

Lattice-Based Group Signatures with Logarithmic Signature Size

Fabien Laguillaumie^{1,3}, Adeline Langlois^{2,3}, Benoît Libert⁴, and Damien Stehlé^{2,3}

¹ Université Claude Bernard Lyon 1

² École Normale Supérieure de Lyon

³ LIP (U. Lyon, CNRS, ENS de Lyon, INRIA, UCBL),
46 Allée d'Italie, 69364 Lyon Cedex 07, France.

{fabien.laguillaumie,adeline.langlois,damien.stehle}@ens-lyon.fr

⁴ Technicolor

975 Avenue des Champs Blancs, 35510 Cesson-Sévigné, France

benoit.libert@technicolor.com

Abstract. Group signatures are cryptographic primitives where users can anonymously sign messages in the name of a population they belong to. Gordon *et al.* (Asiacrypt 2010) suggested the first realization of group signatures based on lattice assumptions in the random oracle model. A significant drawback of their scheme is its linear signature size in the cardinality N of the group. A recent extension proposed by Camenisch *et al.* (SCN 2012) suffers from the same overhead. In this paper, we describe the first lattice-based group signature schemes where the signature and public key sizes are essentially logarithmic in N (for any fixed security level). Our basic construction only satisfies a relaxed definition of anonymity (just like the Gordon *et al.* system) but readily extends into a fully anonymous group signature (*i.e.*, that resists adversaries equipped with a signature opening oracle). We prove the security of our schemes in the random oracle model under the SIS and LWE assumptions.

Keywords. Lattice-based cryptography, group signatures, anonymity, provable security.

1 Introduction

Group signatures are a core cryptographic primitive that paradoxically combines the properties of authenticity and anonymity. They are useful in many real-life applications including trusted computing platforms, auction protocols or privacy-protecting mechanisms for users in public transportation.

Parties involved in such a system are a special entity, called the group manager, and group members. The manager holds a master secret key, generates a system-wide public key, and administers the group members, by providing to each of them an individual secret key that will allow them to anonymously sign on behalf of the group. In case of dispute, the manager (or a separate authority) is able to determine the identity of a signer via an opening operation. This fundamental primitive has been extensively studied, from both theoretical and practical perspectives: It has been enriched with many useful properties, and it has been implemented in the contexts of trusted computing (using privacy-preserving attestation [14]) and of traffic management (e.g., the Vehicle Safety Communications project of the U.S. Dept. of Transportation [29]).

Group signatures were originally proposed by Chaum and van Heyst [20] and made scalable by Ateniese *et al.* in [3]. Proper security models were introduced in [5] and [7, 30] (for dynamic groups), whereas more intricate and redundant properties were considered hitherto. The model of Bellare *et al.* [5] requires two main security properties called *full anonymity* and *full traceability*. The former notion means that signatures do not leak the identities of their originators, whereas the latter implies that no collusion of malicious users can produce a valid signature that cannot be traced to one of them. Bellare *et al.* [5] proved that trapdoor permutations suffice to design group signatures, but their theoretical construction was mostly a proof of concept. Nevertheless, their methodology has been adapted in practical constructions: Essentially, a group member signs a message by verifiably encrypting a valid membership certificate delivered by the authority, while hiding its identity. While numerous schemes (e.g., [3, 15, 17, 9]) rely on the random oracle model (ROM), others are proved secure in the standard model (e.g., [5, 7, 11, 12, 24, 25]). Except theoretical constructions [5, 7], all of these rely on the Groth-Sahai methodology to design non-interactive proof systems for specific languages involving elements in bilinear groups [27]. This powerful tool led to the design of elegant compact group signatures [12, 24, 25] whose security relies on pairing-related assumptions. The resulting signatures typically consist in a constant number of elements of a group admitting a secure and efficient bilinear map.

LATTICES AND GROUP SIGNATURES. Lattices are emerging as a promising alternative to traditional number-theoretic tools like bilinear maps. They lead to asymptotically faster solutions, thanks to the algorithmic simplicity of the involved operations and to the high cost of the best known attacks. Moreover, lattice-based schemes often enjoy strong security guarantees, thanks to worst-case/average-case connections between lattice problems, and to the conjectured resistance to quantum computers.

While numerous works have been (successfully) harnessing the power of lattices for constructing digital signatures (see [35, 22, 19, 32, 10, 33] and references therein), only two works addressed the problem of efficiently realizing lattice-based group signatures. The main difficulty to overcome is arguably the scarcity of efficient and expressive non-interactive proof systems for statements involving lattices, in particular for statements on the witnesses of the hard average-case lattice problems. This state of affairs contrasts with the situation in bilinear groups, where powerful non-interactive proof systems are available [26, 27].

In 2010, Gordon *et al.* [23] described the first group signature based on lattice assumptions using the Gentry *et al.* signature scheme [22] as membership certificate, an adaptation of Regev’s encryption scheme [42] to encrypt it, and a zero-knowledge proof technique due to Micciancio and Vadhan [38]. While elegant in its design principle, their scheme suffers from signatures and public keys of sizes linear in the number of group members, making it utterly inefficient in comparison with constructions based on bilinear maps [9] or the strong RSA assumption [3]. Quite recently, Camenisch *et al.* [18] proposed anonymous attribute token systems, which can be seen as generalizations of group signatures. One of their schemes improves upon [23] in that the group public key has constant size⁵ and the anonymity property is achieved in a stronger model where the adversary is granted access to a signature opening oracle. Unfortunately, all the constructions of [18] inherit the linear signature size of the Gordon *et al.* construction. Thus far, it remained an open problem to break the linear-size barrier. This is an important challenge considering that, as advocated by Bellare *et al.* [5], one should expect practical group signatures not to entail more than poly-logarithmic complexities in the group sizes.

OUR CONTRIBUTIONS. We describe the first lattice-based group signatures featuring sub-linear signature sizes. If t and N denote the security parameter and the maximal group size, the public keys and signatures are $\tilde{O}(t^2 \cdot \log N)$ bit long. Notice that no group signature scheme can provide signatures containing $o(\log N)$ bits (such signatures would be impossible to open), so that the main improvement potential lies in the $\tilde{O}(t^2)$ factor. These first asymptotically efficient (in t and $\log N$) lattice-based group signatures are a first step towards a practical alternative to the pairing-based counterparts. The security proofs hold in the ROM (as for [23, 18]), under the Learning With Error (LWE) and Short Integer Solution (SIS) assumptions.

While our basic system only provides anonymity in a relaxed model (like [23]) where the adversary has no signature opening oracle, we show how to upgrade it into a fully anonymous group signature, in the anonymity model of Bellare *et al.* [5]. This is achieved at a minimal cost in that the signature length is only increased by a constant factor. In contrast, Camenisch *et al.* [18, Se. 5.2] achieve full anonymity at the expense of inflating their basic signatures by a factor proportional to the security parameter.

CONSTRUCTION OVERVIEW. Our basic construction is inspired by the general paradigm from [5] consisting in *encrypting* a membership certificate under the authority’s public key while providing a *non-interactive proof* that the ciphertext encrypts a valid *certificate* belonging to some group member. Nevertheless, our scheme differs from this paradigm in the sense that it is not the certificate itself which is encrypted. Instead, a temporary certificate, produced at each signature generation, is derived from the initial one and encrypted, with a proof of its validity.

We also depart from the approach of [23] at the very core of the design, *i.e.*, when it comes to provide evidence that the encrypted certificate corresponds to a legitimate group member. Specifically, Gordon *et al.* [23] hide their certificate, which is a GPV signature [22, Se. 6], within a set of $N - 1$ (encrypted) GPV pseudo-signatures that satisfy the same verification equation without being short vectors. Here, to avoid the $\mathcal{O}(N)$ factor in the signature size, we take a different approach which is reminiscent of the Boyen-Waters group signature [11]. Each group member is assigned a unique ℓ -bit identifier $\text{id} = \text{id}[1] \dots \text{id}[\ell] \in \{0, 1\}^\ell$, where $\ell = \lceil \log_2 N \rceil$. Its certificate is an extension of a Boyen signature [10] consisting in a *full* short basis of a certain lattice (instead of a single vector), which allows the signer to generate *temporary certificates*

⁵ This can also be achieved with [23] by replacing the public key by a hash thereof, and appending the key to the signature.

composed of a pair $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{Z}^m$ of discrete Gaussian vectors such that

$$\mathbf{x}_1^T \cdot \mathbf{A} + \mathbf{x}_2^T \cdot (\mathbf{A}_0 + \sum_{1 \leq i \leq \ell} \text{id}[i] \cdot \mathbf{A}_i) = \mathbf{0} \pmod{q}. \quad (1)$$

Here, q is a small bit length integer and $\mathbf{A}, \mathbf{A}_0, \dots, \mathbf{A}_\ell \in \mathbb{Z}_q^{m \times n}$ are part of the group public key. Our choice of Boyen’s signature [10] as membership certificate is justified by it being the most efficient known lattice-based signature proven secure in the standard model, and enjoying a simple verification procedure corresponding to a relation for which we can easily design a proof of knowledge. A signature proven secure in the standard model allows us to obtain an easy-to-prove relation that does not involve a random oracle. Indeed, as noted for example in [3, 16, 17] signature schemes without random oracles make it easier to prove knowledge of a valid message-signature pair in the design of privacy-preserving protocols.

We then encrypt $\mathbf{x}_2 \in \mathbb{Z}^m$ as in [23], using a variant of the dual-Regev encryption scheme [22, Se. 7]: the resulting ciphertext is $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s} + \mathbf{x}_2$, where $\mathbf{B}_0 \in \mathbb{Z}_q^{m \times n}$ is a public matrix and \mathbf{s} is uniform in \mathbb{Z}_q^n . We prove that \mathbf{c}_0 hides a short vector using the Micciancio-Vadhan proof system [38], as in [23]. Then, for each $i \in [1, \ell]$, we also compute a proper dual-Regev encryption \mathbf{c}_i of $\text{id}[i] \cdot \mathbf{x}_2$ and generate a non-interactive OR proof that \mathbf{c}_i encrypts either the same vector as \mathbf{c}_0 or the $\mathbf{0}$ vector.

It remains to prove that the encrypted vectors \mathbf{x}_2 are part of a signature satisfying (1) without giving away the $\text{id}[i]$ ’s. To this end, we choose the signing matrices \mathbf{A}_i orthogonally to the encrypting matrices \mathbf{B}_i , as suggested in [23]. Contrarily to the case of [23], the latter technique does not by itself suffice to guarantee the well-formedness of the \mathbf{c}_i ’s. Indeed, we also need to prove properties about the noise vectors used in the dual-Regev ciphertexts $\{\mathbf{c}_i\}_{1 \leq i \leq \ell}$. This is achieved using a modification of Lyubashevsky’s protocol [31, 33] to prove knowledge of a solution to the Inhomogeneous Short Integer Solution problem (ISIS). This modification leads to a Σ -protocol which is zero-knowledge when the transcript is conditioned on the protocol not aborting. As the challenge space of this Σ -protocol is binary, we lowered the abort probability so that we can efficiently apply the Fiat-Shamir heuristic to a parallel repetition of the basic protocol. In the traceability proof, the existence of a witness extractor will guarantee that a successful forger will either yield a forgery for Boyen’s signature or a short non-zero vector in the kernel of one of the matrices $\{\mathbf{A}_i\}_{1 \leq i \leq \ell}$. In either case, the forger allows the simulator to solve a SIS instance.

In the fully anonymous variant of our scheme, the difficulty is to find a way to open adversarially-chosen signatures. This is achieved by implicitly using a “chosen-ciphertext-secure” variant of the signature encryption technique of Gordon *et al.* [23]. While Camenisch *et al.* [18] proceed in a similar way using Peikert’s technique [39], we use a much more economical method borrowed from the Agrawal *et al.* [1] identity-based cryptosystem. In our basic system, each \mathbf{c}_i is of the form $\mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}_i + \text{id}[i] \cdot \mathbf{x}_2$, where p is an upper bound on \mathbf{x}_2 ’s coordinates, and can be decrypted using a short basis \mathbf{S}_i such that $\mathbf{S}_i \cdot \mathbf{B}_i = \mathbf{0} \pmod{q}$. Our fully anonymous system replaces each \mathbf{B}_i by a matrix $\mathbf{B}_{i, \text{VK}}$ that depends on the verification key VK of a one-time signature. In the proof of full anonymity, the reduction will be able to compute a trapdoor for all matrices $\mathbf{B}_{i, \text{VK}}$, except for one specific verification key VK^* that will be used in the challenge phase. This will provide the reduction with a backdoor allowing it to open all adversarially-generated signatures.

OPEN PROBLEMS. The schemes we proposed should be viewed as proofs of concept, since instantiating them with practical parameters would most likely lead to large keys and signature sizes. It is an interesting task to replace the SIS and LWE problems by their ring variants [34, 40, 36], to attempt to save linear factors in the security parameter t . The main hurdle in that direction seems to be the design of appropriate zero-knowledge proofs of knowledge for the LWE and ISIS relations (see Section 2.2).

As opposed to many pairing-based constructions, the security of our scheme is only proven in the random oracle model: We rely on the Fiat-Shamir heuristic to remove the interaction in the interactive proof systems. This is because very few lattice problems are known to belong to NIZK. The problems considered in the sole work on this topic [41] seem ill-fitted to devise group signatures. As a consequence, the security proofs of all known lattice-based group signatures are conducted in the random oracle model. Recently suggested multi-linear maps [21] seem like a possible direction towards solving this problem. However, the only known instantiation so far relies on stronger assumptions than LWE or SIS.

2 Background and Definitions

We first recall standard notations. All vectors will be denoted in bold lower-case letters, whereas bold upper-case letters will be used for matrices. If \mathbf{b} and \mathbf{c} are two vectors of compatible dimensions and base rings, then their inner product will be denoted by $\langle \mathbf{b}, \mathbf{c} \rangle$. Further, if $\mathbf{b} \in \mathbb{R}^n$, its euclidean norm will be denoted by $\|\mathbf{b}\|$. This notation is extended to any matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$ with columns $(\mathbf{b}_i)_{i \leq n}$ by $\|\mathbf{B}\| = \max_{i \leq n} \|\mathbf{b}_i\|$. If \mathbf{B} is full column-rank, we let $\widetilde{\mathbf{B}}$ denote the Gram-Schmidt orthogonalisation of \mathbf{B} .

If D_1 and D_2 are two distributions over the same countable support S , then their statistical distance is defined as $\Delta(D_1, D_2) = \frac{1}{2} \sum_{x \in S} |D_1(x) - D_2(x)|$. A function $f(n)$ is said negligible if $f(n) = n^{-\omega(1)}$. Finally, the acronym PPT stands for probabilistic polynomial-time.

2.1 Lattices

A (full-rank) lattice L is the set of all integer linear combinations of some linearly independent basis vectors $(\mathbf{b}_i)_{i \leq n}$ belonging to some \mathbb{R}^n . For a lattice L and a real $\sigma > 0$, we define the Gaussian distribution of support L and standard deviation σ by $\forall \mathbf{b} \in L : D_{L, \sigma}[\mathbf{b}] \sim \exp(-\pi \|\mathbf{b}\|^2 / \sigma^2)$. We will extensively use the fact that samples from $D_{L, \sigma}$ are short with overwhelming probability.

Lemma 1 ([4, Le. 1.5]). *For any lattice $L \subseteq \mathbb{R}^n$ and $\sigma > 0$, we have $\Pr_{\mathbf{b} \leftarrow D_{L, \sigma}}[\|\mathbf{b}\| \leq \sqrt{n}\sigma] \geq 1 - 2^{-\Omega(n)}$.*

Further, as shown by Gentry *et al.* [22], Gaussian distributions with lattice support can be sampled from efficiently, given a sufficiently short basis of the lattice.

Lemma 2 ([13, Le. 2.3]). *There exists a PPT algorithm `GPVSample` that takes as inputs a basis \mathbf{B} of a lattice $L \subseteq \mathbb{Z}^n$ and a rational $\sigma \geq \|\widetilde{\mathbf{B}}\| \cdot \Omega(\sqrt{\log n})$, and outputs vectors $\mathbf{b} \in L$ with distribution $D_{L, \sigma}$.*

Cash *et al.* [19] showed how to use `GPVSample` to randomize the basis of a given lattice. The following statement is obtained by using [13, Le. 2.3] in the proof of [19].

Lemma 3 (Adapted from [19, Le. 3.3]). *There exists a PPT algorithm `RandBasis` that takes as inputs a basis \mathbf{B} of a lattice $L \subseteq \mathbb{Z}^n$ and a rational $\sigma \geq \|\widetilde{\mathbf{B}}\| \cdot \Omega(\sqrt{\log n})$, and outputs a basis \mathbf{C} of L satisfying $\|\widetilde{\mathbf{C}}\| \leq \sqrt{n}\sigma$ with probability $\geq 1 - 2^{-\Omega(n)}$. Further, the distribution of \mathbf{C} is independent of the input basis \mathbf{B} .*

Let $m \geq n \geq 1$ and $q \geq 2$. For a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, we define the lattice $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \cdot \mathbf{A} = \mathbf{0} \pmod{q}\}$. We will use an algorithm that jointly samples a uniform \mathbf{A} and a short basis of $\Lambda_q^\perp(\mathbf{A})$.

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

Lemma 4 is often combined with the sampler from Lemma 2. Micciancio and Peikert [37] recently proposed a more efficient approach for this combined task, that should be preferred in practice but, for the sake of simplicity, we present our schemes using `TrapGen`.

Lemma 4 was later extended by Gordon *et al.* [23] so that the columns of \mathbf{A} lie within a prescribed linear vector subspace of \mathbb{Z}_q^n (for q prime). For the security proof of our fully anonymous scheme, we will use an extension where the columns of the sampled \mathbf{A} lie within a prescribed *affine* subspace of \mathbb{Z}_q^n . A proof is given in Appendix C.

Lemma 5. *There exists a PPT algorithm `SuperSamp` that takes as inputs integers $m \geq n \geq 1$ and $q \geq 2$ prime such that $m \geq \Omega(n \log q)$, as well as matrices $\mathbf{B} \in \mathbb{Z}_q^{m \times n}$ and $\mathbf{C} \in \mathbb{Z}_q^{n \times n}$ such that the rows of \mathbf{B} span \mathbb{Z}_q^n . It outputs $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and a basis $\mathbf{T}_\mathbf{A}$ of $\Lambda_q^\perp(\mathbf{A})$ such that \mathbf{A} is within statistical distance $2^{-\Omega(n)}$ to $U(\mathbb{Z}_q^{m \times n})$ conditioned on $\mathbf{B}^T \cdot \mathbf{A} = \mathbf{C}$, and $\|\widetilde{\mathbf{T}_\mathbf{A}}\| \leq \mathcal{O}(\sqrt{mn \log q \log m})$.*

Finally, we also make use of an algorithm that extends a trapdoor for $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ to a trapdoor of any $\mathbf{B} \in \mathbb{Z}_q^{m' \times n}$ whose top $m \times n$ submatrix is \mathbf{A} .

Lemma 6 ([19, Le. 3.2]). *There exists a PPT algorithm ExtBasis that takes as inputs a matrix $\mathbf{B} \in \mathbb{Z}_q^{m' \times n}$ whose first m rows span \mathbb{Z}_q^n , and a basis $\mathbf{T}_\mathbf{A}$ of $\Lambda_q^\perp(\mathbf{A})$ where \mathbf{A} is the top $m \times n$ submatrix of \mathbf{B} , and outputs a basis $\mathbf{T}_\mathbf{B}$ of $\Lambda_q^\perp(\mathbf{B})$ with $\|\widetilde{\mathbf{T}}_\mathbf{B}\| \leq \|\widetilde{\mathbf{T}}_\mathbf{A}\|$.*

For the sake of simplicity, we will assume that when the parameter conditions are satisfied, the distributions of the outputs of TrapGen and SuperSamp are *exactly* those they are meant to approximate, and the probabilistic norm bounds of Lemmas 1 and 3 *always* hold.

2.2 Computational problems

The security of our schemes provably relies (in the random oracle model) on the assumption that both algorithmic problems below are hard, *i.e.*, cannot be solved in polynomial time with non-negligible probability and non-negligible advantage, respectively.

Definition 1. *Let m, q, β be functions of a parameter n . The Short Integer Solution problem $\text{SIS}_{m,q,\beta}$ is as follows: Given $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{m \times n})$, find $\mathbf{x} \in \Lambda_q^\perp(\mathbf{A})$ with $0 < \|\mathbf{x}\| \leq \beta$.*

Definition 2. *Let q, α be functions of a parameter n . For $\mathbf{s} \in \mathbb{Z}_q^n$, the distribution $A_{q,\alpha,\mathbf{s}}$ over $\mathbb{Z}_q^n \times \mathbb{Z}_q$ is obtained by sampling $\mathbf{a} \leftarrow U(\mathbb{Z}_q^n)$ and (a noise) $e \leftarrow D_{\mathbb{Z},\alpha q}$, and returning $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e)$. The Learning With Errors problem $\text{LWE}_{q,\alpha}$ is as follows: For $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$, distinguish between arbitrarily many independent samples from $U(\mathbb{Z}_q^n \times \mathbb{Z}_q)$ or the same number of independent samples from $A_{q,\alpha,\mathbf{s}}$.*

If $q \geq \sqrt{n}\beta$ and $m, \beta \leq \text{poly}(n)$, then standard worst-case lattice problems with approximation factors $\gamma = \mathcal{O}(\beta\sqrt{n})$ reduce to $\text{SIS}_{m,q,\beta}$ (see, e.g., [22, Se. 9]). Similarly, if $\alpha q = \Omega(\sqrt{n})$, then standard worst-case lattice problems with approximation factors $\gamma = \mathcal{O}(\alpha/n)$ quantumly reduce to $\text{LWE}_{q,\alpha}$ (see [42], and also [39, 13] for partial dequantizations). Note that we use the discrete noise variant of LWE from [23].

We will make use of a non-interactive zero-knowledge proof of knowledge (NIZPoK) protocol, which can be rather directly derived from [31, 33], for the following relation corresponding to an inhomogenous variant of the SIS relation:

$$R_{\text{SIS}} = \{(\mathbf{A}, \mathbf{y}, \beta; \mathbf{x}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^n \times \mathbb{Q} \times \mathbb{Z}^m : \mathbf{x}^T \cdot \mathbf{A} = \mathbf{y}^T \wedge \|\mathbf{x}\| \leq \beta\}.$$

The protocol, detailed in Appendix 2.3, is derived from the parallel repetition of a Σ protocol with binary challenges. We call $\text{Prove}_{\text{SIS}}$ and $\text{Verify}_{\text{SIS}}$ the PPT algorithms run by the Prover and the Verifier when the scheme is rendered non-interactive using the Fiat-Shamir heuristic (*i.e.*, the challenge is implemented using the random oracle $H(\cdot)$). Algorithm $\text{Prove}_{\text{SIS}}$ takes $(\mathbf{A}, \mathbf{y}, \beta; \mathbf{x})$ as inputs, and generates a transcript $(\text{Comm}, \text{Chall}, \text{Resp})$. Algorithm $\text{Verify}_{\text{SIS}}$ takes $(\mathbf{A}, \mathbf{y}, \beta)$ and such a transcript as inputs, and returns 0 or 1. The scheme has completeness error $2^{-\Omega(n)}$: if $\text{Prove}_{\text{SIS}}$ is given as input an element of R_{SIS} , then given as input the output of $\text{Prove}_{\text{SIS}}$, $\text{Verify}_{\text{SIS}}$ replies 1 with probability $\geq 1 - 2^{-\Omega(m)}$ (over the randomness of Prove). Also, there exists a PPT algorithm $\text{Simulate}_{\text{SIS}}$ that, by reprogramming the random oracle $H(\cdot)$, takes $(\mathbf{A}, \mathbf{y}, \beta)$ as input and generates a transcript $(\text{Comm}, \text{Chall}, \text{Resp})$ whose distribution is within statistical distance $2^{-\Omega(m)}$ of the genuine transcript distribution. Finally, there also exists a PPT algorithm $\text{Extract}_{\text{SIS}}$ that given access to a time T algorithm \mathcal{A} that generates transcripts accepted by $\text{Verify}_{\text{SIS}}$ with probability ε , produces, in time $\text{Poly}(T, 1/\varepsilon)$ a vector \mathbf{x}' such that $(\mathbf{A}, \mathbf{y}, \mathcal{O}(\beta \cdot m^2); \mathbf{x}') \in R_{\text{SIS}}$.

We will also need a NIZKPoK protocol for the following language:

$$R_{\text{LWE}} = \{(\mathbf{A}, \mathbf{b}, \alpha; \mathbf{s}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m \times \mathbb{Q} \times \mathbb{Z}_q^n : \|\mathbf{b} - \mathbf{A} \cdot \mathbf{s}\| \leq \alpha q \sqrt{m}\}.$$

As noted in [33], we may multiply \mathbf{b} by a parity check matrix $\mathbf{G} \in \mathbb{Z}_q^{(m-n) \times m}$ of \mathbf{A} and prove the existence of small $\mathbf{e} \in \mathbb{Z}^m$ such that $\mathbf{e}^T \cdot \mathbf{G}^T = \mathbf{b}^T \cdot \mathbf{G}^T$. This may be done with the above NIZKPoK protocol for R_{SIS} . We call $\text{Prove}_{\text{LWE}}$, $\text{Verify}_{\text{LWE}}$, $\text{Simulate}_{\text{LWE}}$ and $\text{Extract}_{\text{LWE}}$ the obtained PPT algorithms.

2.3 Proof of Knowledge of an ISIS Solution

In [31], Lyubashevsky described an identification scheme whose security relies on the hardness of the SIS problem. Given a public vector $\mathbf{y} \in \mathbb{Z}_q^n$ and a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, the prover holds a short secret \mathbf{x} and generates an interactive witness indistinguishable proof of knowledge of a short vector $\mathbf{x}^T \in \mathbb{Z}^m$ such that $\mathbf{x}^T \cdot \mathbf{A} = \mathbf{y}^T \pmod q$. A variant was later proposed in [33], which enjoys the property of being zero-knowledge (when the distribution of the transcript is conditioned on the prover not aborting). We present an adaptation of [33, Fig. 1] (still enjoying the same zero-knowledgedness property): the secret is a single vector, the challenges are binary (which we use for the extraction vector), and we increase the standard deviation of the committed vector to lower the rejection probability (we use a parallel repetition of the basic scheme, and want the probability that there is a reject among all the parallel iterations to be sufficiently away from 1).

Assume the prover P wishes to prove knowledge of an \mathbf{x} such that $\mathbf{y}^T = \mathbf{x}^T \cdot \mathbf{A} \pmod q$ and $\|\mathbf{x}\| \leq \beta$, where \mathbf{y} and \mathbf{A} are public. The protocol takes place between the prover P and the verifier V and proceeds by the parallel repetition of a basic Σ protocol with binary challenges. We set $\sigma = \Theta(\beta m^{3/2})$ and M_L as specified by [33, Th. 4.6]. Thanks to our larger value of σ , we obtain (by adapting [33, Le. 4.5]) have that M_L is now $1 - \Omega(1/m)$.

1. The prover P generates a commitment $\text{Comm} = (\mathbf{w}_i)_{i \leq t}$ where, for each $i \leq t$, $\mathbf{w}_i \in \mathbb{Z}_q^n$ is obtained by sampling $\mathbf{y}_i \leftarrow D_{\mathbb{Z}^m, \sigma}$ and computing $\mathbf{w}_i^T = \mathbf{y}_i^T \cdot \mathbf{A} \pmod q$. The message Comm is sent to V .
2. The verifier V sends a challenge $\text{Chall} \leftarrow \{0, 1\}^t$ to P .
3. For $i \leq t$, the prover P does the following.
 - a. Compute $\mathbf{z}_i = \mathbf{y}_i + \text{Chall}[i] \cdot \mathbf{x}$, where $\text{Chall}[i]$ denotes the i th bit of Chall .
 - b. Set \mathbf{z}_i to \perp with probability $\min(1, \frac{\exp(-\pi \|\mathbf{z}_i\|^2 / \sigma^2)}{M_L \cdot \exp(-\pi \|\text{Chall}[i] \cdot \mathbf{x} - \mathbf{z}_i\|^2 / \sigma^2)})$.

Then P sends the response $\text{Resp} = (\mathbf{z}_i)_{i \leq t}$ to V .

4. The verifier V checks the transcript $(\text{Comm}, \text{Chall}, \text{Resp})$ as follows:
 1. For $i \leq t$, set $d_i = 1$ if $\|\mathbf{z}_i\| \leq 2\sigma\sqrt{m}$ and $\mathbf{z}_i^T \cdot \mathbf{A} = \mathbf{w}_i^T + \text{Chall}[i] \cdot \mathbf{y}^T$. Otherwise, set $d_i = 0$.
 2. If $\sum_{i \leq t} d_i \geq 0.65t$, return 1 (and accept the transcript). Otherwise, return 0.

The protocol has completeness error $2^{-\Omega(t)}$. Further, by [33, Th. 4.6], the distribution of the transcript conditioned on $\mathbf{z}_i \neq \perp$ can be simulated efficiently. Note that if we implement the challenge phase with a random oracle, we can compute the \mathbf{z}_i 's for increasing value of i , and repeat the whole procedure if $\mathbf{z}_i = \perp$ for some i . Thanks to our choice of σ , for any $t \leq \mathcal{O}(m)$, the probability that $\mathbf{z}_i = \perp$ for some i is $\leq c$, for some constant $c < 1$. Thanks to this random oracle enabled rejection, the simulator produces a distribution that is within statistical distance $2^{-\Omega(m)}$ to the transcript distribution.

Finally, the modified protocol provides special soundness in that there is a simple extractor that takes as input two valid transcripts $(\text{Comm}, \text{Chall}, \text{Resp})$, $(\text{Comm}, \text{Chall}', \text{Resp}')$ with distinct challenges $\text{Chall} \neq \text{Chall}'$ and obtains a witness \mathbf{x}' such that $\mathbf{x}'^T \cdot \mathbf{A} = \mathbf{y}^T \pmod q$ and $\|\mathbf{x}'\| \leq \mathcal{O}(\sigma\sqrt{m}) \leq \mathcal{O}(\beta m^2)$.

2.4 Group Signatures

This section recalls the model of Bellare, Micciancio and Warinschi [5], which assumes static groups. A group signature scheme \mathcal{GS} consists of a tuple of four PPT algorithms $(\text{Keygen}, \text{Sign}, \text{Verify}, \text{Open})$ with the following specifications:

- **Keygen** takes 1^n and 1^N as inputs, where $n \in \mathbb{N}$ is the security parameter, and $N \in \mathbb{N}$ is the maximum number of group members. It returns a tuple $(\mathbf{gpk}, \mathbf{gmsk}, \mathbf{gsk})$ where \mathbf{gpk} is the *group public key*, \mathbf{gmsk} is the group manager secret key, and \mathbf{gsk} is an N -vector of secret keys: $\mathbf{gsk}[j]$ is the signing key of the j -th user, for $j \in \{0, \dots, N-1\}$.
- **Sign** takes the group public key \mathbf{gpk} , a signing key $\mathbf{gsk}[j]$ and a message $M \in \{0, 1\}^*$ as inputs. Its output is a signature $\Sigma \in \{0, 1\}^*$ on M under $\mathbf{gsk}[j]$.

- **Verify** is deterministic and takes the group public key \mathbf{gpk} , a message M and a putative signature Σ of M as inputs. It outputs either 0 (reject) or 1 (accept).
- **Open** is deterministic and takes as input the group public key \mathbf{gpk} , the group manager secret key \mathbf{gmsk} , a message M and a valid group signature Σ w.r.t. \mathbf{gpk} . It returns an index $j \in \{0, \dots, N - 1\}$ or a special symbol \perp in case of opening failure.

The group signature scheme must be *correct*, i.e., for all integers n and N , all $(\mathbf{gpk}, \mathbf{gmsk}, \mathbf{gsk})$ obtained from **Keygen** with $(1^n, 1^N)$ as input, all $j \in \{0, \dots, N - 1\}$ and $M \in \{0, 1\}^*$:

$$\text{Verify}(\mathbf{gpk}, M, \text{Sign}(\mathbf{gpk}, \mathbf{gsk}[j], M)) = 1 \text{ and } \text{Open}(\mathbf{gpk}, \mathbf{gmsk}, M, \text{Sign}(\mathbf{gpk}, \mathbf{gsk}[j], M)) = j,$$

with probability negligibly close to 1 over the internal randomness of **Keygen** and **Sign**.

Beyond the correctness requirement, the security model of Bellare *et al.* [5] considers two notions called traceability and anonymity. The former asks that no coalition of group members be able to create a signature that cannot be traced to one of them. The latter implies that, even if all group members' private keys are given to the adversary, signatures generated by two distinct members should be computationally indistinguishable. Formal definitions of these properties are provided in Appendix A.

3 An Asymptotically Shorter Lattice-Based Group Signature

At a high level, our key generation is based on the variant of Boyen's lattice signatures [10] described in [37, Se. 6.2]: Boyen's secret and verification keys respectively become our secret and public keys, whereas Boyen's message space is mapped to the users' identity space. There are however several additional twists in **Keygen**. First, each group member is given a *full* short basis of the public lattice associated to its identity, instead of a single short lattice vector. The reason is that, for anonymity and unlinkability purposes, the user has to generate each group signature using a *fresh* short lattice vector. Second, we sample our public key matrices $(\mathbf{A}_i)_{i \leq \ell}$ orthogonally to publicly known matrices \mathbf{B}_i , similarly to the group signature scheme from [23]. These \mathbf{B}_i 's will be used to publicly verify the validity of the signatures. They are sampled along with short trapdoor bases, using algorithm **SuperSamp**, which become part of the group signature secret key. These trapdoor bases will be used by the group authority to open signatures.

To anonymously sign a message M , the user samples a Boyen signature $(\mathbf{x}_1, \mathbf{x}_2)$ with its identity as message, which is a temporary certificate of its group membership. It does so using its full trapdoor matrix for the corresponding lattice. The user then encrypts \mathbf{x}_2 , in a fashion that resembles [23], using Regev's dual encryption scheme from [22, Se. 7.1] with the \mathbf{B}_i 's as encryption public keys. Note that in all cases but one (\mathbf{c}_0 at Step 2), the signature is not embedded in the encryption noise as in [23], but as proper plaintext. The rest of the signing procedure consists in proving in zero-knowledge that these are valid ciphertexts and that the underlying plaintexts indeed encode a Boyen signature under the group public key. These ZKPoKs are all based on the interactive protocols recalled in Section 2.2. These were made non-interactive via the Fiat-Shamir heuristic with random oracle $H(\cdot)$ taking values in $\{0, 1\}^t$. The message M is embedded in the very last application of the Fiat-Shamir transform at Step 6 of the signing algorithm.

The verification algorithm merely consists in verifying all proofs of knowledge concerning the Boyen signature embedded in the plaintexts of the ciphertexts.

Finally, the group manager can open any signature by decrypting the ciphertexts (using the group manager secret key) and then recovering the underlying Boyen signature within the plaintexts: this reveals which public key matrices \mathbf{A}_i have been considered by the signer, and therefore its identity.

The scheme depends on several functions m, q, p, α and σ of the security parameter n and the group size $N (= 2^\ell)$. They are set so that all algorithms can be implemented in polynomial time and are correct (Theorem 1), and so that the security properties (Theorems 2 and 3) hold, in the ROM, under the SIS and LWE hardness assumptions for parameters for which these problems enjoy reductions from standard worst-case lattice problems with polynomial approximation factors. More precisely, we require that:

- parameter m is $\Omega(n \log q)$,
- parameter σ is $\Omega(m^{3/2} \sqrt{\ell n \log q \log m})$ and $\leq n^{\mathcal{O}(1)}$,
- parameter p is $\Omega((\alpha q + \sigma)m^{5/2})$,

- parameter α is set so that $\alpha^{-1} \geq \Omega(pm^3 \log m)$ and $\leq n^{\mathcal{O}(1)}$,
- parameter q is prime and $\Omega(\ell + \alpha^{-1}\sqrt{n\ell})$ and $\leq n^{\mathcal{O}(1)}$.

For example, we may set $m = \tilde{\mathcal{O}}(n)$, $\sigma = \tilde{\mathcal{O}}(n^2\sqrt{\ell})$, $p = \tilde{\mathcal{O}}(n^{9/2}\sqrt{\ell})$, $\alpha^{-1} = \tilde{\mathcal{O}}(n^{15/2}\sqrt{\ell})$ and $q = \tilde{\mathcal{O}}(\ell + n^8\sqrt{\ell})$.

Keygen($1^n, 1^N$): Given a security parameter $n > 0$ and the desired number of group members $N = 2^\ell \in \text{poly}(n)$, choose parameters q, m, p, α and σ as specified above and make them public. Choose a hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^t$ for some $t \in [\Omega(n), n^{\mathcal{O}(1)}]$, which will be modeled as a random oracle. Then, proceed as follows.

1. Run $\text{TrapGen}(1^n, 1^m, q)$ to get $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ with a short basis $\mathbf{T}_{\mathbf{A}}$ of $\Lambda_q^\perp(\mathbf{A})$.
2. For $i = 0$ to ℓ , sample $\mathbf{A}_i \leftarrow U(\mathbb{Z}_q^{m \times n})$ and compute $(\mathbf{B}_i, \mathbf{S}'_i) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{A}_i, \mathbf{0})$. Then, randomize \mathbf{S}'_i as $\mathbf{S}_i \leftarrow \text{RandBasis}(\mathbf{S}'_i, \Omega(\sqrt{mn \log q \log m}))$.⁶
3. For $j = 0$ to $N - 1$, let $\text{id}_j = \text{id}_j[1] \dots \text{id}_j[\ell] \in \{0, 1\}^\ell$ be the binary representation of id_j and define:

$$\mathbf{A}_{\text{id}_j} = \left[\frac{\mathbf{A}}{\mathbf{A}_0 + \sum_{i=1}^{\ell} \text{id}_j[i] \mathbf{A}_i} \right] \in \mathbb{Z}_q^{2m \times n}.$$

Then, run $\mathbf{T}'_{\text{id}_j} \leftarrow \text{ExtBasis}(\mathbf{A}_{\text{id}_j}, \mathbf{T}_{\mathbf{A}})$ to get a short delegated basis $\mathbf{T}'_{\text{id}_j}$ of $\Lambda_q^\perp(\mathbf{A}_{\text{id}_j})$. Eventually, run $\mathbf{T}_{\text{id}_j} \leftarrow \text{RandBasis}(\mathbf{T}'_{\text{id}_j}, \Omega(m\sqrt{\ell n \log q \log m}))$.⁶ The j -th member's private key is $\text{gsk}[j] := \mathbf{T}_{\text{id}_j}$.

4. The group manager's private key is $\text{gmsk} := \{\mathbf{S}_i\}_{i=0}^{\ell}$ and the group public key is defined to be $\text{gpk} := (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_{i=0}^{\ell})$. The algorithm outputs $(\text{gpk}, \text{gmsk}, \{\text{gsk}[j]\}_{j=0}^{N-1})$.

Sign($\text{gpk}, \text{gsk}[j], M$): To sign a message $M \in \{0, 1\}^*$ using the private key $\text{gsk}[j] = \mathbf{T}_{\text{id}_j}$, proceed as follows.

1. Run $\text{GPVSample}(\mathbf{T}_{\text{id}_j}, \sigma)$ to obtain a vector $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T \in \Lambda_q^\perp(\mathbf{A}_{\text{id}_j})$ of norm $\leq \sigma\sqrt{2m}$.
2. Sample $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$ and encrypt $\mathbf{x}_2 \in \mathbb{Z}_q^m$ by computing $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^m$.
3. Sample $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$. For $i = 1$ to ℓ , sample $\mathbf{e}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and compute $\mathbf{c}_i = \mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}_i + \text{id}_j[i] \cdot \mathbf{x}_2$, which encrypts $\mathbf{x}_2 \in \mathbb{Z}_q^m$ (resp. $\mathbf{0}$) if $\text{id}_j[i] = 1$ (resp. $\text{id}_j[i] = 0$).
4. Generate a NIZKPoK π_0 of \mathbf{s}_0 so that $(\mathbf{B}_0, \mathbf{c}_0, \sqrt{2}\sigma/q; \mathbf{s}_0) \in R_{\text{LWE}}$ (see Section 2.2).
5. For $i = 1$ to ℓ , generate a NIZKPoK $\pi_{\text{OR},i}$ of \mathbf{s} and \mathbf{s}_0 so that either:
 - (i) $((\mathbf{B}_i | \mathbf{B}_0), p^{-1}(\mathbf{c}_i - \mathbf{c}_0), \sqrt{2}\alpha; (\mathbf{s}^T | -\mathbf{s}_0^T)^T) \in R_{\text{LWE}}$ (the vectors \mathbf{c}_i and \mathbf{c}_0 encrypt the same \mathbf{x}_2 , so that $p^{-1}(\mathbf{c}_i - \mathbf{c}_0)$ is close to the \mathbb{Z}_q -span of $(\mathbf{B}_i | \mathbf{B}_0)$);
 - (ii) or $(\mathbf{B}_i, p^{-1}\mathbf{c}_i, \alpha; \mathbf{s}) \in R_{\text{LWE}}$ (the vector \mathbf{c}_i encrypts $\mathbf{0}$, so that $p^{-1}\mathbf{c}_i$ is close to the \mathbb{Z}_q -span of \mathbf{B}_i). This can be achieved by OR-ing two proofs for R_{LWE} , and making the resulting protocol non-interactive with the Fiat-Shamir heuristic.⁷
6. For $i = 1$ to ℓ , set $\mathbf{y}_i = \text{id}_j[i] \mathbf{x}_2 \in \mathbb{Z}^m$ and generate a NIZKPoK π_K of $(\mathbf{e}_i)_{1 \leq i \leq \ell}, (\mathbf{y}_i)_{1 \leq i \leq \ell}, \mathbf{x}_1$ s.t.:

$$\mathbf{x}_1^T \mathbf{A} + \sum_{i=0}^{\ell} \mathbf{c}_i^T \mathbf{A}_i = \sum_{i=1}^{\ell} \mathbf{e}_i^T (p \mathbf{A}_i) \quad \text{and} \quad \mathbf{e}_i^T (p \mathbf{A}_i) + \mathbf{y}_i^T \mathbf{A}_i = \mathbf{c}_i^T \mathbf{A}_i \quad \text{for } i \in [1, \ell] \quad (2)$$

with $\|\mathbf{e}_i\|, \|\mathbf{y}_i\|, \|\mathbf{x}_1\| \leq \max(\sigma, \alpha q)\sqrt{m}$ for all i . This is achieved using $\text{Prove}_{\text{ISIS}}$, producing a triple $(\text{Comm}_K, \text{Chall}_K, \text{Resp}_K)$, where $\text{Chall}_K = H(M, \text{Comm}_K, (\mathbf{c}_i)_{0 \leq i \leq \ell}, \pi_0, (\pi_{\text{OR},i})_{1 \leq i \leq \ell})$.

The signature consists of

$$\Sigma = ((\mathbf{c}_i)_{0 \leq i \leq \ell}, \pi_0, (\pi_{\text{OR},i})_{1 \leq i \leq \ell}, \pi_K). \quad (3)$$

Verify(gpk, M, Σ): Parse Σ as in (3). Then, return 1 if $\pi_0, (\pi_{\text{OR},i})_{1 \leq i \leq \ell}, \pi_K$ properly verify. Else, return 0.

Open($\text{gpk}, \text{gmsk}, M, \Sigma$): Parse gmsk as $\{\mathbf{S}_i\}_{i=0}^{\ell}$ and Σ as in (3). Compute \mathbf{x}_2 by decrypting \mathbf{c}_0 using \mathbf{S}_0 . For $i = 1$ to ℓ , use \mathbf{S}_i to determine which one of the vectors $p^{-1}\mathbf{c}_i$ and $p^{-1}(\mathbf{c}_i - \mathbf{x}_2)$ is close to the \mathbb{Z}_q -span of \mathbf{B}_i . Set $\text{id}[i] = 0$ in the former case and $\text{id}[i] = 1$ in the latter. Eventually, output $\text{id} = \text{id}[1] \dots \text{id}[\ell]$.

All steps of the scheme above can be implemented in polynomial-time as a function of the security parameter n , assuming that $q \geq 2$ prime, $m \geq \Omega(n \log q)$, $\sigma \geq \Omega(m^{3/2}\sqrt{\ell n \log q \log m})$ (using Lemmas 2 and 3), $\alpha q \geq \Omega(1)$ (using Lemma 2 in dimension 1). Under some mild conditions on the parameters, the scheme above is correct, *i.e.*, the verifier accepts honestly generated signatures, and the group manager successfully opens honestly generated signatures. In particular, multiplying the ciphertexts by the \mathbf{S}_i modulo q should reveal $p \cdot \mathbf{e}_i + \text{id}_j[i] \cdot \mathbf{x}_2$ over the integers, and $\text{id}_j[i] \cdot \mathbf{x}_2$ should always be smaller than p .

⁶ These randomisation steps are not needed for the correctness of the scheme but are important in the traceability proof.

⁷ Recall that HVZK is stable under disjunctions: The OR of two HVZK proofs enjoying special soundness is itself HVZK.

Theorem 1. Assume that $q \geq 2$ is prime and that $m \geq \Omega(n \log q)$, $\sigma \geq \Omega(m^{3/2} \sqrt{\ell n \log q} \log m)$, $\alpha^{-1} \geq \Omega(pm^{5/2} \log m \sqrt{n \log q})$ and $q \geq \Omega(\alpha^{-1} + \sigma m^{5/2} \log m \sqrt{n \log q})$. Then the group signature scheme above can be implemented in time polynomial in n , is correct, and the bit-size of the generated signatures is $\mathcal{O}(\ell t m \log q)$. (The proof is given in Appendix B).

4 Security

4.1 Anonymity

Like in [23, 9], we use a relaxation of the anonymity definition, called weak anonymity and recalled in Definition 3 (in Appendix A). Analogously to the notion of IND-CPA security for public-key encryption, the adversary does not have access to a signature opening oracle.

We show that the two versions (for $b = 0, 1$) of the anonymity security experiment recalled in Figure 2 (in Appendix A) are indistinguishable under the LWE assumption. We use several intermediate hybrid experiments called $G_b^{(i)}$, and show that each of these experiments is indistinguishable from the next one. At each step, we only change one element of the game (highlighted by an arrow in Figure 1), to finally reach the experiment $G^{(4)}$ where the signature scheme does not depend on the identity of the user anymore.

Theorem 2. In the random oracle model, the scheme provides weak anonymity in the sense of Definition 3 under the $\text{LWE}_{q,\alpha}$ assumption. Namely, for any PPT adversary \mathcal{A} with advantage ε , there exists an algorithm \mathcal{B} solving the $\text{LWE}_{q,\alpha}$ problem with the same advantage.

Proof. We define by G_0 the experiment of Definition 3 with $b = 0$ and by G_1 the same experiment with $b = 1$. To show the anonymity of the scheme, we prove that G_0 and G_1 are indistinguishable. We use several hybrid experiments named $G_b^{(1)}$, $G_b^{(2)}$, $G_b^{(3)}$ and $G^{(4)}$ (described in Figure 1), where b is either 0 or 1.

Lemma 7. For each $b \in \{0, 1\}$, the experiments G_b and $G_b^{(1)}$ are statistically indistinguishable.

We only change the way we generate $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T$, by using the fact that one way to generate it is to first sample \mathbf{x}_2 from $D_{\mathbb{Z}^m, \sigma}$ and then generate \mathbf{x}_1 from $D_{\mathbb{Z}^m, \sigma}$ such that $(\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \mathbf{A}_{id_{j_b}} = 0 \pmod q$ (by using the trapdoor \mathbf{T}_A). This change is purely conceptual and the vector $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T$ has the same distribution anyway. The two experiments are thus identical from \mathcal{A} 's view and \mathbf{x}_2 is chosen independently of the signer's identity in the challenge phase.

Lemma 8. For each $b \in \{0, 1\}$, the experiments $G_b^{(1)}$ and $G_b^{(2)}$ are statistically indistinguishable.

The differences between the two experiments are simply: Instead of generating the proofs $\{\pi_{\text{OR}, i}\}_{i=1}^{\ell}$ and π_K using the witnesses, we simulate them (see Section 2.2).

Lemma 9. For each $b \in \{0, 1\}$, if the $\text{LWE}_{q,\alpha}$ problem is hard, then the experiments $G_b^{(2)}$ and $G_b^{(3)}$ are computationally indistinguishable.

Proof. This proof uses the same principle as the proof of [23, Claim 1]: We use the adversary \mathcal{A} to construct a PTT algorithm \mathcal{B} for the $\text{LWE}_{q,\alpha}$ problem. We consider an LWE instance $(\mathbf{B}', \mathbf{z}) \in \mathbb{Z}_q^{m\ell \times (n+1)}$ such that $\mathbf{B}' = (\mathbf{B}'_1, \dots, \mathbf{B}'_{\ell})$ and $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_{\ell})$ with $\mathbf{B}'_i \in \mathbb{Z}_q^{m \times n}$ and $\mathbf{z}_i \in \mathbb{Z}_q^m$. The component \mathbf{z} is either uniform in $\mathbb{Z}_q^{m\ell}$, either of the form $\mathbf{z} = \mathbf{s} \cdot \mathbf{B}'^T + \mathbf{e}$ where \mathbf{e} is sampled from $D_{\mathbb{Z}^{m\ell}, \alpha q}$.

We construct a modified Keygen algorithm using this LWE instance: It generates the matrix \mathbf{A} with a basis \mathbf{T}_A of $\Lambda_q^{\perp}(\mathbf{A})$. Instead of generating the \mathbf{B}_i 's genuinely, we pick \mathbf{B}_0 uniformly in $\mathbb{Z}^{m \times n}$ and set $\mathbf{B}_i = \mathbf{B}'_i$ for $1 \leq i \leq \ell$. For $0 \leq i \leq \ell$, we compute $(\mathbf{A}_i, \mathbf{T}_i) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{B}_i, \mathbf{0})$. Then, for each $j \in [1, N - 1]$, we define \mathbf{A}_{id_j} as in the original Keygen algorithm, and compute a trapdoor \mathbf{T}_{id_j} using \mathbf{T}_A . The adversary \mathcal{A} is given gpk and $\{\text{gsk}_j\}_j$. In the challenge phase, it outputs j_0, j_1 and a message M . By [23], this Keygen algorithm and the one in all the experiments are statistically indistinguishable. Then, the signature is created on behalf of the group member j_b . Namely, \mathcal{B} first chooses $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^m, \sigma}$ and finds \mathbf{x}_1 such that $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T \cdot \mathbf{A}_{id_{j_b}} = \mathbf{0} \pmod q$. Then it chooses $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$ and computes $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^m$.

Experiment G_b

- Run Keygen; give $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_i)$ and $\text{gsk} = \{\mathbf{T}_{id_j}\}_j$ to \mathcal{A} .
- \mathcal{A} outputs j_0, j_1 and a message M .
- The signature of user j_b is computed as follows:
 1. $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T \leftarrow \text{GPVSample}(\mathbf{T}_{id_{j_b}}, \sigma)$;
we have $(\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \mathbf{A}_{id_{j_b}} = \mathbf{0} \pmod q$.
 2. Choose $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$, compute $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^m$.
 3. Choose $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$, and for $i = 1$ to ℓ , choose $\mathbf{e}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and compute $\mathbf{c}_i = \mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}_i + \text{id}_{j_b}[i] \cdot \mathbf{x}_2$.
 4. Generate π_0 .
 5. Generate $\{\pi_{\text{OR}, i}\}_i$.
 6. Generate π_K .

Experiment $G_b^{(2)}$

- Run Keygen; give $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_i)$ and $\text{gsk} = \{\mathbf{T}_{id_j}\}_j$ to \mathcal{A} .
- \mathcal{A} outputs j_0, j_1 and a message M .
- The signature of user j_b is computed as follows:
 1. Sample $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^m, \sigma}$; sample $\mathbf{x}_1 \leftarrow D_{\mathbb{Z}^m, \sigma}$, conditioned on $(\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \mathbf{A}_{id_{j_b}} = \mathbf{0} \pmod q$.
 2. Choose $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$ and compute $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^m$,
 3. Choose $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$, and for $i = 1$ to ℓ , choose $\mathbf{e}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and compute $\mathbf{c}_i = \mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}_i + \text{id}_{j_b}[i] \cdot \mathbf{x}_2$.
 4. Generate π_0 .
 - 5. Simulate $\{\pi_{\text{OR}, i}\}_i$.
 - 6. Simulate π_K .

Experiment $G^{(4)}$

- Run Keygen; give $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_i)$ and $\text{gsk} = \{\mathbf{T}_{id_j}\}_j$ to \mathcal{A} .
- \mathcal{A} outputs j_0, j_1 and a message M .
- The signature of user j_b is computed as follows:
 - 1. Sample $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^m, \sigma}$.
 2. Choose $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$ and compute $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^m$,
 - 3. For $i = 1$ to ℓ , choose $\mathbf{z}_i \leftarrow U(\mathbb{Z}_q^m)$ and set $\mathbf{c}_i = \mathbf{z}_i$.
 4. Generate π_0 .
 5. Simulate $\{\pi_{\text{OR}, i}\}_i$.
 6. Simulate π_K .

Experiment $G_b^{(1)}$

- Run Keygen; give $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_i)$ and $\text{gsk} = \{\mathbf{T}_{id_j}\}_j$ to \mathcal{A} .
- \mathcal{A} outputs j_0, j_1 and a message M .
- The signature of user j_b is computed as follows:
 - 1. Sample $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^m, \sigma}$ and, using \mathbf{T}_A , sample $\mathbf{x}_1 \leftarrow D_{\mathbb{Z}^m, \sigma}$ conditioned on $(\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \mathbf{A}_{id_{j_b}} = \mathbf{0} \pmod q$.
 2. Choose $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$, compute $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^m$,
 3. Choose $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$, and for $i = 1$ to ℓ , choose $\mathbf{e}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and compute $\mathbf{c}_i = \mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}_i + \text{id}_{j_b}[i] \cdot \mathbf{x}_2$.
 4. Generate π_0 .
 5. Generate $\{\pi_{\text{OR}, i}\}_i$.
 6. Generate π_K .

Experiment $G_b^{(3)}$

- Run Keygen; give $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_i)$ and $\text{gsk} = \{\mathbf{T}_{id_j}\}_j$ to \mathcal{A} .
- \mathcal{A} outputs j_0, j_1 and a message M .
- The signature of user j_b is computed as follows:
 1. Sample $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^m, \sigma}$ Sample $\mathbf{x}_1 \leftarrow D_{\mathbb{Z}^m, \sigma}$ conditioned on $(\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \mathbf{A}_{id_{j_b}} = \mathbf{0} \pmod q$.
 2. Choose $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$ and compute $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^m$,
 - 3. For $i = 1$ to ℓ , choose $\mathbf{z}_i \leftarrow U(\mathbb{Z}_q^m)$ and compute $\mathbf{c}_i = \mathbf{z}_i + \text{id}_{j_b}[i] \cdot \mathbf{x}_2$.
 4. Generate π_0 .
 5. Simulate $\{\pi_{\text{OR}, i}\}_i$.
 6. Simulate π_K .

Fig. 1. Experiments $G_b, G_b^{(1)}, G_b^{(2)}, G_b^{(3)}$ and $G^{(4)}$.

Third, it computes $\mathbf{c}_i = p \cdot \mathbf{z}_i + \text{id}_{j_b}[i] \cdot \mathbf{x}_2$ (with the \mathbf{z}_i of the LWE instance). Then it generates π_0 and simulates the $\pi_{\text{OR}, i}$'s and π_K proofs.

We let \mathcal{D}_{LWE} denote this experiment when $\mathbf{z} = \mathbf{s} \cdot \mathbf{B}^T + \mathbf{e}$: This experiment is statistically close to $G_b^{(2)}$. Then, we let $\mathcal{D}_{\text{rand}}$ denote this experiment when \mathbf{z} is uniform: It is statistically close to $G_b^{(3)}$. As a consequence, if the adversary \mathcal{A} can distinguish between the experiments $G_b^{(2)}$ and $G_b^{(3)}$ with non-negligible advantage, then we can solve the $\text{LWE}_{q, \alpha}$ problem with the same advantage.

Lemma 10. *For each $b \in \{0, 1\}$, the experiments $G_b^{(3)}$ and $G^{(4)}$ are indistinguishable.*

Between these two experiments, we change the first and third steps. In the former, we no longer generate \mathbf{x}_1 and, in the latter, \mathbf{c}_i is uniformly sampled in \mathbb{Z}_q^m . These changes are purely conceptual. Indeed, in experiment $G_b^{(3)}$, \mathbf{x}_1 is not used beyond Step 1. In the same experiment, we also have $\mathbf{c}_i = \mathbf{z}_i + \text{id}_{j_b}[i]$. Since the \mathbf{z}_i 's are uniformly sampled in \mathbb{Z}_q^m , the \mathbf{c}_i 's are also uniformly distributed in \mathbb{Z}_q^m . As a consequence, the \mathbf{c}_i 's of $G_b^{(3)}$ and the \mathbf{c}_i 's of $G^{(4)}$ have the same distribution. In $G_b^{(4)}$, we conclude that \mathcal{A} 's view is exactly the same as in experiments $G_b^{(3)}$. Since the experiment $G^{(4)}$ no longer depends on the bit $b \in \{0, 1\}$ that determines the signer's identity, the announced result follows. \square

4.2 Traceability

The proof of traceability relies on the technique of [1, 10] and a refinement from [28, 37], which is used in order to allow for a smaller modulus q .

A difference with the proof of [23] is that we need to rely on the knowledge extractor of a proof of knowledge π_K . Depending on whether the extracted witnesses $\{\mathbf{e}_i, \mathbf{y}_i\}_{i=1}^\ell$ of relation (2) satisfy $\mathbf{y}_i = \text{id}_j[i]\mathbf{x}_2$ for all i or not, we need to distinguish two cases. The strategy of the reduction and the way it uses its given $\text{SIS}_{q,\beta}$ instance will depend on which case is expected to occur.

Theorem 3. *Assume that $q > \log N$, $p \geq \Omega((\alpha q + \sigma)m^{5/2})$ and $\beta \geq \Omega(\sigma m^{7/2}\sqrt{\log N} + p\alpha q m^{5/2})$. Then for any PPT traceability adversary \mathcal{A} with success probability ε , there exists a PPT algorithm \mathcal{B} solving $\text{SIS}_{m,q,\beta}$ with probability $\varepsilon'' \geq \frac{\varepsilon'}{2N} \cdot (\frac{\varepsilon'}{q_H} - 2^{-t}) + \frac{\varepsilon'}{2\log N}$, where $\varepsilon' = \varepsilon - 2^{-t} - 2^{-\Omega(n)}$ and q_H is the number of queries to the random oracle $H : \{0, 1\}^* \rightarrow \{0, 1\}^t$.*

Proof. Let \mathcal{A} be a PPT adversary that can defeat the traceability of the scheme with non-negligible success probability ε in the game of Definition 4. We construct a PPT algorithm \mathcal{B} that emulates \mathcal{A} 's challenger and attacks $\text{SIS}_{q,\beta}$: It takes as input $\bar{\mathbf{A}} \in \mathbb{Z}_q^{m \times n}$ with the task of finding $\mathbf{v} \in \Lambda_q^\perp(\bar{\mathbf{A}})$ with $0 < \|\mathbf{v}\| \leq \beta$.

Initialization. Before starting its interaction with \mathcal{A} , algorithm \mathcal{B} samples $\text{coin} \leftarrow U(\{0, 1\})$. Depending on the latter, the group public key is prepared in two different ways.

- If $\text{coin} = 0$, algorithm \mathcal{B} first calls $\text{TrapGen}(1^n, 1^m, q)$ to obtain $\mathbf{C} \in \mathbb{Z}_q^{m \times n}$ and a basis $\mathbf{T}_{\mathbf{C}}$ of $\Lambda_q^\perp(\mathbf{C})$ with $\|\widetilde{\mathbf{T}_{\mathbf{C}}}\| \leq \mathcal{O}(\sqrt{n \log q})$. Then, it samples $\ell + 1$ matrices $\mathbf{Q}_k \in \mathbb{Z}^{m \times m}$, with each matrix entry sampled independently from $D_{\mathbb{Z}, \sqrt{m}}$ (as in [10, Th. 25], with a larger standard deviation to get exponentially small statistical distances later on). Also, algorithm \mathcal{B} samples $j^* \leftarrow U([0, N - 1])$, a guess that \mathcal{A} 's forgery will open to user j^* . Let $\text{id}_{j^*} = \text{id}_{j^*}[1] \dots \text{id}_{j^*}[\ell] \in \{0, 1\}^\ell$ denote the binary expansion of id_{j^*} . The reduction \mathcal{B} defines the matrices $\{\mathbf{A}_i\}_{i=0}^\ell$ as

$$\begin{cases} \mathbf{A}_0 = \mathbf{Q}_0 \cdot \bar{\mathbf{A}} + (\sum_{i=1}^\ell \text{id}_{j^*}[i]) \cdot \mathbf{C} \\ \mathbf{A}_i = \mathbf{Q}_i \cdot \bar{\mathbf{A}} + (-1)^{\text{id}_{j^*}[i]} \cdot \mathbf{C}, & \text{for } i \in [1, \ell]. \end{cases}$$

It also sets $\mathbf{A} = \bar{\mathbf{A}}$. Next, it runs $\text{SuperSamp}(1^n, 1^m, q, \mathbf{A}_i, \mathbf{0})$ to obtain $\mathbf{B}_i \in \mathbb{Z}_q^{m \times n}$ along with short bases \mathbf{S}'_i of $\Lambda_q^\perp(\mathbf{B}_i)$, and then computes $\mathbf{S}_i \leftarrow \text{RandBasis}(\mathbf{S}'_i, \Omega(\sqrt{mn \log q} \log m))$, as in Step 2 of the genuine key generation algorithm. The group public key $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_{i=0}^\ell)$ is finally given to \mathcal{A} .

We note that, for each $j \neq j^*$, we have

$$\begin{aligned} \mathbf{A}_{\text{id}_j} &= \left[\frac{\bar{\mathbf{A}}}{\mathbf{A}_0 + \sum_{i=1}^\ell \text{id}_j[i]\mathbf{A}_i} \right] = \left[\frac{\bar{\mathbf{A}}}{(\mathbf{Q}_0 + \sum_{i=1}^\ell \text{id}_j[i]\mathbf{Q}_i) \cdot \bar{\mathbf{A}} + (\sum_{i=1}^\ell \text{id}_{j^*}[i] + (-1)^{\text{id}_{j^*}[i]}\text{id}_j[i]) \cdot \mathbf{C}} \right] \\ &= \left[\frac{\bar{\mathbf{A}}}{(\mathbf{Q}_0 + \sum_{i=1}^\ell \text{id}_j[i]\mathbf{Q}_i) \cdot \bar{\mathbf{A}} + h_{\text{id}_j} \cdot \mathbf{C}} \right] \end{aligned}$$

where $h_{\text{id}_j} \in [1, \ell]$ stands for the Hamming distance between the identifiers id_j and id_{j^*} . Since $q > \ell$, we have $h_{\text{id}_j} \neq 0 \pmod q$ whenever $\text{id}_j \neq \text{id}_{j^*}$, so that algorithm \mathcal{B} is able to compute (see [1, Se. 4.2], using the basis $\mathbf{T}_{\mathbf{C}}$ of $\Lambda_q^\perp(\mathbf{C})$ and the refined GPVSample of Lemma 2) a basis $\mathbf{T}'_{\text{id}_j}$ of $\Lambda_q^\perp(\mathbf{A}_{\text{id}_j})$ with $\|\widetilde{\mathbf{T}'_{\text{id}_j}}\| \leq \Omega(m\sqrt{\ell n \log q})$. Then algorithm \mathcal{B} runs $\mathbf{T}_{\text{id}_j} \leftarrow \text{RandBasis}(\mathbf{T}'_{\text{id}_j}, \Omega(m\sqrt{\ell n \log q} \log m))$. Algorithm \mathcal{B} is thus able to compute a trapdoor \mathbf{T}_{id_j} for each $j \neq j^*$. In contrast, algorithm \mathcal{B} lacks a trapdoor for $\mathbf{A}_{\text{id}_{j^*}}$ as the latter only depends on \mathbf{A} and $\{\mathbf{Q}_k\}_{k=0}^\ell$.

Observe that since the rows of the \mathbf{Q}_k 's are sampled from $D_{\mathbb{Z}^m, \sqrt{m}}$, the matrices $\mathbf{A}, \mathbf{A}_0, \dots, \mathbf{A}_\ell$ are within statistical distance $2^{-\Omega(m)}$ of $U(\mathbb{Z}_q^{m \times n})$ (this is a consequence of [22, Le. 5.2]). Further, by Lemma 3, the distribution of the \mathbf{T}_{id_j} 's generated by \mathcal{B} is statistically close to that of the real scheme.

- If $\text{coin} = 1$, algorithm \mathcal{B} samples $i^* \leftarrow U([1, \ell])$ and embeds its $\text{SIS}_{q,\beta}$ instance in the matrix \mathbf{A}_{i^*} that will be part of gpk . It calls $\text{TrapGen}(1^n, 1^m, q)$ to obtain $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and a basis $\mathbf{T}_{\mathbf{A}}$ of $\Lambda_q^\perp(\mathbf{A})$ with $\|\widetilde{\mathbf{T}_{\mathbf{A}}}\| \leq \mathcal{O}(\sqrt{n \log q})$. Next, it independently samples $\mathbf{A}_0, \dots, \mathbf{A}_{i^*-1}, \mathbf{A}_{i^*+1}, \dots, \mathbf{A}_\ell \leftarrow U(\mathbb{Z}_q^{m \times n})$ and defines $\mathbf{A}_{i^*} = \mathbf{A}$. Then, algorithm \mathcal{B} computes $(\mathbf{B}_i, \mathbf{S}'_i) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{A}_i, \mathbf{0})$ and $\mathbf{S}_i \leftarrow \text{RandBasis}(\mathbf{S}'_i, \Omega(\sqrt{mn \log q} \log m))$, as in Step 2 of Keygen. The group public key $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_{i=0}^\ell)$, which is distributed as in the real scheme, is given to the adversary \mathcal{A} . Since it knows $\mathbf{T}_{\mathbf{A}}$, algorithm \mathcal{B} is able to sample a trapdoor \mathbf{T}_{id_j} for all users, with exactly the same distribution as in the real scheme.

In either case, \mathcal{B} runs the adversary \mathcal{A} on inputs $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, \mathbf{B}_i\}_{i=0}^\ell)$ and $\text{gmsk} = \{\mathbf{S}_i\}_{i=0}^\ell$.

Queries. Algorithm \mathcal{B} then starts interacting with \mathcal{A} and handles \mathcal{A} 's queries depending on *coin*.

- If *coin* = 0, it aborts in the event that \mathcal{A} queries the unavailable secret key $\text{gsk}[j^*]$. When \mathcal{A} queries a secret key $\text{gsk}[j]$ for $j \neq j^*$, algorithm \mathcal{B} reveals the short basis \mathbf{T}_{id_j} that was computed in the initialization phase. When it comes to answer signing queries, algorithm \mathcal{B} faithfully runs the signing algorithm whenever the involved user j differs from j^* . As for signing queries involving the expected target user j^* , the reduction \mathcal{B} samples $\mathbf{s}_0, \mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$, $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^m, \sigma}$ and $\mathbf{e}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$ for each $i \in [1, \ell]$. It then computes $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2$ as well as $\mathbf{c}_i = \mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}_i + \text{id}_{j^*}[i] \mathbf{x}_2$ for each $i \in [1, \ell]$. The proof π_0 is then generated using the actual witness \mathbf{x}_2 whereas the other non-interactive proofs $\{\pi_{\text{OR}, i}\}_{i=1}^\ell$ and π_K are simulated using the appropriate NIZK simulator (exactly as in experiment $G_b^{(2)}$ in the proof of anonymity). By the statistical zero-knowledge property of the simulator, the signature Σ will be statistically indistinguishable from a genuine signature.
- If *coin* = 1, algorithm \mathcal{B} knows \mathbf{T}_A and can answer \mathcal{A} 's queries by running the real signing algorithm or returning the queried secret keys $\text{gsk}[j]$ (all of which are available).

Regardless of the value of *coin*, queries to the random oracle H are handled by returning a uniformly chosen value in $\{0, 1\}^t$. For each $\kappa \leq q_H$, we let r_κ denote the answer to the κ -th H -query. Of course, if the adversary makes a given query more than once, then \mathcal{B} consistently returns the previously defined value.

Forgery. When \mathcal{A} terminates, it outputs a signature $\Sigma^* = ((\mathbf{c}_i^*)_{0 \leq i \leq \ell}, \pi_0^*, (\pi_{\text{OR}, i}^*)_{1 \leq i \leq \ell}, \pi_K^*)$ on some message M^* with probability $\geq \varepsilon - 2^{-\Omega(n)}$. If we parse π_K^* as $(\text{Comm}_K^*, \text{Chall}_K^*, \text{Resp}_K^*)$, with overwhelming probability, the adversary \mathcal{A} must have queried H on the input $(M^*, \text{Comm}_K^*, (\mathbf{c}_i^*)_{0 \leq i \leq \ell}, \pi_0^*, (\pi_{\text{OR}, i}^*)_{1 \leq i \leq \ell})$. Indeed, otherwise, the probability to have the equality $\text{ch}_K^* = H(M^*, \text{Comm}_K^*, (\mathbf{c}_i^*)_{0 \leq i \leq \ell}, \pi_0^*, (\pi_{\text{OR}, i}^*)_{1 \leq i \leq \ell})$ is at most 2^{-t} . With probability $\geq \varepsilon' := \varepsilon - 2^{-t} - 2^{-\Omega(n)}$, the tuple $(M^*, \text{Comm}_K^*, (\mathbf{c}_i^*)_{0 \leq i \leq \ell}, \pi_0^*, (\pi_{\text{OR}, i}^*)_{1 \leq i \leq \ell})$ thus coincides with the κ^* -th hash query for some $\kappa^* \leq q_H$.

At this stage, the reduction \mathcal{B} runs a second execution of the adversary \mathcal{A} with the *same* random tape and input as in the original execution. All queries are answered as previously with only one difference in the treatment of random oracle queries. Namely, the first $\kappa^* - 1$ hash queries – which are identical to those of the first execution since \mathcal{A} is run with the same random tape as before – receive the same answers $r_1, \dots, r_{\kappa^* - 1}$ as in the initial run. This implies that the κ^* -th query will involve the tuple $(M^*, \text{Comm}_K^*, (\mathbf{c}_i^*)_{0 \leq i \leq \ell}, \pi_0^*, (\pi_{\text{OR}, i}^*)_{1 \leq i \leq \ell})$ as in the first execution. However, from the κ^* -th query onwards, \mathcal{A} obtains fresh random oracle values $r'_{\kappa^*}, \dots, r'_{q_H}$ which depart from the sequence of answers in the first execution. The General Forking Lemma of [6] implies that, with probability $\geq \varepsilon'(\varepsilon'/q_H - 2^{-t})$, \mathcal{A} 's forgery also involves $(M^*, \text{Comm}_K^*, (\mathbf{c}_i^*)_{0 \leq i \leq \ell}, \pi_0^*, (\pi_{\text{OR}, i}^*)_{1 \leq i \leq \ell})$ in the second run and we also have $r'_{\kappa^*} \neq r_{\kappa^*}$. In this case, using Extract, algorithm \mathcal{B} can obtain vectors $\mathbf{e}_1, \dots, \mathbf{e}_\ell, \mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{y}_\ell \in \mathbb{Z}^m$ satisfying

$$\mathbf{x}_1^T \mathbf{A} + \sum_{i=0}^{\ell} \mathbf{c}_i^T \mathbf{A}_i = \sum_{i=1}^{\ell} \mathbf{e}_i^T (p \mathbf{A}_i) \quad \text{and} \quad \mathbf{e}_i^T (p \mathbf{A}_i) + \mathbf{y}_i^T \mathbf{A}_i = \mathbf{c}_i^T \mathbf{A}_i \quad \forall i \in [1, \ell] \quad (4)$$

with $\|\mathbf{e}_i\|, \|\mathbf{y}_i\|, \|\mathbf{x}_1\| \leq \mathcal{O}((\alpha q + \sigma)m^{5/2})$ for all $i \in [1, \ell]$ (see Section 2.2).

The reduction \mathcal{B} then opens one of the two forgeries using $\{\mathbf{S}_i\}_{i=0}^\ell$ (note that both signatures necessarily open to the same identity id). At this point, \mathcal{B} aborts and reports failure if the opening algorithm does not point to user j^* . However, with probability $\geq 1/N$, \mathcal{B} 's initial choice for j^* turns out to be correct and the opening algorithm reveals id_{j^*} .

We now assume that Σ^* indeed traces to user j^* . We let $\mathbf{x}_2 \in \mathbb{Z}^m$ denote the vector obtained by decrypting \mathbf{c}_0^* using \mathbf{S}_0 . Algorithm \mathcal{B} considers the following two situations:

- If $\mathbf{y}_i = \text{id}_{j^*}[i] \mathbf{x}_2$ for all $i \in [1, \ell]$, then \mathcal{B} aborts if *coin* = 1 and continues if *coin* = 0. The relations (4) and the fact that \mathbf{c}_0^* is of the form $\mathbf{c}_0^* = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2 \pmod q$ with $\mathbf{B}_0^T \cdot \mathbf{A}_0 = \mathbf{0} \pmod q$ imply that (modulo q):

$$\begin{aligned} \mathbf{0} &= \mathbf{x}_1^T \mathbf{A} + \mathbf{c}_0^{*T} \mathbf{A}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \mathbf{x}_2^T \cdot \mathbf{A}_i = (\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \left[\frac{\mathbf{A}}{\mathbf{A}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \mathbf{A}_i} \right] \\ &= (\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \left[\frac{\bar{\mathbf{A}}}{(\mathbf{Q}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \mathbf{Q}_i) \cdot \bar{\mathbf{A}}} \right], \end{aligned}$$

by construction of the matrices $\mathbf{A}, \mathbf{A}_0, \dots, \mathbf{A}_\ell$. It comes that $\mathbf{v}^T = \mathbf{x}_1^T + \mathbf{x}_2^T \cdot (\mathbf{Q}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \mathbf{Q}_i) \in \Lambda^\perp(\bar{\mathbf{A}})$. The same analysis as in [10] shows that $0 < \|\mathbf{v}\| \leq \mathcal{O}((\alpha q + \sigma)m^{7/2} \sqrt{\ell})$ holds with probability $1 - 2^{-\Omega(m)}$.

• If there exists $i \in [1, \ell]$ such that $\mathbf{y}_i \neq \text{id}_{j^*}[i]\mathbf{x}_2$, then \mathcal{B} aborts if $\text{coin} = 0$ and continues if $\text{coin} = 1$. The non-interactive proofs π_0 and $\pi_{\text{OR},i}$ imply that $\mathbf{c}_i = \mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}'_i + \text{id}_{j^*}[i]\mathbf{x}_2 \pmod q$ for some $\mathbf{s}_0, \mathbf{s} \in \mathbb{Z}_q^n$ and $\mathbf{x}_2, \mathbf{e}'_i \in \mathbb{Z}^m$ such that $\|\mathbf{x}_2\| \leq \mathcal{O}(\sigma m^{5/2})$ and $\|\mathbf{e}'_i\| \leq \mathcal{O}(\alpha q m^{5/2})$. If we multiply \mathbf{c}'_i by \mathbf{A}_i , we find

$$\mathbf{c}'_i \mathbf{A}_i = p \mathbf{e}'_i{}^T \cdot \mathbf{A}_i + \text{id}_{j^*}[i]\mathbf{x}_2^T \cdot \mathbf{A}_i.$$

By subtracting the latter equation from the second equation of (4), we find (still modulo q):

$$(p(\mathbf{e}_i^T - \mathbf{e}'_i{}^T) + (\mathbf{y}_i^T - \text{id}_{j^*}[i]\mathbf{x}_2^T)) \cdot \mathbf{A}_i = \mathbf{0}.$$

If $p(\mathbf{e}_i - \mathbf{e}'_i) + (\mathbf{y}_i - \text{id}_{j^*}[i]\mathbf{x}_2) \neq \mathbf{0}$, it is a non-zero vector in $\Lambda^\perp(\mathbf{A}_i)$ of norm $\leq \mathcal{O}((\sigma + p\alpha q)m^{5/2})$. Given that we have $\mathbf{A}_i = \bar{\mathbf{A}}$ with probability $1/\ell$, we solved the given SIS instance with the same probability. Finally, if $p(\mathbf{e}_i - \mathbf{e}'_i) + (\mathbf{y}_i - \text{id}_{j^*}[i]\mathbf{x}_2) = \mathbf{0}$, the relative norms of the vectors $\mathbf{e}_i, \mathbf{e}'_i, \mathbf{y}_i, \mathbf{x}_2$ with respect to p imply $\mathbf{e}_i = \mathbf{e}'_i$ and $\mathbf{y}_i = \text{id}_{j^*}[i]\mathbf{x}_2$ (over the integers), which is in contradiction with $\mathbf{y}_i \neq \text{id}_{j^*}[i]\mathbf{x}_2$.

The lower bound on \mathcal{B} 's advantage is obtained by combining the probability of obtaining a successful forking, the fact that \mathcal{B} 's choice for $j^* \in U([0, N-1])$ is independent of \mathcal{A} 's view when $\text{coin} = 0$ and the observation that \mathcal{B} 's choice for coin is also independent of \mathcal{A} 's view. \square

5 A Variant with Full (CCA-)Anonymity

We modify in the sequel our basic group signature scheme to reach the strongest possible anonymity level (detailed in Definition 3 of Appendix A), in which the attacker is authorized to query an opening oracle. This implies the simulation of an oracle which opens adversarially-chosen signatures in the proof of anonymity. To this end, we essentially replace each \mathbf{B}_i from our previous scheme by a matrix $\mathbf{B}_{i,\text{VK}}$ that depends on the verification key VK of a strongly unforgeable one-time signature. The reduction will be able to compute a trapdoor for all these matrices, except for one specific verification key VK^* that will be used in the challenge phase. This will provide the reduction with a backdoor allowing it to open all adversarially-generated signatures.

It is assumed that the one-time verification keys VK belong to \mathbb{Z}_q^n (note that this condition can always be enforced by hashing VK). Following Agrawal *et al.* [1], we rely on a full-rank difference function $H_{vk} : \mathbb{Z}_q^n \rightarrow \mathbb{Z}_q^{n \times n}$ such that, for any two distinct $\mathbf{u}, \mathbf{v} \in \mathbb{Z}_q^n$, the difference $H_{vk}(\mathbf{u}) - H_{vk}(\mathbf{v})$ is a full rank matrix. This function H_{vk} can be obtained as in [1, Se. 5]. Below is the description of our new scheme.

Keygen($1^n, 1^N$): given a security parameter $n > 0$ and the desired number of members $N = 2^\ell \in \text{poly}(n)$, choose parameters as in Section 3 and make them public. Choose a hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^t$, that will be seen as a random oracle, and the specification of a one-time signature $\Pi^{\text{ots}} = (\mathcal{G}, \mathcal{S}, \mathcal{V})$. Then, proceed as follows.

1. Run $\text{TrapGen}(1^n, 1^m, q)$ to get $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ with a short basis $\mathbf{T}_\mathbf{A} \in \mathbb{Z}^{m \times m}$ of $\Lambda_q^\perp(\mathbf{A})$.
2. For $i = 0$ to ℓ , repeat the following steps.
 - a. Choose uniformly random matrices $\mathbf{A}_{i,1}, \mathbf{B}_{i,0}, \mathbf{B}_{i,1} \in \mathbb{Z}_q^{m \times n}$.
 - b. Compute $(\mathbf{A}_{i,2}, \mathbf{T}_{i,2}) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{B}_{i,1}, 0^{n \times n})$ such that $\mathbf{B}_{i,1}^T \cdot \mathbf{A}_{i,2} = 0 \pmod q$ and discard $\mathbf{T}_{i,2}$, which will not be needed. Define $\mathbf{A}_i = \begin{bmatrix} \mathbf{A}_{i,1} \\ \mathbf{A}_{i,2} \end{bmatrix} \in \mathbb{Z}_q^{2m \times n}$.
 - c. Run $(\mathbf{B}_{i,-1}, \mathbf{S}'_i) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{A}_{i,1}, -\mathbf{A}_{i,2}^T \cdot \mathbf{B}_{i,0})$ to obtain $\mathbf{B}_{i,-1} \in \mathbb{Z}_q^{m \times n}$ such that $\mathbf{B}_{i,-1}^T \cdot \mathbf{A}_{i,1} + \mathbf{B}_{i,0}^T \cdot \mathbf{A}_{i,2} = 0 \pmod q$.
 - d. Finally, compute a re-randomized trapdoor $\mathbf{S}_i \leftarrow \text{RandBasis}(\mathbf{S}'_i)$ for $\mathbf{B}_{i,-1}$.

For any string VK , if the matrix $H_{vk}(\text{VK})$ is used to define $\mathbf{B}_{i,\text{VK}} = \begin{bmatrix} \mathbf{B}_{i,-1} \\ \mathbf{B}_{i,0} + \mathbf{B}_{i,1}H_{vk}(\text{VK}) \end{bmatrix} \in \mathbb{Z}_q^{2m \times n}$, we have $\mathbf{B}_{i,\text{VK}}^T \cdot \mathbf{A}_i = 0 \pmod q$ for all i .

3. For $j = 0$ to $N-1$, let $\text{id}_j = \text{id}_j[1] \dots \text{id}_j[\ell] \in \{0, 1\}^\ell$ be the binary representation of id_j and define:

$$\mathbf{A}_{\text{id}_j} = \begin{bmatrix} \mathbf{A} \\ \mathbf{A}_0 + \sum_{i=1}^{\ell} \text{id}_j[i]\mathbf{A}_i \end{bmatrix} \in \mathbb{Z}_q^{3m \times n}.$$

Then run $\mathbf{T}_{\text{id}_j} \leftarrow \text{ExtBasis}(\mathbf{T}_\mathbf{A}, \mathbf{A}_{\text{id}_j})$ to get a short delegated basis $\mathbf{T}_{\text{id}_j} \in \mathbb{Z}^{3m \times 3m}$ of $\Lambda_q^\perp(\mathbf{A}_{\text{id}_j})$ and define $\text{gsk}[j] := \mathbf{T}_{\text{id}_j}$.

4. Define $\text{gpk} := (\mathbf{A}, \{\mathbf{A}_i, (\mathbf{B}_{i,-1}, \mathbf{B}_{i,0}, \mathbf{B}_{i,1})\}_{i=0}^{\ell}, H, \Pi^{\text{ots}})$ and $\text{gmsk} := \{\mathbf{S}_i\}_{i=0}^{\ell}$. The algorithm outputs $(\text{gpk}, \text{gmsk}, \{\text{gsk}[j]\}_{j=0}^{N-1})$.

Sign($\text{gpk}, \text{gsk}[j], M$): To sign a message $M \in \{0, 1\}^*$ using the private key $\text{gsk}[j] = \mathbf{T}_{\text{id}_j}$, generate a one-time signature key pair $(\mathbf{VK}, \mathbf{SK}) \leftarrow \mathcal{G}(1^n)$ for Π^{ots} and proceed as follows.

1. Run $\text{GPVSample}(\mathbf{T}_{\text{id}_j}, \sigma)$ to obtain a vector $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T \in \Lambda_q^\perp(\mathbf{A}_{\text{id}_j})$ of norm $\leq \sigma\sqrt{3m}$.
2. Sample $\mathbf{s}_0 \leftarrow U(\mathbb{Z}_q^n)$ and encrypt $\mathbf{x}_2 \in \mathbb{Z}_q^{2m}$ by computing $\mathbf{c}_0 = \mathbf{B}_{0, \mathbf{VK}} \cdot \mathbf{s}_0 + \mathbf{x}_2 \in \mathbb{Z}_q^{2m}$.
3. Sample $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$. For $i = 1$ to ℓ , sample $\mathbf{e}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$ and a random matrix $\mathbf{R}_i \in \mathbb{Z}^{m \times m}$ whose columns are sampled from $D_{\mathbb{Z}^m, \sigma}$. Then, compute $\mathbf{c}_i = \mathbf{B}_{i, \mathbf{VK}} \cdot \mathbf{s} + p \cdot [\mathbf{e}_i | \mathbf{e}_i \cdot \mathbf{R}_i] + \text{id}_j[i] \cdot \mathbf{x}_2$, which encrypts $\mathbf{x}_2 \in \mathbb{Z}_q^{2m}$ (resp. 0^{2m}) if $\text{id}_j[i] = 1$ (resp. $\text{id}_j[i] = 0$).
4. Generate a NIZKPoK π_0 of \mathbf{s}_0 so that $(\mathbf{B}_0, \mathbf{c}_0, \sqrt{2}\sigma/q; \mathbf{s}_0) \in R_{\text{LWE}}$ (see Section 2.2).
5. For $i = 1$ to ℓ , generate a NIZKPoK $\pi_{\text{OR}, i}$ of \mathbf{s} and \mathbf{s}_0 so that either:
 - (i) $((\mathbf{B}_{i, \mathbf{VK}} | \mathbf{B}_{0, \mathbf{VK}}), p^{-1}(\mathbf{c}_i - \mathbf{c}_0), \sqrt{2}\alpha; (\mathbf{s}^T | -\mathbf{s}_0^T)^T) \in R_{\text{LWE}}$ (the vectors \mathbf{c}_i and \mathbf{c}_0 encrypt the same \mathbf{x}_2 , so that $p^{-1}(\mathbf{c}_i - \mathbf{c}_0)$ is close to the \mathbb{Z}_q -span of $(\mathbf{B}_{i, \mathbf{VK}} | \mathbf{B}_{0, \mathbf{VK}})$);
 - (ii) or $(\mathbf{B}_{i, \mathbf{VK}}^T, p^{-1}\mathbf{c}_i, \alpha; \mathbf{s}) \in R_{\text{LWE}}$ (\mathbf{c}_i encrypts $\mathbf{0}$, so that $p^{-1}\mathbf{c}_i$ is close to the \mathbb{Z}_q -span of $\mathbf{B}_{i, \mathbf{VK}}^T$).
6. For $i = 1$ to ℓ , set $\mathbf{y}_i = \text{id}_j[i]\mathbf{x}_2 \in \mathbb{Z}^{2m}$ and generate a NIZKPoK π_K of $(\mathbf{e}_i)_{1 \leq i \leq \ell}, (\mathbf{y}_i)_{1 \leq i \leq \ell}, \mathbf{x}_1$ s.t:

$$\mathbf{x}_1^T \mathbf{A} + \sum_{i=0}^{\ell} \mathbf{c}_i^T \mathbf{A}_i = \sum_{i=1}^{\ell} \mathbf{e}_i^T (p \cdot \mathbf{A}_i) \quad \text{and} \quad \mathbf{e}_i^T (p \cdot \mathbf{A}_i) + \mathbf{y}_i^T \mathbf{A}_i = \mathbf{c}_i^T \mathbf{A}_i, \text{ for } i \in \{1, \dots, \ell\} \quad (5)$$

with $\|\mathbf{x}_1\| \leq \sigma\sqrt{m}$ and $\|\mathbf{y}_i\| \leq \sigma\sqrt{2m}$ for all i . This is achieved using $\text{Prove}_{\text{SIS}}$, producing a triple $(\text{Comm}_K, \text{Chall}_K, \text{Resp}_K)$, where $\text{Chall}_K = H(M, \text{Comm}_K, (\mathbf{c}_i)_{0 \leq i \leq \ell}, \pi_0, (\pi_{\text{OR}, i})_{1 \leq i \leq \ell})$.

7. Compute a one-time signature $\text{sig} = \mathcal{S}(\mathbf{SK}, (\mathbf{c}_i)_{0 \leq i \leq \ell}, \pi_0, (\pi_{\text{OR}, i})_{1 \leq i \leq \ell}, \pi_K)$.

The signature consists of

$$\Sigma = (\mathbf{VK}, (\mathbf{c}_i)_{0 \leq i \leq \ell}, \pi_0, (\pi_{\text{OR}, i})_{1 \leq i \leq \ell}, \pi_K, \text{sig}). \quad (6)$$

Verify(gpk, M, Σ): Parse Σ as in (6). Return 0 if $\mathcal{V}(\mathbf{VK}, \text{sig}, (\mathbf{c}_i)_{0 \leq i \leq \ell}, \pi_0, (\pi_{\text{OR}, i})_{1 \leq i \leq \ell}, \pi_K) = 0$. Then, return 1 if all proofs $\pi_0, (\pi_{\text{OR}, i})_{1 \leq i \leq \ell}, \pi_K$ properly verify. Otherwise, return 0.

Open($\text{gpk}, \text{gmsk}, M, \Sigma$): Parse gmsk as $\{\mathbf{S}_i\}_{i=0}^{\ell}$ and Σ as in (6). For $i = 0$ to ℓ , compute a trapdoor $\mathbf{S}_{i, \mathbf{VK}} \leftarrow \text{ExtBasis}(\mathbf{S}_i, \mathbf{B}_{i, \mathbf{VK}})$ for $\mathbf{B}_{i, \mathbf{VK}}$. Using the delegated basis $\mathbf{S}_{0, \mathbf{VK}} \in \mathbb{Z}^{2m \times 2m}$ (for which we have $\mathbf{S}_{0, \mathbf{VK}} \cdot \mathbf{B}_{0, \mathbf{VK}} = 0 \pmod{q}$), compute \mathbf{x}_2 by decrypting \mathbf{c}_0 . Then, using $\mathbf{S}_{i, \mathbf{VK}} \in \mathbb{Z}^{2m \times 2m}$, determine which vector among $p^{-1}\mathbf{c}_i \pmod{q}$ and $p^{-1}(\mathbf{c}_i - \mathbf{x}_2) \pmod{q}$ is close to the \mathbb{Z}_q -span of $\mathbf{B}_{i, \mathbf{VK}}$. Set $\text{id}[i] = 0$ in the former case and $\text{id}[i] = 1$ in the latter case. Eventually, output $\text{id} = \text{id}[1] \dots \text{id}[\ell]$.

In Appendix D, we prove the following theorems which provide evidence that the above scheme provides full anonymity and full traceability in the random oracle model.

Theorem 4. *The scheme provides full anonymity in the random oracle model if the $\text{LWE}_{q, \alpha}$ assumption holds and if the one-time signature is strongly unforgeable.*

Theorem 5. *Assuming that $q > \log N$, the scheme is fully traceable in the random oracle model under the $\text{SIS}_{q, \beta}$ assumption. More precisely, for any PPT traceability adversary \mathcal{A} with success probability ε , there exists an algorithm \mathcal{B} solving the $\text{SIS}_{q, \beta}$ problem with probability at least $\frac{1}{2N} \cdot \left(\varepsilon - \frac{1}{2^t}\right) \cdot \left(\frac{\varepsilon - 1/2^t}{q_H} - \frac{1}{2^t}\right)$, where q_H is the number of queries to $H : \{0, 1\}^* \rightarrow \{0, 1\}^t$ and N stands for the maximal number of users.*

Acknowledgements. We thank Daniele Micciancio and Khoa Nguyen for helpful discussions. Part of this work was done while the third author was an FNRS scientific collaborator (at Université catholique de Louvain, Belgium) and visiting ENS de Lyon, under the INRIA invited researcher scheme. The last author was partly supported by the Australian Research Council Discovery Grant DP110100628.

References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 553–572. Springer, 2010.
2. J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. *Theor. Comput. Science*, 48(3):535–553, 2011.
3. G. Ateniese, J. Camenisch, M. Joye, and G. Tsudik. A practical and provably secure coalition-resistant group signature scheme. In *Proceedings of Crypto 2000*, number 1880 in *LNCS*, pages 255–270. Springer, 2000.
4. W. Banaszczyk. New bounds in some transference theorems in the geometry of number. *Math. Ann.*, 296:625–635, 1993.
5. M. Bellare, D. Micciancio, and B. Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In *Eurocrypt'03*, volume 2656 of *LNCS*, pages 614–629, 2003.
6. M. Bellare and G. Neven. Multi-signatures in the plain public-key model and a general forking lemma. In *ACM Conference on Computer and Communications Security (ACM-CCS) 2006*, pages 390–399. ACM Press, 2006.
7. M. Bellare, H. Shi, and C. Zhang. Foundations of group signatures: The case of dynamic groups. In *Proceedings of CT-RSA 2005*, volume 3376 of *LNCS*, pages 136–153. Springer, 2005.
8. D. Boneh and X. Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *Proceedings of Eurocrypt 2004*, volume 3027 of *LNCS*, pages 223–238. Springer, 2004.
9. D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In *Proceedings of Crypto 2004*, volume 3152 of *LNCS*, pages 41–55. Springer, 2004.
10. X. Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Proc. of PKC*, volume 6056 of *LNCS*, pages 499–517. Springer, 2010.
11. X. Boyen and B. Waters. Compact group signatures without random oracles. In *Proceedings of Eurocrypt 2006*, volume 4004 of *LNCS*, pages 427–444. Springer, 2006.
12. X. Boyen and B. Waters. Full-domain subgroup hiding and constant-size group signatures. In *Proceedings of PKC 2007*, volume 4450 of *LNCS*, pages 1–15. Springer, 2007.
13. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. On the classical hardness of learning with errors. To appear in the proceedings of STOC'13.
14. E. Brickell. An efficient protocol for anonymously providing assurance of the container of a private key, apr. 2003., 2003. Submitted to the Trusted Computing Group.
15. J. Camenisch and A. Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In *Crypto'02*, volume 2442 of *LNCS*, pages 61–76. Springer, 2002.
16. J. Camenisch and A. Lysyanskaya. A signature scheme with efficient protocols. In *Security in Communication Networks (SCN 2002)*, volume 2576 of *LNCS*, pages 268–289. Springer, 2002.
17. J. Camenisch and A. Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In *Crypto'04*, volume 3152 of *LNCS*, pages 56–72. Springer, 2004.
18. J. Camenisch, G. Neven, and M. Rückert. Fully anonymous attribute tokens from lattices. In *Proceedings of SCN 2012*, volume 7485 of *LNCS*, pages 57–75. Springer, 2012.
19. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 523–552. Springer, 2010.
20. D. Chaum and E. van Heyst. Group signatures. In *Eurocrypt'91*, volume 547 of *LNCS*, pages 257–265. Springer, 1991.
21. S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices and applications. In *Proc. of EUROCRYPT*, LNCS. Springer, 2013.
22. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008. Full version available at <http://eprint.iacr.org/2007/432.pdf>.
23. S. D. Gordon, J. Katz, and V. Vaikuntanathan. A group signature scheme from lattice assumptions. In *Proc. of ASIACRYPT*, volume 2647 of *LNCS*, pages 395–412. Springer, 2010.
24. J. Groth. Simulation-sound nzk proofs for a practical language and constant size group signatures. In *Asiacrypt'06*, volume 4284 of *LNCS*, pages 444–459. Springer, 2006.
25. J. Groth. Fully anonymous group signatures without random oracles. In *Asiacrypt07*, volume 4833 of *LNCS*, pages 164–180. Springer, 2007.
26. J. Groth, Rafail Ostrovsky, and A. Sahai. Perfect non-interactive zero knowledge for NP. In *Eurocrypt'06*, volume 4004 of *LNCS*, pages 339–358. Springer, 2006.
27. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt'08*, volume 4965 of *LNCS*, pages 415–432. Springer, 2008.
28. S. Hohenberger and B. Waters. Short and stateless signatures from the RSA assumption. In *Proc. of CRYPTO*, volume 5677 of *LNCS*, pages 654–670. Springer, 2009.
29. VSC Project IEEE P1556 Working Group. Dedicated short range communications (dsrc), 2003.
30. A. Kiayias and M. Yung. Secure scalable group signature with dynamic joins and separable authorities. *International Journal of Security and Networks (IJSN)*, 1(1/2):24–45, 2006.
31. V. Lyubashevsky. Lattice-based identification schemes secure unde active attacks. In *Proceedings of PKC 2008*, volume 4939 of *LNCS*, pages 162–179. Springer, 2008.
32. V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *Proc. of ASIACRYPT*, volume 5912 of *LNCS*, pages 598–616. Springer, 2009.
33. V. Lyubashevsky. Lattice signatures without trapdoors. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 738–755. Springer, 2012.

34. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. of ICALP*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
35. V. Lyubashevsky and D. Micciancio. Asymptotically efficient lattice-based digital signatures. In *Proc. of TCC*, volume 4948 of *LNCS*, pages 37–54. Springer, 2008.
36. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 1–23. Springer, 2010.
37. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 700–718. Springer, 2012. Full version available at <http://www.cc.gatech.edu/~cpeikert/pubs/efftrap.pdf>.
38. D. Micciancio and S. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *CRYPTO2003*, pages 282–298, 2003.
39. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
40. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proc. of TCC*, volume 3876 of *LNCS*, pages 145–166. Springer, 2006.
41. C. Peikert and V. Vaikuntanathan. Non-interactive statistical zero-knowledge proofs for lattice problems. In *Proceedings of Crypto'08*, volume LNCS, pages 536–553. Springer, 2008.
42. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.

A Security Definitions for Group Signatures

Bellare *et al.* [5] gave a unified security model for group signatures in static groups. The two main security requirements are called *anonymity* and *traceability*.




Fig. 2. Random experiments for anonymity and full traceability

Anonymity. Anonymity requires that an adversary who does not know the group manager secret key cannot recognize the identity of a user given its signature. More formally, the attacker, modeled as a two-stage adversary (*choose* and *guess*), is engaged in the first random experiment depicted in Figure 2. The *advantage* of such an adversary \mathcal{A} against a group signature \mathcal{GS} with N members is defined as

$$\text{Adv}_{\mathcal{GS}, \mathcal{A}}^{\text{anon}}(n, N) = |\Pr[\text{Exp}_{\mathcal{GS}, \mathcal{A}}^{\text{anon-1}}(n, N) = 1] - \Pr[\text{Exp}_{\mathcal{GS}, \mathcal{A}}^{\text{anon-0}}(n, N) = 1]|.$$

In our first scheme, we consider a *weak anonymity* scenario in which the adversary is not allowed to query an opening oracle. This relaxed model is precisely the one considered in [23], and was firstly introduced in [9]. Nonetheless, we provide in Section 5 a variant of our scheme enjoying chosen-ciphertext security. The adversary is then granted an access to an opening oracle that can be called on any string except the challenge signature Σ^* .

Definition 3 (Weak and full anonymity, [5, 9]). A group signature scheme \mathcal{GS} is said to be weakly anonymous (*resp.* fully anonymous) if for all polynomial $N(\cdot)$ and all PPT adversaries \mathcal{A} (*resp.* PPT adversaries \mathcal{A} with access to an opening oracle except for the challenge signature), its advantage $\text{Adv}_{\mathcal{GS}, \mathcal{A}}^{\text{anon}}(n, N)$ is negligible in the security parameter n .

Full traceability. Full traceability ensures that all signatures, even those created by a coalition of users *and* the group manager, pooling their secret keys together, can be traced to a member of the forging coalition. Once again, the attacker is modeled as a two-stage adversary who is run within the second experiment described in Figure 2. Its success probability against \mathcal{GS} is defined as

$$\text{Succ}_{\mathcal{GS}, \mathcal{A}}^{\text{trace}}(n, N) = \Pr[\text{Exp}_{\mathcal{GS}, \mathcal{A}}^{\text{trace}}(n, N) = 1].$$

Definition 4 (Full traceability, [5]). A group signature scheme \mathcal{GS} is said to be fully traceable if for all polynomial $N(\cdot)$ and all PPT adversaries \mathcal{A} , its success probability $\text{Succ}_{\mathcal{GS}, \mathcal{A}}^{\text{trace}}(n, N)$ is negligible in the security parameter n .

B Proof of Theorem 1

Proof. Setting $m = \Omega(n \log q)$ allows us to use algorithms `TrapGen` and `SuperSamp` from Lemmas 4 and 5, at Steps 1 and 2 of algorithm `Keygen`. Also, the rows of the matrix \mathbf{A} sampled at Step 1 span \mathbb{Z}_q^n with probability $\geq 1 - 2^{-\Omega(n)}$. At Steps 2 and 3, the second inputs to the calls to `RandBasis` are sufficiently large for the assumption of Lemma 3 to hold (note that in the second case, it is much larger than needed, but this choice is important for the simulation in the traceability proof). At the end of the execution of `Keygen`, we have $\|\widetilde{\mathbf{S}}_i\| \leq \mathcal{O}(m \log m \sqrt{n \log q})$ for all $i \in [0, \ell]$ and $\|\widetilde{\mathbf{T}}_{\text{id}_j}\| \leq \mathcal{O}(m^{3/2} \sqrt{\ell n \log q \log m})$ for all $j \in [0, N-1]$.

At Step 1 of algorithm `Sign`, the parameter σ is sufficiently large for applying Lemma 2 and obtain a distribution within statistical distance $2^{-\Omega(n)}$ from $D_{\Lambda_q^+(\mathbf{A}_{\text{id}_j}, \sigma)}$. The same holds for all \mathbf{e}_i 's of Step 3.

Correctness of algorithm `Verify` follows the completeness property of the underlying proof systems. Now, consider algorithm `Open`. We have $\mathbf{S}_0 \cdot \mathbf{c}_0 = \mathbf{S}_0 \cdot \mathbf{x}_2 \pmod q$. But on the other hand $\|\mathbf{S}_0 \cdot \mathbf{x}_2\| \leq \sqrt{m} \|\mathbf{S}_0\| \|\mathbf{x}_2\| \leq m \|\widetilde{\mathbf{S}}_0\| \|\mathbf{x}_2\|$, which is itself $\mathcal{O}(\sigma m^{5/2} \log m \sqrt{n \log q})$ with probability $\geq 1 - 2^{-\Omega(n)}$, by Lemma 1. As q has been set sufficiently large, we obtain that $\mathbf{S}_0 \cdot \mathbf{x}_2$ is known over the integers: Multiplying by \mathbf{S}_0^{-1} over the rationals allows the group manager to recover \mathbf{x}_2 . The argument is similar for the other \mathbf{c}_i 's, except that $\|\mathbf{S}_i \cdot \mathbf{c}_i \pmod q\| \leq \mathcal{O}(p \alpha q m^{5/2} \log m \sqrt{n \log q})$. Again, α has been set sufficiently small to allow the group manager to recover $p \cdot \mathbf{e}_i + \text{id}_j[i] \cdot \mathbf{x}_2$.

Finally, the total bit-size of all proofs is $\mathcal{O}(\ell t m \log q)$. The same bound holds for the ciphertexts. \square

C Proof of Lemma 5

The algorithm is a simple extension of the one in [23]. It first partitions \mathbf{B} into matrices $\mathbf{B}_1 \in \mathbb{Z}_q^{m_1 \times n}$ and $\mathbf{B}_2 \in \mathbb{Z}_q^{n \times n}$, with $m_1 = m - n$, such that \mathbf{B}_2 is invertible over \mathbb{Z}_q and $\mathbf{B}^T = [\mathbf{B}_1^T | \mathbf{B}_2^T]$. Such a partition can always be found by re-arranging the rows of \mathbf{B} if necessary. The execution of `GenSuperSamp`($1^n, 1^m, q, \mathbf{B}, \mathbf{C}$) then proceeds with the following steps.

1. Generate $(\mathbf{A}_1, \mathbf{T}_1) \leftarrow \text{TrapGen}(1^n, 1^{m_1}, q)$. Return \perp if the rows of $\mathbf{A}_1 \in \mathbb{Z}_q^{m_1 \times n}$ do not span \mathbb{Z}_q^n .
2. Compute $\mathbf{A}_2 = \mathbf{B}_2^{-T} \cdot (\mathbf{C} - \mathbf{B}_1^T \cdot \mathbf{A}_1) \pmod q$. Note that $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}$ satisfies $\mathbf{B}^T \cdot \mathbf{A} = \mathbf{C} \pmod q$.
3. Extend $\mathbf{T}_1 \in \mathbb{Z}^{m_1 \times m_1}$ to have a basis $\mathbf{T} \in \mathbb{Z}^{m \times m}$ for \mathbf{A} using the basis delegation algorithm from Lemma 6. Then, re-randomize \mathbf{T}' to obtain \mathbf{T}'' using the basis randomization algorithm `RandBasis`.

The rest of the proof is exactly identical to the proof of Lemma 4 in [23]. \square

D Security Proofs for the Fully Anonymous Construction

D.1 Proof of Theorem 4 (full anonymity)

We now prove the full anonymity of the scheme in an attack game which is exactly the one of Definition 3 with the difference that the adversary is granted access to a signature opening oracle. Namely, before and after the challenge phase, the latter oracle can be invoked for adversarially-chosen signatures as long as these do not coincide with the challenge signature Σ^* . The proof of Theorem 4 relies on the all-but-one simulation technique [8] in the same way as in the Agrawal-Boneh-Boyen IBE [1].

Proof. Like the proof of Theorem 2, the proof proceeds via a sequence of hybrid experiments. For each i , we define W_i to be the event that experiment $G_i^{(b)}$ outputs 1.

Experiment $G_0^{(b)}$. This experiment is the real attack game. Namely, the challenger performs the setup of the system by following the specification of the Keygen algorithm. The adversary \mathcal{A} is given \mathbf{gpk} and $\{\mathbf{gsk}[j]\}_{j=0}^{N-1}$ at the beginning of the game. All opening queries are answered faithfully, by returning the uncovered identity $\text{id} \in \{0, 1\}^\ell$. At the challenge phase, the adversary chooses a message M as well as indexes $j_0, j_1 \in \{0, \dots, N-1\}$ and obtains a challenge $\Sigma^* = (\mathbf{VK}^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \dots, \mathbf{c}_\ell^*, \pi_0^*, \pi_{\text{OR},1}^*, \dots, \pi_{\text{OR},\ell}^*, \pi_K^*, \text{sig}^*) \leftarrow \text{Sign}(\mathbf{gpk}, \mathbf{gsk}[j_b], M)$. The experiment ends with the adversary \mathcal{A} outputting a bit $b' \in \{0, 1\}$. At this point, the experiment returns 1 if $b' = b$ and 0 otherwise. The probability $\Pr[W_0]$ is thus the probability to have $b' = b$.

Experiment $G_1^{(b)}$. We make a simple conceptual change to the generation of the challenge signature Σ^* . Namely, instead of sampling $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T \in \mathbb{Z}^{3m}$ in $\Lambda^\perp(\mathbf{A}_{\text{id}})$, Experiment $G_1^{(b)}$ first samples $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^{2m}, \sigma}$ and uses the trapdoor \mathbf{T}_A to compute $\mathbf{x}_1 \in D_{\mathbb{Z}^{m}, \sigma}$ such that $(\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \mathbf{A}_{\text{id}} = 0 \pmod q$. This change is purely conceptual since the vector $(\mathbf{x}_1^T | \mathbf{x}_2^T)^T$ has the same distribution either way. Clearly, it holds that $\Pr[W_1] = \Pr[W_2]$.

Experiment $G_2^{(b)}$. We introduce a slight modification w.r.t. Experiment $G_1^{(b)}$. At the outset of the game, the challenger generates a one-time signature key pair $(\mathbf{VK}^*, \mathbf{SK}^*) \leftarrow \mathcal{G}(1^n)$. If \mathcal{A} queries the opening oracle with a valid signature $\Sigma = (\mathbf{VK}, \mathbf{c}_0, \mathbf{c}_1, \dots, \mathbf{c}_\ell, \pi_0, \pi_{\text{OR},1}, \dots, \pi_{\text{OR},\ell}, \pi_K, \text{sig})$ such that $\mathbf{VK} = \mathbf{VK}^*$, the experiment halts and outputs a random bit. The assumed strong security of the one-time signature implies that Experiment $G_2^{(b)}$ cannot depart from Experiment $G_1^{(b)}$. Indeed, if a valid opening query is made after the challenge phase, the adversary is able to break the strong unforgeability of the one-time signature (the proof is straightforward and omitted). Moreover, before the challenge phase, the one-time verification key \mathbf{VK}^* is independent of \mathcal{A} 's view. As long as no one-time verification key is produced by the one-time key generation algorithm with too high probability (which is implied by the strong unforgeability property), the chance of \mathbf{VK}^* to show up in a valid pre-challenge opening query is negligible. There thus exists a PPT forger \mathcal{B}^{ots} against the one-time signature for which $|\Pr[W_2] - \Pr[W_1]| \leq \mathbf{Adv}^{\text{suf-ots}}(\mathcal{B}^{\text{ots}})$. In the following, we henceforth assume that no opening query involves \mathbf{VK}^* .

Experiment $G_3^{(b)}$. We bring a first modification to the generation of the group public key \mathbf{gpk} in the setup phase. Namely, for each $i \in \{0, \dots, \ell\}$, the experiment first runs $(\mathbf{B}_{i,1}, \mathbf{T}_{i,1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ to obtain a matrix $\mathbf{B}_{i,1} \in \mathbb{Z}_q^{m \times n}$ with a short basis $\mathbf{T}_{i,1} \in \mathbb{Z}^{m \times m}$. Note that the distribution of $\mathbf{B}_{i,1}$ is statistically close to the uniform distribution over $\mathbb{Z}_q^{m \times n}$. Next, the experiment sets $\mathbf{B}_{i,0} = \mathbf{R}_i \cdot \mathbf{B}_{i,1} - \mathbf{B}_{i,1} \cdot H_{vk}(\mathbf{VK}^*)$, where $\mathbf{R}_i \in \mathbb{Z}^{m \times m}$ is a matrix whose rows are vectors sampled from the distribution $D_{\mathbb{Z}^m, \sigma}$. The result of [22, Lemma 5.2] implies that matrices $\{\mathbf{B}_{i,0}\}_{i=0}^\ell$ will be statistically close to the uniformly distributed matrices produced by the real key generation algorithm. We can write $|\Pr[W_3] - \Pr[W_2]| \in \text{negl}(1^n)$.

Experiment $G_4^{(b)}$. In this experiment, we modify the signature opening oracle in the following way. Recall that, due to the modification introduced in Experiment $G_2^{(b)}$, each opening query involves a signature $\Sigma = (\mathbf{VK}, \mathbf{c}_0, \mathbf{c}_1, \dots, \mathbf{c}_\ell, \pi_0, \pi_{\text{OR},1}, \dots, \pi_{\text{OR},\ell}, \pi_K, \text{sig})$ for which $\mathbf{VK} \neq \mathbf{VK}^*$ unless the one-time signature is not strongly unforgeable. For this reason, each matrix $\mathbf{B}_{i, \mathbf{VK}}$ can be written as

$$\mathbf{B}_{i, \mathbf{VK}} = \begin{bmatrix} \mathbf{B}_{i,1} \\ \mathbf{B}_{i,0} + \mathbf{B}_{i,1} H_{vk}(\mathbf{VK}) \end{bmatrix} = \begin{bmatrix} \mathbf{B}_{i,1} \\ \mathbf{R}_i \cdot \mathbf{B}_{i,1} + \mathbf{B}_{i,1} \cdot (H_{vk}(\mathbf{VK}) - H_{vk}(\mathbf{VK}^*)) \end{bmatrix},$$

where $H_{vk}(\mathbf{VK}) - H_{vk}(\mathbf{VK}^*)$ is a non-singular $n \times n$ matrix over \mathbb{Z}_q . This implies that the trapdoor $\mathbf{T}_{i,1} \in \mathbb{Z}^{m \times m}$ of $\mathbf{B}_{i,1}$ – which was defined in Experiment $G_3^{(b)}$ – can be used to generate a short basis for the lattice $\Lambda^\perp(\mathbf{B}_{i, \mathbf{VK}})$ as in step 2 of the SampleRight algorithm of [1, Section 4.2]. The obtained short basis $\mathbf{T}_{i, \mathbf{VK}} \in \mathbb{Z}^{2m \times 2m}$ satisfies $\mathbf{T}_{i, \mathbf{VK}} \in \mathbb{Z}^{2m \times 2m} \cdot \mathbf{B}_{i, \mathbf{VK}} = 0 \pmod q$ and it can be used exactly in the same way as the delegated bases $\mathbf{S}_{i, \mathbf{VK}}$ of the actual opening algorithm to identify the signer. This modification is thus purely conceptual and we thus have $\Pr[W_4] = \Pr[W_3]$. We remark that, in this experiment, the trapdoors $\{\mathbf{S}_i\}_{i=0}^\ell$ of matrices $\{\mathbf{B}_{i,1}\}_{i=0}^\ell$ are not used any longer.

Experiment $G_5^{(b)}$. This experiment is identical to Experiment $G_4^{(b)}$ but we slightly modify the setup phase in step c of the key generation algorithm. Recall that Experiment $G_4^{(b)}$ generates $(\mathbf{B}_{i,1}, \mathbf{S}'_i) \leftarrow$

SuperSamp($1^n, 1^m, q, \mathbf{A}_{i,1}, -\mathbf{A}_{i,2}^T \cdot \mathbf{B}_{i,0}$) so as to obtain a matrix $\mathbf{B}_{i,-1} \in \mathbb{Z}_q^{m \times n}$ satisfying the equality

$$\mathbf{B}_{i,-1}^T \cdot \mathbf{A}_{i,1} + \mathbf{B}_{i,0}^T \cdot \mathbf{A}_{i,2} = 0 \pmod{q} \quad (7)$$

at step c of Keygen. In contrast, Experiment $G_5^{(b)}$ proceeds by generating $(\mathbf{B}_{i,1}, \mathbf{T}_{i,1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ and choosing $\mathbf{B}_{i,-1}, \mathbf{B}_{i,0}$ uniformly in $\mathbb{Z}_q^{m \times n}$. Then, it generates

$$(\mathbf{A}_{i,1}, \mathbf{T}'_i) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{B}_{i,-1}, -\mathbf{B}_{i,0}^T \cdot \mathbf{A}_{i,2}),$$

which satisfies (7). The same arguments as in [23, Lemma 5] imply that $\{\mathbf{B}_{i,-1}, \mathbf{B}_{i,0}, \mathbf{B}_{i,1}, \mathbf{A}_i\}_{i=0}^\ell$ have a distribution which is negligibly far apart from their distribution in Experiment $G_4^{(b)}$.

The setup phase is completed by using \mathbf{T}_A to compute group member's private keys $\{\text{gsk}[j]\}_{j=0}^{N-1}$. Since \mathcal{A} 's view is not noticeably affected by this modification, we have $|\Pr[W_5] - \Pr[W_4]| \in \text{negl}(1^n)$.

Experiment $G_6^{(b)}$. Here, we modify the generation of the challenge signature Σ^* as follows. At step 5 of the signing algorithm, instead of computing the NIZK proofs $\{\pi_{\text{OR},i}^*\}_{i=1}^\ell$ using the actual witnesses, the experiment generates a simulated non-interactive proof by programming the random oracle. The statistical zero-knowledge property of the Micciancio-Vadhan proof system [38] guarantees that the distribution of $\{\pi_{\text{OR},i}^*\}_{i=1}^\ell$ remains statistically unchanged (note that $\{\pi_{\text{OR},i}^*\}_{i=1}^\ell$ are simulated proofs for true statements). Therefore it comes that $|\Pr[W_6] - \Pr[W_5]| \in \text{negl}(1^n)$. Note that $\text{negl}(1^n)$ incorporates the small probability that the NIZK simulator fails because it accidentally has to program the random oracle on an input where it was previously defined.

Experiment $G_7^{(b)}$. In this experiment, we bring a new modification to the generation of Σ^* . The real proof of knowledge π_K^* is replaced by a simulated proof which is obtained by programming the random oracle H at step 6 of the signing algorithm. Similarly to the previous transition, we can write $|\Pr[W_7] - \Pr[W_6]| \in \text{negl}(1^n)$, where $\text{negl}(1^n)$ encompasses the tiny probability that the NIZK simulator fails.

Experiment $G_8^{(b)}$. We introduce yet another change in the generation of Σ^* . For each $i \in \{1, \dots, \ell\}$, instead of computing $\mathbf{c}_i^* = \mathbf{B}_{i, \text{VK}^*} \cdot \mathbf{s} + p \cdot \mathbf{e}_i + \text{id}[j_b] \cdot \mathbf{x}_2$, where $\mathbf{x}_2 \in \mathbb{Z}^{2m}$ is the vector encrypted by \mathbf{c}_0 , the experiment sets $\mathbf{c}_i^* = \mathbf{z}_i + \text{id}[j_b] \cdot \mathbf{x}_2$ for a randomly drawn $\mathbf{z}_i \leftarrow U(\mathbb{Z}_q^{2m})$. Under the $\text{LWE}_{q,\alpha}$ assumption, we argue that this change should not significantly affect \mathcal{A} 's view. Concretely, assuming that an adversary can distinguish Experiment $G_8^{(b)}$ from Experiment $G_7^{(b)}$, we can build a distinguisher \mathcal{B}^{lwe} for the $\text{LWE}_{q,\alpha}$. The latter distinguisher is described in the proof of Lemma 11 for completeness. For this reason, we find $|\Pr[W_8] - \Pr[W_7]| \leq \text{Adv}^{\text{LWE}_{q,\alpha}}(\mathcal{B}^{\text{lwe}})$.

Experiment $G_9^{(b)}$. As a final change in the generation of Σ^* , we choose \mathbf{c}_i^* at random in $U(\mathbb{Z}_q^{2m})$ for $i = 1$ to ℓ . This is just a conceptual change since $\{\mathbf{c}_i^*\}_{i=1}^\ell$ have exactly the same distribution as in Experiment $G_8^{(b)}$. This implies $\Pr[W_9] = \Pr[W_8]$. Moreover, in Experiment $G_9^{(b)}$, it is obvious that $\Pr[W_9] = 1/2$ since Σ^* is completely independent of the random bit $b \in_R \{0, 1\}$.

To conclude the proof, we prove the indistinguishability of Experiment $G_8^{(b)}$ and Experiment $G_7^{(b)}$.

Lemma 11. *Under the $\text{LWE}_{q,\alpha}$ assumption, no PPT adversary can distinguish Experiment $G_8^{(b)}$ and Experiment $G_7^{(b)}$.*

Proof. Towards a contradiction, suppose that an adversary \mathcal{A} can tell the two experiments apart with non-negligible advantage. We build the following LWE distinguisher \mathcal{B}^{lwe} . It takes as input a $\text{LWE}_{q,\alpha}$ instance $\{(\mathbf{B}'_i, \mathbf{z}_i)\}_{i=1}^\ell$, where $\mathbf{B}'_i \in \mathbb{Z}_q^{m \times n}$ and $\mathbf{z}_i \in \mathbb{Z}_q^m$ for each $i \in \{1, \dots, \ell\}$. Each component \mathbf{z}_i is either uniform in \mathbb{Z}_q^m or of the form $\mathbf{z}_i = \mathbf{B}'_i \cdot \mathbf{s} + \mathbf{e}_i$, where \mathbf{e}_i is sampled from $D_{\mathbb{Z}^m, \alpha q}$.

In order to prepare the group public key gpk , algorithm \mathcal{B}^{lwe} defines $\mathbf{B}_{i,-1} = \mathbf{B}'_i$ for $i = 1$ to ℓ . For each $i \in \{1, \dots, \ell\}$, it also generates $\mathbf{B}_{i,1}$ by running $(\mathbf{B}_{i,1}, \mathbf{T}_{i,1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ and also sets $\mathbf{B}_{i,0} = \mathbf{R}_i \cdot \mathbf{B}_{i,-1} - \mathbf{B}_{i,1} H_{\text{vk}}(\text{VK}^*)$ as in Experiment $G_7^{(b)}$. By doing so, \mathcal{B}^{lwe} is able to answer all signature opening queries using the trapdoor $\mathbf{T}_{i,1}$ of $\mathbf{B}_{i,1}$ unless the failure event introduced in Experiment $G_2^{(b)}$ occurs.

During the challenge phase, \mathcal{B}^{lwe} samples \mathbf{x}_2 in $D_{\mathbb{Z}^{2m}, \sigma}$ and defines

$$\mathbf{c}_i^* = \left[\frac{p \cdot \mathbf{z}_i}{\mathbf{R}_i \cdot (p \cdot \mathbf{z}_i)} \right] + \text{id}[j_b] \cdot \mathbf{x}_2, \text{ for } i \in \{1, \dots, \ell\},$$

while \mathbf{c}_0^* is obtained by faithfully encrypting \mathbf{x}_2 . The proof π_0^* is generated as a real proof whereas $\{\pi_{\text{OR},i}^*\}_{i=1}^\ell$ and π_K^* are obtained from their respective NIZK simulators.

After the challenge phase, \mathcal{A} is granted further access to the opening oracle and its opening queries are handled as in the first phase. At the end of the experiment, \mathcal{A} outputs a random bit b' and \mathcal{B}^{lwe} outputs 1 if and only if $b' = b$.

We note that each $\mathbf{B}_{i,\text{VK}^*}$ is such that $\mathbf{B}_{i,\text{VK}^*} = \begin{bmatrix} \mathbf{B}_{i,-1} \\ \mathbf{R}_i \cdot \mathbf{B}_{i,-1} \end{bmatrix}$ for $i = 1$ to ℓ . If each \mathbf{z}_i is such that $\mathbf{z}_i = \mathbf{B}'_i \cdot \mathbf{s} + \mathbf{e}_i$, where $\mathbf{e}_i \in D_{\mathbb{Z}^{m\ell}, \alpha q}$, then $\{\mathbf{c}_i^*\}_{i=1}^\ell$ are distributed as in Experiment $G_7^{(b)}$. Indeed, the matrices $\{\mathbf{R}_i\}_{i=1}^\ell$ introduced in Experiment $G_3^{(b)}$ are statistically independent of \mathcal{A} 's view until the challenge phase because the product $\mathbf{R}_i \cdot \mathbf{B}_{i,-1}$ is statistically close to the uniform distribution over $\mathbb{Z}_q^{m \times n}$. In this case, the reduction \mathcal{B}^{lwe} is running Experiment $G_7^{(b)}$ with \mathcal{A} . Now, if each \mathbf{z}_i is uniform in \mathbb{Z}_q^m , we are clearly in Experiment $G_8^{(b)}$. \square

D.2 Proof of Theorem 5 (full traceability)

The traceability property is proved in the same way as in the proof of Theorem 3.

Proof. For the sake of contradiction, let us assume that a full traceability adversary \mathcal{A} has non-negligible success probability ε in the model of Definition 4. In the random oracle model, we build an algorithm \mathcal{B} that solves a given $\text{SIS}_{q,\beta}$ instance with non-negligible probability. Algorithm \mathcal{B} receives as input a matrix $\hat{\mathbf{A}} \in \mathbb{Z}_q^{2m \times n}$ and has to find a vector $\mathbf{v} \in \mathbb{Z}^{2m}$ in $\Lambda_q^\perp(\hat{\mathbf{A}})$ such that $0 < \|\mathbf{v}\| \leq \beta$. Let $\bar{\mathbf{A}} \in \mathbb{Z}_q^{m \times n}$ be the matrix consisting of the first m rows of $\hat{\mathbf{A}}$.

Initialization. As in the proof of Theorem 3, algorithm \mathcal{B} first flips a fair coin $\text{coin} \leftarrow U(\{0,1\})$ that will determine its strategy and the way to set up the group public key. If $\text{coin} = 0$, \mathcal{B} will try to find a non-zero short vector of $\Lambda^\perp(\bar{\mathbf{A}})$ and pad it with zeroes to obtain a short non-zero vector in $\Lambda^\perp(\hat{\mathbf{A}})$. If $\text{coin} = 1$, \mathcal{B} will embed the entire input matrix $\hat{\mathbf{A}}$ in one of the $\{\mathbf{A}_i\}_{i=1}^\ell$.

• If $\text{coin} = 0$, \mathcal{B} first runs $\text{TrapGen}(1^n, 1^{2m}, q)$ to generate $\mathbf{C} \in \mathbb{Z}_q^{2m \times n}$ with a basis $\mathbf{T}_\mathbf{C} \in \mathbb{Z}^{2m \times 2m}$ of $\Lambda_q^\perp(\mathbf{C})$ with $\|\mathbf{T}_\mathbf{C}\| \leq \mathcal{O}(\sqrt{n \log q})$. Next, \mathcal{B} samples a collection of $\ell + 1$ matrices $\mathbf{Q}_0, \dots, \mathbf{Q}_\ell \in \mathbb{Z}^{2m \times m}$, where each matrix entry sampled independently in $D_{\mathbb{Z}, \omega(\sqrt{\log n})}$. Then, \mathcal{B} draws $j^* \leftarrow U([0, N - 1])$, hoping that user j^* will be the one whose identity $\text{id}_{j^*} = \text{id}_{j^*}[1] \dots \text{id}_{j^*}[\ell] \in \{0,1\}^\ell$ will be uncovered by the opening algorithm for \mathcal{A} 's forgery at the end of the game. Also, \mathcal{B} defines $\mathbf{A}_0 = \mathbf{Q}_0 \cdot \bar{\mathbf{A}} + (\sum_{i=1}^\ell \text{id}_{j^*}[i]) \cdot \mathbf{C}$ and $\mathbf{A}_i = \mathbf{Q}_i \cdot \bar{\mathbf{A}} + (-1)^{\text{id}_{j^*}[i]} \cdot \mathbf{C}$ for each $i \in [1, \ell]$. It also sets $\mathbf{A} = \bar{\mathbf{A}}$.

Then, for each $i \in \{0, \dots, \ell\}$, \mathcal{B} chooses $\mathbf{B}_{i,0} \leftarrow U(\mathbb{Z}_q^{m \times n})$ and parses the matrix $\mathbf{A}_i \in \mathbb{Z}_q^{2m \times n}$ as $\mathbf{A}_i^T = [\mathbf{A}_{i,1}^T \mid \mathbf{A}_{i,2}^T]$, where $\mathbf{A}_{i,1}, \mathbf{A}_{i,2} \in \mathbb{Z}_q^{m \times n}$. It runs $(\mathbf{B}_{i,1}, \mathbf{T}_{\mathbf{B}_{i,1}}) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{A}_{i,2}, \mathbf{0})$ to obtain a matrix $\mathbf{B}_{i,1} \in \mathbb{Z}_q^{m \times n}$ such that $\mathbf{A}_{i,2}^T \cdot \mathbf{B}_{i,1} = 0 \pmod q$. It erases $\mathbf{T}_{\mathbf{B}_{i,1}}$, that will not be needed, and generates $(\mathbf{B}_{i,-1}, \mathbf{S}'_i) \leftarrow \text{SuperSamp}(1^n, 1^m, q, \mathbf{A}_{i,1}, -\mathbf{B}_{i,0}^T \cdot \mathbf{A}_{i,2})$ which will satisfy

$$\mathbf{B}_{i,-1}^T \cdot \mathbf{A}_{i,1} + \mathbf{B}_{i,0}^T \cdot \mathbf{A}_{i,2} = 0 \pmod q,$$

as desired. Finally, \mathcal{B} re-randomizes each \mathbf{S}'_i as $\mathbf{S}_i \leftarrow \text{RandBasis}(\mathbf{S}'_i)$ for $i = 0$ to ℓ . We observe that \mathcal{B} notably departs from the real key generation algorithm in that $\mathbf{B}_{i,1}$ is generated from $\mathbf{A}_{i,2}$ (whereas Keygen proceeds the other way around at step 2) using SuperSamp . However, by Lemma 4 in [23], the distribution of the resulting matrices is statistically the same either way.

The group public key $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, (\mathbf{B}_{i,-1}, \mathbf{B}_{i,0}, \mathbf{B}_{i,1})\}_{i=0}^\ell)$ is finally given to \mathcal{A} . As in the proof of Theorem 3, for each $j \neq j^*$, we have

$$\mathbf{A}_{\text{id}_j} = \begin{bmatrix} \bar{\mathbf{A}} \\ \mathbf{A}_0 + \sum_{i=1}^\ell \text{id}_j[i] \mathbf{A}_i \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{A}} \\ (\mathbf{Q}_0 + \sum_{i=1}^\ell \text{id}_j[i] \mathbf{Q}_i) \cdot \bar{\mathbf{A}} + h_{\text{id}_j} \cdot \mathbf{C} \end{bmatrix} \in \mathbb{Z}_q^{2m \times n},$$

where $h_{\text{id}_j} \in [1, \ell]$ denotes the Hamming distance between id_j and id_{j^*} . As in the proof of Theorem 3, for each identifier $\text{id}_j \neq \text{id}_{j^*}$, \mathcal{B} is able to compute a basis $\mathbf{T}'_{\text{id}_j}$ of $\Lambda_q^\perp(\mathbf{A}_{\text{id}_j})$ with $\|\mathbf{T}'_{\text{id}_j}\| \leq \omega(\sqrt{2mn \log q \log n})$

from the basis $\mathbf{T}_{\mathbf{C}}$ of $\Lambda_q^\perp(\mathbf{C})$. The obtained bases $\{\mathbf{T}'_{\text{id}_j}\}_{\text{id}_j \neq \text{id}_{j^*}}$ are then re-randomized as $\mathbf{T}_{\text{id}_j} \leftarrow \text{RandBasis}(\mathbf{T}'_{\text{id}_j}, \omega(\sqrt{2mn \log q \log n}))$. However, the reduction \mathcal{B} is unable to compute a trapdoor for the matrix $\mathbf{A}_{\text{id}_{j^*}}$ corresponding to the expected target group member j^* . Fortunately, \mathcal{B} can derive a trapdoor \mathbf{T}_{id_j} for each $j \neq j^*$.

Since the rows of each \mathbf{Q}_k are sampled from $D_{\mathbb{Z}^m, \omega(\sqrt{\log n})}$, the matrices $\mathbf{A}_0, \dots, \mathbf{A}_\ell \in \mathbb{Z}_q^{2m \times n}$ have a distribution which is statistically close to that of independent and uniformly random matrices over $\mathbb{Z}_q^{2m \times n}$, which are also statistically independent of \mathbf{A} . Also, by Lemma 3, the distribution of $\{\mathbf{T}_{\text{id}_j}\}_{j \neq j^*}$ is statistically close to that of the real system.

- If $\text{coin} = 1$, the reduction \mathcal{B} chooses $i^* \leftarrow U([1, \ell])$ and defines $\hat{\mathbf{A}}$ to be the matrix $\mathbf{A}_{i^*} \in \mathbb{Z}_q^{2m \times n}$ that will be part of gpk . It runs $\text{TrapGen}(1^n, 1^m, q)$ to obtain $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ with a basis $\mathbf{T}_{\mathbf{A}}$ of $\Lambda_q^\perp(\mathbf{A})$ such that $\|\widetilde{\mathbf{T}_{\mathbf{A}}}\| \leq \mathcal{O}(\sqrt{n \log q})$. Next, it independently samples $\mathbf{A}_0, \dots, \mathbf{A}_{i^*-1}, \mathbf{A}_{i^*+1}, \dots, \mathbf{A}_\ell \leftarrow U(\mathbb{Z}_q^{2m \times n})$ and sets $\mathbf{A}_{i^*} = \hat{\mathbf{A}}$. Finally, \mathcal{B} computes $\{(\mathbf{B}_{i,-1}, \mathbf{B}_{i,0}, \mathbf{B}_{i,1})\}_{i=0}^\ell$ in the same way as in the case $\text{coin} = 0$. As in the previous case, \mathcal{B} thus knows a trapdoor \mathbf{S}_i for $\mathbf{B}_{i,-1}$ for each $i \in \{0, \dots, \ell\}$. The group public key $\text{gpk} = (\mathbf{A}, \{\mathbf{A}_i, (\mathbf{B}_{i,-1}, \mathbf{B}_{i,0}, \mathbf{B}_{i,1})\}_{i=0}^\ell)$, which is distributed (statistically) as in the real system, is given as input to \mathcal{A} . Using $\mathbf{T}_{\mathbf{A}}$, the reduction \mathcal{B} is able to compute a delegated basis \mathbf{T}_{id_j} for all users $j \in [0, N-1]$ exactly as in the real scheme.

Regardless of the value of $\text{coin} \in \{0, 1\}$, the adversary \mathcal{A} is run on input of $\text{gmsk} := \{\mathbf{S}_i\}_{i=0}^\ell$ and $\text{gpk} := (\mathbf{A}, \{\mathbf{A}_i, (\mathbf{B}_{i,-1}, \mathbf{B}_{i,0}, \mathbf{B}_{i,1})\}_{i=0}^\ell, H, \Pi^{\text{ots}}, p)$.

Queries. Algorithm \mathcal{B} starts interacting with adversary \mathcal{A} whose queries are handled in a way that depends on $\text{coin} \in \{0, 1\}$.

- If $\text{coin} = 0$, \mathcal{B} aborts if \mathcal{A} ever queries the private key $\text{gsk}[j^*]$ of user j^* . When \mathcal{A} queries a private key $\text{gsk}[j]$ for $j \in \{0, \dots, N-1\} \setminus \{j^*\}$, \mathcal{B} reveals the previously computed short basis \mathbf{T}_{id_j} . When \mathcal{A} queries the signing oracle, \mathcal{B} faithfully runs the signing algorithm whenever the involved user j is not j^* . For each signing query involving the expected target user j^* , \mathcal{B} samples $\mathbf{x}_2 \leftarrow D_{\mathbb{Z}^{2m}, \sigma}$ and $\mathbf{s}_0, \mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$. Then, it computes $\mathbf{c}_0 = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2$ as well as $\mathbf{c}_i = \mathbf{B}_{i, \text{VK}} \cdot \mathbf{s} + p \cdot [e_i | e_i \cdot \mathbf{R}_i] + \text{id}_{j^*}[i^*] \cdot \mathbf{x}_2$ for each $i \in [1, \ell]$. The proof π_0 is computed as a real proof (*i.e.*, using the witness \mathbf{x}_2), whereas all other non-interactive proofs $\{\pi_{\text{OR}, i}\}_{i=1}^\ell$ and π_K are simulated using the appropriate NIZK simulator, by programming the random oracle. Since the simulator is statistically zero-knowledge property, the resulting signature Σ will be statistically indistinguishable from a real signature.

- If $\text{coin} = 1$, \mathcal{B} has all private keys $\{\text{gsk}[j]\}_{j=0}^{N-1}$ at disposal since it knows $\mathbf{T}_{\mathbf{A}}$. It can thus perfectly answer \mathcal{A} 's queries by running the actual signing algorithm or returning the queried private keys $\text{gsk}[j]$.

For each $\text{coin} \in \{0, 1\}$, queries to the random oracle H are handled by returning a uniformly chosen value in $\{0, 1\}^t$. For each $\kappa \in \{1, \dots, q_H\}$, r_κ will stand for the answer to the κ -th H -query. As usual, if a given random oracle query occurs more than once, \mathcal{B} responds by returning the previously defined value.

Forgery. Eventually, \mathcal{A} outputs a signature $\Sigma^* = (\mathbf{c}_0^*, \mathbf{c}_1^*, \dots, \mathbf{c}_\ell^*, \pi_0^*, \pi_{\text{OR}, 1}^*, \dots, \pi_{\text{OR}, \ell}^*, \pi_K^*)$ on some message M^* with probability ε . If we parse the proof of knowledge π_K^* as $(\text{Comm}_K^*, \text{Chall}_K^*, \text{Resp}_K^*)$, w.h.p., \mathcal{A} must have queried H on the input $(M^*, \text{Comm}_K^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \dots, \mathbf{c}_\ell^*, \pi_0^*, \pi_{\text{OR}, 1}^*, \dots, \pi_{\text{OR}, \ell}^*)$. Indeed, otherwise, the probability to have $\text{Chall}_K^* = H(M^*, \text{Comm}_K^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \dots, \mathbf{c}_\ell^*, \pi_0^*, \pi_{\text{OR}, 1}^*, \dots, \pi_{\text{OR}, \ell}^*)$ is at most $1/2^t$. With probability $\varepsilon - 1/2^t$, the tuple $(M^*, \text{Comm}_K^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \dots, \mathbf{c}_\ell^*, \pi_0^*, \pi_{\text{OR}, 1}^*, \dots, \pi_{\text{OR}, \ell}^*)$ was the input of the κ -th random oracle query for some $\kappa^* \in \{1, \dots, q_H\}$.

Now, \mathcal{B} starts a second execution of the adversary \mathcal{A} with the *same* random tape and input as in the first run. All queries are answered as in the latter with a difference in the treatment of random oracle queries. Namely, the first $\kappa^* - 1$ hash queries – which are necessarily the same as in the first execution because \mathcal{A} 's random tape has not changed – receive the same answers $r_1, \dots, r_{\kappa^*-1}$ as in the first run. Consequently, the κ^* -th query will involve the tuple $(M^*, \text{Comm}_K^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \dots, \mathbf{c}_\ell^*, \pi_0^*, \pi_{\text{OR}, 1}^*, \dots, \pi_{\text{OR}, \ell}^*)$ as in the first execution. However, a forking occurs as, from this point forward, \mathcal{A} obtains fresh random oracle values $r'_{\kappa^*}, \dots, r'_{q_H}$ which are independent of the subsequence of answers in the first execution. The General Forking Lemma of Bellare and Neven [6] implies that, with probability at least $(\varepsilon - \frac{1}{2^t}) \left(\frac{\varepsilon - 1/2^t}{q_H} - \frac{1}{2^t} \right)$, it holds that: (1) \mathcal{A} 's forgery also pertains to $(M^*, \text{Comm}_K^*, \mathbf{c}_0^*, \mathbf{c}_1^*, \dots, \mathbf{c}_\ell^*, \pi_0^*, \pi_{\text{OR}, 1}^*, \dots, \pi_{\text{OR}, \ell}^*)$ in the second

run; (2) we also have $r'_{\kappa^*} \neq r_{\kappa^*}$. Hence, using the knowledge extractor of the proof of knowledge $\pi_{K^*}^*$, \mathcal{B} extracts vectors $\mathbf{e}_1, \dots, \mathbf{e}_\ell \in D_{\mathbb{Z}^{2m}, \alpha q}$ and $\mathbf{x}_1 \in \mathbb{Z}^m, \mathbf{y}_1, \dots, \mathbf{y}_\ell \in \mathbb{Z}^{2m}$ satisfying

$$\mathbf{x}_1^T \mathbf{A} + \sum_{i=0}^{\ell} \mathbf{c}_i^T \mathbf{A}_i = \sum_{i=1}^{\ell} \mathbf{e}_i^T (p \cdot \mathbf{A}_i) \quad \text{and} \quad \mathbf{e}_i^T (p \cdot \mathbf{A}_i) + \mathbf{y}_i^T \mathbf{A}_i = \mathbf{c}_i^T \mathbf{A}_i, \text{ for } i \in \{1, \dots, \ell\} \quad (8)$$

with $\|\mathbf{x}_1\| \leq \sigma\sqrt{m}$ and $\|\mathbf{y}_i\| \leq \sigma\sqrt{2m}$ for each $i \in \{1, \dots, \ell\}$.

The reduction \mathcal{B} then opens either of the two forgeries using $\{\mathbf{S}_i\}_{i=0}^{\ell}$ (note that both signatures necessarily open to the same identity id as they involve the same $\{\mathbf{c}_i^*\}_{i=1}^{\ell}$). At this point, \mathcal{B} aborts and declares failure if the opening does not unveil user j^* 's identity. Still, with probability at least $1/N$, \mathcal{B} 's was fortunate in its random choice for j^* and the opening algorithm reveals id_{j^*} .

If this desirable event occurs, \mathcal{B} considers the following situations.

- If $\mathbf{y}_i = \text{id}_{j^*}[i] \cdot \mathbf{x}_2$ for each $i \in \{1, \dots, \ell\}$, where $\mathbf{x}_2 \in \mathbb{Z}^{2m}$ is the vector encrypted by \mathbf{c}_0^* , \mathcal{B} aborts if $\text{coin} = 1$. Otherwise, relations (8) guarantee that

$$\begin{aligned} \mathbf{x}_1^T \cdot \mathbf{A} + \mathbf{c}_0^T \cdot \mathbf{A}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \cdot \mathbf{x}_2^T \cdot \mathbf{A}_i &= \mathbf{x}_1^T \cdot \mathbf{A} + \mathbf{x}_2^T \cdot \mathbf{A}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \cdot \mathbf{x}_2^T \cdot \mathbf{A}_i \\ &= (\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \left[\frac{\mathbf{A}}{\mathbf{A}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \cdot \mathbf{A}_i} \right] \\ &= (\mathbf{x}_1^T | \mathbf{x}_2^T) \cdot \left[\frac{\bar{\mathbf{A}}}{(\mathbf{Q}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \mathbf{Q}_i) \cdot \bar{\mathbf{A}}} \right] = 0 \pmod{q}, \end{aligned}$$

where the first equality follows from the fact that $\mathbf{B}_0^T \cdot \mathbf{A}_0 = 0 \pmod{q}$ and \mathbf{c}_0^* is of the form $\mathbf{c}_0^* = \mathbf{B}_0 \cdot \mathbf{s}_0 + \mathbf{x}_2$. This implies that $\mathbf{v} = \mathbf{x}_1 + \mathbf{x}_2 \cdot (\mathbf{Q}_0 + \sum_{i=1}^{\ell} \text{id}_{j^*}[i] \mathbf{Q}_i)$ is a vector of $\Lambda^\perp(\bar{\mathbf{A}})$. A similar analysis to [10] shows that \mathbf{v} is both short and non-zero with overwhelming probability. As a consequence, \mathcal{B} outputs $(\mathbf{x}_2 | 0^m)^T$ which is a short non-zero vector such that $(\mathbf{x}_2 | 0^m)^T \cdot \hat{\mathbf{A}} = 0 \pmod{q}$.

- If there exists $i \in \{1, \dots, \ell\}$ such that $\mathbf{y}_i \neq \text{id}_{j^*}[i] \cdot \mathbf{x}_2$, where $\mathbf{x}_2 \in \mathbb{Z}^{2m}$ is the vector obtained by decrypting \mathbf{c}_0^* using \mathbf{S}_0 , then \mathcal{B} aborts if $\text{coin} = 0$. Otherwise, the non-interactive proofs $\{\pi_{\text{OR},i}^*\}_i$ imply that $\mathbf{c}_i^* = \mathbf{B}_i \cdot \mathbf{s} + p \cdot \mathbf{e}'_i + \text{id}_{j^*}[i] \cdot \mathbf{x}_2$ for some $\mathbf{x}_2, \mathbf{e}'_1, \dots, \mathbf{e}'_\ell \in \mathbb{Z}^{2m}$ and $\mathbf{s}_0, \mathbf{s} \in \mathbb{Z}_q^n$. By multiplying the latter expression of \mathbf{c}_i^{*T} by \mathbf{A}_i , we find

$$\mathbf{c}_i^T \cdot \mathbf{A}_i = p \cdot (\mathbf{e}'_i^T \cdot \mathbf{A}_i) + \text{id}_{j^*}[i] \cdot \mathbf{x}_2^T \mathbf{A}_i.$$

Subtracting the latter equation from the second equation of (8), we find

$$(p \cdot (\mathbf{e}_i^T - \mathbf{e}'_i^T) + (\mathbf{y}_i^T - \text{id}_{j^*}[i] \cdot \mathbf{x}_2^T)) \cdot \mathbf{A}_i = 0 \pmod{q}.$$

If $p \cdot (\mathbf{e}_i^T - \mathbf{e}'_i^T) + (\mathbf{y}_i^T - \text{id}_{j^*}[i] \cdot \mathbf{x}_2^T) \neq 0 \pmod{q}$, it is a short non-zero vector in $\Lambda^\perp(\mathbf{A}_i)$. Given that $\mathbf{A}_i = \hat{\mathbf{A}}$ with probability $1/\ell$, we solved the given SIS instance with the same probability. Finally, if

$$p \cdot (\mathbf{e}_i^T - \mathbf{e}'_i^T) + (\mathbf{y}_i^T - \text{id}_{j^*}[i] \cdot \mathbf{x}_2^T) = 0 \pmod{q},$$

the relative lengths of vectors $\mathbf{e}_i, \mathbf{e}'_i, \mathbf{y}_i, \mathbf{x}_2$ with respect to p implies $\mathbf{e}_i = \mathbf{e}'_i$ and $\mathbf{y}_i = \text{id}_{j^*}[i] \cdot \mathbf{x}_2$, which contradicts the assumption that $\mathbf{y}_i \neq \text{id}_{j^*}[i] \cdot \mathbf{x}_2$.

The lower bound on the reduction's probability of success is assessed exactly in the same way as in the proof of Theorem 3. \square