} c_{f/g} \cdot f/g$ is a non-zero rational function, and therefore h is a non-zero polynomial. Furthermore, $\deg(h) \leq \deg(m) + \max_{f/g \in \mathcal{F}} \deg(f) = d$. The claim thus follows from the analysis of the polynomial case. \square

A.2 Proof of Lemma 2

Proof. From the definition of strong linear independence, if $1 \notin \mathcal{R}_q\text{-span}(\mathcal{F})$, it holds that, for any $d \in \mathcal{R}_q$ and any not-all-zero $(c_f)_{f \in \mathcal{F}}$,

$$\Pr\left[\sum_{f \in \mathcal{F}} c_f \cdot f(\mathbf{a}^\top) = d \bmod q \mid \mathbf{a} \leftarrow \mathcal{R}_q^k\right] \leq \epsilon_{\mathcal{F}}.$$

Suppose w.l.o.g. that $1 \notin \text{span}_{\mathcal{R}_q}(\mathcal{F})$ (else, swap \mathcal{F} and \mathcal{G}). Unless $\epsilon_{\mathcal{F}} = 1$ or $\epsilon_{\mathcal{G}} = 1$, we know that all elements in \mathcal{F} (resp. in \mathcal{G}) are linearly independent. By assumption $\text{span}_{\mathcal{R}_q}(\mathcal{F}) \cap \text{span}_{\mathcal{R}_q}(\mathcal{G}) = \{0\}$, so $\mathcal{F} \cup \mathcal{G}$ is a basis of $\text{span}_{\mathcal{R}_q}(\mathcal{F} \cup \mathcal{G})$. Hence, for any $h \in \text{span}_{\mathcal{R}_q}(\mathcal{F} \cup \mathcal{G})$, there is a unique coefficient vector $\mathbf{c} \in \mathcal{R}_q^{\mathcal{F} \cup \mathcal{G}}$ such that h is a linear combination of $\mathcal{F} \cup \mathcal{G}$ with coefficient \mathbf{c} . Since \mathcal{F}, \mathcal{G} are disjoint, \mathbf{c} can be split into $\mathbf{c}_{\mathcal{F}} \in \mathcal{R}^{\mathcal{F}}$ and $\mathbf{c}_{\mathcal{G}} \in \mathcal{R}^{\mathcal{G}}$. Therefore

$$\begin{aligned} & \Pr[h(\mathbf{a}^\top) = 0 \bmod q \mid \mathbf{a} \leftarrow \mathcal{R}_q^k] \\ &= \Pr\left[\sum_{f \in \mathcal{F}} c_f \cdot f(\mathbf{a}^\top) + \sum_{g \in \mathcal{G}} c_g \cdot g(\mathbf{a}^\top) = 0 \bmod q \mid \mathbf{a} \leftarrow \mathcal{R}_q^k\right] \\ &= \Pr\left[\sum_{f \in \mathcal{F}} c_f \cdot f(\mathbf{a}^\top) = -\sum_{g \in \mathcal{G}} c_g \cdot g(\mathbf{a}^\top) \bmod q \mid \mathbf{a} \leftarrow \mathcal{R}_q^k\right] \\ &\leq \max_{d \in \mathcal{R}_q} \Pr\left[\sum_{f \in \mathcal{F}} c_f \cdot f(\mathbf{a}^\top) = d \bmod q \mid \mathbf{a} \leftarrow \mathcal{R}_q^k\right] \leq \epsilon_{\mathcal{F}}. \end{aligned}$$

On the other hand, if $(c_f)_{f \in \mathcal{F}}$ in the above are all-zero, then clearly the probability is bounded by $\epsilon_{\mathcal{G}}$. Thus, we find $\epsilon_{\mathcal{F} \cup \mathcal{G}} \leq \max(\epsilon_{\mathcal{F}}, \epsilon_{\mathcal{G}})$. \square

A.3 Proof of Theorem 7

Proof. The detailed descriptions of both concerned experiments are provided in Figure 6. For ease of comparison we incorporate syntactical changes to the sEUF-RMA experiment of Σ_{vSIS} , in particular introducing OSign for sampling Q random messages and generating their signatures. It is straightforward to verify that the stated experiment is functionally equivalent to that in Definition 4 and Figure 3.

We define the following hybrids:

Hyb₀: The sEUF-RMA experiment of Σ_{vSIS} in Figure 6.

Hyb₁: Same as Hyb₀, except that we sample $\mathbf{s} \leftarrow \mathcal{D}_{A_q^\top(\mathbf{A}),s}$ inefficiently from the discrete Gaussian distribution.

Hyb₂: Same as Hyb₁, except that we sample $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$ uniformly randomly.

sEUF-RMA $_{\Sigma_{\text{vSIS}}, \mathcal{A}, Q}(1^\lambda)$	ExpGenISISf $_{\mathcal{B}}(1^\lambda)$
$\mathbf{B} \leftarrow \mathcal{R}_q^{n \times \ell}$	$S := \emptyset; \kappa \leftarrow \mathcal{K}$
$(\mathbf{A}, \text{td}_{\mathbf{A}}) \leftarrow \text{TrapGen}(\mathcal{R}, 1^n, 1^m, q)$	$\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$
$(\text{sig}_1, \dots, \text{sig}_Q) \leftarrow \text{OSign}()$	
$(\mu^*, (\chi^*, \mathbf{s}^*)) \leftarrow \mathcal{A}(\mathbf{B}, \mathbf{A}, \text{sig}_1, \dots, \text{sig}_Q)$	$(\mathbf{s}^*, \mu^*, \chi^*) \leftarrow \mathcal{B}^{\text{OPre}}(\kappa, \mathbf{A})$
$\mathbf{t}^* := h_{\mu^*, \chi^*}(\mathbf{B}) \bmod q$	$\mathbf{t}^* := f(\kappa, \mu^*, \chi^*)$
$b_0 := (\mathbf{A} \cdot \mathbf{s}^* = \mathbf{t}^* \bmod q)$	$b_0 := (\mathbf{A} \cdot \mathbf{s}^* = \mathbf{t}^* \bmod q)$
$b_1 := (0 < \ \mathbf{s}^*\ \leq \beta)$	$b_1 := (0 < \ \mathbf{s}^*\ \leq \beta)$
$b_2 := (\mu^* \in \mathcal{M}_{\mathcal{X}, g, \mathbf{B}}) \wedge (\chi^* \in \mathcal{X})$	
$b_3 := ((\text{msg}^*, \text{sig}^*) \notin \{(\mu_i, \text{sig}_i)_{i=1}^Q\})$	$b_2 := ((\mathbf{s}^*, \mu^*, \chi^*) \notin S)$
return $b_0 \wedge b_1 \wedge b_2 \wedge b_3$	return $b_0 \wedge b_1 \wedge b_2$
<hr/>	<hr/>
OSign ()	OPre ()
for $i \in [Q]$:	assert $ S < Q$
$\mu_i \leftarrow \mathcal{M}; \chi_i \leftarrow \mathcal{X}$	$\mu \leftarrow \mathcal{M}; \chi \leftarrow \mathcal{X}$
$\mathbf{t}_i := h_{\mu_i, \chi_i}(\mathbf{B}) \bmod q$	$\mathbf{t} = f(\kappa, \mu, \chi)$
$\mathbf{s}_i \leftarrow \text{SampPre}(\text{td}_{\mathbf{A}}, \mathbf{t}_i, s)$	$\mathbf{s} \leftarrow \mathcal{D}_{\Lambda_q^{\dagger}(\mathbf{A}), s}$
$\text{sig}_i := (\chi_i, \mathbf{s}_i)$	$S := S \cup \{(\mathbf{s}, \mu, \chi)\}$
return $(\text{sig}_1, \dots, \text{sig}_Q)$	return (\mathbf{s}, μ, χ)

Fig. 6: Experiments of sEUF-RMA security of Σ_{vSIS} and GenISISf.

We have that $\text{Hyb}_0, \text{Hyb}_1$ and Hyb_2 are computationally indistinguishably directly by the lattice trapdoor properties (Definition 1). At the point, we observe that Hyb_2 is almost identical to the ExpGenISISf experiment in Figure 6, since $f(\kappa, \mu, \chi) = h_{\mu, \chi}(\mathbf{B})$ for any $\mu \in \mathcal{M}, \chi \in \mathcal{X}$, and the sampling of $\kappa \leftarrow \mathcal{K}$ is identical to $\mathbf{B} \leftarrow \mathcal{R}_q^{n \times \ell}$ by definition.

Finally, we argue that there is a PPT \mathcal{B} winning the ExpGenISISf experiment with non-negligible probability if and only if there is a PPT \mathcal{A} winning Hyb_2 with non-negligible probability, hence completing the proof.

For the “if” direction: Given \mathcal{A} , we let \mathcal{B} query the OPre oracle exactly Q times to obtain Q tuples of random messages and their signatures, pass all of them together with $\kappa = \mathbf{B}$ and \mathbf{A} to \mathcal{A} , and observe that a valid forgery of \mathcal{A} in sEUF-RMA $_{\Sigma_{\text{vSIS}}, \mathcal{A}, Q}$ is immediately also a valid forgery in ExpGenISISf $_{\mathcal{B}}$.

For the “only if” direction: Given \mathcal{B} , we let \mathcal{A} proceed as follows. \mathcal{A} obtains \mathbf{B}, \mathbf{A} and Q tuples of random messages and their signatures from the sEUF-RMA experiment upfront, passes \mathbf{B}, \mathbf{A} to \mathcal{B} , and upon each OPre query made by \mathcal{B} , it answers with one of the Q message-signature tuples. By design of ExpGenISISf $_{\mathcal{B}}$, \mathcal{B} can query at most Q times so \mathcal{A} can answer all queries. Then, we observe that a valid forgery of \mathcal{B} in ExpGenISISf $_{\mathcal{B}}$ is also a valid forgery in sEUF-RMA $_{\Sigma_{\text{vSIS}}, \mathcal{A}, Q}$, since Condition b_2 in sEUF-RMA $_{\Sigma_{\text{vSIS}}, \mathcal{A}, Q}$ is satisfied only if $h_{\mu, \chi}(\mathbf{B}) \neq \perp$, i.e. if $f(\kappa, \mu, \chi) \neq \perp$. If this condition is not satisfied, then b_0 in ExpGenISISf $_{\mathcal{B}}$ is not. \square

A.4 Proof of Theorem 8

Proof. Suppose there exists a PPT solver \mathcal{A} for s- \mathcal{H} int-vSIS, we construct a PPT solver \mathcal{B} for GenISISf. On input (κ, \mathbf{A}) , where $\kappa = \mathbf{B} \leftarrow \mathcal{R}_q^{n \times \ell}$ and $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$, our reduction \mathcal{B} queries OPre Q times to obtain $(\mathbf{s}_i, \mu_i, \chi_i)$ for $i \in [Q]$, and passes $([\mathbf{A}|\mathbf{B}], \{(\mathbf{s}_i, \mu_i, \chi_i)\}_{i \in [Q]})$ to \mathcal{A} . Upon receiving (g^*, \mathbf{s}^*) from \mathcal{A} , where $g^* = h_{\mu^*, \chi^*}$ for some (μ^*, χ^*) , \mathcal{B} returns $(\mathbf{s}^*, \mu^*, \chi^*)$.

It is clear that \mathcal{B} runs in PPT, and simulates a s- \mathcal{H} int-vSIS instance for \mathcal{A} faithfully. Denote $\mathcal{Q} := \{h_{\mu_i, \chi_i} : i \in [Q]\}$. By assumption, with non-negligible probability, the output of \mathcal{A} satisfies $\mathbf{A} \cdot \mathbf{s}^* = h_{\mu^*, \chi^*}(\mathbf{B}) \bmod q$, $0 < \|\mathbf{s}^*\| \leq \beta$, $g^* \in \mathcal{G} \setminus \mathcal{Q}$, and $\mathcal{P}(\mathcal{F} \cup \mathcal{Q} \cup (\{g^*\} \setminus \{0\})) = 1$. The condition $g^* \in \mathcal{G} \setminus \mathcal{Q}$

translates to $(\mathbf{s}^*, \mu^*, \chi^*) \notin \{(\mathbf{s}_1, \mu_1, \chi_1), \dots, (\mathbf{s}_Q, \mu_Q, \chi_Q)\}$. Our reduction \mathcal{B} therefore solves the given GenSIS_f instance with non-negligible probability. \square