G+G: A Fiat-Shamir Lattice Signature Based on Convolved Gaussians

Julien Devevey[1], Alain Passelègue[1,2,3], and Damien Stehlé[1,3]

[1] ENS de Lyon, France
julien.devevey@ens-lyon.fr
[2] INRIA, France
alain.passelegue@cryptolab.co.kr
[3] CryptoLab Inc., Lyon, France
damien.stehle@cryptolab.co.kr

Abstract. We describe an adaptation of Schnorr's signature to the lattice setting, which relies on Gaussian convolution rather than flooding or rejection sampling as previous approaches. It does not involve any abort, can be proved secure in the ROM and QROM using existing analyses of the Fiat-Shamir transform, and enjoys smaller signature sizes (both asymptotically and for concrete security levels).

1 Introduction

Schnorr's identification protocol [Sch91] allows secure authentication between a prover and a verifier based on the hardness on the discrete logarithm problem in a cyclic group of order p, generated by an element g. The prover's public verification key is simply a group element g^s, whose discrete logarithm s forms the prover's signing key. The identification protocol proceeds as follows: the prover first commits to some uniform $y \hookleftarrow U(\mathbb{Z}_p)$ by sending g^y to a verifier. The latter returns some challenge $c \in \mathbb{Z}_p$, to which the prover replies with a response z, namely $z = y + cs \bmod p$. Here, no information about s is revealed as z is still uniform modulo p. However, a verifier is convinced that the prover knows s as it can verify $g^z = g^y (g^s)^c$. This can be compiled into a signature scheme by using the Fiat-Shamir heuristic [FS86].

Adapting this protocol to the lattice setting has proved challenging. At a high-level, the approach adopted in [Lyu09,Lyu12] and subsequent works proceeds as follows. The discrete logarithms s is replaced with a short, tall matrix $\mathbf{S}$ in $\mathbb{Z}^{k \times m}$, whereas y and z are replaced with elements $\mathbf{y}$ and $\mathbf{z}$ of $\mathbb{Z}^k$ and the generator g is replaced with a uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times k}$. The challenge vector $\mathbf{c}$ belongs to a finite subset of $\mathbb{Z}^m$, typically designed to have the shortest possible vectors under the constraint that the challenge has sufficiently high min-entropy to prevent guessing. For security, one needs $\mathbf{z}$ and hence $\mathbf{y}$ to be short. Leaving things as they are described so far would make signatures leak the secret matrix $\mathbf{S}$, as $\mathbf{z}$ is centered around $\mathbb{E}[\mathbf{y}] + \mathbf{Sc}$ (see [ASY22] for a detailed key recovery). A solution could be to take a large enough standard deviation to "flood" this center (this is considered for example in [DPSZ12, Appendix A.1] in the context of zero-knowledge proofs), but this results in very large signatures as the modulus then needs to grow exponentially with the security parameter λ (see the discussion in [ASY22]). The most efficient approach so far, introduced by Lyubashevsky [Lyu09,Lyu12] and notably leading to Dilithium [DKL$^+$18], relies on rejection sampling to erase the center from $\mathbf{z}$. This comes at the cost of restarting the protocol multiple times before finally outputting an appropriately distributed response $\mathbf{z}$. This strategy still allows the identification protocol to be compiled into a signature, using a variant of the Fiat-Shamir heuristic called Fiat-Shamir with Aborts. To obtain shorter signatures, Ducas *et al.* [DDLL13] suggested to reject a bimodal Gaussian distribution against a Gaussian distribution. This was later argued in [DFPS22] to be essentially optimal among pairs of source and target distributions. Finally, we note that Fiat-Shamir with Aborts turns out to be complex to analyze, and flaws in many analyses have been recently discovered [DFPS23,BBD$^+$23].

Removing rejection sampling while keeping similar signature sizes has been a long-standing open problem. Steps in this direction were made in [BCM21] for instance. The authors noticed that in the setting where $\mathbf{y}$ is sampled uniformly in a hypercube and one uses signature truncation [BG14], one rejection condition out of two is superfluous. They however argue that removing the second one is difficult.

Contribution. We introduce a new paradigm for adapting Schnorr's identification protocol to the lattice setting. It relies on Gaussian convolution, rather than flooding or rejection sampling. Our $\mathsf{G}+\mathsf{G}$ (Gaussian Plus Gaussian) identification protocol can be compiled into a signature using the Fiat-Shamir heuristic (without aborts), in the Quantum Random Oracle Model (QROM). The resulting signature is asymptotically more compact than those based on rejection sampling and its analysis relies on the well-understood properties of the standard Fiat-Shamir transform. Finally, we provide concrete parameters which show that $\mathsf{G}+\mathsf{G}$ is competitive with the state-of-the-art optimizations of Lyubashevsky's signature.

Technical Overview. $\mathsf{G}+\mathsf{G}$ involves two Gaussians that are being summed. The first one is $\mathbf{y}$ and the second one corresponds to $\mathbf{Sc}$. The first difficulty that we face is that $\mathbf{S}$ is fixed and $\mathbf{c}$ is publicly known as part of the resulting signature and hence cannot be assumed random for the sake of studying the distribution of $\mathbf{z}$.

To introduce the required new randomness, we start from BLISS [DDLL13]. The verification key $\mathbf{A} \in \mathbb{Z}_{2q}^{m\times k}$ and the signing key $\mathbf{S} \in \mathbb{Z}^{k\times m}$ satisfy the relation $\mathbf{AS} = q\mathbf{I}_m \bmod 2q$. Among the variants of Lyubashesvky's signature, it is a specificity of BLISS to work modulo $2q$, which is particularly useful in our case. The commitment of the prover is $\mathbf{w} = \mathbf{Ay} \bmod 2q$, and upon receiving $\mathbf{c} \in \{0,1\}^m$, the prover replies with either $\mathbf{z} = \mathbf{y} + \mathbf{Sc}$ or $\mathbf{z} = \mathbf{y} - \mathbf{Sc}$ with probability $1/2$ each. The verifier checks that $\mathbf{z}$ is short and $\mathbf{Az} = \mathbf{w} + q\mathbf{c} \bmod 2q$. This check works for both values of $\mathbf{z}$ that the prover chose from. This can be explained by observing that the verification views $\mathbf{c}$ modulo 2, i.e., as a coset of $\mathbb{Z}^m/2\mathbb{Z}^m$, and negating it does not change the coset. This observation was used in [Duc14] to take negations of individual coordinates of $\mathbf{c}$ to minimize the Euclidean norm of $\mathbf{Sc}$ and hence decrease the standard deviation of $\mathbf{y}$ necessary to hide $\mathbf{Sc}$ via rejection sampling. We go further and let the prover extend the coset $\mathbf{c}$ sent by the verifier to a Gaussian sample with support $2\mathbb{Z}^m + \mathbf{c}$ and center $\mathbf{0}$. The verification equation above still holds, and we now have our second Gaussian.

At this stage, the prover samples a Gaussian $\mathbf{y}$ over $\mathbb{Z}^k$, receives a uniform coset $\mathbf{c} \in \mathbb{Z}^m/2\mathbb{Z}^m$ from the verifier, produces a Gaussian sample $\mathbf{x}$ with support $2\mathbb{Z}^m + \mathbf{c}$ and computes $\mathbf{z} = \mathbf{y} + \mathbf{Sx}$. Equivalently, it samples $\mathbf{k}$ Gaussian with support $2\mathbf{S}\mathbb{Z}^m$ and center $-\mathbf{Sc}$, which will be used to cancel the center $\mathbf{Sc}$, and returns $\mathbf{z} = \mathbf{y} + \mathbf{k} + \mathbf{Sc}$. In order to obtain the zero-knowledge property (i.e., be able to simulate signatures without knowing the signing key), we aim to prove that the distribution of the Gaussian convolution $\mathbf{z}$ can be sampled from publicly. If $\mathbf{y}$ and $\mathbf{k}$ were continuous Gaussians, we would set their covariance matrices $\mathbf{\Sigma_y}$ and $\mathbf{\Sigma_k}$ such that $\mathbf{\Sigma_y} + \mathbf{\Sigma_k} = \mathbf{\Sigma_z}$ for a known covariance matrix $\mathbf{\Sigma_z}$ for $\mathbf{z}$. To fix the ideas, we could set $\mathbf{\Sigma_z} = \sigma^2\mathbf{I}$ for some $\sigma > 0$, i.e., the distribution of $\mathbf{z}$ is a spherical Gaussian, and set $\mathbf{\Sigma_y} = \sigma^2\mathbf{I} - \mathbf{\Sigma_k}$. If we sample $\mathbf{x}$ from a spherical Gaussian with standard deviation $s > 0$, then $\mathbf{\Sigma_k} = s^2\mathbf{SS}^\top$ and $\mathbf{\Sigma_y} = \sigma^2\mathbf{I} - s^2\mathbf{SS}^\top$ (by taking σ sufficiently large, the latter is indeed definite positive). This is the choice we actually make for $\mathsf{G} + \mathsf{G}$, but there is flexibility.

The above over-simplifies the situation as the Gaussians we manipulate are discrete rather than continuous. Further, their supports do not have the same dimensions. Indeed, the support of $\mathbf{y}$ is $\mathbb{Z}^k$ whereas the support of $\mathbf{k}$ is exactly $2\mathbf{S}\mathbb{Z}^m + \mathbf{Sc}$ whose span has dimension $m < k$: the second Gaussian lives in a smaller dimension and its support is sparser. This is illustrated in Figure 1.

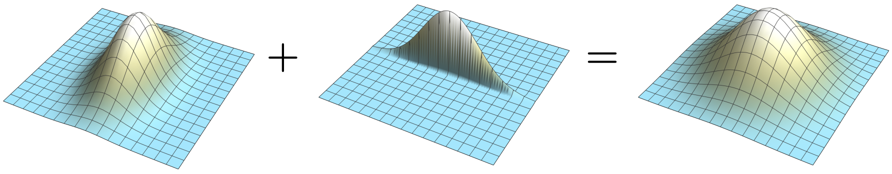

Fig. 1. The sum of two Gaussians with compensating covariance matrices is a spherical Gaussian, even when the second Gaussian is rank-deficient. In the $\mathsf{G} + \mathsf{G}$ identification protocol and signature, the first Gaussian corresponds to $\mathbf{y}$, the second Gaussian is associated to $\mathbf{Sc}$ and the resulting one corresponds to $\mathbf{z}$.

Thanks to the above, if the covariance matrices are set appropriately, then $\mathsf{G} + \mathsf{G}$ is honest-verifier zero-knowledge (HVZK). The proofs of completeness and soundness are adapted from [DDLL13].

Our final goal is to apply the Fiat-Shamir heuristic on this protocol to get a signature scheme. This heuristic replaces the uniform challenge with one derived from a hash function called on input the commitment and the message to be signed. The signature is then the whole transcript. As the commitment of $\mathsf{G}+\mathsf{G}$ can be recomputed from the challenge and its response, we actually exclude it from the signature for compactness. Then, as long as $\mathsf{G}+\mathsf{G}$ is complete, the resulting signature is correct. Moreover, the security reduction proceeds in two steps. First, it is shown that the $\mathsf{EU\text{-}CMA}$ security of the signature can be reduced to the $\mathsf{EU\text{-}NMA}$ security of the signature, where no signature query can be made. To do so, one shows that signatures queries can be answered with simulated ones (up to reprogramming the random oracle) from the HVZK property, as long as the commitment $\mathbf{Ay}$ has sufficiently high min-entropy. This is technically more complex than for Lyubashevsky's signatures as $\mathbf{y}$ is distributed from a skewed Gaussian. Second, computational soundness (resp. lossy-soundness) implies security against no-message attacks for different parametrizations.

We stress that convolution of discrete Gaussian distributions is the core technical idea to make the signature distribution independent from the signing key. Exploiting Gaussian convolution in lattice-based signatures dates back to [Pei10], which used it to simplify the message-dependent component of the signing algorithm of the GPV signature scheme [GPV08]. At a high level, our contribution can be summarized as applying the Gaussian convolution technique in the context of Fiat-Shamir lattice signatures.

Comparison with BLISS. Among variants of Lyubashevsky's signatures, BLISS provides the smallest $\mathbf{z}$: its expected norm can be as small as $\sigma_1(\mathbf{S})m/\sqrt{\log M}$ (up to a constant factor), where $\sigma_1(\mathbf{S})$ is the largest singular value of $\mathbf{S}$ and M is the expected number of repetitions (see [DFPS22, Appendix C]). Further, an argument is made in [DFPS22] that this is essentially optimal for Lyubashevsky's signatures, even if we allow to optimize over the choice of source and target distributions. In the case of $\mathsf{G}+\mathsf{G}$, the strongest constraint on parameters is essentially that the standard deviation σ of $\mathbf{z}$ be sufficiently large to "smooth out" the lattice $2\mathbf{S}\mathbb{Z}^m$. By using a variant of the HVZK property based on the Rényi divergence rather than the statistical distance, which suffices for the signature application, it suffices that σ be above $\sigma_1(\mathbf{S})\sqrt{\log Q_S}$, up to a constant factor, where Q_S is the maximum number of signature queries that the adversary is allowed to make. As a result, the expected norm of $\mathbf{z}$ in $\mathsf{G}+\mathsf{G}$ is $\sigma_1(\mathbf{S})\sqrt{m\log Q_S}$. We conclude by observing that $\log Q_S$ is typically much smaller than m, and that the $\sqrt{\log M}$ term from BLISS cannot grow sufficiently to compensate for the difference. More concretely, if we set $M=\lambda^{\Theta(1)}$, $Q_S=\lambda^{\Theta(1)}$ and $m=\Theta(\lambda)$, where λ is the security parameter, then the expected norms of $\mathbf{z}$ in BLISS and $\mathsf{G}+\mathsf{G}$ respectively grow as $\sigma_1(\mathbf{S})\cdot\lambda/\sqrt{\log\lambda}$ and $\sigma_1(\mathbf{S})\cdot\sqrt{\lambda\log\lambda}$.

Optimization and concrete parameters. While all key generation techniques presented in [DDLL13] can be used with our $\mathsf{G}+\mathsf{G}$ protocol, we present alternative versions which offer more flexibility. A first improvement is that we can set $\mathbf{AS}=q\mathbf{J} \bmod 2q$, where $\mathbf{J}\in\mathbb{Z}_q^{m\times\ell}$ is only rectangular and full column-rank rather than set to the identity. When instantiating $\mathsf{G}+\mathsf{G}$ with the MLWE and

MSIS hardness assumptions [BGV12,LS15] over a ring $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$ with n a power of 2, we take $\mathbf{j} = (x^{n/2} + 1, 0, \ldots, 0)$. This allows us to replace the lattice $2\mathbf{s}\mathcal{R}$ with $(x^{n/2} - 1)\mathbf{s}\mathcal{R}$, and to decrease the standard deviation of $\mathbf{z}$ by a factor $\sqrt{2}$. Overall, we obtain signature sizes that are between 12% and 20% smaller than those in [DFPS22], or 30% to 40% smaller than Dilithium [DKL⁺18] for typical target levels of security.

Related Work. As pointed out in [CLMQ21], GPV signatures [GPV08] can be seen as a special case of the lattice-based Fiat-Shamir signatures by considering a specific instance of the hash function and adapting parameters. This analysis can be extended to $\mathsf{G} + \mathsf{G}$, and we then recover the hash-and sign scheme described in [YJW23]. More details are provided in Appendix B.

2 Preliminaries

For any integers $k \geq m$, we let $\mathbf{I}_k$ denote the $k \times k$ identity matrix as well as $\mathbf{J}_{k,m} = (\mathbf{I}_m | \mathbf{0}^{m \times (k-m)})^\top$ denote the $k \times m$ matrix whose first m diagonal elements are 1 and all others are 0. The notations log and ln respectively refer to the base-2 and natural logarithms. The notation $\|\cdot\|$ refers to the Euclidean norm, while $\|\cdot\|_\infty$ refers to the infinity norm.

2.1 Probabilities

Let P, Q be two discrete random variables. The min-entropy of P is defined as

$$H_\infty(P) = -\log \max_{x \in \mathrm{Supp}(P)} \Pr[P = x] \enspace .$$

The conditional min-entropy of P on Q is defined as

$$H_\infty(P|Q) = -\log \sum_{y \in \mathrm{Supp}(Q)} \Pr[Q = y] \cdot \max_{x \in \mathrm{Supp}(P)} \Pr[P = x | Q = y] \enspace .$$

Let $\mathit{\Omega} = \mathrm{Supp}(P) \cup \mathrm{Supp}(Q)$. The statistical distance between P and Q is defined as $\mathit{\Delta}(P, Q) = \sum_{x \in \mathit{\Omega}} |\Pr[P = x] - \Pr[Q = x]|/2$.

If $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, the Rényi divergence of infinite order between P and Q is defined as

$$R_\infty(P \| Q) = \sup_{x \in \mathrm{Supp}(P)} \frac{\Pr[P = x]}{\Pr[Q = x]} \in [1, +\infty] \enspace .$$

We will use the following properties of the Rényi divergence.

Lemma 1 ([vEH14]). *Let P and Q be two discrete random variables such that* $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$. *Let $f : \mathrm{Supp}(Q) \to \mathcal{X}$ be a (possibly probabilistic) function. Let $E \subseteq \mathrm{Supp}(P)$ be an event. The Rényi divergence satisfies the probability preservation property:*

$$\Pr[P \in E] \leq R_\infty(P \| Q) \cdot \Pr[Q \in E] \tag{1}$$

and the data processing inequality:

$$R_\infty(f(P)\|f(Q)) \leq R_\infty(P\|Q) \ . \tag{2}$$

We will also use the following result.

Lemma 2. *Let $\varepsilon < 1$. Let P and Q be two random variables taking values in some countable set Ω. Let $c \in \mathbb{R}$ be a constant such that*

$$\forall a \in \Omega : \ \Pr[Q = a] = c(1 - \delta(a)) \Pr[P = a] \ ,$$

for some function $\delta : \Omega \to [0, \varepsilon]$. Then it holds that:

$$R_\infty(P\|Q) \leq \frac{1}{1-\varepsilon} \ , \quad R_\infty(Q\|P) \leq \frac{1}{1-\varepsilon} \quad \text{and} \quad \Delta(P,Q) \leq \frac{\varepsilon}{1-\varepsilon} \ .$$

Proof. Let us first note that $(1-\varepsilon)c \leq 1 \leq c$, by summing the above equality over all $a \in \Omega$ and applying the bounds on $\delta(a)$. Then we have

$$R_\infty(P\|Q) = \sup_{a\in\Omega} \frac{\Pr[P=a]}{\Pr[Q=a]} = \sup_{a\in\Omega} \frac{1}{c(1-\delta(a))} \leq \frac{1}{1-\varepsilon} \ .$$

We also have

$$R_\infty(Q\|P) = \sup_{a\in\Omega} \frac{\Pr[Q=a]}{\Pr[P=a]} = \sup_{a\in\Omega} c(1-\delta(a)) \leq c \leq \frac{1}{1-\varepsilon} \ .$$

Finally, we refer to [BF11, Lemma A.2] for the third bound. □

2.2 Lattice Gaussian Distributions

Let $k > 0, \mathbf{c} \in \mathbb{R}^k$ and $\mathbf{\Sigma} \in \mathbb{R}^{k\times k}$ be a positive-definite symmetric matrix. The Gaussian function with covariance parameter $\mathbf{\Sigma}$ and center parameter $\mathbf{c}$ is defined as

$$\rho_{\mathbf{\Sigma},\mathbf{c}} : \mathbf{x} \mapsto \exp\left(-\pi(\mathbf{x}-\mathbf{c})^\top \mathbf{\Sigma}^{-1}(\mathbf{x}-\mathbf{c})\right) \ .$$

The Gaussian distribution over the lattice $\Lambda \subseteq \operatorname{span}(\mathbf{\Sigma})$ with covariance parameter $\mathbf{\Sigma}$ and center parameter $\mathbf{c}$ is the distribution with support Λ and probability mass function

$$D_{\Lambda,\mathbf{\Sigma},\mathbf{c}} : \mathbf{x} \mapsto \frac{\rho_{\mathbf{\Sigma},\mathbf{c}}(\mathbf{x})}{\sum_{\mathbf{y}\in\Lambda} \rho_{\mathbf{\Sigma},\mathbf{c}}(\mathbf{y})} \ .$$

If $\mathbf{\Sigma} = \sigma^2 \mathbf{I}_k$, we write $\rho_{\sigma,\mathbf{c}}$ and $D_{\Lambda,\sigma,\mathbf{c}}$. We omit $\mathbf{c}$ when it is $\mathbf{0}$. We also define $D_{\Lambda+\mathbf{c},\mathbf{\Sigma}} = D_{\Lambda,\mathbf{\Sigma},-\mathbf{c}} + \mathbf{c}$. For convenience, we let $\rho_{\mathbf{\Sigma},\mathbf{c}}(S)$ denote the quantity $\sum_{\mathbf{y}\in S} \rho_{\mathbf{\Sigma},\mathbf{c}}(\mathbf{y})$ for any countable set S.

For spherical Gaussians, the upper and lower part of a vector are statistically independent. This is not the case anymore for general covariance matrices. The following lemma give the conditional distribution of the lower part of a Gaussian vector, given the upper part. The proof is adaptated from the continuous setting and relies on writing the covariance as a 2×2 block matrix and inverting it using the Schur complement of the upper left matrix.

Lemma 3 (Conditional distribution). *Let $k \geq m > 0$, $\mathbf{\Sigma} \in \mathbb{R}^{k \times k}$ be a symmetric positive-definite matrix and $\mathbf{c} \in \mathbb{R}^k$. Write*

$$\mathbf{c} = \begin{pmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{pmatrix} \quad \text{and} \quad \mathbf{\Sigma} = \begin{pmatrix} \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} \end{pmatrix} ,$$

where $\mathbf{c}_1 \in \mathbb{R}^{k-m}$ and $\mathbf{\Sigma}_{11} \in \mathbb{R}^{(k-m)\times(k-m)}$. Let $(Y_1^\top | Y_2^\top) \hookleftarrow D_{\mathbb{Z}^k, \mathbf{\Sigma}, \mathbf{c}}$, where Y_1 takes values in $\mathbb{Z}^{k-m}$. Given any $\mathbf{y}_1 \in \mathbb{Z}^{k-m}$, the conditional distribution of Y_2 conditioned on $Y_1 = \mathbf{y}_1$ is $D_{\mathbb{Z}^m, \overline{\mathbf{\Sigma}}, \overline{\mathbf{c}}}$, where

$$\overline{\mathbf{c}} = \mathbf{c}_2 + \mathbf{\Sigma}_{21}\, \mathbf{\Sigma}_{11}^{-1}(\mathbf{y}_1 \ - \mathbf{c}_1) \quad \text{and} \quad \overline{\mathbf{\Sigma}} = \mathbf{\Sigma}_{22} \ - \mathbf{\Sigma}_{21}\, \mathbf{\Sigma}_{11}^{-1} \mathbf{\Sigma}_{12} \, .$$

Proof. As $\mathbf{\Sigma}$ is symmetric and positive-definite, both $\mathbf{\Sigma}_{11}$ and $\mathbf{\Sigma}_{22}$ are also symmetric and positive-definite and thus invertible. This is shown by considering vectors of the form $(\mathbf{x}^\top | (\mathbf{0}^m)^\top)^\top$ or $((\mathbf{0}^{k-m})^\top | \mathbf{y}^\top)^\top$. Let us write the block inverse of $\mathbf{\Sigma}$ as follows:

$$\mathbf{\Sigma}^{-1} = \left(\begin{array}{c|c} \mathbf{\Sigma}_{11}^{-1} + \mathbf{\Sigma}_{11}^{-1}\mathbf{\Sigma}_{12}^{-1}\overline{\mathbf{\Sigma}}^{-1}\mathbf{\Sigma}_{21}\, \mathbf{\Sigma}_{11}^{-1} & -\mathbf{\Sigma}_{11}^{-1}\mathbf{\Sigma}_{12}^{-1}\overline{\mathbf{\Sigma}}^{-1} \\ \hline -\overline{\mathbf{\Sigma}}^{-1}\mathbf{\Sigma}_{21}\, \mathbf{\Sigma}_{11}^{-1} & \overline{\mathbf{\Sigma}}^{-1} \end{array}\right) = \begin{pmatrix} \mathbf{S}_{11} & \mathbf{S}_{12} \\ \mathbf{S}_{21} & \mathbf{S}_{22} \end{pmatrix} .$$

This formula also ensures that $\overline{\mathbf{\Sigma}}$ is invertible, as it is a diagonal block of the positive definite symmetric matrix $\mathbf{\Sigma}^{-1}$.

Let $\mathbf{y}_2 \in \mathbb{Z}^m$. The probability that $Y_2 = \mathbf{y}_2$ conditioned on $Y_1 = \mathbf{y}_1$ is

$$\rho_{\mathbf{\Sigma},\mathbf{c}} \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} \Big/ \sum_{\mathbf{y} \in \mathbb{Z}^m} \rho_{\mathbf{\Sigma},\mathbf{c}} \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y} \end{pmatrix} .$$

Let us then study $\rho_{\mathbf{\Sigma},\mathbf{c}}((\mathbf{y}_1^\top | \mathbf{y}^\top)^\top)$ by expanding it and completing the square.

$$\rho_{\mathbf{\Sigma},\mathbf{c}} \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y} \end{pmatrix} \sim \exp\left(-\pi \left((\mathbf{y} - \overline{\mathbf{c}})^\top \mathbf{S}_{22} (\mathbf{y} - \overline{\mathbf{c}})\right)\right) = \rho_{\overline{\mathbf{\Sigma}}, \overline{\mathbf{c}}}(\mathbf{y}) \ ,$$

where the notation $\sim$ hides terms that do not depend on $\mathbf{y}$. Using the fact that the probability mass sums to 1, we obtain that the distribution of Y_2 conditioned on $Y_1 = \mathbf{y}_1$ is $D_{\mathbb{Z}^m, \overline{\mathbf{\Sigma}}, \overline{\mathbf{c}}}$. □

As showed in [GPV08], Gaussian distributions can be sampled from by using Klein's algorithm [Kle00]. We will rely on the following variant.

Lemma 4 (Adapted from [BLP$^+$13, Lemma 2.3]). *There is a* ppt *algorithm that, given a basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_\ell)$ of a full-rank ℓ-dimensional lattice Λ, a positive definite symmetric matrix $\mathbf{\Sigma}$ and $\mathbf{c} \in \mathbb{R}^\ell$ returns a sample from $D_{\Lambda, \mathbf{\Sigma}, \mathbf{c}}$, assuming that $\sqrt{\ln(2\ell + 4)/\pi} \cdot \max_i \|\mathbf{\Sigma}^{-1/2}\mathbf{b}_i\| \leq 1$.*

2.3 Smoothing Parameter

Given a k-dimensional lattice $\Lambda \subseteq \mathbb{R}^k$, its dual lattice Λ^* is defined as the set $\Lambda^* = \{\mathbf{x} \in \operatorname{span}(\Lambda) \mid \mathbf{x}^\top \mathbf{y} \in \mathbb{Z}, \forall \mathbf{y} \in \Lambda\}$. If $\mathbf{B}$ is a basis of Λ, then $(\mathbf{B}^\dagger)^\top$ is a basis of Λ^*.

Given a lattice $\Lambda \subseteq \mathbb{R}^k$ and $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\Lambda)$ of the lattice Λ is defined as the smallest σ such that $\rho_{1/\sigma}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \varepsilon$. The smoothing parameter satisfies the following two properties.

Lemma 5 ([ZXZ18, Theorem 2]). *Let $k > 1$ and $\varepsilon < 0.086k$. Let $\Lambda \subseteq \mathbb{R}^k$ be a full-rank lattice with basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_k)$. It holds that*

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(k - 1 + 2k/\varepsilon)}{\pi}} \cdot \max_{i \leq k} \|\mathbf{b}_i\| \ .$$

Lemma 6 ([MR07]). *Let Λ be a k-dimensional full-rank lattice. Let $\varepsilon > 0$ and $\mathbf{\Sigma} \in \mathbb{R}^{k \times k}$ be a definite positive symmetric matrix with all singular values larger than $\eta_\varepsilon(\Lambda)$ and $\mathbf{c} \in \mathbb{R}^k$. We have*

$$\rho_{\mathbf{\Sigma},\mathbf{c}}(\Lambda) \ \in \ \frac{\sqrt{\det \mathbf{\Sigma}}}{\det \Lambda} \cdot [1 - \varepsilon, 1 + \varepsilon] \quad \text{ and } \quad \frac{\rho_{\mathbf{\Sigma},\mathbf{c}}(\Lambda)}{\rho_{\mathbf{\Sigma}}(\Lambda)} \ \in \ \left[\frac{1 - \varepsilon}{1 + \varepsilon}, 1\right] \ .$$

The last upper bound holds for all $\mathbf{\Sigma}$.

The following lemma is adapted from [BMKMS22, Lemma 1] (but could also be obtained from [GMPW20, Theorem 3.1]). It is at the core of the completeness and zero-knowledge proofs. While [BMKMS22] does not give explicit statistical bounds, we note that Lemma 6 above, which is applied at the end of the proof from [BMKMS22], allows us to do so when combined with Lemma 2. A further adaptation is the use of the smoothing parameter bound from Lemma 5. Note that the dimension involved for this condition is ℓ rather than k, as this is the small-rank lattice that needs to be smoothed out (the corresponding condition from [BMKMS22, Lemma 1] is stronger than needed).

Lemma 7 (Gaussian decomposition). *Let $k \geq \ell$, $\varepsilon \in (0, 1)$ and $\mathbf{S} \in \mathbb{Z}^{k \times \ell}$. Let $s \geq \sqrt{2 \ln(\ell - 1 + 2\ell/\varepsilon)/\pi}$ and $\sigma \geq \sqrt{8}\sigma_1(\mathbf{S}) \cdot s$. Define*

$$\mathbf{\Sigma}(\mathbf{S}) = \sigma^2 \mathbf{I}_k - 4s^2 \mathbf{S}\mathbf{S}^\top \ ,$$

and let $\mathbf{y} \hookleftarrow D_{\mathbb{Z}^k, \mathbf{\Sigma}(\mathbf{S})}$ and $\mathbf{k} \hookleftarrow D_{\mathbb{Z}^\ell, s, -\mathbf{c}/2}$ for any $\mathbf{c} \in \mathbb{Z}^\ell$. Then $\mathbf{\Sigma}(\mathbf{S})$ is positive definite and the distribution $P_\mathbf{z}$ of $\mathbf{z} = \mathbf{y} + \mathbf{S}(2\mathbf{k} + \mathbf{c})$ satisfies

$$R_\infty(P_\mathbf{z} \| D_{\mathbb{Z}^k, \sigma}) \leq \frac{1 + \varepsilon}{1 - \varepsilon} \quad \text{ and } \quad \Delta(P_\mathbf{z}, D_{\mathbb{Z}^k, \sigma}) \leq \frac{2\varepsilon}{1 - \varepsilon} \ .$$

Note that the matrix $\mathbf{\Sigma}(\mathbf{S})$ is positive definite since $\sigma \geq \sqrt{8}\sigma_1(\mathbf{S}) \cdot s$ ensures that all singular values of $\sigma^2 \mathbf{I}_k$ are larger than those of $4s^2 \mathbf{S}\mathbf{S}^\top$.

2.4 Cryptographic Definitions

We recall the definition of an identification scheme and how such a scheme can be transformed into a digital signature via the Fiat-Shamir transform (see Figure 6, p.27). For an identification scheme ID and a hash function H (modeled as a random oracle in the analysis), we let $\mathsf{FS}[\mathsf{ID}, H]$ denote the resulting signature scheme. Details about correctness and security of $\mathsf{FS}[\mathsf{ID}, H]$ are provided in Appendix A.

Definition 1 (Identification Scheme). *An identification scheme is a tuple of PPT algorithms* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *such that:*

- Igen : *On input the security parameter* 1^λ*, algorithm* Igen *outputs a verification key* vk *and a signing key* sk*. We assume that* vk *defines the challenge space* $\mathcal{C}$*.*
- P : *The prover* $\mathsf{P} = (\mathsf{P}_1, \mathsf{P}_2)$ *is split into two algorithms: given* sk*, algorithm* P_1 *produces a* commitment w *(first message sent to the verifier) and a state* st*; algorithm* P_2*, on input* $(\mathsf{sk}, w, \mathsf{st})$ *and a uniformly random* challenge $c \in \mathcal{C}$ *sent by the verifier in response to commitment* w*, outputs an* answer z*.*
- V : *On input* (vk, w, c, z)*, the deterministic verifier* V *outputs* 1 *or* 0*.*

We let $\mathsf{P}(\mathsf{sk}, \mathsf{vk}) \leftrightarrow \mathsf{V}(\mathsf{vk})$ *denote the transcript* (w, c, z) *of an interaction between the prover and the verifier, as illustrated in Figure 2.*

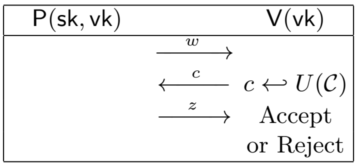

Fig. 2. Interaction Between P and V

We further define the following properties of identification schemes and recall their roles in the analysis of the signature obtained by applying the Fiat-Shamir transform to an identification protocol. We first recall *completeness* and *commitment-recoverability*, which allow to prove correctness of $\mathsf{FS}[\mathsf{ID}, H]$.

Definition 2 (Completeness and commitment-recoverability). *An identification scheme* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *is* ε-complete *for some* $\varepsilon > 0$ *if for any* $(\mathsf{vk}, \mathsf{sk}) \leftarrow \mathsf{Igen}(1^\lambda)$*, for any challenge* $c \in \mathcal{C}$*, we have:*

$$\Pr\Big[\mathsf{V}(\mathsf{vk}, (w, c, z)) = 0 \mid (w, c, z) \leftarrow (\mathsf{P}(\mathsf{sk}, \mathsf{vk}) \leftrightarrow \mathsf{V}(\mathsf{vk}))\Big] \leq \varepsilon \ ,$$

where the randomness is taken over the random coins of P*.*

In addition, ID *satisfies* commitment-recoverability *if for any public key* vk*, challenge* $c \in \mathcal{C}$*, and answer* z*, there is at most one commitment* w *such that the transcript* (w, c, z) *is valid, and there exists a PPT algorithm* Rec *such that* $w = \mathsf{Rec}(\mathsf{vk}, c, z)$*.*

We then recall the definitions of *honest-verifier zero-knowledge* and *commitment min-entropy*, which allow to reduce $\mathsf{EU\text{-}CMA}$ security of $\mathsf{FS}[\mathsf{ID}, H]$ to its $\mathsf{EU\text{-}NMA}$ security.

Definition 3 (HVZK and commitment min-entropy). *An identification scheme* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *is* Honest-Verifier Zero-Knowledge *if there exists a PPT simulator* Sim *such that one of the following holds:*

- $\Delta((w,c,z) \leftarrow (\mathsf{P}(\mathsf{sk},\mathsf{vk}) \leftrightarrow \mathsf{V}(\mathsf{vk}))\ ,\ \mathsf{Sim}(c,\mathsf{vk})) \leq \varepsilon$*. In this case, we say that* ID *is* ε*-HVZK.*
- $R_\infty((w,c,z) \leftarrow (\mathsf{P}(\mathsf{sk},\mathsf{vk}) \leftrightarrow \mathsf{V}(\mathsf{vk})) \parallel \mathsf{Sim}(c,\mathsf{vk})) \leq 1+\varepsilon$*. In this case, we say that* ID *is* $(1+\varepsilon)$*-divergence HVZK.*

Furthermore, we say that ID *satisfies* α-Min Entropy *or has* α *bits of commitment min-entropy if for any* $(\mathsf{vk},\mathsf{sk})$ *in the range of* IGen*:*

$$H_\infty\Big(w|(w,c,z) \leftarrow (\mathsf{P}(\mathsf{sk},\mathsf{vk}) \leftrightarrow \mathsf{V}(\mathsf{vk}))\Big) \geq \alpha\ \ .$$

Finally, we recall the notions of *key-indistinguishability* and *lossy-soundness*, which allow to prove $\mathsf{EU\text{-}NMA}$ security of $\mathsf{FS}[\mathsf{ID},H]$ in the QROM.

Definition 4 (Lossy identification scheme). *An identification scheme* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *is* lossy *if it satisfies both following properties.*

- key-indistinguishability*: there exists a PPT lossy key generation algorithm* $\mathsf{LossyIGen}$ *that, on input a security parameter, outputs a verification key* $\mathsf{vk}_{\mathsf{ls}}$*, such that for any (possibly quantum) adversary* $\mathcal{A}$*:*

  $$\mathsf{Adv}^{\mathsf{key\text{-}ind}}(\mathcal{A}) := |\Pr(\mathcal{A}(\mathsf{vk}_{\mathsf{ls}}) = 1) - \Pr(\mathcal{A}(\mathsf{vk}) = 1)| = \mathsf{negl}(\lambda),$$

  where $\mathsf{vk}_{\mathsf{ls}} \leftarrow \mathsf{LossyIGen}(1^\lambda)$ *and* $(\mathsf{vk},\mathsf{sk}) \leftarrow \mathsf{IGen}(1^\lambda)$ *and the probability is taken over the random coins of* $\mathcal{A}$ *and* $\mathsf{LossyIGen}$ *and* IGen*.*
- $\varepsilon_{\mathsf{ls}}$-lossy-soundness *for some* $\varepsilon_{\mathsf{ls}} > 0$*: for any (unbounded)* P^* *interacting with* V*, we have:*

  $$\Pr\Big[\mathsf{V}(\mathsf{vk}_{\mathsf{ls}},(w,c,z)) = 1 \mid (w,c,z) \leftarrow (\mathsf{P}^*(\mathsf{vk}_{\mathsf{ls}}) \leftrightarrow \mathsf{V}(\mathsf{vk}_{\mathsf{ls}}))\Big]\ \leq\ \varepsilon_{\mathsf{ls}}\ \ .$$

If we only consider classical adversaries, $\mathsf{EU\text{-}NMA}$ security of $\mathsf{FS}[\mathsf{ID},H]$ can be argued by relying on the simpler notion of *lossy special soundness.*

Definition 5 (Lossy special soundness). *Let* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *be an identification scheme with key-indistinguishability and* $\mathsf{LossyIGen}$ *be the corresponding lossy key generation algorithm.The scheme* ID *is* lossy-special-sound *if for any PPT adversary* $\mathcal{A}$*, the quantity*

$$\Pr\Big[c_0 \neq c_1 \wedge \mathsf{V}(\mathsf{vk},(w,c_0,z_0)) = \mathsf{V}(\mathsf{vk},(w,c_1,z_1)) = 1\,|(w,c_0,z_0,c_1,z_1) \leftarrow \mathcal{A}(\mathsf{vk})\Big]$$

is $\mathsf{negl}(\lambda)$*, where the probability is over the choice of* $\mathsf{vk} \leftarrow \mathsf{LossyIGen}(1^\lambda)$ *and the coins of* $\mathcal{A}$*.*

An identification scheme has unique response if, for any (w,c), there is at most one z such that the transcript (w,c,z) is accepted by the verifier. While this is not true by design for the $\mathsf{G}+\mathsf{G}$ identification scheme, we relax this property to computational unique response (CUR), which states that it is difficult to come up with two different responses for (w,c).

Definition 6 (Computational unique responses). *Let* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *be an identification scheme. It has* computational unique responses *if, for any PPT adversary* $\mathcal{A}$*, the quantity*

$$\Pr\Big[z_0 \neq z_1 \wedge \mathsf{V}(\mathsf{vk}, (w, c, z_0)) = \mathsf{V}(\mathsf{vk}, (w, c, z_1)) = 1 \mid (w, c, z_0, z_1) \leftarrow \mathcal{A}(\mathsf{vk})\Big]$$

is $\mathsf{negl}(\lambda)$*, where the probability is over the choice of* vk *and the coins of* $\mathcal{A}$*.*

We now briefly recall the formalism of digital signatures.

Definition 7. *A signature scheme is a tuple* $(\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ *of PPT algorithms with the following specifications:*

- $\mathsf{KeyGen} : 1^\lambda \to (\mathsf{vk}, \mathsf{sk})$ *takes as input a security parameter* λ *and outputs a verification key* vk *and a signing key* sk*.*
- $\mathsf{Sign} : (\mathsf{sk}, \mu) \to \sigma$ *takes as inputs a signing key* sk *and a message* μ *and outputs a signature* σ*.*
- $\mathsf{Verify} : (\mathsf{vk}, \mu, \sigma) \to b \in \{0, 1\}$ *takes as inputs a verification key* vk*, a message* μ *and a signature* σ *and accepts* $(b = 1)$ *or rejects* $(b = 0)$*.*

We say that it is ε*-correct if for any pair* $(\mathsf{vk}, \mathsf{sk})$ *in the range of* KeyGen *and* μ*,*

$$\Pr\Big[\mathsf{Verify}(\mathsf{vk}, \mu, \mathsf{Sign}(\mathsf{sk}, \mu)) = 1\Big] \geq 1 - \mathsf{negl}(\lambda),$$

where the probability is taken over the random coins of Sign*.*

Finally, we recall the weak and strong Existential Unforgeability under Chosen Message Attack (EU-CMA and sEU-CMA) and the Existential Unforgeability under No Message Attack (EU-NMA) security game for digital signatures.

Definition 8. *Let* $\delta > 0$*. A signature scheme* $(\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ *is said to be* δ*-EU-CMA (resp.* δ*-EU-NMA) secure if no ppt adversary* $\mathcal{A}$ *given* vk *and access to a signing oracle (resp. without access to a signing oracle) has probability* $\geq \delta$ *over the choice of the signing and verification keys* $(\mathsf{vk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$ *and its random coins of outputting* (μ^*, σ^*) *such that*

1. μ^* *was not queried to the signing oracle,*
2. $\mathsf{Verify}(\mathsf{vk}, \mu^*, \sigma^*) = 1$*, i.e., the forged signature must be accepted.*

The scheme is said δ*-EU-CMA secure in the ROM if the above holds when the adversary can also make queries to a random oracle that models some hash function used in the scheme. The probability of forging a signature is also called the advantage of* $\mathcal{A}$*. If condition 1 is replaced with* σ^* *is not an answer of a signature query for* μ^**, the scheme is instead said* δ*-sEU-CMA.*

2.5 Hardness Assumptions

The security of our constructions relies on the hardness of two lattice problems, namely the decisional Learning with Errors problem and the Short Integer Solution problem.

Definition 9 (Learning With Errors). *Let $m, k > 0$ and $q \geq 2$. Let χ be a distribution over $\mathbb{Z}$. The $\mathsf{LWE}_{m,k,\ell,q,\chi}$ assumption states that no (quantum) adversary has non-negligible advantage in distinguishing $(\mathbf{A}, \mathbf{AS}+\mathbf{E})$ from $(\mathbf{A}, \mathbf{U})$, where $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{m\times k})$, $\mathbf{U} \hookleftarrow U(\mathbb{Z}_q^{m\times \ell})$ and $(\mathbf{S}^\top|\mathbf{E}^\top)^\top \hookleftarrow \chi^{k+m\times\ell}$.*

Definition 10 (Short Integer Solution). *Let $m, k, \gamma > 0$ and $q \geq 2$ be a modulus. The $\mathsf{SIS}_{m,k,q,\gamma}$ assumption states that no (quantum) adversary has non-negligible probability of finding $\mathbf{s} \in \mathbb{Z}^k$ such that*

$$\mathbf{As} = \mathbf{0} \bmod q \quad \textit{and} \quad 0 < \|\mathbf{s}\| \leq \gamma \ ,$$

when given $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{m\times k})$ as input.

3 The $\mathsf{G}+\mathsf{G}$ Identification Protocol

In this section, we first describe the $\mathsf{G}+\mathsf{G}$ identification protocol, then prove the required properties to compile it into a signature using the Fiat-Shamir heuristic, and then discuss asymptotic parameters.

3.1 Description of the Scheme

Let us first introduce the parameters of the scheme as well as some notations. Let $m \geq \ell > 0$, $k > m + \ell$ and $\mathbf{J} = \mathbf{J}_{m,\ell}$. Let χ be a distribution over $\mathbb{Z}$. Let $\mathcal{C} \subseteq \mathbb{Z}_2^\ell$ be the challenge space, which we assume to be finite. Let $\sigma, s \geq 0$ and define $\mathbf{\Sigma} : \mathbb{Z}^{k\times\ell} \to \mathbb{R}^{k\times k}$ as

$$\mathbf{\Sigma} : \mathbf{S} \mapsto \sigma^2\mathbf{I}_k - 4s^2\mathbf{S}\mathbf{S}^\top.$$

The scheme is also parametrized by an odd modulus q and an acceptance bound γ.

The $\mathsf{G}+\mathsf{G}$ identification protocol is described in Figure 3. The instance generation algorithm samples a verification key $\mathbf{A} \in \mathbb{Z}_{2q}^{m\times k}$ and a signing key $\mathbf{S} \in \mathbb{Z}^{k\times\ell}$ with small-magnitude coefficients such that $\mathbf{A}\cdot\mathbf{S} = q\mathbf{J} \bmod 2q$. In the first phase of the interaction, the prover samples a vector $\mathbf{y}$ with well-crafted covariance matrix, and sends the commitment $\mathbf{w} = \mathbf{Ay} \bmod 2q$ to the verifier. The protocol is public-coin, i.e., the verifier just samples $\mathbf{c}$ uniformly in the challenge space and sends it to the prover. After receiving $\mathbf{c}$, the prover samples a Gaussian vector $\mathbf{k}$ over the lattice coset $2\mathbf{S}\mathbb{Z}^\ell + \mathbf{c}$. The covariance matrices of $\mathbf{y}$ and $\mathbf{k}$ are set so that the Gaussian plus Gaussian sum is statistically close to a spherical Gaussian distribution.

The first sampling that the prover has to perform is well-defined only if $\mathbf{\Sigma}(\mathbf{S})$ is definite positive, which we show thanks to Lemma 7. The first sampling is implemented using Lemma 4, which requires $\sigma^2 - s^2\sigma_1(\mathbf{S})^2 \geq \sqrt{\ln(2\ell+4)/\pi}$, where we let $\sigma_1(\mathbf{S})$ denote the largest singular value of $\mathbf{S}$. The protocol can then be executed in polynomial time.

```
$\underline{\mathsf{IGen}(1^\lambda)}$:
$\mathbf{A}_1 \hookleftarrow U(\mathbb{Z}_q^{m\times(k-m-\ell)})$
$(\mathbf{S}_1, \mathbf{S}_2) \hookleftarrow \chi^{(k-m-\ell)\times\ell} \times \chi^{m\times\ell}$
$\mathbf{B} \leftarrow \mathbf{A}_1\mathbf{S}_1 + \mathbf{S}_2 \bmod q$
$\mathbf{A} \leftarrow (q\mathbf{J} - 2\mathbf{B}|2\mathbf{A}_1|2\mathbf{I}_m) \in \mathbb{Z}_{2q}^{m\times k}$
$\mathbf{S} \leftarrow (\mathbf{I}_\ell|\mathbf{S}_1^\top|\mathbf{S}_2^\top)^\top \in \mathbb{Z}^{k\times\ell}$
$\mathsf{vk} \leftarrow \mathbf{A}, \mathsf{sk} \leftarrow \mathbf{S}$
return $(\mathsf{vk}, \mathsf{sk})$
```

$\mathsf{P}(\mathbf{A}, \mathbf{S})$		$\mathsf{V}(\mathbf{A})$
$\mathbf{y} \hookleftarrow D_{\mathbb{Z}^k, \mathbf{\Sigma}(\mathbf{S})}$		
$\mathbf{w} \leftarrow \mathbf{A}\mathbf{y} \bmod 2q$	$\xrightarrow{\mathbf{w}}$	
	$\xleftarrow{\mathbf{c}}$	$\mathbf{c} \hookleftarrow U(\mathcal{C})$
$\mathbf{k} \hookleftarrow D_{\mathbb{Z}^\ell, s, -\mathbf{c}/2}$		
$\mathbf{z} \leftarrow \mathbf{y} + 2\mathbf{S}\mathbf{k} + \mathbf{S}\mathbf{c}$	$\xrightarrow{\mathbf{z}}$	Accept if $\mathbf{A}\mathbf{z} = \mathbf{w} + q\mathbf{J}\mathbf{c} \bmod 2q$ and $\|\mathbf{z}\| \leq \gamma$

Fig. 3. The $\mathsf{G}+\mathsf{G}$ Identification Protocol.

Combining this identification protocol with the Fiat-Shamir (without aborts) paradigm, we then obtain a lattice-based signature $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$, as stated in the following Theorem. The correctness and security of the scheme are inherited from the properties of the underlying identification protocol.

Theorem 1. *Let $m \geq \ell > 0$, $k > m+\ell$, $\varepsilon \in (0, 1/2]$, $s \geq \sqrt{2\ln(\ell - 1 + 2\ell/\varepsilon)/\pi}$ and $\sigma \geq \sqrt{8}\sigma_1(\mathbf{S}) \cdot s$ for all $\mathbf{S} \in \mathbb{Z}^{k\times\ell}$ in the range of* IGen. *Let γ and ε_c be such that $\Pr_{\mathbf{z}\hookleftarrow D_{\mathbb{Z}^k,\sigma}}[\|\mathbf{z}\| > \gamma] \leq \varepsilon_c/3$. Let $q > \max(2\gamma, \sigma\cdot\eta_\varepsilon(\mathbb{Z}^m))$ be an odd modulus.*

Then the signature scheme $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ is ε_c-correct and:

- *EU-CMA-secure in the ROM under the $\mathsf{SIS}_{m,k,q,2\gamma}$ assumption. Namely, for any adversary $\mathcal{A}$ against the EU-CMA security of $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ making at most Q_S sign queries and at most Q_H hash queries, there exists an adversary $\mathcal{B}$ against the $\mathsf{SIS}_{m,k,q,2\gamma}$ assumption and an adversary $\mathcal{B}$ against the*

$\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ *assumption such that:*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{EU\text{-}CMA}}(\mathcal{A}) \leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{Q_S} &\left[Q_H \cdot \left(\sqrt{\mathsf{Adv}^{\mathsf{SIS}_{m,k,q,2\gamma}}(\mathcal{B})} + \frac{2}{|\mathcal{C}|}\right)\right] \\ &+ 3Q_S/2 \cdot \sqrt{(Q_H + Q_S + 1) \cdot (s/3)^{-m}} \\ &+ \mathsf{Adv}^{\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}') \ ;\end{aligned}$$

- *EU-CMA-secure in the QROM under the* $\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ *assumption, assuming that* $1/|\mathcal{C}| + (|\mathcal{C}|^2(2\gamma+1)^{2k})/q^m$ *is negligible. Namely, for any quantum adversary* $\mathcal{A}$ *against the EU-CMA security of* $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ *making at most* Q_S *classical sign queries and at most* Q_H *quantum hash queries, there exists an adversary* $\mathcal{B}'$ *against the* $\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ *assumption such that:*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{EU\text{-}CMA}}(\mathcal{A}) \leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{Q_S} &\mathsf{Adv}^{\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}') \\ + \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{Q_S} &8(Q_H+1)^2 \cdot \left(\frac{1}{|\mathcal{C}|} + \frac{|\mathcal{C}|^2(2\gamma+1)^{2k}}{q^m}\right) \\ + 3Q_S/2 \cdot &\sqrt{(Q_H + Q_S + 1) \cdot (s/3)^{-m}} \ .\end{aligned}$$

Moreover, these two bounds holds when $\mathcal{A}$ *is an adversary against the sEU-CMA security of the scheme by adding an extra* $+Q_S \cdot (s/3)^{-m} + \mathsf{Adv}^{\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}')$ $+\mathsf{Adv}^{\mathsf{SIS}_{m,k,q,2\gamma}}(\mathcal{B})$ *term on the right hand side.*

The proof of Theorem 1 follows from Corollaries 1, 2, 3, and 4, which are derived from the properties of the underlying identification protocol proved in Sections 3.2, 3.3, and 3.4, by applying the Fiat-Shamir transform. The Fiat-Shamir transform results are reminded in Appendix A.

3.2 Completeness and Commitment Recoverability

We first show that the $\mathsf{G}+\mathsf{G}$ protocol is complete and commitment recoverable. As a corollary, we obtain that the resulting Fiat-Shamir signature scheme $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ is correct.

Theorem 2. *Let* $m \geq \ell > 0$, $k > m+\ell$, $\varepsilon \in (0, 1/2]$, $s \geq \sqrt{2\ln(\ell - 1 + 2\ell/\varepsilon)/\pi}$ *and* $\sigma \geq \sqrt{8}\sigma_1(\mathbf{S}) \cdot s$ *for all* $\mathbf{S} \in \mathbb{Z}^{k\times\ell}$ *in the range of* IGen*. Let* γ *and* ε_c *be such that* $\Pr_{\mathbf{z}\hookleftarrow D_{\mathbb{Z}^k,\sigma}}[\|\mathbf{z}\| > \gamma] \leq \varepsilon_c/3$*. Let* $q > 2\gamma$ *be an odd modulus. Then the* $\mathsf{G}+\mathsf{G}$ *identification protocol is* ε_c*-complete and achieves commitment-recoverability.*

Proof. First, we note that $\mathbf{AS} = q\mathbf{J} \bmod 2q$ holds for any matrix pair output by IGen. Then, in order to pass the first verification step, a transcript $(\mathbf{w}, \mathbf{c}, \mathbf{z})$ must satisfy:

$$\mathbf{Az} = \mathbf{A}(\mathbf{y} + 2\mathbf{Sk} + \mathbf{Sc}) = \mathbf{w} + \mathbf{0} + q\mathbf{Jc} \bmod 2q \ . \tag{3}$$

In particular, this defines a unique commitment $\mathbf{w} = \mathbf{A}\mathbf{z} - q\mathbf{J}\mathbf{c} \bmod 2q$ such that $(\mathbf{w}, \mathbf{c}, \mathbf{z})$ can be a valid transcript, and $\mathbf{w}$ is efficiently recoverable, by defining Rec as $\mathsf{Rec}(\mathbf{A}, \mathbf{c}, \mathbf{z}) := \mathbf{A}\mathbf{z} - q\mathbf{J}\mathbf{c} \bmod 2q$.

Now, we note that an honestly generated transcript $(\mathbf{w}, \mathbf{c}, \mathbf{z})$ always satisfies Equation (3). The probability preservation property of the Rényi divergence (Equation (1)) and Lemma 7 give the following bound:

$$\begin{aligned}\Pr_{(\mathbf{w},\mathbf{c},\mathbf{z})}[\|\mathbf{z}\| > \gamma] &\leq R_\infty(P_\mathbf{z} \| D_{\mathbb{Z}^k,\sigma}) \cdot \Pr_{\mathbf{z} \hookleftarrow D_{\mathbb{Z}^k,\sigma}}[\|\mathbf{z}\| > \gamma] \\ &\leq \frac{1+\varepsilon}{1-\varepsilon} \cdot \Pr_{\mathbf{z} \hookleftarrow D_{\mathbb{Z}^k,\sigma}}[\|\mathbf{z}\| > \gamma] \\ &\leq \frac{1+1/2}{1-1/2} \cdot \Pr_{\mathbf{z} \hookleftarrow D_{\mathbb{Z}^k,\sigma}}[\|\mathbf{z}\| > \gamma].\end{aligned}$$

Then the probability that an honest transcript $(\mathbf{w}, \mathbf{c}, \mathbf{z})$ be rejected at most $\leq 3 \cdot \Pr_{\mathbf{z} \hookleftarrow D_{\mathbb{Z}^k,\sigma}}[\|\mathbf{z}\| > \gamma]$. □

We then obtain the following corollary.

Corollary 1. *Using the same assumptions as in Theorem 2, the resulting signature scheme* $\mathsf{FS}[\mathsf{G} + \mathsf{G}, H]$ *is* ε_c*-correct.*

Note that correctness of $\mathsf{FS}[\mathsf{G} + \mathsf{G}, H]$ does not require to assume that H is modeled as a random oracle, as Lemma 7 holds without relying on the randomness of $\mathbf{c}$. This is in contrast to Lemma 8 that generically considers completeness of signatures obtained using the Fiat-Shamir transform.

3.3 Honest-Verifier Zero-Knowledge and Commitment Min-Entropy

We now show that the $\mathsf{G} + \mathsf{G}$ protocol is HVZK and has large commitment min-entropy. As a corollary, we obtain that the signature scheme $\mathsf{FS}[\mathsf{G} + \mathsf{G}, H]$ is EU-CMA-secure provided it is EU-NMA-secure.

Theorem 3. *Let* $m \geq \ell > 0$, $k > m + \ell$, $\varepsilon \in (0, 1/2]$, $s \geq \sqrt{2\ln(\ell - 1 + 2\ell/\varepsilon)/\pi}$ *and* $\sigma \geq \sqrt{8}\sigma_1(\mathbf{S}) \cdot s$ *for all* $\mathbf{S} \in \mathbb{Z}^{k\times\ell}$ *in the range of* IGen. *Let* $q > \sigma \cdot \eta_\varepsilon(\mathbb{Z}^m)$ *be an odd modulus. Then the* $\mathsf{G} + \mathsf{G}$ *identification protocol satisfies:*

- $(1 + \varepsilon)/(1 - \varepsilon)$*-divergence* HVZK*,*
- $2\varepsilon/(1 - \varepsilon)$*-*$\mathsf{HVZK}$*.*

In addition, its commitment min-entropy is $\geq m \cdot \log(s/3)$.

Proof. We prove both properties separately. We start by proving HVZK, which is inherited from Lemma 7 and then focus on commitment min-entropy.

HVZK. The simulator on input a challenge $\mathbf{c} \in \mathcal{C}$ and a public matrix $\mathbf{A}$ samples $\mathbf{z} \hookleftarrow D_{\mathbb{Z}^k,\sqrt{2}\sigma}$, sets $\mathbf{w} = \mathbf{A}\mathbf{z} - q\mathbf{J}\mathbf{c}$ and returns $(\mathbf{w}, \mathbf{c}, \mathbf{z})$ as a transcript. As everything here is a function of $\mathbf{z}$ and $\mathbf{c}$, we can rely on Lemma 7. The bounds from the above claim are immediately inherited from the latter lemma by applying the data processing inequalities (which we recall in Equation (2) for the Rényi divergence – the same inequality holds replacing the Rényi divergence by the statistical distance). This completes the zero-knowledge analysis.

Commitment Min-Entropy. Let $\mathbf{w} \in \mathbb{Z}_{2q}^m$ and $(Y_1^\top, Y_2^\top)^\top \hookleftarrow D_{\mathbb{Z}^k, \boldsymbol{\Sigma}(\mathbf{S})}$, where Y_1 takes values in $\mathbb{Z}^{k-m}$. Given a matrix $\mathbf{A} = (\mathbf{A}_0 | 2\mathbf{I}_m) \in \mathbb{Z}_{2q}^{m \times k}$, it holds that

$$\Pr_{(Y_1,Y_2)}[\mathbf{A}_0 Y_1 + 2Y_2 = \mathbf{w} \bmod 2q] = \Pr_{(Y_1,Y_2)}[2Y_2 = \mathbf{w} - \mathbf{A}_0 Y_1 \bmod 2q]$$
$$\leq \Pr_{(Y_1,Y_2)}[Y_2 = (\mathbf{w} - \mathbf{A}_0 Y_1)\zeta \bmod q] \ ,$$

where ζ is the modular inverse of 2 mod q. Hence, the min-entropy of the commitment is $\geq H_\infty(Y_2 \bmod q | Y_1)$ and we move on to bounding the latter quantity from below. Note that there exist $\sigma \geq \sigma_1 \geq \cdots \geq \sigma_m \geq (\sigma^2 - s^2\sigma_1(\mathbf{S})^2)^{1/2}$ and $\mathbf{Q} \in \mathbb{R}^{m \times m}$ orthogonal such that

$$\boldsymbol{\Sigma}(\mathbf{S}) = \mathbf{Q} \begin{pmatrix} \sigma_1^2 & & \\ & \ddots & \\ & & \sigma_m^2 \end{pmatrix} \mathbf{Q}^\top .$$

Let $\mathbf{y}_1 \in \mathbb{Z}^{k-m}$ be fixed. The distribution of Y_2 conditioned on $Y_1 = \mathbf{y}_1$ is exactly $D_{\mathbb{Z}^m, \overline{\boldsymbol{\Sigma}}, \overline{\mathbf{c}}}$, as defined in Lemma 3 (with $\mathbf{c} = \mathbf{0}$). Let $\overline{\sigma}_1^2$ (resp. $\overline{\sigma}_m^2$) be the largest (resp. smallest) eigenvalue of $\overline{\boldsymbol{\Sigma}}$ and $\overline{\mathbf{c}} = (\overline{c}_1, \ldots, \overline{c}_m)^\top$. We are interested in obtaining an upper bound on $\rho_{\overline{\boldsymbol{\Sigma}}, \overline{\mathbf{c}}}(\mathbf{z} + q\mathbb{Z}^m)/\rho_{\overline{\boldsymbol{\Sigma}}, \overline{\mathbf{c}}}(\mathbb{Z}^m)$ for all $\mathbf{z} \in (-q/2, q/2]^m$. Indeed, this quantity corresponds to all values taken by the probability mass function of the random variable Y_2 mod q conditioned on $Y_1 = \mathbf{y}_1$, namely $\Pr_{Y_2|Y_1=\mathbf{y}_1}(Y_2 = \mathbf{z} \bmod q) = \sum_{\mathbf{u} \in q\mathbb{Z}^m} \rho_{\overline{\boldsymbol{\Sigma}}, \overline{\mathbf{c}}}(\mathbf{z} + \mathbf{u})/\rho_{\overline{\boldsymbol{\Sigma}}, \overline{\mathbf{c}}}(\mathbb{Z}^m)$.

As $\overline{\boldsymbol{\Sigma}}^{-1}$ is the bottom right submatrix of $\boldsymbol{\Sigma}^{-1}$ of size $m \times m$, it holds that for any $\mathbf{y} \in \mathbb{R}^m$, we have $\mathbf{y}^\top \overline{\boldsymbol{\Sigma}}^{-1} \mathbf{y} \in \|\mathbf{y}\|^2 \cdot [1/\sigma_1^2, 1/\sigma_m^2]$. Hence all singular values $\overline{\sigma}_i$ of $\overline{\boldsymbol{\Sigma}}$ lie in $[(\sigma^2 - s^2\sigma_1(\mathbf{S})^2)^{1/2}, \sigma]$. Thanks to the theorem assumptions, we obtain that all $\overline{\sigma}_i$'s are above $\eta_\varepsilon(\mathbb{Z}^m)$. Using Lemma 6, it holds that

$$\rho_{\overline{\boldsymbol{\Sigma}}, \overline{\mathbf{c}}}(\mathbb{Z}^m) \geq (1 - \varepsilon) \cdot \sqrt{\det \overline{\boldsymbol{\Sigma}}} \geq (1 - \varepsilon) \cdot \left(\sigma^2 - s^2\sigma_1(\mathbf{S})^2\right)^{m/2} .$$

The latter is $\geq (1 - \varepsilon) \cdot (s\sigma_1(\mathbf{S}))^m$, by assumption on σ. For the numerator, we first use Lemma 6 once more, to obtain:

$$\rho_{\overline{\boldsymbol{\Sigma}}, \overline{\mathbf{c}}}(\mathbf{z} + q\mathbb{Z}^m) \leq \rho_{\overline{\boldsymbol{\Sigma}}}(q\mathbb{Z}^m) = 1 + \rho_{\overline{\boldsymbol{\Sigma}}}(q\mathbb{Z}^m \setminus \{\mathbf{0}\}) \leq 1 + \rho_\sigma(q\mathbb{Z}^m \setminus \{\mathbf{0}\}) \ .$$

Rewriting the assumption on q we have $1/\sigma > \eta_\varepsilon((1/q)\mathbb{Z}^m)$. Note that the dual lattice of $(1/q)\mathbb{Z}^m$ is $q\mathbb{Z}^m$. Hence, we have $\rho_\sigma(q\mathbb{Z}^m \setminus \{\mathbf{0}\}) \leq \varepsilon$ by definition of the smoothing parameter. The result follows by noting that for any $\mathbf{S}$ in the range of IGen, we have $\sigma_1(\mathbf{S}) \geq 1$ as $\mathbf{S}$ includes an identity matrix. We note also that $(1 + \varepsilon)/(1 - \varepsilon) \leq 3$ for any $\varepsilon \in [0, 1)$. □

We then obtain the following corollary as an application of Theorem 7.

Corollary 2. *Using the same assumptions as in Theorem 3 the resulting signature scheme* $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ *is* EU-CMA*-secure in the QROM, provided it is* EU-NMA*-secure. Namely, for any (possibly quantum) adversary* $\mathcal{A}$ *against the* EU-CMA *security of* $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ *making at most* Q_S *(classical) sign queries and at most* Q_H *(possibly quantum) hash queries, there exists an adversary* $\mathcal{B}$ *against the* EU-NMA *security of* $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ *such that:*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{EU\text{-}CMA}}(\mathcal{A}) \leq &\left(1+\frac{2\varepsilon}{1-\varepsilon}\right)^{Q_S} \mathsf{Adv}^{\mathsf{EU\text{-}NMA}}(\mathcal{B}) \\ &+ 3Q_S/2 \cdot \sqrt{(Q_H+Q_S+1)\cdot(s/3)^{-m}} \ .\end{aligned}$$

3.4 Lossy Identification Scheme and Lossy Special Soundness

To complete the analysis, we show that (i) $\mathsf{G}+\mathsf{G}$ is a lossy identification scheme and that (ii) $\mathsf{G}+\mathsf{G}$ is lossy-special sound. As a corollary, we obtain that the signature scheme $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ is EU-NMA-secure in the ROM, and in the QROM under some parameters constraint.

Theorem 4. *Let* $m \geq \ell > 0$*,* $k > m+\ell$*,* $\varepsilon \in (0, 1/2]$*,* $s \geq \sqrt{2\ln(\ell-1+2\ell/\varepsilon)/\pi}$ *and* $\sigma \geq \sqrt{8}\sigma_1(\mathbf{S})\cdot s$ *for all* $\mathbf{S} \in \mathbb{Z}^{k\times\ell}$ *in the range of* IGen*. Let* $\gamma > 0$ *and* $q > 2\gamma$ *be an odd modulus. Then the* $\mathsf{G}+\mathsf{G}$ *identification protocol:*

- *has key-indistinguishability, under the* $\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ *assumption,*
- *is* $\varepsilon_{\mathsf{ls}}$*-lossy sound for*

$$\varepsilon_{\mathsf{ls}} \ = \ \frac{1}{|\mathcal{C}|} + \frac{|\mathcal{C}|^2(2\gamma+1)^{2k}}{q^m} \quad ,$$

  and is thus a lossy identification scheme,
- *is lossy-special-sound, under the* $\mathsf{SIS}_{m,k,q,2\gamma}$ *assumption,*

Proof. We first set $\mathsf{G}+\mathsf{G}$ in lossy mode and then show that it achieves the two flavours of soundness.

Key-indistinguishability. Compared to IGen, the lossy key generation algorithm $\mathsf{LossyIGen}$ only modifies the generation of $\mathbf{B}$. Recall that in IGen, the latter is defined as $\mathbf{B} \leftarrow \mathbf{A}_1\mathbf{S}_1 + \mathbf{S}_2$, with $\mathbf{A}_1 \hookleftarrow U(\mathbb{Z}_q^{m\times(k-m-\ell)})$ and the components of the signing key $(\mathbf{S}_1, \mathbf{S}_2) \hookleftarrow \chi_\eta^{(k-m-\ell)\times\ell} \times \chi_\eta^{m\times\ell}$. The lossy key generation algorithm $\mathsf{LossyIGen}$ samples it as $\mathbf{B} \hookleftarrow U(\mathbb{Z}_q^{m\times\ell})$. Lossy verification keys are computationally indistinguishable from non-lossy ones, under the $\mathsf{LWE}_{k-m-\ell,m,\ell,\eta,q}$ assumption.

$\varepsilon_{\mathsf{ls}}$**-Lossy Soundness.** First note that, if the lossy verification key $\mathbf{A}$ is such that, for all commitment $\mathbf{w}$, there exists at most one challenge $\mathbf{c}$ such that there exists $\mathbf{z}$ with $(\mathbf{w}, \mathbf{c}, \mathbf{z})$ passing verification, then, as the challenge is sampled uniformly and independently of $\mathbf{w}$, an (unbounded) prover cannot pass verification, except with probability at most $1/|\mathcal{C}|$.

We then focus on proving that the above holds with overwhelming probability over the choice of the lossy key $\mathbf{A}$. By contradiction, assume there exists $\mathbf{w}, \mathbf{c}_0, \mathbf{c}_1, \mathbf{z}_0, \mathbf{z}_1$ with $\|\mathbf{z}_0\|, \|\mathbf{z}_1\| \leq \gamma$ and $\mathbf{c}_0 \neq \mathbf{c}_1 \in \mathcal{C}$, such that we have both $\mathbf{A}\mathbf{z}_0 = \mathbf{w} + q\mathbf{J}\mathbf{c}_0 \bmod 2q$ and $\mathbf{A}\mathbf{z}_1 = \mathbf{w} + q\mathbf{J}\mathbf{c}_1 \bmod 2q$. Then, we have:

$$\mathbf{A}(\mathbf{z}_0 - \mathbf{z}_1) = q\mathbf{J}(\mathbf{c}_1 - \mathbf{c}_0) \bmod 2q \ .$$

Recall that $\mathbf{A}$ is of the form $(q\mathbf{J} - 2\mathbf{B}|2\mathbf{A}_1|2\mathbf{I}_m)$, with $\mathbf{A}_1, \mathbf{B}$ uniform over $\mathbb{Z}_q$. Hence, the matrix $\mathbf{A} \bmod q$ is of the form $(\mathbf{B}|\mathbf{A}_1|\mathbf{I}_m)$, since q is odd. Then the above implies that $(\mathbf{B}|\mathbf{A}_1|\mathbf{I}_m)(\mathbf{z}_0 - \mathbf{z}_1) = \mathbf{0} \bmod q$ with $\mathbf{z}_0 - \mathbf{z}_1 \neq \mathbf{0} \bmod q$. This happens with probability at most $1/q^m$.

To conclude, note that there are at most $(2\gamma+1)^{2k} \cdot |\mathcal{C}|^2$ choices for $\mathbf{z}_0, \mathbf{z}_1, \mathbf{c}_0$ and $\mathbf{c}_1$. A union bound therefore implies that the probability over $\mathbf{A}$ that there is a commitment with at least two challenges permitting valid transcripts is at most $|\mathcal{C}|^2(2\gamma+1)^{2k}/q^m$. Our lossy identification scheme is then $\varepsilon_{\mathsf{ls}}$-lossy-sound, with

$$\varepsilon_{\mathsf{ls}} \ \leq \ \frac{1}{|\mathcal{C}|} + \frac{|\mathcal{C}|^2(2\gamma+1)^{2k}}{q^m} \ ,$$

which shows combined with key-indistinguishability, that it is lossy.

Lossy Special Soundness. Assume there exists a PPT adversary $\mathcal{A}$ which, given a lossy verification key $\mathsf{vk} = \mathbf{A}$, produces two valid transcripts $(\mathbf{w}, \mathbf{c}_0, \mathbf{z}_0)$ and $(\mathbf{w}, \mathbf{c}_1, \mathbf{z}_1)$ with $\mathbf{c}_0 \neq \mathbf{c}_1$. It can be turned into an $\mathsf{SIS}_{m,k,q,2\gamma}$ solver. Indeed, by definition, such transcripts satisfy $\mathbf{A}(\mathbf{z}_0 - \mathbf{z}_1) = q\mathbf{J}(\mathbf{c}_1 - \mathbf{c}_0) \bmod 2q$.

Notice that we have $\mathbf{A}(\mathbf{z}_0 - \mathbf{z}_1) = \mathbf{0} \bmod q$, which implies that $\mathbf{z}_0 - \mathbf{z}_1$ is a solution to the (uniformly sampled) SIS instance defined by $\mathbf{A}$. In addition, when reducing modulo 2, we also have $\mathbf{A}(\mathbf{z}_0 - \mathbf{z}_1) = \mathbf{J}(\mathbf{c}_1 - \mathbf{c}_0) \bmod 2$, which implies that $\mathbf{z}_0 \neq \mathbf{z}_1$. Finally, note that the condition on γ implies that $\|\mathbf{z}_0 - \mathbf{z}_1\| \leq 2\gamma$ (as transcript validity implies $\|\mathbf{z}\| \leq \gamma$), and that $\mathbf{z}_0 - \mathbf{z}_1 \neq \mathbf{0} \bmod q$.

Hence, there exists an adversary $\mathcal{B}$ against the $\mathsf{SIS}_{m,k,q,2\gamma}$ problem such that:

$$\mathsf{Adv}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{SIS}_{m,k,q,2\gamma}}(\mathcal{B}) \ ,$$

which completes the proof of the theorem. □

We then obtain the following corollary as an application of Lemma 10.

Corollary 3. *Using the same assumptions as in Theorem 4, the resulting signature scheme* $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ *is EU-NMA-secure, in the ROM. Namely, for any adversary* $\mathcal{A}$ *against the EU-NMA security of* $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$*, there exists an adversary* $\mathcal{B}$ *against the* $\mathsf{SIS}_{m,k,q,2\gamma}$ *assumption as well as an adversary* $\mathcal{B}'$ *against the* $\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ *assumption such that:*

$$\mathsf{Adv}^{\mathsf{EU\text{-}NMA}}(\mathcal{A}) \leq Q_H \cdot \left(\sqrt{\mathsf{Adv}^{\mathsf{SIS}_{m,k,q,2\gamma}}(\mathcal{B})} + \frac{2}{|\mathcal{C}|}\right) + \mathsf{Adv}^{\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}') \ .$$

We also obtain the following corollary as an application of Theorem 8.

Corollary 4. *Using the same assumptions as in Theorem 4, and if $\varepsilon_{\mathsf{ls}}$ is negligible, the signature scheme $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ is EU-NMA-secure, in the QROM. Namely, for any (possibly quantum) adversary $\mathcal{A}$ against the EU-NMA security of $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ making at most Q_H (possibly quantum) hash queries, there exists a quantum adversary $\mathcal{B}$ against the $\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ assumption such that:*

$$\mathsf{Adv}^{\textsf{EU-NMA}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}) + 8(Q_H+1)^2 \cdot \left(\frac{1}{|\mathcal{C}|} + \frac{|\mathcal{C}|^2(2\gamma+1)^{2k}}{q^m}\right) .$$

To conclude this section, we introduce an additional assumption of a similar flavour as the $\mathsf{SelfTargetMSIS}$ assumption [KLS18], which allows to directly prove EU-NMA-security of $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ in the QROM as it is (up to LWE) the EU-NMA security game of the resulting signature. As for $\mathsf{SelfTargetMSIS}$, this problem can be related in the ROM to SIS, using the special soundness property of the scheme.

Definition 11 ($\mathsf{GpGSelfTargetSIS}$). *Let $m \geq \ell > 0$, $k > m+\ell$. Let $\gamma > 0$ and $q > 2\gamma$ be an odd modulus. The $\mathsf{GpGTargetSIS}_{m,k,\ell,\gamma,q}$ states that given a matrix $\mathbf{A} := (q\mathbf{J} - 2\mathbf{B}|2\mathbf{A}_1|2\mathbf{I}_m) \in \mathbb{Z}_{2q}^{m\times k}$, where $\mathbf{A}_1 \hookleftarrow U(\mathbb{Z}_q^{m\times(k-m-\ell)})$ and $\mathbf{B} \hookleftarrow U(\mathbb{Z}_q^{m\times \ell})$, and oracle access to a hash function H, it is computationally hard to find $\mathbf{c} \in \mathcal{C}$, $\mathbf{z} \in \mathbb{Z}^k$ and $\mu \in \{0,1\}^\star$ such that $H(\mathbf{A}\mathbf{z} - q\mathbf{J}\mathbf{c}, \mu) = \mathbf{c}$ and $\|\mathbf{z}\| \leq \gamma$.*

3.5 Computational Unique Responses

Finally, in order to show the strong unforgeability of the scheme, we show that $\mathsf{G}+\mathsf{G}$ has Computational Unique Response (CUR).

Theorem 5. *Let $m \geq \ell > 0$, $k > m+\ell$, $\varepsilon \in (0, 1/2]$, $s \geq \sqrt{2\ln(\ell - 1 + 2\ell/\varepsilon)/\pi}$ and $\sigma \geq \sqrt{8}\sigma_1(\mathbf{S}) \cdot s$ for all $\mathbf{S} \in \mathbb{Z}^{k\times\ell}$ in the range of IGen. Let $\gamma > 0$ and $q > 2\gamma$ be an odd modulus. Then the $\mathsf{G}+\mathsf{G}$ identification protocol satisfies the CUR property under both SIS and LWE assumptions. Namely, for any (possibly quantum) adversary $\mathcal{A}$ against the CUR property, there exists an adversary $\mathcal{B}$ against $\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ and an adversary $\mathcal{B}'$ against $\mathsf{SIS}_{m,k,q,2\gamma}$ such that:*

$$\mathsf{Adv}^{\mathsf{CUR}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}) + \mathsf{Adv}^{\mathsf{SIS}_{m,k,q,2\gamma}}(\mathcal{B}').$$

Proof. We first change the verification key to the lossy one, as defined in the proof of Theorem 4. This means that the matrix $2^{-1}\mathbf{A}$ is now uniform modulo q among matrices in Hermite Normal Form. Now, given elements $(\mathbf{w}, \mathbf{c}, \mathbf{z}_0, \mathbf{z}_1)$ such that $\mathbf{A}\mathbf{z}_0 - q\mathbf{J}\mathbf{c} = \mathbf{A}\mathbf{z}_1 - q\mathbf{J}\mathbf{c} \bmod 2q$, we see that $\mathbf{A}(\mathbf{z}_0 - \mathbf{z}_1) = \mathbf{0} \bmod q$. This means that breaking the CUR property of the scheme implies finding a SIS solution for the instance defined by $2^{-1}\mathbf{A} \bmod q$. □

Corollary 5. *Using the same assumptions as in Theorem 5 the resulting signature scheme $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ is sEU-CMA-secure in the QROM, provided it is EU-NMA-secure. Namely, for any (possibly quantum) adversary $\mathcal{A}$ against the*

EU-CMA security of $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$ *making at most* Q_S *(classical) sign queries and at most* Q_H *(possibly quantum) hash queries, there exist an adversary* $\mathcal{B}$ *against the EU-NMA security of* $\mathsf{FS}[\mathsf{G}+\mathsf{G}, H]$*, an adversary* $\mathcal{B}'$ *against* $\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}$ *and an adversary* $\mathcal{B}''$ *against* $\mathsf{SIS}_{m,k,q,2\gamma}$ *such that:*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{EU\text{-}CMA}}(\mathcal{A}) \leq & \left(1+\frac{2\varepsilon}{1-\varepsilon}\right)^{Q_S} \mathsf{Adv}^{\mathsf{EU-NMA}}(\mathcal{B}) \\ & + 3Q_S/2 \cdot \sqrt{(Q_H+Q_S+1)\cdot(s/3)^{-m}} + Q_S \cdot (s/3)^{-m} \\ & + \mathsf{Adv}^{\mathsf{LWE}_{k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}') + \mathsf{Adv}^{\mathsf{SIS}_{m,k,q,2\gamma}}(\mathcal{B}'') \ .\end{aligned}$$

3.6 Asymptotic Parameters Analysis

Our analysis above is applicable to the following instantiation of parameters, as a function of the security parameter λ and the number of signature queries Q_S. We assume Q_S to be a large polynomial in λ. We consider k, ℓ, m linear in λ. We set χ as $D_{\mathbb{Z},\sqrt{k}}$ with tailcutting to get samples in $\{-k,\ldots,0,\ldots,k\}$ with overwhelming probability. We let $\varepsilon = 1/Q_S$.

We make the security of the $\mathsf{G}+\mathsf{G}$ scheme rely on the following two assumptions. First, the $\mathsf{LWE}_{k-m-\ell,k,\ell,q,\chi}$ assumption, where $\sqrt{k} = \alpha q$. This LWE parametrization is compatible with the reduction from worst-case lattice problems from [Reg09]. Second, the $\mathsf{SIS}_{m,k,\beta}$ assumption, where $\beta = O(\sqrt{k}\sigma)$. The SIS parametrization is compatible with the reductions from worst-case lattice problems from [MR07,GPV08] when $q \geq \Omega(\sqrt{k}\beta)$. The hardness of both problems is balanced out when $\alpha \approx 1/\beta$.

Further, the distribution of $\mathbf{z}$ is centered Gaussian with standard deviation $\sigma = 4\sigma_1(\mathbf{S})\sqrt{\ln(\ell-1+2\ell/\varepsilon)/\pi}$, which is $O(\sigma_1(\mathbf{S})\sqrt{\log(Q_S\lambda)})$. Moreover as $\sigma_1(\mathbf{S}) = O(\lambda)$, the norm of $\mathbf{z}$ is at most $\beta = O(\lambda^{3/2}\log^{1/2} Q_S)$. Finally, we set $q = \Theta(\lambda^2 \log^{1/2} Q_S)$.

The verification key and a signature respectively have bit-sizes $O(\lambda^2\log\lambda)$ and $O(\lambda\log\lambda)$.

4 Optimizations and Concrete Parameters

In order to decrease the sizes of a lattice-based scheme, a common approach is to replace $\mathbb{Z}$ with a cyclotomic polynomial ring of the form $\mathcal{R} = \mathbb{Z}[x]/(1+x^n)$, where n is a power of 2, and to rely on the intractability of the module versions of SIS and LWE [BGV12,LS15]. Gaussian distributions are extended by considering the coefficients of the polynomials.

The parameters and data provided in this section are computed using scripts available at `https://github.com/jdevevey/GplusG`.

4.1 Description of the Module-Based Scheme

In this section, we propose parameters for an optimized, module version of the $\mathsf{G}+\mathsf{G}$ signature, that we present in Figure 4.

As in Section 3, let $m > 0$, $k > m+1$ and $\ell = 1$. Let $\mathbf{j} = (\zeta^*, 0, \ldots, 0) \in \mathcal{R}^m$, where $\zeta = 1 + x^{n/2}$ and $\zeta^* = 1 - x^{n/2}$ satisfy $\zeta^*\zeta = 2 \bmod 1 + x^n$. The challenge space is $\mathcal{R}/\zeta^*\mathcal{R}$. We let $\eta > 0$ and $\chi_\eta = U(\{y \in \mathcal{R} |\ \|y\|_\infty \leq \eta\})$. Given $s \in \mathcal{R}$, we define $\mathsf{rot}(\zeta s)$ as the $n \times n$ matrix whose (i,j)-th entry is the coefficient of degree $n-1-j$ of $x^i \cdot \zeta s \bmod 1+x^n$. This matrix maps the coefficient embedding of a polynomial c to the coefficient embedding of sc. We extend this definition to vectors coordinate-wise and we define $\mathbf{\Sigma}(\mathbf{s}) = \sigma^2 \mathbf{I}_{nk} - s^2 \mathbf{S}\mathbf{S}^\top$, where $\mathbf{S} = \mathsf{rot}(\zeta \mathbf{s})$.

Following the key generation algorithm from [CCD+23], we introduce a function Decomp. On input a vector $\mathbf{b}$ and a bit d, it returns $(\mathbf{0}, \mathbf{b})$ if $d = 0$. Otherwise, it computes $\mathbf{b}_1$ and $\mathbf{b}_0$ as follows. For each coordinate b_i of $\mathbf{b}$, if $b_i = 0 \bmod 2$ then the corresponding coordinate of $\mathbf{b}_1$ is set to b_i and the one of $\mathbf{b}_0$ is set to 0. Otherwise, the coordinate of $\mathbf{b}_0$ is set to the nearest multiple of 4 and the one of $\mathbf{b}_1$ is the element in $\{-1, 1\}$ such that $\mathbf{b}_0 + \mathbf{b}_1 = \mathbf{b}$. In the case where $d = 0$, the vector $\mathbf{b}_0$ can be omitted for the verification key, decreasing its size, at the cost of increasing the largest singular value of $\zeta\mathbf{s}$. However, having balanced probability of getting $+1$'s and -1's in $\mathbf{b}_0$ mitigates this increase as opposed to taking it as the lowest bit of $\mathbf{b}$. Finally, the vector $\mathbf{a}$ is used to re-randomize the first column, and is actually useless when $d = 0$.

The above leads to the signature scheme presented in Figure 4.

```
KeyGen(1^λ) :
  $(\mathbf{a}|\mathbf{A}_0) \hookleftarrow U(\mathcal{R}_q^{m \times k-m})$
  do $(\mathbf{s}_1, \mathbf{s}_2) \hookleftarrow \chi_\eta^{k-m-1} \times \chi_\eta^m$
      $\mathbf{b} \leftarrow \mathbf{a} + \mathbf{A}_0\mathbf{s}_1 + \mathbf{s}_2 \bmod q$
      $(\mathbf{b}_0, \mathbf{b}_1) \leftarrow \mathsf{Decomp}(\mathbf{b}, d)$
      $\mathbf{s} \leftarrow (1|\mathbf{s}_1^\top|\mathbf{s}_2^\top - \mathbf{b}_0^\top)^\top \in \mathcal{R}_{2q}^k$
  while $\sigma_1(\mathsf{rot}(\zeta\mathbf{s})) \geq S$
  $\mathbf{A} \leftarrow (2(\mathbf{a} - \mathbf{b}_1) + q\mathbf{j}|2\mathbf{A}_0|2\mathbf{I}_m)$
  return $(\mathsf{vk}, \mathsf{sk}) = (\mathbf{A}, \mathbf{s})$
```

```
Sign(A, s, μ) :
  $\mathbf{y} \hookleftarrow D_{\mathcal{R}^k, \mathbf{\Sigma}(\mathbf{s})}$
  $\mathbf{w} \leftarrow \mathbf{A}\mathbf{y} \bmod 2q$
  $c \leftarrow H(\mathbf{w}, \mu)$
  $u \hookleftarrow D_{\mathcal{R}, s, -\zeta^* \cdot c/2}$
  $\mathbf{z} \leftarrow \mathbf{y} + (\zeta u + c)\mathbf{s}$
  return $(\mathbf{z}, c)$
```

```
Verify(A, μ, z, c) :
  $\mathbf{w} \leftarrow \mathbf{A}\mathbf{z} - qc\mathbf{j} \bmod 2q$
  if $\mathbf{c} = H(\mathbf{w}, \mu)$
  and $\|\mathbf{z}\| \leq \gamma$ then
      return 1
  end if
  return 0
```

Fig. 4. The Module $\mathsf{G} + \mathsf{G}$ Signature Scheme.

Beyond relying on polynomial rings and Decomp, we consider various improvements and optimizations, which we discuss now.

- **KeyGen:** The key generation step includes a rejection sampling step. The threshold S will be set such that about 50% of the keys will be rejected. This helps controlling the upper bound on the smoothing parameter of the secret lattice.
- **Sign:** Instead of computing $\mathbf{z} = \mathbf{y} + (2u + c)\mathbf{s}$, we compute $\mathbf{z} = \mathbf{y} + (\zeta u + c)\mathbf{s}$. As $\mathbf{A}\mathbf{s} = \mathbf{j} \bmod 2q$, we have $\zeta\mathbf{A}\mathbf{s} = \mathbf{0} \bmod 2q$ by definition of $\mathbf{j}$. Thus, the identity $\mathbf{A}\mathbf{z} - qc\mathbf{j} = \mathbf{A}\mathbf{y} \bmod 2q$ still holds. The main advantage of this modification is that the secret lattice is now $\zeta\mathbf{s}\mathcal{R}$ instead of $2\mathbf{s}\mathcal{R}$, whose smoothing parameter is a factor $\sqrt{2}$ smaller.

Verify: The verification bound is set to $\gamma = 1.01 \cdot \sqrt{nk}\sigma$, and the signer may verify that its signature is accepted before outputting it, up to restarting in the somewhat rare event that it is not.

An analysis similar to the one from the previous section would bring the following result. We omit the QROM analysis relying on the lossy-soundness, as the concrete parameters we propose in the next section are outside of the parameters range required for this analysis to hold.

Theorem 6. *Let $n > 0$ be a power of two and $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$. Let $m > 0$, $k > m + 1$, $\varepsilon \in (0, 1/2]$, $s \geq \sqrt{2 \ln(n - 1 + 2n/\varepsilon)/\pi}$ and $\sigma \geq \sqrt{2}\sigma_1(\mathbf{S}) \cdot s$ for all $\mathbf{S} = \mathsf{rot}(\zeta\mathbf{s}) \in \mathbb{Z}^{kn \times n}$ such that $(\mathbf{A}, \mathbf{s})$ is in the range of* KeyGen. *Let γ and ε_c be such that $\Pr_{\mathbf{z} \hookleftarrow D_{\mathcal{R}^k,\sigma}}[\|\mathbf{z}\| > \gamma] \leq \varepsilon_c/3$. Let $q > \max(2\gamma, \sigma \cdot \eta_\varepsilon(\mathbb{Z}^{mn}))$ be an odd modulus.*

Then the signature scheme from Figure 4 is ε_c-correct and EU-CMA-secure in the ROM under the $\mathsf{MSIS}_{n,m,k,q,2\gamma}$ and $\mathsf{MLWE}_{n,k-m-\ell,m,\ell,\chi,q}$ assumptions. Namely, for any adversary $\mathcal{A}$ against the EU-CMA security of $\mathsf{FS}[\mathsf{G} + \mathsf{G}, H]$ making at most Q_S sign queries and at most Q_H hash queries, there is an adversary $\mathcal{B}$ against the $\mathsf{MSIS}_{n,m,k,q,2\gamma}$ assumption and an adversary $\mathcal{B}'$ against the $\mathsf{MLWE}_{n,k-m-\ell,m,\ell,\chi,q}$ assumption such that:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{EU\text{-}CMA}}(\mathcal{A}) \leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{Q_S} &\left[Q_H \cdot \left(\sqrt{\mathsf{Adv}^{\mathsf{MSIS}_{n,m,k,q,2\gamma}}(\mathcal{B})} + \frac{2}{|\mathcal{C}|}\right)\right] \\ &+ \frac{3}{2} Q_S \cdot \sqrt{(Q_H + Q_S + 1) \cdot (s/3)^{-mn}} + \mathsf{Adv}^{\mathsf{MLWE}_{n,k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}').\end{aligned}$$

*This holds when $\mathcal{A}$ is an adversary against the **sEU-CMA** security of the scheme by adding "$+Q_S \cdot (s/3)^{-m} + \mathsf{Adv}^{\mathsf{MLWE}_{n,k-m-\ell,m,\ell,\chi,q}}(\mathcal{B}') + \mathsf{Adv}^{\mathsf{MSIS}_{n,m,k,q,2\gamma}}(\mathcal{B})$" on the right hand side.*

4.2 Concrete Parameters

We now give concrete parameters and estimates of the public key and signature sizes resulting from these optimizations in Table 1. This gives rise to the following estimates. The script we used is derived from the one provided with Dilithium [DKL+18] and is available as supplementary material. We made the following additional assumptions:

- We use the compression technique from [BG14] to get rid of the lower $\log \alpha$ bits of the signature, except the lowest.[4] The hint resulting from the compression technique is assumed to follow a Gaussian distribution whose standard deviation is $\sqrt{2}\sigma/\alpha$. The technique presented in [DKL+18] can be readily adapted to the mod $2q$ setting. This comes at the cost of increasing the verification bound to $\gamma = 1.01 \cdot \sqrt{nk}\sigma + \sqrt{nm}(1 + \alpha/4)$ to take into account the inaccuracy of the commitment recovered by the verifier.

[4] As our key generation algorithm outputs a $\mathbf{A}$ with $2\mathbf{I}_m$, what we cut is cyclically bit-shifted.

- The final signature is compressed using range Asymmetric Numeral System, as explained in [ETWY22]. For simplicity, we assume that this gives expected bitsizes equal to the entropy of the compressed vector.

Target Security	120	180	260
n	256	256	256
q	64513	50177	202753
S	82.74	90.65	79.75
Keygen Acceptance Rate	0.5	0.5	0.5
s	14.22	14.22	14.22
σ	664.18	727.68	640.14
γ	31972.19	39405.92	38437.36
$(m, k-m)$	(3,4)	(4,5)	(4,7)
η	1	1	1
α	512	512	512
d	1	1	0
BKZ block-size b to break SIS	553 (461)	752 (633)	891 (751)
Best Known Classical bit-cost	161 (134)	219 (185)	260 (219)
Best Known Quantum bit-cost	142 (118)	193 (162)	228 (193)
BKZ block-size b to break LWE	415	610	895
Best Known Classical bit-cost	121	178	261
Best Known Quantum bit-cost	106	156	230
Signature size with rANS	1677	2143	2804
Expected public key size	1472	1952	2336
Sum	3149	4095	5140
Signature size [DFPS22]	1903	2473	3461
Public Key size	800	1056	1760
Sum	2703	3529	5221
Signature size [CCD$^+$23]	1474	2349	2908
Public Key size	992	1472	2080
Sum	2455	3809	4988

Table 1. Parameter sets for the Module G + G signature scheme. Numbers in parentheses for SIS security are for strong unforgeability.

For comparison, we include in Table 1 a reminder on estimated sizes of optimized Lyubashevsky signatures from [DFPS22], in the hyperball setting, as well as the experimental sizes of Haetae [CCD$^+$23], which implements the bimodal hyperball. As far as we are aware of, these are the lowest signatures and key sizes provided in the literature for Lyubashevsky's signatures (when using the core-SVP hardness methodology to estimate security). We note that the resulting signature sizes are 12% to 20% smaller than those from [DFPS22]. The asymptotic gain of our signature is observable when comparing the signature sizes with Haetae, as the tradeoff is first in their favor but ends up in our favor for the higher security levels. However, the sum of the public key and the signature sizes is somewhat similar across the three signatures. This is due to the fact that in the non-bimodal setting, a practical optimization due to [DKL+18] consists in truncating the low bits of the public key, at the cost of increasing the verification bound. While the technique is present in both Haetae and G+G, its efficiency is moderate in the G+G setting.

4.3 Optimized NTRU Key Generation Algorithm

We can alternatively use the NTRU-based key generation algorithm described in [DDLL13]. In our setting, it is possible to improve it, by relying on the aforementioned technique based on the divisibility of 2 by $(1 + x^{n/2})$. This leads to the key generation algorithm presented in Figure 5.

```
KeyGen($1^\lambda$) :
do $(f, g) \hookleftarrow U(\{\mathbf{x} \in \mathbb{R}^2[\|\mathbf{x}\|_\infty \leq \eta\})$
while $\sigma_1(\mathsf{rot}(\zeta f|2x^{n/2}g+\zeta)) \geq S$ or $f$ non-invertible mod $q$
$\mathbf{h} \leftarrow [\zeta g + 1]/f \bmod q$
$\mathbf{A} \leftarrow (\zeta^*(q-1)h \mid \zeta^*) \bmod 2q$
$\mathbf{s} \leftarrow (f \mid \zeta g + 1)^\top$
return $\mathsf{vk} = \mathbf{A}$ and $\mathsf{sk} = (\mathbf{A}, \mathbf{s})$
```

Fig. 5. NTRU KeyGen for $\mathsf{G} + \mathsf{G}$

Target Security	85	180
n	512	1024
q	32257	50177
S	99.60	94.36
Keygen Acceptance Rate	0.1	0.5
s	14.32	14.42
σ	804.94	767.76
B	27486.44	43316.29
$(m, k-m)$	(1,1)	(1,1)
η	2	1
α	256	1024
BKZ block-size b to break SIS	293 (220)	735 (619)
Best Known Classical bit-cost	85 (64)	214 (181)
Best Known Quantum bit-cost	75 (56)	188 (159)
BKZ block-size b to break LWE	305	610
Best Known Classical bit-cost	89	178
Best Known Quantum bit-cost	78	156
Signature size with rANS	1021	1769
Expected public key size	992	2080
Sum	2013	3849

Table 2. Parameter Sets for NTRU $\mathsf{G} + \mathsf{G}$.

The algorithm outputs keys $\mathbf{A}$ and $(\mathbf{A}, \mathbf{s})$ satisfying $\mathbf{A}\mathbf{s} = \zeta^* q \bmod 2q$ as it holds that $(q-1)hf = (q-1)(\zeta g + 1) \bmod 2q$ since $(q-1)$ is even. This implies that $\zeta \mathbf{A}\mathbf{s} = 0 \bmod 2q$, and the lattice that needs to be smoothed out is $\zeta \mathbf{s} \mathcal{R}$ where $\zeta \mathbf{s}^\top = (\zeta f | 2x^{n/2} g + \zeta)$. We then propose two sets of parameters in Table 2, for ring dimensions 512 and 1024. The former leads to only around 85 bits of security, but the latter allows to reach NIST security level III. While the sum $|\mathsf{vk}| + |\mathsf{sig}|$ is similar to those of the other schemes, we note that the signature size is further decreased, compared to module $\mathsf{G} + \mathsf{G}$. The resulting signature is 30% smaller than [DFPS22] and 47% smaller than Dilithium.

Acknowledgments. We thank Chris Peikert, Keita Xagawa and Shiduo Zhang for useful feedback that helped improve and fix errors in a prior version of this work. This work was supported by the France 2030 ANR Project ANR-22-PECY-003 SecureCompute, the France 2030 ANR Project ANR-22-PETQ-0008 PQ-TLS and the AMIRAL ANR grant (ANR-21-ASTR-0016).

References

ASY22. S. Agrawal, D. Stehlé, and A. Yadav. Round-optimal lattice-based threshold signatures, revisited. In *ICALP*, 2022.

BBD$^+$23. M. Barbosa, G. Barthe, C. Doczkal, J. Don, Serge Fehr, B. Grégoire, Y.-H. Huang, A. Hülsing, Y. Lee, and X. Wu. Fixing and mechanizing the security proof of Fiat-Shamir with aborts and Dilithium. In *CRYPTO*, 2023.

BCM21. R. Behnia, Y. Chen, and D. Masny. On removing rejection conditions in practical lattice-based signatures. In *PQCRYPTO*, 2021.

BF11. D. Boneh and D. M. Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In *PKC*, 2011.

BG14. S. Bai and S. D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, 2014.

BGV12. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, 2012.

BLP$^+$13. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, 2013.

BMKMS22. J. M. Bermudo Mera, A. Karmakar, T. Marc, and A. Soleimanian. Efficient lattice-based inner-product functional encryption. In *PKC*, 2022.

BP02. M. Bellare and A. Palacio. GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In *CRYPTO*, 2002.

CCD$^+$23. J. H. Cheon, H. Choe, J. Devevey, T. Güneysu, D. Hong, M. Krausz, G. Land, M. Möller, D. Stehlé, and M. Yi. HAETAE: Shorter lattice-based Fiat-Shamir signatures. Cryptology ePrint Archive, 2023. `https://ia.cr/2023/624`.

CLMQ21. Y. Chen, A. Lombardi, F. Ma, and W. Quach. Does Fiat-Shamir require a cryptographic hash function? In *CRYPTO*, 2021.

DDLL13. L. Ducas, A. Durmus, T. Lepoint, and V. Lyubashevsky. Lattice signatures and bimodal Gaussians. In *CRYPTO*, 2013.

DFPS22. J. Devevey, O. Fawzi, A. Passelègue, and D. Stehlé. On rejection sampling in Lyubashevsky's signature scheme. In *ASIACRYPT*, 2022.

DFPS23. J. Devevey, P. Fallahpour, A. Passelègue, and D. Stehlé. A detailed analysis of Fiat-Shamir with aborts. In *CRYPTO*, 2023.

DKL$^+$18. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-DILITHIUM: A lattice-based digital signature scheme. *IACR TCHES*, 2018.

DPSZ12. I. Damgård, V. Pastro, N. P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, 2012.

Duc14. L. Ducas. Accelerating Bliss: the geometry of ternary polynomials. Cryptology ePrint Archive, 2014. `https://ia.cr/2014/874`.

ETWY22. T. Espitau, M. Tibouchi, A. Wallet, and Y. Yu. Shorter hash-and-sign lattice-based signatures. In *CRYPTO*, 2022.

FS86. A. Fiat and A. Shamir. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *CRYPTO*, 1986.

GHHM21. A. B. Grilo, K. Hövelmanns, A. Hülsing, and C. Majenz. Tight adaptive reprogramming in the QROM. In *ASIACRYPT*, 2021.

GMPW20. N. Genise, D. Micciancio, C. Peikert, and M. Walter. Improved discrete gaussian and subgaussian analysis for lattice cryptography. In *PKC*, 2020.

GPV08. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.

Kle00. P. N. Klein. Finding the closest lattice vector when it's unusually close. In *SODA*, 2000.

KLS18. E. Kiltz, V. Lyubashevsky, and C. Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In *EUROCRYPT*, 2018.

LS15. A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.*, 2015.

Lyu09. V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, 2009.

Lyu12. V. Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, 2012.

MR07. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 2007.

Pei10. C. Peikert. An efficient and parallel gaussian sampler for lattices. In *CRYPTO*, 2010.

Reg09. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 2009.

Sch91. C.-P. Schnorr. Efficient Signature Generation by Smart Cards. *J. Cryptol.*, 1991.

vEH14. T. van Erven and P. Harremos. Rényi Divergence and Kullback-Leibler Divergence. *IEEE T. Inform. Theory*, 2014.

YJW23. Y. Yu, H. Jia, and X. Wang. Compact lattice gadget and its applications to hash-and-sign signatures. In *CRYPTO*, 2023.

ZXZ18. Z. Zheng, G. Xu, and C. Zhao. Discrete Gaussian measures and new bounds of the smoothing parameter for lattices. Cryptology ePrint Archive, 2018. `https://ia.cr/2018/786`.

A The Fiat-Shamir Transform

In this section, we recall the Fiat-Shamir transform, which allows to transform an identification scheme into a digital signature. It removes interaction by sampling the challenge as a hash function evaluation $H(w, \mu)$ with w being the prover's commitment and μ the signed message. The hash function is then modeled as a random oracle in the analysis. The signature is the pair (w, z), which is verified by checking validity of the transcript $(w, H(w, \mu), z)$.

As the challenge c being typically much shorter than w, it is desirable to replace w by c in the signature. This is possible if the underlying identification scheme is commitment-recoverable (see Definition 2). Verification simply starts by recovering $w \leftarrow \mathsf{Rec}(\mathsf{vk}, c, z)$. Our protocol satisfies this property, thus we describe the signature obtained applying this version of the Fiat-Shamir transform. See Figure 6.

```
$\underline{\mathsf{KeyGen}(1^\lambda):}$
$(\mathsf{vk}, \mathsf{sk}) \leftarrow \mathsf{IGen}(1^\lambda)$
return $\mathsf{vk}$ and $\mathsf{sk}$
```

```
$\underline{\mathsf{Sign}(\mathsf{sk}, \mu):}$
$(w, \mathsf{st}) \leftarrow \mathsf{P}_1(\mathsf{sk})$
$c \leftarrow H(w, \mu)$
$z \leftarrow \mathsf{P}_2(\mathsf{sk}, \mathsf{st}, w, c)$
return $(c, z)$
```

```
$\underline{\mathsf{Verify}(\mathsf{vk}, (c, z), \mu):}$
$w \leftarrow \mathsf{Rec}(\mathsf{vk}, c, z)$
if $c \neq H(w, \mu)$ then
    return 0
end if
return $\mathsf{V}(\mathsf{vk}, (w, c, z))$
```

Fig. 6. Fiat-Shamir Signature $\mathsf{FS}[\mathsf{ID}, H]$.

For the sake of completeness, we state the following lemma arguing correctness of the signature scheme $\mathsf{FS}[\mathsf{ID}, H]$, which immediately follows from the completeness and commitment-recoverability of the underlying identification scheme.

Lemma 8. *Let* $\mathsf{ID} = (\mathsf{IGen}, \mathsf{P}, \mathsf{V})$ *denote an identification scheme. Further assume that* ID *is ε-complete and commitment-recoverable. Then the signature scheme* $\mathsf{FS}[\mathsf{ID}, H]$ *described in Figure 6 is ε-correct in the ROM.*

Security of $\mathsf{FS}[\mathsf{ID}, H]$ can be proven by successive claims. First, one can reduce $\mathsf{EU\text{-}CMA}$ security of $\mathsf{FS}[\mathsf{ID}, H]$ to its $\mathsf{EU\text{-}NMA}$ security assuming ID has large commitment min-entropy and is honest-verifier zero-knowledge (see Definition 3). This can be shown by relying on the following theorem.

Theorem 7 (Adapted from [GHHM21, Theorem 3]). *Let* ID *be an identification scheme which has α-min-entropy and satisfies ε-statistical* HVZK. *Let H a hash function modeled as a random oracle. Then, for any (possibly quantum) adversary $\mathcal{A}$ against the* $\mathsf{EU\text{-}CMA}$ *security of* $\mathsf{FS}[\mathsf{ID}, H]$ *making at most Q_S (classical) sign queries and at most Q_H (possibly quantum) hash queries, there exists an adversary $\mathcal{B}$ against the* $\mathsf{EU\text{-}NMA}$ *security of* $\mathsf{FS}[\mathsf{ID}, H]$ *such that:*

$$\mathsf{Adv}^{\mathsf{EU\text{-}CMA}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{EU\text{-}NMA}}(\mathcal{B}) + Q_S \varepsilon + 3\frac{Q_S}{2} \cdot \sqrt{(Q_H + Q_S + 1) \cdot 2^{-\alpha}} \ .$$

Furthermore, if ID *is* $(1+\varepsilon)$*-divergence* HVZK*, the following bound applies:*

$$\mathsf{Adv}^{\mathsf{EU\text{-}CMA}}(\mathcal{A}) \leq (1+\varepsilon)^{Q_S} \mathsf{Adv}^{\mathsf{EU\text{-}NMA}}(\mathcal{B}) + 3Q_S/2 \cdot \sqrt{(Q_H + Q_S + 1) \cdot 2^{-\alpha}} \ .$$

The result can be adapted to sEU-CMA *security by adding* $Q_S 2^{-\alpha} + \mathsf{Adv}^{\mathsf{CUR}}(\mathcal{B}')$ *to the bounds, for some adversary* $\mathcal{B}'$ *against the CUR property.*

The last statement, while not present in [GHHM21], can be found in [DFPS23] by considering that the abort probability is 0 and the number of iterations $B = 1$.

It remains to prove EU-NMA-security to conclude the security analysis, which can be argued via the following statement for lossy identification schemes (see Definition 4).

Theorem 8 ([KLS18, Theorem 3.4]). *Let* ID *be a lossy identification scheme satisfying* $\varepsilon_{\mathsf{ls}}$*-lossy soundness for some* $\varepsilon_{\mathsf{ls}} > 0$*. Let* H *a hash function modeled as a random oracle. For any (possibly quantum) adversary* $\mathcal{A}$ *against the* EU-NMA *security of* $\mathsf{FS}[\mathsf{ID}, H]$ *making at most* Q_H *(possibly quantum) hash queries, there exists a quantum adversary* $\mathcal{B}$ *against the key-indistinguishability of* ID *such that*

$$\mathsf{Adv}^{\mathsf{EU\text{-}NMA}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{key\text{-}ind}}(\mathcal{B}) + 8(Q_H + 1)^2 \cdot \varepsilon_{\mathsf{ls}} \ .$$

Finally, we describe a reduction in the (classical) ROM which relies on weaker properties compared to the above QROM reduction. Various folklore reductions are known in this setting, and we consider a variant based on special soundness (see Definition 5), which is first reduced to the soundness as recalled below.

Definition 12 (Soundness). *Let* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *be an identification scheme. It is* sound *if for any PPT adversary* $\mathcal{A}$*, the quantity*

$$\Pr\Big[\mathsf{V}(\mathsf{vk}, (w, c, z)) = 1 \mid (w, c, z) \leftarrow \mathcal{A}(\mathsf{vk})\Big]$$

is $\mathsf{negl}(\lambda)$*, where the probability is over the choice of* vk *and the coins of* $\mathcal{A}$*.*

We recall the Reset Lemma, which is a standard reduction between soundness and special soundness.

Lemma 9 (Reset Lemma [BP02]). *Let* $\mathsf{ID} = (\mathsf{Igen}, \mathsf{P}, \mathsf{V})$ *be an identification scheme. Given any adversary* $\mathcal{A}$ *against the soundness of* ID*, there exists an adversary* $\mathcal{B}$ *against the special soundness of* ID *such that*

$$\mathsf{Adv}^{\mathsf{special\text{-}sound}}(\mathcal{B}) \geq \left(\mathsf{Adv}^{\mathsf{sound}}(\mathcal{A}) - \frac{1}{|\mathcal{C}|}\right)^2 .$$

This can also be adapted to work with lossy soundness and lossy special soundness. While this result is folklore, we finally show that lossy special soundness implies EU-NMA security in the ROM.

Lemma 10. *Let* ID *be an identification scheme and* H *a hash function modeled as a random oracle. For any adversary* $\mathcal{A}$ *against the* $\mathsf{EU\text{-}NMA}$ *security of* $\mathsf{FS}[\mathsf{ID}, H]$ *making* Q_H *classical hash queries, there exists an adversary* $\mathcal{B}$ *against the lossy special soundness of* ID *and an adversary* $\mathcal{B}'$ *against key-indistinguishability of the scheme such that:*

$$\mathsf{Adv}^{\mathsf{EU\text{-}NMA}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{key\text{-}ind}}(\mathcal{B}') + Q_H \cdot \left(\sqrt{\mathsf{Adv}^{\mathsf{special\text{-}sound}}(\mathcal{B})} + \frac{2}{|\mathcal{C}|} \right) \ .$$

Proof. We first reduce the soundness of ID to the $\mathsf{EU\text{-}NMA}$ security of $\mathsf{FS}[\mathsf{ID}, H]$. We start by turning the scheme into its lossy mode, leading to the first term in the bound. Then, if $\mathcal{A}$ outputs a forgery $(\mu^*, (c^*, z^*))$ such that $H(\mathsf{Rec}(\mathsf{vk}, c^*, z^*), \mu^*)$ was never queried, it has probability at most $1/|\mathcal{C}|$ of outputting a valid forgery.

The reduction $\mathcal{B}'$ guesses the hash query $H(w^*, \mu^*)$ made by $\mathcal{A}$ which is used in $\mathcal{A}$'s forgery. When this query is made, $\mathcal{B}'$ answers it by running sending w^* as commitments to its challenger. The latter replies with a challenge c^* and $\mathcal{B}'$ programs $H(w^*, \mu^*)$ as c^*. With probability $1/Q_H$, $\mathcal{B}'$'s guess is correct and the adversary $\mathcal{A}$ halts with a forgery $(\mu^*, (c^*, z^*))$ with $\mathsf{Rec}(\mathsf{vk}, c^*, z^*) = w^*$. We then have

$$\mathsf{Adv}^{\mathsf{sound}}(\mathcal{B}') \geq \frac{1}{Q_H} \cdot \mathsf{Adv}^{\mathsf{EU\text{-}NMA}}(\mathcal{A}) - 1/|\mathcal{C}| \ .$$

Finally, Lemma 9 gives an adversary $\mathcal{B}$ against the special soundness such that

$$\mathsf{Adv}^{\mathsf{special\text{-}sound}}(\mathcal{B}) \geq \left(\mathsf{Adv}^{\mathsf{sound}}(\mathcal{B}') - \frac{1}{|\mathcal{C}|} \right)^2 ,$$

which completes the proof. □

B Related Work

In Figure 7, we give a simplified version of the Eagle signature scheme described in [YJW23] (with our notations from Section 4 and an extra parameter $\gamma' > 0$). Minor differences with the scheme from Figure 4 include the facts that Eagle works in the ring setting as opposed to the module setting, that a parameterizable integer p is considered while we work with $p = 2$, and that the RLWE sample from Eagle is computed modulo $Q = pq$, while we use MLWE samples computed modulo q. The exact signing algorithm from [YJW23] is omitting some elements of the final vector $\mathbf{z}$ to optimize compactness, but we do not consider this optimization to better illustrate the relationship with $\mathsf{G} + \mathsf{G}$. Moreover, as usual in hash-and-sign schemes, the message is padded using some salt, chosen as a uniform 320-bit long bitstring.

We now explain how to decompose Eagle as an instance of $\mathsf{G} + \mathsf{G}$ with a specific hash function, as well as the differences that arise during verification due to this hash function, following the steps of [CLMQ21]. The instance of the hash function H that turns the signing algorithm of $\mathsf{G} + \mathsf{G}$ into a simplified version of Eagle is described in Steps 3, 4 and 5 of the signing algorithm from Figure 7.

```
$\mathsf{KeyGen}(1^\lambda)$:
    $\mathbf{a}_0 \hookleftarrow U(\mathcal{R}_q)$
    $(\mathbf{s}_1, \mathbf{s}_2) \hookleftarrow \chi_\eta \times \chi_\eta$
    $\mathbf{s} \leftarrow (1|\mathbf{s}_1|\mathbf{s}_2)^\top \in \mathcal{R}_{2q}^3$
    $\mathbf{b} \leftarrow \mathbf{a}_0\mathbf{s}_1 + \mathbf{s}_2 \bmod Q$
    $\mathbf{A} \leftarrow (q - \mathbf{b}|\mathbf{a}_0|1)$
    return $(\mathsf{vk}, \mathsf{sk}) = (\mathbf{A}, \mathbf{s})$
```

```
$\mathsf{Sign}(\mathbf{A}, \mathbf{s}, \mu)$:
    $\mathsf{salt} \hookleftarrow U(\{0,1\}^{320})$
    $\mathbf{S} = \mathsf{rot}(\mathbf{s})$
    $\mathbf{y} \hookleftarrow D_{\mathcal{R}^3, \sigma^2\mathbf{I}_{3n} - 4s^2\mathbf{S}\mathbf{S}^\top}$
    $u \leftarrow H'(\mu, \mathsf{salt})$
    $u' \leftarrow u - \mathbf{A}\mathbf{y} \bmod Q$
    $c \leftarrow \lfloor u' \rceil_q$
    $k \leftarrow D_{\mathcal{R}, s, -c/2}$
    $\mathbf{z} \leftarrow \mathbf{y} + \mathbf{S}c + p\mathbf{S}k$
    return $(\mathsf{salt}, \mathbf{z})$
```

```
$\mathsf{Verify}(\mathbf{A}, \mu, \sigma)$
    $\sigma = (\mathsf{salt}, \mathbf{z})$
    $u \leftarrow H(\mu, \mathsf{salt})$
    $z' \leftarrow u - \mathbf{A}\mathbf{z}$
    Accept if $\|\mathbf{z}\| \leq \gamma$
        and $\|z'\| \leq \gamma'$
```

Fig. 7. Simplified Eagle Signature Scheme.

It proceeds as follows. On input $w \in \mathcal{R}$, μ and salt, the function H computes a target $u = H'(\mu, \mathsf{salt})$ using another hash function H' and sets $u' = u - w$. The challenge is then $\lfloor u' \rceil_q$, i.e., a rounding of u' to the $q\mathcal{R}$ lattice.

The verification algorithm differs substantially due to the fact that Verify is aware of the inner workings of the hash function. It knows in particular that $\mathbf{A}\mathbf{z} = \mathbf{A}\mathbf{y} + qc \bmod Q \approx u$. However, the challenge c is omitted from the signature and instead of checking that $H(\mathbf{A}\mathbf{z} - qc, \mu, \mathsf{salt}) = c$, it checks that $u - \mathbf{A}\mathbf{z}$ is sufficiently short, i.e., has norm smaller than γ'. While this check is less accurate than recomputing the hash value, it allows one to omit c in the signature, hence reducing its size. Finally, the verification algorithm also checks that $\mathbf{z}$ has norm $\leq \gamma$, as in Figure 4.