Optimized G+G Signature

Renjie Jin, Shuoqu Jian, and Longjiang Qu

College of Science, National University of Defense Technology, Changsha, 410073, China
{jinrj97,jian_souqi}@163.com, ljqu_happy@hotmail.com

Abstract. At ASIACRYPT 2023, Devevey, Passelègue, and Stehlé introduced the G+G signature, a Fiat–Shamir lattice signature based on Gaussian convolution that avoids rejection sampling during signing. Although G+G achieves competitive signature sizes, its spherical Gaussian response distribution leaves room for further optimization.
In this work, we propose an optimized G+G signature by incorporating asymmetric secret-key distributions into key generation and introducing a three-block non-spherical Gaussian response distribution. We develop a refined spectral-norm estimation method for the secret-key matrix, enabling more flexible selection of Gaussian parameters across different blocks. Combined with signature truncation and compression techniques, this approach reduces the signature size while preserving the underlying G+G signing paradigm. For two representative parameter sets, our entropy-based estimates show signature-size reductions of 9.61% and 5.53%, respectively, compared with G+G under matched correctness requirements and a common concrete attack-cost model.

Keywords: Lattice signature · Asymmetric LWE · Discrete Gaussian.

Note on this ePrint revision. Following the November 18, 2026 revision of G+G [DPS26a], we correct the Gaussian convention, effective challenge spaces, assumption formalizations, and parameter estimates. This revision supersedes the former concrete parameter tables. Supplementary details on key-generation conditioning and compression are given in the appendices. For this revision, the authors relied on some AI tools.

1 Introduction

Lattice-based cryptography stands as one of the most pivotal cryptographic primitives in the post-quantum era. Among its core components, lattice-based signatures have emerged as a central research topic. Currently, existing lattice-based signature schemes can be broadly categorized into two classes based on their design paradigms: Hash-and-Sign [GPV08] and the Fiat-Shamir transform [Lyu09]. The former heavily relies on discrete Gaussian sampling over lattice cosets, with major advancements typically arising from improvements in lattice trapdoors and optimizations of Gaussian samplers [CGM19, ETWY22, YJW23]. The latter, primarily employing the rejection sampling technique [DDLL13, Lyu12] to support the compactness of signature size, achieves optimization mainly through enhancements in the underlying source distribution [DFPS23] and the application of rejection strategies [Gär25]. In 2023, Devevey et al. [DPS23] proposed a Fiat-Shamir signature without signing-time rejection sampling, which is regarded as a promising new direction in this branch, particularly as the so-called "noise flooding" technique fails to produce signatures with compact parameters. However, the construction of G+G signatures remains high-level. While its signature size is competitive compared to known schemes [DDLL13, DKL$^+$18, CGM19, YJW23], there is significant room for improvement since the underlying spherical Gaussian distribution is not the optimal distribution for constructing signatures with compact size.

1.1 Our Contributions

In this work, we integrate the *Asymmetric Learning with Errors* (ALWE) problem into the key setting of the G+G signature, leveraging the asymmetry of key distribution combined with signature truncation techniques to output signatures with smaller size. Specifically, we propose a new method for estimating the largest singular values of the secret key and introduce a three-block non-spherical Gaussian allocation to achieve fine-grained control over Gaussian parameters. The revised experimental results compare the size of our optimized signature with the corrected G+G baseline [DPS26a] under a common correctness target and the stated concrete attack model. Moreover, we analyze the security of the ALWE problem using an explicit security evaluation model. The updated parameter estimates and their limitations are reported in Section 3.2.

1.2 Technique Overview

Our goal is to optimize the size of the G+G signature. Specifically, we seek to do this without modifying the idea of the signing process, so as to retain the original proof strategy, with the corrections and conditioning stated below. To this end, we transform the signature distribution from a spherical Gaussian to a non-spherical one. In the G+G signature [DPS23], the secret key $\mathbf{S} \leftarrow (\mathbf{I}_\ell \mid \mathbf{S}_1^\top \mid \mathbf{S}_2^\top)^\top \in \mathbb{Z}^{k\times\ell}$, where $(\mathbf{S}_1, \mathbf{S}_2) \hookleftarrow \chi^{(k-m-\ell)\times\ell} \times \chi^{m\times\ell}$ and χ is a probability distribution. Let $\sigma, s > 0$ and define $\Sigma : \mathbb{Z}^{k\times\ell} \to \mathbb{R}^{k\times k}$ as

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

The ideal response follows the distribution $\mathcal{D}_{\mathbb{Z}^k,\sigma}$. The width σ is chosen so that the perturbation with shape $\Sigma(\mathbf{S})$ can be efficiently sampled by the signer and the convolution satisfies the smoothing condition used for signature simulation. Appendix B gives the resulting pointwise bound for the actual response distribution. Therefore, the lower bound for σ is critically determined by the largest singular value of $\mathbf{S}$. Consequently, optimizing the structure of $\mathbf{S}$ to reduce this bound serves as an effective approach for minimizing the signature size.

Step 1: Asymmetric secret key setting. We sample $(\mathbf{S}_1, \mathbf{S}_2) \hookleftarrow \chi_1^{(k-m-\ell)\times\ell} \times \chi_2^{m\times\ell}$ where χ_1, χ_2 belong to the same distribution family but have different standard deviations, with χ_1 being the narrower one. At this point, the key pair is set as follows

$$\mathsf{pk} := \mathbf{A} \leftarrow (q\mathbf{J} - 2\mathbf{B} \mid 2\mathbf{A}_1 \mid 2\mathbf{I}_m) \in \mathbb{Z}_{2q}^{m\times k},$$
$$\mathsf{sk} := \mathbf{S} \leftarrow (\mathbf{I}_\ell \mid \mathbf{S}_1^\top \mid \mathbf{S}_2^\top)^\top \in \mathbb{Z}^{k\times\ell},$$

where $\mathbf{B} \leftarrow \mathbf{A}_1\mathbf{S}_1 + \mathbf{S}_2 \bmod q$ and $\mathbf{J} = (\mathbf{I}_\ell \mid \mathbf{0}^{\ell\times(m-\ell)})^\top \in \mathbb{Z}^{m\times\ell}$. The security of the keys is based on computational indistinguishability between an ALWE distribution [ZYF+20] and the uniform distribution. For the concrete module instantiation, the actual key-generation acceptance event is included in the assumption specified in Appendix B, whose comparison distribution retains the accepted public marginal.

Step 2: New signature distribution: three-block Gaussian allocation. We compress the signature distribution from a spherical Gaussian into a non-spherical one with a well-defined shape. The primary goal is to eliminate the redundancy in the Gaussian width inherent in the original spherical Gaussian distribution of the G+G signature, enabling further compression of the signature size along certain dimensions. Specifically, our improvement lies in the construction of the Gaussian shape matrix $\Sigma(\mathbf{S})$. We set

$$\Sigma(\mathbf{S}) = \begin{bmatrix} \sigma_1^2 \cdot \mathbf{I}_\ell & & \\ & \sigma_2^2 \cdot \mathbf{I}_{k-m-\ell} & \\ & & \sigma_3^2 \cdot \mathbf{I}_m \end{bmatrix} - 4s^2 \left[\mathbf{I}_\ell, \mathbf{S}_1^\top, \mathbf{S}_2^\top\right]^\top \left[\mathbf{I}_\ell, \mathbf{S}_1^\top, \mathbf{S}_2^\top\right].$$

Our goal is to find lower bounds for σ_i $(i = 1, 2, 3)$ ensuring the positive definiteness of $\Sigma(\mathbf{S})$ and the required smoothing condition[1]. Notably, the structure of $\mathbf{S}\mathbf{S}^\top$ allows for a significant reduction in σ_1, which is tied to the identity matrix component. This analysis requires examining the impact of $\mathbf{S}_1$ and $\mathbf{S}_2$—sampled from the same distribution with varying parameters—on the singular values.

Remark 1. From the perspective of reducing the value of σ_1 in $\Sigma(\mathbf{S})$, the ellipsoidal Gaussian distribution proposed in [JHT22] has already achieved this. The purpose of employing an asymmetrically distributed secret key in this work is to further optimize the ellipsoidal Gaussian distribution, which requires integrating the signature truncation technique introduced by Bai et al. [BG14]. By narrowing the distribution width of $\mathbf{S}_1$ while increasing the distribution parameter of $\mathbf{S}_2$, we obtain a different ALWE parameter choice whose security must be evaluated for the actual distributions and sample count. Meanwhile, by leveraging the signature truncation technique, we can effectively reduce the size of the G+G signature. This will be discussed in more detail later.

Step 3: Bound on the largest singular value of $\mathbf{S}$. The structure of our secret key $\mathbf{S}$ makes directly estimating the largest singular value of $\mathbf{S}\mathbf{S}^\top$ inconvenient. Instead, we analyze the simulatability condition through a rescaled matrix. Define

$$\mathbf{D} := \mathrm{diag}(\sigma_1\mathbf{I}_\ell, \sigma_2\mathbf{I}_{k-m-\ell}, \sigma_3\mathbf{I}_m), \qquad \widetilde{\mathbf{S}} := \mathbf{D}^{-1}\mathbf{S}.$$

[1] A sufficient condition is $\mathbf{D}^2 \succeq 8s^2\mathbf{S}\mathbf{S}^\top$ together with $s/\sqrt{2} \geq \eta_\epsilon(\mathbb{Z}^\ell)$, as proved in Appendix B.

Then the requirement

$$\begin{bmatrix} \sigma_1^2\mathbf{I}_\ell & & \\ & \sigma_2^2\mathbf{I}_{k-m-\ell} & \\ & & \sigma_3^2\mathbf{I}_m \end{bmatrix} - 8s^2\mathbf{S}\mathbf{S}^\top \geq 0$$

is equivalent to

$$\sin_1(\widetilde{\mathbf{S}})^2 \leq (8s^2)^{-1}.$$

To exploit the block structure of $\mathbf{S}$, we set $t = \sigma_2/\sigma_3$ and rewrite

$$\widetilde{\mathbf{S}} = \begin{pmatrix} \sigma_1^{-1}\mathbf{I}_\ell \\ \sigma_2^{-1}\mathbf{S}_1 \\ \sigma_2^{-1}t\mathbf{S}_2 \end{pmatrix} = \begin{pmatrix} \sigma_1^{-1}\mathbf{I}_\ell \\ \sigma_2^{-1}\mathbf{S}_{\text{new}} \end{pmatrix}, \qquad \mathbf{S}_{\text{new}} := \begin{pmatrix} \mathbf{S}_1 \\ t\mathbf{S}_2 \end{pmatrix}.$$

Consequently,

$$\widetilde{\mathbf{S}}^\top\widetilde{\mathbf{S}} = \sigma_1^{-2}\mathbf{I}_\ell + \sigma_2^{-2}\mathbf{S}_{\text{new}}^\top\mathbf{S}_{\text{new}},$$

which implies

$$\sin_1(\widetilde{\mathbf{S}})^2 = \sigma_1^{-2} + \sigma_2^{-2}\sin_1(\mathbf{S}_{\text{new}})^2.$$

Therefore, instead of estimating $\mathbf{S}\mathbf{S}^\top$ directly, we first determine the scaling parameter t, and then estimate $\sin_1(\mathbf{S}_{\text{new}})$ numerically. This provides a stable way to evaluate the admissible region of the non-spherical Gaussian parameters $(\sigma_1, \sigma_2, \sigma_3)$.

Step 4: Truncate the signature. Based on the public key setting of the G+G signature, the response component corresponding to $\mathbf{S}_2$ is multiplied by the public identity block. We can truncate this response component, storing its information in the form of a hint, and subsequently recover a nearby representative in the verification algorithm. According to the structure of our secret key, $\mathbf{S}_1$ has a narrower distribution, while $\mathbf{S}_2$ has a wider distribution. Truncating the signature segment corresponding to $\mathbf{S}_2$ can further reduce the signature size relative to a spherical response distribution. The complete hint and the increase in the verification bound must be included in this comparison. Appendix C specifies the recovery relation and its effect on correctness and the security reduction.

1.3 Related Work

Lattice signatures. The first secure lattice-based signature scheme following the Hash-and-Sign paradigm was introduced by Gentry, Peikert, and Vaikuntanathan [GPV08]. Their seminal work systematically resolved the issue of secret key leakage in prior trapdoor sampling techniques. Subsequently, Micciancio and Peikert [MP12] proposed the notion of gadget-based trapdoors, which substantially enhanced the efficiency of lattice-based trapdoor mechanisms while offering practical implementability and inherent parallelism. Chen et al. [CGM19] generalized exact gadget trapdoors to an approximate variant, enabling its use in practical lattice signature schemes. Espitau et al. [ETWY22] used the range Asymmetric Numeral Systems (rANS) encoding mechanism as a means of signature compression, tackling the size reduction problem from a distinct perspective. Jia et al. [JHT22] and Yu et al. [YJW23] further optimized gadget trapdoors and preimage samplers to achieve near-optimal performance, with the latter presenting a compact construction under this paradigm.

Lattice-based signatures following the Fiat-Shamir paradigm were first introduced by Lyubashevsky [Lyu09], who proposed the signature scheme on lattices using rejection sampling techniques. Later, Lyubashevsky [Lyu12] further refined this approach by employing Gaussian distributions, establishing a general framework for signature simulation based on rejection sampling mechanisms. Ducas et al. [DDLL13] proposed an efficient bimodal Gaussian distribution, which reduced rejection rates while optimizing signature size. The work of Devevey et al. [DFPS23, DFPS22] not only revealed potential issues in the security analysis of rejection sampling in Lyubashevsky-type signatures but also demonstrated that the bimodal Gaussian distribution is the optimal source distribution within the class analyzed there. Subsequently, they proposed a Fiat-Shamir signature without signing-time rejection sampling [DPS23] based on Gaussian convolution techniques, while a comprehensive security analysis opens up new possibilities for efficiency improvements in this signature paradigm. Recently, Gärtner [Gär25] achieved significant signature size reduction through an iterative rejection sampling strategy, resulting in a compact Fiat-Shamir signature construction.

Roadmap. This paper is organized as follows. Section 2 introduces the necessary definitions and foundational results. Section 3 presents an analysis of the G+G signature scheme and introduces its optimized construction. Section 4 concludes the paper. Related details for evaluating the hardness of asymmetric problems are provided in the Summary of the Hardness Estimation for the Asymmetric Problems. The supplementary convolution and security arguments, and the precise compression procedure, are given in Appendices B and C, respectively.

2 Notations and Preliminaries

We fix notation throughout the paper. Let $\mathbb{R}$ denote the field of real numbers and $\mathbb{Z}$ the ring of integers. For a positive integer q, $\mathbb{Z}/q\mathbb{Z}$ denotes the ring of integers modulo q. Let $[n] := \{1, \cdots, n\}$. A function $f(n)$ is negligible, denoted by $\mathsf{negl}(n)$, if it decreases faster than the reciprocal of any polynomial function $g(n)$. We use $\log n$ to denote the logarithm to the base 2. Matrices and vectors are represented using bold fonts, for example, matrix $\mathbf{A}$ and vector $\mathbf{a}$. The notations $(\mathbf{A} \mid \mathbf{B})$ and $(\mathbf{A}, \mathbf{B})$ both denote matrix concatenation. We use $\sin_1(\mathbf{A}) = \|\mathbf{A}\|_{\mathrm{op}}$ to denote the largest singular value of $\mathbf{A}$, where $\|\mathbf{A}\|_{\mathrm{op}} = \sup_{\|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$. For vectors $\mathbf{a}, \mathbf{b}$, denote by $\mathbf{a}^\top \mathbf{b}$ (or $\langle \mathbf{a}, \mathbf{b} \rangle$) the Euclidean inner product of $\mathbf{a}$ and $\mathbf{b}$. For a matrix $\mathbf{B} = (\mathbf{b}_1, \cdots, \mathbf{b}_t)$ with columns in $\mathbb{R}^m$, denote by $\|\mathbf{b}_i\|$ the ℓ_2-norm of $\mathbf{b}_i$ $(1 \le i \le t)$ and

$$\|\mathbf{B}\|_{\mathrm{col}} := \max_{i \in [t]} \|\mathbf{b}_i\|_2$$

the maximum column norm. This is distinct from the operator norm, and $\|\mathbf{B}\|_{\mathrm{col}} \le \|\mathbf{B}\|_{\mathrm{op}} \le \sqrt{t}\|\mathbf{B}\|_{\mathrm{col}}$. Unless indicated otherwise, vector norms are Euclidean.

Let $\mathcal{C}$ be a PPT (probabilistic polynomial time) algorithm. $\mathbf{t} \leftarrow \mathcal{C}(\mathbf{s})$ means that the random variable $\mathbf{t}$ is the output of $\mathcal{C}$ with input $\mathbf{s}$ and some random values. Write vector $\mathbf{s} \hookleftarrow \mathcal{D}$ if $\mathbf{s}$ is sampled from $\mathcal{D}$ at random with distribution $\mathcal{D}$. $\mathbf{s} \sim \mathcal{D}$ denotes that $\mathbf{s}$ follows the distribution $\mathcal{D}$. We use $\mathcal{U}(\mathbb{Z}_q)$ to represent the uniform distribution over $\mathbb{Z}_q$. We denote $\Sigma > 0$ ($\Sigma \ge 0$) when a symmetric matrix $\Sigma \in \mathbb{R}^{m \times m}$ is positive definite (positive semidefinite), i.e., $\mathbf{x}^\top \Sigma \mathbf{x} > 0$ ($\mathbf{x}^\top \Sigma \mathbf{x} \ge 0$) for all nonzero $\mathbf{x} \in \mathbb{R}^m$. We write $\Sigma_1 \ge \Sigma_2$ if $\Sigma_1 - \Sigma_2 \ge 0$.

2.1 Lattice and Discrete Gaussian

Given a set of linearly independent vectors $\mathbf{S} := \{\mathbf{s}_1, \mathbf{s}_2, \cdots, \mathbf{s}_n\} \subset \mathbb{R}^m$, define the lattice $\Lambda(\mathbf{S})$ generated by $\mathbf{S}$ as

$$\Lambda(\mathbf{S}) := \left\{ \mathbf{x} \in \mathbb{R}^m \mid \mathbf{x} = \sum_{i=1}^{n} d_i \mathbf{s}_i \text{ where } d_i \in \mathbb{Z} \right\}.$$

When $n = m$, we call the lattice a full-rank lattice.

For a positive integer q and public matrix $\mathbf{S} \in \mathbb{Z}_q^{n \times m}$, define the following q-ary lattices and, for $\mathbf{u} \in \mathrm{im}(\mathbf{S}) \subseteq \mathbb{Z}_q^n$, a lattice coset:

$$\Lambda_q(\mathbf{S}) := \{\mathbf{x} \in \mathbb{Z}^m \mid \exists \mathbf{t} \in \mathbb{Z}_q^n \; s.t. \; \mathbf{S}^\top \mathbf{t} = \mathbf{x} \pmod q\},$$
$$\Lambda_q^\perp(\mathbf{S}) := \{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{S}\mathbf{x} = \mathbf{0} \pmod q\},$$
$$\Lambda_q^{\mathbf{u}}(\mathbf{S}) := \{\mathbf{x} \in \mathbb{Z}^m \mid \mathbf{S}\mathbf{x} = \mathbf{u} \pmod q\}.$$

The Gaussian function $\rho(\mathbf{x}) = \exp(-\pi \langle \mathbf{x}, \mathbf{x} \rangle)$ maps from $\mathbb{R}^n$ to $(0, 1]$ where $\mathbf{x} \in \mathbb{R}^n$. Given an invertible matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, the Gaussian function associated with $\mathbf{A}$ is defined as

$$\rho_{\mathbf{A}}(\mathbf{x}) = \exp(-\pi \mathbf{x}^\top \Sigma^{-1} \mathbf{x}),$$

where $\Sigma = \mathbf{A}\mathbf{A}^\top$ is the Gaussian shape matrix. For a full-rank lattice $\Lambda \subset \mathbb{R}^n$, the discrete Gaussian distribution over Λ is defined as:

$$\mathcal{D}_{\Lambda, \sqrt{\Sigma}, \mathbf{c}}(\mathbf{x}) := \frac{\rho_{\sqrt{\Sigma}}(\mathbf{x} - \mathbf{c})}{\rho_{\sqrt{\Sigma}}(\Lambda - \mathbf{c})}, \qquad \mathbf{x} \in \Lambda,$$

where $\mathbf{c} \in \mathrm{span}(\Lambda)$ and $\rho_{\sqrt{\Sigma}}(T) := \sum_{\mathbf{x} \in T} \rho_{\sqrt{\Sigma}}(\mathbf{x})$ for a discrete set T. The probability is zero outside Λ. In this work, we use $\sigma > 0$ to denote the Gaussian width, so that $\rho_\sigma(\mathbf{x}) = \exp(-\pi \|\mathbf{x}\|_2^2 / \sigma^2)$, and omit $\mathbf{c}$ if $\mathbf{c} = \mathbf{0}$. The

corresponding continuous Gaussian has standard deviation $\sigma/\sqrt{2\pi}$ and covariance $\Sigma/(2\pi)$ in the matrix case. A discrete Gaussian does not in general have exactly this covariance. Thus widths and standard deviations are not interchangeable. Unless otherwise specified, let $\mathcal{D}_{\mathbb{Z}^{n\times m},\sigma}$ denote the distribution obtained by independently sampling m integer vectors from $\mathcal{D}_{\mathbb{Z}^n,\sigma}$ and viewing them as the columns of a matrix in $\mathbb{Z}^{n\times m}$.

Definition 1 (Smoothing Parameter). *For a lattice Λ and any real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest real $\sigma > 0$ such that $\rho(\sigma \cdot \Lambda^*) \leq 1 + \epsilon$ where Λ^* is the dual lattice of Λ.*

2.2 Lattice Assumptions

The Learning With Errors (LWE) problem was introduced by Regev [Reg05]. We adopt the normal-form formulation of LWE, in which the secret and error are independently sampled from the same distribution, following the standard treatment in [LP11].

Definition 2 (LWE). *Let $n, m > 0$ and $q \geq 2$ be integers, and let χ_α be a specified distribution over $\mathbb{Z}$. Sample*

$$\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{m\times n}), \qquad \mathbf{s} \hookleftarrow \chi_\alpha^n, \qquad \mathbf{e} \hookleftarrow \chi_\alpha^m$$

independently, and set $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod q$. Here, the secret $\mathbf{s}$ is sampled once and shared by all m samples. We define the following two problems:

- *Search-LWE$_{n,m,q,\chi_\alpha}$: Given $(\mathbf{A}, \mathbf{b})$, recover $\mathbf{s} \bmod q$.*
- *Decision-LWE$_{n,m,q,\chi_\alpha}$: Distinguish $(\mathbf{A}, \mathbf{b})$ from $(\mathbf{A}, \mathbf{u})$, where $\mathbf{u} \leftarrow \mathcal{U}(\mathbb{Z}_q^m)$.*

The asymmetric variant, denoted by $\mathsf{ALWE}_{n,m,q,\alpha_1,\alpha_2}$, is defined in the same way, except that the secret and error are independently sampled from $\chi_{\alpha_1}^n$ and $\chi_{\alpha_2}^m$, respectively.

Definition 3 ([ZYF+20], Asymmetric-LWE (ALWE)). *Let $n, m > 0$ and $q \geq 2$ be integers, and let $\chi_{\alpha_1}, \chi_{\alpha_2}$ be specified distributions on $\mathbb{Z}$. The parameters $\alpha_1, \alpha_2 > 0$ index this family and are integers only when the family requires integer parameters. Suppose $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{m\times n})$, $\mathbf{s} \hookleftarrow \chi_{\alpha_1}^n$ and $\mathbf{e} \hookleftarrow \chi_{\alpha_2}^m$ independently.*

- *Search-ALWE$_{n,m,q,\alpha_1,\alpha_2}$: Given the sample $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$, find $\mathbf{s} \bmod q$.*
- *Decision-ALWE$_{n,m,q,\alpha_1,\alpha_2}$: Distinguish $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q)$ from the uniform distribution over $\mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$.*

We next introduce the Module-LWE problem, which extends LWE to polynomial quotient rings.

Definition 4 (Module-LWE (MLWE)). *Let N be a power of two and define $R = \mathbb{Z}[x]/(x^N + 1)$ and $R_q = R/qR$. For a distribution χ_α over $\mathbb{Z}$, let $\chi_{\alpha,R}$ denote its coefficient-wise extension to R. Let $d, m > 0$ be the module rank and the number of samples, respectively. Sample*

$$\mathbf{A} \leftarrow \mathcal{U}(R_q^{m\times d}), \qquad \mathbf{s} \hookleftarrow \chi_{\alpha,R}^d, \qquad \mathbf{e} \hookleftarrow \chi_{\alpha,R}^m$$

independently, and set $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod q$. We define the following two problems:

- *Search-MLWE$_{N,d,m,q,\chi_\alpha}$: Given $(\mathbf{A}, \mathbf{b})$, recover $\mathbf{s} \bmod q$.*
- *Decision-MLWE$_{N,d,m,q,\chi_\alpha}$: Distinguish $(\mathbf{A}, \mathbf{b})$ from $(\mathbf{A}, \mathbf{u})$, where $\mathbf{u} \leftarrow \mathcal{U}(R_q^m)$.*

The asymmetric Module-LWE (AMLWE) problem is defined analogously, where the secret and error are independently sampled from $\chi_{\alpha_1,R}^d$ and $\chi_{\alpha_2,R}^m$, respectively. When $\alpha_1 = \alpha_2$, AMLWE reduces to MLWE.

We introduce the Short Integer Solution (SIS) and Inhomogeneous SIS (ISIS) problems.

Definition 5 (SIS and ISIS). *Let $n, m > 0$ and $q \geq 2$ be integers, and let $\beta > 0$ be real. We introduce the following two problems with Euclidean norms:*

- *SIS$_{n,m,q,\beta}$: Given a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, find a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ which satisfies $\mathbf{A}\mathbf{x} = \mathbf{0} \bmod q$ and $\|\mathbf{x}\| \leq \beta$.*

- *ISIS$_{n,m,q,\beta}$: Given a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, a uniformly random $\mathbf{u} \in \mathbb{Z}_q^n$, find a vector $\mathbf{x} \in \mathbb{Z}^m$ which satisfies $\mathbf{Ax} = \mathbf{u} \mod q$ and $\|\mathbf{x}\| \le \beta$.*

Definition 6 ([ZYF+20], Asymmetric-SIS (ASIS) and Asymmetric-ISIS (AISIS)). *Let $n, m_1, m_2 > 0$ and $q \ge 2$ be integers and $\beta_1, \beta_2 > 0$ be real numbers. We introduce the following two problems with Euclidean norms:*

- *ASIS$_{n,m_1,m_2,q,\beta_1,\beta_2}$: Given a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times(m_1+m_2)}$, find a non-zero vector $\mathbf{x} := (\mathbf{x}_1^\top, \mathbf{x}_2^\top)^\top \in \mathbb{Z}^{m_1+m_2}$ which satisfies $\mathbf{Ax} = \mathbf{0} \mod q$ and $\|\mathbf{x}_1\| \le \beta_1$, $\|\mathbf{x}_2\| \le \beta_2$.*
- *AISIS$_{n,m_1,m_2,q,\beta_1,\beta_2}$: Given a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times(m_1+m_2)}$, an independent uniformly random $\mathbf{u} \in \mathbb{Z}_q^n$, find a vector $\mathbf{x} := (\mathbf{x}_1^\top, \mathbf{x}_2^\top)^\top \in \mathbb{Z}^{m_1+m_2}$ which satisfies $\mathbf{Ax} = \mathbf{u} \mod q$ and $\|\mathbf{x}_1\| \le \beta_1$, $\|\mathbf{x}_2\| \le \beta_2$.*

We can similarly define the asymmetric Module-SIS (AMSIS) problem by extending this asymmetric setting to the module lattice case and using coefficient Euclidean norms. Writing $P \preceq Q$ when a solver for Q yields a solver for P, we have

$$\mathrm{SIS}_{n,m_1+m_2,q,\sqrt{\beta_1^2+\beta_2^2}} \preceq \mathrm{ASIS}_{n,m_1,m_2,q,\beta_1,\beta_2} \preceq \mathrm{SIS}_{n,m_1+m_2,q,\min(\beta_1,\beta_2)}.$$

Indeed, $\|\mathbf{x}\|_2^2 = \|\mathbf{x}_1\|_2^2 + \|\mathbf{x}_2\|_2^2$, and each reduction returns the same solution. The corresponding statement with $\max(\beta_1, \beta_2)$ as the left-hand radius applies to infinity norms, not to the Euclidean norms used here. The same inclusions hold for the inhomogeneous problems. In a signature scheme, unequal sampling widths do not by themselves impose separate bounds on adversarial response blocks. A verifier with a single global Euclidean bound leads to an ordinary MSIS problem. Its precise matrix distribution, including any normal-form identity block, is specified in Appendix B.

2.3 Digital Signature

A *digital signature scheme* is defined by a triple of probabilistic polynomial-time (PPT) algorithms $(\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$. For rejection-based sampling, we allow expected polynomial time and include its expected cost. A fixed time limit must account for the resulting abort probability. The algorithms have the following inputs and outputs:

- $\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$: Takes a security parameter λ as input and outputs a public key pk and a secret key sk.
- $\mathsf{Sign}(\mathsf{sk}, \mu) \to \sigma$: Takes the secret key sk and a message μ as input, and outputs a signature σ.
- $\mathsf{Verify}(\mathsf{pk}, \mu, \sigma) \to \{0, 1\}$: Takes the public key pk, a message μ, and a signature σ as input, and outputs 1 if the signature is valid, or 0 otherwise.

Security Model: Existential Unforgeability under Chosen-Message Attacks (EUF-CMA). A digital signature scheme achieves *EUF-CMA security* if, for every PPT adversary $\mathcal{A}$, the advantage $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{EUF-CMA}}(\lambda)$ in the following game is negligible in the security parameter λ:

- **Setup:** The challenger runs $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$ and gives pk to $\mathcal{A}$.
- **Signing Queries:** $\mathcal{A}$ may adaptively query the signing oracle on messages $\mu_1, \ldots, \mu_Q$. For each query μ_i, the challenger returns $\sigma_i \leftarrow \mathsf{Sign}(\mathsf{sk}, \mu_i)$.
- **Forgery:** $\mathcal{A}$ outputs a forgery (μ^*, σ^*).
- **Success:** $\mathcal{A}$ wins if:
  - $\mathsf{Verify}(\mathsf{pk}, \mu^*, \sigma^*) = 1$,
  - and μ^* was not previously queried to the signing oracle.

The scheme is EUF-CMA secure if:

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{EUF-CMA}}(\lambda) = \Pr[\mathcal{A} \text{ wins}] \le \mathsf{negl}(\lambda).$$

Modulus Rounding. We utilize rounding to diminish the bit lengths of public keys and signatures. Let $v \in \mathbb{N} \setminus \{0\}$. Any integer $x \in \mathbb{Z}$ can be uniquely expressed as:

$$x = 2^v \cdot x_+ + x_-, \quad (x_+, x_-) \in \mathbb{Z} \times [-2^{v-1}, 2^{v-1} - 1],$$

which amounts to separating the lower-order bits from the higher-order ones. We define the function $\lfloor\cdot\rceil_v : \mathbb{Z} \to \mathbb{Z}$ such that $\lfloor x\rceil_v = \lfloor x/2^v + 1/2\rfloor = x_+$, where $\lfloor\cdot\rfloor$ is the floor function. Thus the "rounding half-up" convention is used, and half-way values are rounded up. With a slight abuse of notation, when $q > 2^v$, we extend $\lfloor\cdot\rceil_v$ to inputs in $\mathbb{Z}_q$, with output in $\mathbb{Z}_{q_v}$, where $q_v = \lfloor q/2^v\rfloor$. Formally, we define

$$\lfloor\cdot\rceil_v : \mathbb{Z}_q \longrightarrow \mathbb{Z}_{q_v}, \qquad \lfloor x\rceil_v = \lfloor x/2^v + 1/2\rfloor \bmod q_v,$$

where $x \in \mathbb{Z}_q$ is represented by its canonical integer in $\{0, \ldots, q-1\}$. The function $\lfloor\cdot\rceil_v$ naturally extends to vectors component-wise. Write $\mathrm{cent}_q(x)$ for the representative of $x \bmod q$ in $[-q/2, q/2) \cap \mathbb{Z}$, also applied component-wise. Reconstruction errors below use this centered modular distance, since the ordinary integer difference can be large at the wrap-around point.

We state an elementary bound for this rounding convention.

Lemma 1. *Let v, q be positive integers such that $q > 2^v$, and set $q_v = \lfloor q/2^v\rfloor$. Assume in addition that the remainder $q_\perp = q - 2^v q_v$ satisfies $0 \le q_\perp \le 2^{v-1} - 1$, and put $E = 2^{v-1} + q_\perp$. Then, for any $x \in \mathbb{Z}_q$, we have*

$$\left|\mathrm{cent}_q(x - 2^v\lfloor x\rceil_v)\right| \le E \le 2^v - 1.$$

Moreover, for any $\mathbf{x}, \boldsymbol{\delta} \in \mathbb{Z}_q^n$, we have

$$\left\|\mathrm{cent}_q\big(2^v\big(\lfloor\mathbf{x}+\boldsymbol{\delta}\rceil_v - \lfloor\mathbf{x}\rceil_v\big)\big)\right\|_2 \le \|\,\mathrm{cent}_q(\boldsymbol{\delta})\|_2 + 2E\sqrt{n},$$

where the high-bit values are represented in $\{0, \ldots, q_v - 1\}$.

Proof. Put $a = 2^v$. Before reduction modulo q_v, half-up rounding has error in $[-a/2, a/2 - 1]$. If the rounded value is q_v, its reduced value is zero and the canonical input lies in $[aq_v - a/2, q - 1]$. Subtracting q then bounds the absolute reconstruction error by $a/2 + q_\perp = E$. Centering can only decrease this absolute value. For the vector bound, write $\mathbf{e}(\mathbf{x}) = \mathrm{cent}_q(\mathbf{x} - a\lfloor\mathbf{x}\rceil_v)$, so that $\|\mathbf{e}(\mathbf{x})\|_2 \le E\sqrt{n}$. The identity

$$a\big(\lfloor\mathbf{x}+\boldsymbol{\delta}\rceil_v - \lfloor\mathbf{x}\rceil_v\big) \equiv \boldsymbol{\delta} - \mathbf{e}(\mathbf{x}+\boldsymbol{\delta}) + \mathbf{e}(\mathbf{x}) \pmod q$$

and the triangle inequality give the claimed bound. □

The specific rounding and hint reconstruction used by the compressed scheme are given in Appendix C.

3 Asymmetric G+G: More Compact Signature Scheme Based on Asymmetric LWE

In this section, we propose a variant of the G+G signature that permits smaller signature encodings by using the signature truncation technique from [BG14]. Specifically, we introduce ALWE in the key generation algorithm. Through an analysis of the largest singular value of a scaled secret-key matrix, we derive admissible Gaussian width parameters and experimentally optimize their allocation. The revised parameter comparison below evaluates the resulting sizes against the corrected G+G baseline.

Description of the G+G scheme. We first introduce the G+G protocol as well as some notation. Let $m \ge \ell > 0$, $k > m + \ell$, and $\mathbf{J} := (\mathbf{I}_\ell \mid \mathbf{0}^{\ell\times(m-\ell)})^\top \in \mathbb{Z}^{m\times\ell}$. Let χ be a probability distribution over $\mathbb{Z}$. Define the challenge space $\mathcal{C} \subseteq \mathbb{Z}_2^\ell$. Let $\sigma, s > 0$ and define $\Sigma : \mathbb{Z}^{k\times\ell} \to \mathbb{R}^{k\times k}$ as

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

The scheme is also parametrized by an odd modulus q and an acceptance bound γ. The G+G identification protocol is described in Fig. 1. Its Gaussian parameters must satisfy the convolution condition for every key in the stated key-generation support. The modified generator in Fig. 2 enforces the corresponding weighted condition, with $\mathbf{D}$ defined below.

Note that the Gaussian width of the signature distribution depends on the shape matrix $\Sigma(\mathbf{S})$. We begin by analyzing the largest singular value of the secret key $\mathbf{S}$. For this block display, write $\mathbf{S} = (\mathbf{T}_1^\top \mid \mathbf{T}_2^\top \mid \mathbf{T}_3^\top)^\top$, where $\mathbf{T}_1 = \mathbf{I}_\ell$, $\mathbf{T}_2 = \mathbf{S}_1$ and $\mathbf{T}_3 = \mathbf{S}_2$. The shape matrix can be written as

$$\Sigma(\mathbf{S}) = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} & \Sigma_{13} \\ \Sigma_{21} & \Sigma_{22} & \Sigma_{23} \\ \Sigma_{31} & \Sigma_{32} & \Sigma_{33} \end{pmatrix} \quad \text{where } \Sigma_{ij} = \begin{cases} \sigma^2 \mathbf{I}_\ell - 4s^2 \mathbf{T}_1 \mathbf{T}_1^\top & i = j = 1; \\ \sigma^2 \mathbf{I}_{k-m-\ell} - 4s^2 \mathbf{T}_2 \mathbf{T}_2^\top & i = j = 2; \\ \sigma^2 \mathbf{I}_m - 4s^2 \mathbf{T}_3 \mathbf{T}_3^\top & i = j = 3; \\ -4s^2 \mathbf{T}_i \mathbf{T}_j^\top & i \neq j. \end{cases}$$

The Gaussian parameter σ critically influences both signature size and verification bound—the former being a key performance metric in signature schemes, and the latter a decisive factor for security against forgery attacks. To enhance both security and efficiency, we propose optimizing the Gaussian sampling by replacing the spherical Gaussian with a non-spherical distribution under otherwise identical conditions.

$\mathsf{IGen}(1^\lambda)$

1 : $\mathbf{A}_1 \hookleftarrow \mathcal{U}(\mathbb{Z}_q^{m \times (k-m-\ell)})$
2 : $(\mathbf{S}_1, \mathbf{S}_2) \hookleftarrow \chi^{(k-m-\ell)\times \ell} \times \chi^{m \times \ell}$
3 : $\mathbf{B} \leftarrow \mathbf{A}_1 \mathbf{S}_1 + \mathbf{S}_2 \bmod q$
4 : $\mathbf{A} \leftarrow (q\mathbf{J} - 2\mathbf{B} | 2\mathbf{A}_1 | 2\mathbf{I}_m) \in \mathbb{Z}_{2q}^{m \times k}$
5 : $\mathbf{S} \leftarrow (\mathbf{I}_\ell | \mathbf{S}_1^\top | \mathbf{S}_2^\top)^\top \in \mathbb{Z}^{k \times \ell}$
6 : $\mathrm{vk} \leftarrow \mathbf{A}, \mathrm{sk} \leftarrow \mathbf{S}$
7 : **return** $(\mathrm{vk}, \mathrm{sk})$

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

Fig. 1. G+G Protocol

Asymmetric LWE in the key generation. Our idea, initially proposed in [ZYF+20], is simple and intuitive, though its application to the analysis of the G+G signature is somewhat more involved. In simple terms, we need to modify the distributions of $\mathbf{S}_1, \mathbf{S}_2$ in the second step of the $\mathsf{IGen}(1^\lambda)$ algorithm such that they differ—namely, by *narrowing*

the distribution of $\mathbf{S}_1$ while *widening* that of $\mathbf{S}_2$. For specific modifications, refer to Fig. 2. The following analysis identifies when this adjustment permits smaller response widths. Its security and size effects must be checked for the selected distributions and parameters.

$\mathsf{IGen}'(1^\lambda)$

$\mathbf{A}_1 \hookleftarrow \mathcal{U}(\mathbb{Z}_q^{m\times(k-m-\ell)})$

$(\mathbf{S}_1, \mathbf{S}_2) \hookleftarrow \chi_1^{(k-m-\ell)\times\ell} \times \chi_2^{m\times\ell}$

$\mathbf{B} \leftarrow \mathbf{A}_1\mathbf{S}_1 + \mathbf{S}_2 \bmod q$

$\mathbf{A} \leftarrow (q\mathbf{J} - 2\mathbf{B}|2\mathbf{A}_1|2\mathbf{I}_m) \in \mathbb{Z}_{2q}^{m\times k}$

$\mathbf{S} \leftarrow (\mathbf{I}_\ell|\mathbf{S}_1^\top|\mathbf{S}_2^\top)^\top \in \mathbb{Z}^{k\times\ell}$

if $\sin_1(\mathbf{D}^{-1}\mathbf{S}) > (\sqrt{8}s)^{-1}$ **then**

  return $\mathsf{IGen}'(1^\lambda)$

end if

$\mathrm{vk} \leftarrow \mathbf{A}, \quad \mathrm{sk} \leftarrow \mathbf{S}$

return $(\mathrm{vk}, \mathrm{sk})$

Fig. 2. Modified key generation algorithm

3.1 Analysis of Distribution Convolution

In this section, we analyze the general case where the last two components of the secret key $\mathbf{S} = (\mathbf{I}_\ell \mid \mathbf{S}_1^\top \mid \mathbf{S}_2^\top)^\top$ are drawn from a distribution family with different parameters. The matrix analysis applies to each fixed secret key and does not require Gaussian secret coefficients. In particular, the concrete instances below retain bounded-uniform coefficient distributions.

In the following analysis, let $\mathbf{S}_1 \in \mathbb{Z}^{(k-m-\ell)\times\ell}$ and $\mathbf{S}_2 \in \mathbb{Z}^{m\times\ell}$. The distribution of the perturbation vector $\mathbf{y}$ changes, and its modified shape matrix is given by

$$\Sigma(\mathbf{S}) = \begin{bmatrix} \sigma_1^2 \cdot \mathbf{I}_\ell & & \\ & \sigma_2^2 \cdot \mathbf{I}_{k-m-\ell} & \\ & & \sigma_3^2 \cdot \mathbf{I}_m \end{bmatrix} - 4s^2\mathbf{S}\mathbf{S}^\top.$$

The ideal response distribution is $\mathcal{D}_{\mathbb{Z}^k,\sqrt{\Sigma}}$, where $\Sigma = \sigma_1^2\mathbf{I}_\ell \oplus \sigma_2^2\mathbf{I}_{k-m-\ell} \oplus \sigma_3^2\mathbf{I}_m$. The real response is close to this distribution under the conditions below, with an explicit likelihood-ratio bound proved in Appendix B. We analyze the relationship between the largest singular values of $\mathbf{S}_1, \mathbf{S}_2$ and the Gaussian widths $\sigma_1, \sigma_2, \sigma_3$ to select an admissible allocation. Throughout, a width σ_i corresponds to continuous standard deviation $\sigma_i/\sqrt{2\pi}$. To justify the convolution, it is sufficient to take $s/\sqrt{2} \geq \eta_\epsilon(\mathbb{Z}^\ell)$ and require the following matrix to be positive semidefinite. This condition makes the perturbation shape positive definite and supplies the smoothing bound needed after completing the square, as proved in Appendix B:

$$\Sigma_1 = \begin{bmatrix} \sigma_1^2 \cdot \mathbf{I}_\ell & & \\ & \sigma_2^2 \cdot \mathbf{I}_{k-m-\ell} & \\ & & \sigma_3^2 \cdot \mathbf{I}_m \end{bmatrix} - 8s^2\mathbf{S}\mathbf{S}^\top \geq 0.$$

In the following, we address the problem of estimating the largest singular value of the secret key matrix. This is essential to ensure the positive definiteness condition holds while achieving more optimal parameters σ_1, σ_2, and σ_3.

Bounds on Singular Values of Secret Key $\mathbf{S}$. Due to the nonlinearity of singular values, deriving an upper bound for the largest singular value of $\mathbf{S} = (\mathbf{I}_\ell \mid \mathbf{S}_1^\top \mid \mathbf{S}_2^\top)^\top$ would be non-trivial. Intuitively, the largest singular value of $\mathbf{S}$ determines the lower bound for the Gaussian width of the signature distribution. To make the analysis tractable, we apply a diagonal rescaling analogous to that used for the ellipsoidal Gaussian setting studied by Jia et al. [JHT22].

Our core idea is to introduce a scaling factor t between σ_2 and σ_3 to reduce the number of free variables. This approach is motivated by the inherent proportionality we observe between σ_2 and σ_3. By exploring t within a bounded range, we obtain a sufficiently accurate estimate for our experimental goals, despite its potential suboptimality. A routine calculation shows that

$$\begin{bmatrix} \sigma_1^2\mathbf{I}_\ell & & \\ & \sigma_2^2\mathbf{I}_{k-m-\ell} & \\ & & \sigma_3^2\mathbf{I}_m \end{bmatrix} - 8s^2\mathbf{S}\mathbf{S}^\top \geq 0$$

is equivalent to

$$\mathbf{I}_k - 8s^2\mathbf{D}^{-1}\mathbf{S}\mathbf{S}^\top\mathbf{D}^{-1} \geq 0, \qquad \mathbf{D} := \mathrm{diag}(\sigma_1\mathbf{I}_\ell, \sigma_2\mathbf{I}_{k-m-\ell}, \sigma_3\mathbf{I}_m).$$

Hence,

$$\lambda_{\max}\big(\mathbf{D}^{-1}\mathbf{S}\mathbf{S}^\top\mathbf{D}^{-1}\big) \leq (8s^2)^{-1},$$

where $\lambda_{\max}(\cdot)$ denotes the largest eigenvalue. Now define

$$\widetilde{\mathbf{S}} := \mathbf{D}^{-1}\mathbf{S} = \begin{pmatrix} \sigma_1^{-1}\mathbf{I}_\ell \\ \sigma_2^{-1}\mathbf{S}_1 \\ \sigma_3^{-1}\mathbf{S}_2 \end{pmatrix}.$$

Then

$$\mathbf{D}^{-1}\mathbf{S}\mathbf{S}^\top\mathbf{D}^{-1} = \widetilde{\mathbf{S}}\,\widetilde{\mathbf{S}}^\top,$$

and therefore

$$\|\widetilde{\mathbf{S}}\|_{\mathrm{op}}^2 = \lambda_{\max}\big(\mathbf{D}^{-1}\mathbf{S}\mathbf{S}^\top\mathbf{D}^{-1}\big) \leq (8s^2)^{-1}.$$

Setting $t = \sigma_2/\sigma_3$, we rewrite

$$\widetilde{\mathbf{S}} = \begin{pmatrix} \sigma_1^{-1}\mathbf{I}_\ell \\ \sigma_2^{-1}\mathbf{S}_1 \\ \sigma_2^{-1}t\mathbf{S}_2 \end{pmatrix} = \begin{pmatrix} \sigma_1^{-1}\mathbf{I}_\ell \\ \sigma_2^{-1}\mathbf{S}_{\mathrm{new}} \end{pmatrix}, \qquad \mathbf{S}_{\mathrm{new}} := \begin{pmatrix} \mathbf{S}_1 \\ t\mathbf{S}_2 \end{pmatrix}.$$

Consequently,

$$\widetilde{\mathbf{S}}^\top\widetilde{\mathbf{S}} = \sigma_1^{-2}\mathbf{I}_\ell + \sigma_2^{-2}\mathbf{S}_{\mathrm{new}}^\top\mathbf{S}_{\mathrm{new}},$$

which implies

$$\sin_1(\widetilde{\mathbf{S}})^2 = \sigma_1^{-2} + \sigma_2^{-2}\sin_1(\mathbf{S}_{\mathrm{new}})^2.$$

To estimate the cost of spectral rejection, we generate random samples and evaluate the distribution of $\sin_1(\mathbf{S}_{\mathrm{new}})$ for each scaling parameter $t = \sigma_2/\sigma_3$. These samples guide parameter selection and estimate the acceptance probability. The inequality above is then checked for every generated key. An empirical median is not a guarantee that an individual key satisfies the required spectral bound.

Comparison with known Gaussian-type distributions. At a high level, we compare three allocations for the ideal distribution of the signature $\mathbf{z}$: a spherical Gaussian, a two-width ellipsoidal Gaussian, and the three-block non-spherical Gaussian considered in this work. These three cases correspond to the G+G scheme, the hash-and-sign signature scheme in [JHT22], and our G+G variant, respectively. Moreover, the three-block diagonal shape offers more parametric flexibility than these two restricted cases. All of them are ellipsoidal Gaussian distributions in the usual geometric sense. By setting $\sigma_1 \neq \sigma_2 = \sigma_3$, the distribution reduces to the two-width case; if all σ_i are identical, it simplifies to a spherical Gaussian.

This highlights the potential of the distribution proposed in this work, as it allows us to adjust the *allocation* of Gaussian widths to the encoding costs of different response blocks. In particular, signature truncation [BG14] permits the last response block to be reconstructed from a hint. This makes it useful to increase σ_3 while reducing σ_1 and σ_2,

provided that the spectral condition and key-security estimates remain satisfied. The hint cost and its effect on the verification bound must also be included. Appendix C gives the explicit reconstruction and its security accounting.

Indeed, if we aim to find the "narrowest" possible distribution for $\mathbf{z}$, the spherical restriction can be suboptimal, since the sufficient convolution condition requires the Gaussian width σ of the signature distribution to satisfy $\sigma > \sqrt{8} \cdot \sin_1(\mathbf{S}) \cdot s$, as demonstrated in the G+G scheme. Allowing two widths relaxes this restriction, while a third width separates the retained response from the block reconstructed using a hint. In the aforementioned calculations regarding σ_i, our objective is to minimize σ_1 and σ_2 at the cost of a slight increase in σ_3, which is more compatible with signature truncation techniques. The resulting improvement depends on the concrete key distribution, correctness target, and complete encoding cost.

At this point, an intuitive question remains: *Why use three distinct Gaussian widths?* The public key has three algebraic blocks, namely a column containing the public-key sample, a uniformly generated matrix, and an identity matrix. The secret and response have corresponding blocks, so three widths provide a natural allocation for this construction. This observation does not prove global optimality, and finer block partitions could be studied separately.

Security of the ALWE problem. In [ZYF+20, HDD21], the authors experimentally discussed the computational complexity of solving ALWE and LWE under known attacks using the secret and error standard deviations. Such estimates depend on the actual distributions, dimensions, and available samples. In this work, key recovery is evaluated for the bounded-uniform distributions $\chi_{\zeta_1}, \chi_{\zeta_2}$. The response widths σ_2, σ_3 are optimized separately and do not determine the entropy of those secret distributions.

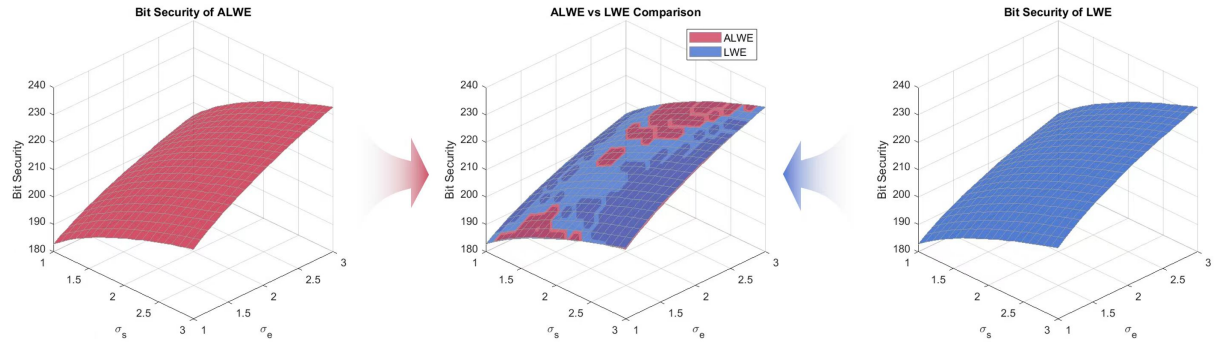


Fig. 3. Historical illustrative comparison of LWE and ALWE attack estimates from the original version. This figure is not used to select or validate the revised parameters.

Figure 3 retains the illustrative attack-model comparison from the original version. Its axes give the secret and error standard deviations σ_s, σ_e and the estimated attack cost. One surface used a direct primal-attack calculation, whereas the other used the historical geometric-mean proxy

$$\sigma_{\text{proxy}} = \sqrt{\sigma_s \sigma_e}.$$

The superimposed view illustrates agreement for the inputs and attack model used there. It is not a reduction from ALWE to uniform-secret LWE or short-secret LWE, and it does not establish equality of their security against other attacks. The revised parameter tables use the actual asymmetric distributions and do not use this figure or proxy as security evidence. The normal-form LWE convention used here is stated in Section 2.2.

Impact on the provable security. The existential unforgeability of the G+G signature relies on completeness, commitment recoverability, the HVZK property, key indistinguishability, and lossy special soundness. Our modification changes the secret distributions and the Gaussian shape, so these properties should be checked for the modified algorithms. Appendix B gives the convolution and security (in the ROM) arguments, including the acceptance event of key generation. As in the corrected Theorem 6 of G+G [DPS26a], key indistinguishability is assumed for the actual

accepted public-key distribution, denoted $\mathsf{AMLWE}^{\mathsf{Acc}}$. The verification rule below uses one Euclidean norm bound, so ordinary normal-form MSIS with radius 2γ suffices for extraction. The compressed commitment and reconstructed response are analyzed in Appendix C.

3.2 Optimized G+G: An Experimental Perspective

In this section, we evaluate the efficiency and security of an optimized variant of G+G by considering key recovery attacks and signature forgery attacks, which are typically examined in signature security assessments. We first present the optimized version of the G+G signature. Subsequently, we conduct an evaluation of the variant's signatures and provide a detailed comparison with the original G+G scheme.

We introduce the module version of the new G+G signature variant in Fig. 4 as well as some notation. Let $n > 1$ be a power of two, $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$, $m > 0$, $k > m + 1$, and $\ell = 1$. The module modulus is an odd prime $q \equiv 1 \pmod 4$. Put $r = k - m - 1$ and $\mathbf{j} = (\eta^*, 0, \ldots, 0)^\top \in \mathcal{R}^m$. We take $(\eta, \eta^*) = (1 + x^{n/2}, 1 - x^{n/2})$ when a $2^{n/2}$-element challenge space suffices, and $(\eta, \eta^*) = (2, 1)$ when a larger space is required, as in the revised Set-II. In both cases $\eta\eta^* = 2$. The challenge space is $\mathcal{C} = \mathcal{R}/\eta\mathcal{R}$, with canonical binary representatives of degree below $n/2$ and n, respectively.

For an integer $\zeta \geq 0$, let χ_ζ be the distribution on $\mathcal{R}$ whose coefficients are independently uniform in $\{-\zeta, \ldots, \zeta\}$. Thus its coefficient variance is $\zeta(\zeta + 1)/3$, and the revised instances retain the original bounded-uniform distribution family. The values ζ_1, ζ_2 are secret-distribution bounds and are distinct from the challenge factor η.

In coefficient embedding, define

$$\mathbf{D} = \mathrm{diag}(\sigma_1\mathbf{I}_n, \sigma_2\mathbf{I}_{rn}, \sigma_3\mathbf{I}_{mn}), \quad \mathbf{T} = \mathrm{rot}(\eta\mathbf{s}), \quad \Sigma(\mathbf{s}) = \mathbf{D}^2 - s^2\mathbf{T}\mathbf{T}^\top.$$

Here $\mathrm{rot}(f)$ is the coefficient matrix of multiplication by f modulo $x^n + 1$, extended to vectors by vertical stacking. The module key generator checks $\|\mathbf{D}^{-1}\mathbf{T}\|_{\mathrm{op}} \leq 1/(\sqrt{2}s)$ for the actual secret, including its low-bit correction. The routines Decomp, PackCommit, Compress and Recover are specified in Appendix C. They make both public-key and response compression explicit. The oracle G expands a 256-bit public seed and is domain-separated from H, under the seeded model stated in Appendix B. In Fig. 4, Sign and Verify use the matrix $\mathbf{A}$ reconstructed from $\mathsf{vk} = (\tau, \mathbf{b}_1/2)$, and the signer obtains this verification key from $\mathsf{sk} = (\mathsf{vk}, \mathbf{s})$.

Comparison of parameter size. In Table 1, we propose the revised parameter selections for the G+G variant signature. We retain the original bounded-uniform secret family and the pairs $(\zeta_1, \zeta_2) = (1, 2)$ and $(1, 3)$. The module dimensions are aligned with the corrected G+G signature of November 18, 2026 [DPS26a]. The Gaussian convention is $\rho_s(x) = \exp(-\pi\|x\|^2/s^2)$, with $s = 5.68$ and $\epsilon = 2^{-64}$. In particular, the former value $s = 14.22$ cannot be reused as a width under this convention. The challenge spaces have sizes 2^{128} and 2^{256}, respectively. Set-II uses $\eta = 2$, since $\eta = 1 + x^{n/2}$ would only give 2^{128} challenges at $n = 256$.

The widths were selected using separate training samples, followed by the independent validation reported in Table 1. For every generated key, the algorithm checks the actual weighted spectral condition, including $\mathbf{s}_2 - \mathbf{b}_0$. Signing makes a single attempt. We choose the verification bound using Lemma 4 with $t = 41 \ln 2$, add the compression allowance $\sqrt{mn}(\alpha/4 + 1)$, and round the result upwards. The contribution of fresh oracle inputs to the failure probability is at most $R_\epsilon 2^{-41} < 2^{-40}$. Oracle-input collisions and sampler implementation error are accounted for separately in Appendix B.

In Table 2, we compare our two sets with symmetric G+G controls using the corrected reference moduli 64513 and 114689, respectively. These controls use $(\zeta_1, \zeta_2) = (1, 1)$, spherical widths 586.23 and 995.50, and truncation parameters 256 and 512. In particular, the first control uses $\alpha = 256$ to retain the 120-bit direct MSIS model target under the stricter tail bound. For a common comparison, we recompute their verification bounds and entropy estimates with the same correctness target and cyclic-hint calculation used for our sets. Thus the G+G columns below are matched recalculations, rather than verbatim copies of the published size table. The nominal 120- and 180-bit targets refer to the stated classical direct-attack model, not proven concrete EUF-CMA security levels. Appendix A specifies all attack-model inputs.

The signature-size estimate includes the challenge, the retained response, and the complete hint. It is the ideal entropy quantity defined in Appendix C, not a measured rANS output length. An implementation must additionally

KeyGen(1^λ)

repeat
  $\tau \hookleftarrow \mathcal{U}(\{0,1\}^{256})$, $(\mathbf{a} \mid \mathbf{A}_1) \leftarrow G(\tau)$
  $(\mathbf{s}_1, \mathbf{s}_2) \hookleftarrow \chi_{\zeta_1}^r \times \chi_{\zeta_2}^m$
  $\mathbf{b} \leftarrow \mathbf{a} + \mathbf{A}_1\mathbf{s}_1 + \mathbf{s}_2 \bmod q$
  $(\mathbf{b}_0, \mathbf{b}_1) \leftarrow \mathsf{Decomp}(\mathbf{b})$
  $\mathbf{s} \leftarrow (1 \mid \mathbf{s}_1^\top \mid (\mathbf{s}_2 - \mathbf{b}_0)^\top)^\top$
until $\|\mathbf{D}^{-1}\,\mathrm{rot}(\eta\mathbf{s})\|_{\mathrm{op}} \le 1/(\sqrt{2}s)$
$\mathbf{A} \leftarrow (2(\mathbf{a} - \mathbf{b}_1) + q\mathbf{j} \mid 2\mathbf{A}_1 \mid 2\mathbf{I}_m) \bmod 2q$
$\mathsf{vk} \leftarrow (\tau, \mathbf{b}_1/2)$, $\mathsf{sk} \leftarrow (\mathsf{vk}, \mathbf{s})$
return $(\mathsf{vk}, \mathsf{sk})$

Sign($\mathbf{A}, \mathbf{s}, \mu$)

$\mathbf{y} \hookleftarrow \mathcal{D}_{\mathcal{R}^k, \sqrt{\Sigma(\mathbf{s})}}$
$\mathbf{w} \leftarrow \mathbf{A}\mathbf{y} \bmod 2q$, $\overline{\mathbf{w}} \leftarrow \mathsf{PackCommit}(\mathbf{w})$
$c \leftarrow H(\overline{\mathbf{w}}, \mu)$
$u \hookleftarrow \mathcal{D}_{\mathcal{R}, s, -\eta^* c/2}$
$\mathbf{z} \leftarrow \mathbf{y} + (c + \eta u)\mathbf{s}$
return $\mathsf{Compress}(\mathbf{A}, \mathbf{z}, c, \overline{\mathbf{w}})$

Verify($\mathbf{A}, \mu, \mathbf{z}_0, c, \mathbf{h}$)

Reject malformed or noncanonical encodings
$(\overline{\mathbf{w}}, \widehat{\mathbf{z}}) \leftarrow \mathsf{Recover}(\mathbf{A}, \mathbf{z}_0, c, \mathbf{h})$
if $c = H(\overline{\mathbf{w}}, \mu)$ and $\|\widehat{\mathbf{z}}\| \le \gamma$ **then**
  return 1
end if
return 0

Fig. 4. The new G+G variant signature, with the revised spectral test and explicit compression routines.

account for coder state, framing, finite-precision probabilities, and any tail escape encoding. The public-key count includes the 32-byte expansion seed and every coefficient of $\mathbf{b}_1/2$. The larger moduli are needed to reach the direct MSIS model targets with the revised widths and verification bounds. They increase the public keys by 64 and 96 bytes, respectively. With these costs included, the decreases in public-key-plus-signature entropy are approximately 77.61 and 33.58 bytes. We therefore replace the former 25% signature-size claim by the more limited comparison in Table 2. This is a comparison to the specified corrected G+G parameters, not a global optimization over all symmetric parameter choices.

The acceptance rates are measured rather than fixed at 0.5. For the specified sampler and low-bit-corrected spectral test, the G+G-I control accepts approximately 26.2% of keys, whereas the other three rates are close to one half. Each row has an independent 100000-trial validation. The observed rates determine expected key-generation attempts, but do not certify a running-time speedup for the Gaussian sampler.

Table 3 separates the direct-attack comparison from the radius 2γ used in Theorem 1. The theorem also has the displayed query and resetting losses. The key-cost figures in Table 2 concern unconditioned (A)MLWE. Using the lower endpoints of the empirical 99% acceptance intervals in Equation (5) gives key-recovery sensitivity values of 122.75 and 225.73 bits for our sets. These values do not prove the conditioned decisional assumption. The quantum

Table 1. Revised parameter settings for the G+G variant signature.

Parameter	Set-I	Set-II
q	125441	260609
n	256	256
$(m, k-m)$	$(2, 4)$	$(3, 6)$
(ζ_1, ζ_2)	$(1, 2)$	$(1, 3)$
(η, η^*)	$(1+x^{128}, 1-x^{128})$	$(2, 1)$
$\log_2 \lvert\mathcal{C}\rvert$	128	256
s	5.68	5.68
$(\sigma_1, \sigma_2, \sigma_3)$	$(37.07, 588.49, 995.43)$	$(54.98, 955.75, 2202.16)$
α	512	1024
Public-key low-bit parameter d	1	1
Verification bound γ	16090.13	39642.01
Tail-error bound ϵ_c	$< 2^{-40}$	$< 2^{-40}$
Accepted keys / independent trials	51390/100000	52287/100000
Acceptance probability, 99% interval	$[0.50982, 0.51798]$	$[0.51880, 0.52694]$

Table 2. Comparison with the corrected G+G parameter sets under matched accounting.

	G+G-I	Set-I	G+G-II	Set-II
Public key (bytes)	992	1056	1568	1664
Signature entropy (bytes)	1473.14	1331.53	2342.38	2212.80
Public key + signature entropy	2465.14	2387.53	3910.38	3876.80
Signature entropy decrease	—	9.61%	—	5.53%
Observed key acceptance	0.26197	0.51390	0.48862	0.52287
Estimated key-generation attempts	3.82	1.95	2.05	1.91
Verification bound γ	11957.79	16090.13	24860.74	39642.01
Direct MSIS block size at γ	413	411	635	616
Classical MSIS cost (bits)	120.79	120.21	185.73	180.17
Quantum lattice MSIS cost (bits)	106.14	105.63	163.20	158.31
Best (A)MLWE block size	428	423	769	775
Classical (A)MLWE cost (bits)	125.18	123.72	225.14	226.67
Quantum lattice (A)MLWE cost (bits)	110.00	108.71	197.85	199.18

Table 3. Reduction radii and challenge-space limits for the revised sets.

	G+G-I	Set-I	G+G-II	Set-II
MSIS block size at 2γ	336	336	528	516
Classical cost at 2γ (bits)	98.27	98.27	154.43	150.92
Quantum lattice cost at 2γ	86.35	86.35	135.70	132.61
$\log_2 \lvert\mathcal{C}\rvert$	128	128	256	256
Generic quantum challenge-search exponent	64	64	128	128

lattice columns do not include generic challenge search, which has exponents 64 and 128. They are not full QROM security claims.

The original version also reported parameters evaluated with the two-step mode (TSM) methodology of Xia et al. [XWW+24]. Those numbers used the former parameter sets and are superseded here. The present numerical comparison is restricted to the explicitly reproduced primal, dual, and MSIS models. TSM, hybrid, and other distribution-sensitive attacks require a separate evaluation before claiming a comprehensive concrete security level.

4 Conclusion

This paper proposed an optimized G+G signature by incorporating asymmetric LWE and a three-block non-spherical Gaussian allocation. We provided an experimental largest singular value analysis of the secret key matrix, which enables a more compact response distribution. The revised parameter comparison uses the corrected G+G baseline and explicitly accounts for challenge cardinality, Gaussian widths, key-generation conditioning, and the stated concrete attack model. A question worth studying is: Is there a more specific analytical approach to the hardness of asymmetric LWE/SIS problems?

Acknowledgments

We thank the anonymous reviewers of PKC 2026 for their insightful comments and constructive suggestions. This work was supported by the National Natural Science Foundation of China (Grant No. 12525115) and the Natural Science Foundation of Hunan Province (Grant Nos. 2026JJ60006 and 2026JJ90267). We also thank Damien Stehlé for informing us of the corrections to G+G, and Jan Niklas Siemer for pointing out the norm and LWE-definition issues.

References

BG14. Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *Topics in Cryptology - CT-RSA 2014*, volume 8366 of *Lecture Notes in Computer Science*, pages 28–47. Springer, 2014.

CGM19. Yilei Chen, Nicholas Genise, and Pratyay Mukherjee. Approximate trapdoors for lattices and smaller hash-and-sign signatures. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019*, volume 11923 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2019.

DDLL13. Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013*, volume 8042 of *Lecture Notes in Computer Science*, pages 40–56. Springer, 2013.

DFPS22. Julien Devevey, Omar Fawzi, Alain Passelègue, and Damien Stehlé. On rejection sampling in lyubashevsky's signature scheme. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology - ASIACRYPT 2022*, volume 13794 of *Lecture Notes in Computer Science*, pages 34–64. Springer, 2022.

DFPS23. Julien Devevey, Pouria Fallahpour, Alain Passelègue, and Damien Stehlé. A detailed analysis of fiat-shamir with aborts. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023*, volume 14085 of *Lecture Notes in Computer Science*, pages 327–357. Springer, 2023.

DKL+18. Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.

DPS23. Julien Devevey, Alain Passelègue, and Damien Stehlé. G+G: A fiat-shamir lattice signature based on convolved gaussians. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology - ASIACRYPT 2023*, volume 14444 of *Lecture Notes in Computer Science*, pages 37–64. Springer, 2023.

DPS26a. Julien Devevey, Alain Passelègue, and Damien Stehlé. G+G: A Fiat–Shamir lattice signature based on convolved gaussians. Cryptology ePrint Archive, Report 2023/1477, 2026. Revision of November 18, 2026.

DPS26b. Julien Devevey, Alain Passelègue, and Damien Stehlé. G+G parameter estimation scripts. GitHub repository, jdevevey/GplusG, 2026. Commit f145c4f21e0f2746c98487bbca2d09250d692ba2, November 18, 2026.

ETWY22. Thomas Espitau, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. Shorter hash-and-sign lattice-based signatures. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology - CRYPTO 2022*, volume 13508 of *Lecture Notes in Computer Science*, pages 245–275. Springer, 2022.

Gär25. Joel Gärtner. Compact lattice signatures via iterative rejection sampling. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology - CRYPTO 2025*, volume 16000 of *Lecture Notes in Computer Science*, pages 514–547. Springer, 2025.

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.

HDD21. Yupu Hu, Siyue Dong, and Xingting Dong. Analysis on aigis-enc: Asymmetrical and symmetrical. *IET Inf. Secur.*, 15(2):147–155, 2021.

JHT22. Huiwen Jia, Yupu Hu, and Chunming Tang. Lattice-based hash-and-sign signatures using approximate trapdoor, revisited. *IET Inf. Secur.*, 16(1):41–50, 2022.

LP11. Richard Lindner and Chris Peikert. Better key sizes (and attacks) for lwe-based encryption. In *Topics in Cryptology – CT-RSA 2011*, pages 319–339. Springer, 2011.

Lyu09. Vadim Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 598–616. Springer, 2009.

Lyu12. Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012.

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718. Springer, 2012.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *ACM Symposium on Theory of Computing*, pages 84–93. ACM, 2005.

XWW+24. Wenwen Xia, Leizhang Wang, Geng Wang, Dawu Gu, and Baocang Wang. A refined hardness estimation of LWE in two-step mode. In Qiang Tang and Vanessa Teague, editors, *Public-Key Cryptography - PKC 2024*, volume 14603 of *Lecture Notes in Computer Science*, pages 3–35. Springer, 2024.

YJW23. Yang Yu, Huiwen Jia, and Xiaoyun Wang. Compact lattice gadget and its applications to hash-and-sign signatures. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023*, volume 14085 of *Lecture Notes in Computer Science*, pages 390–420. Springer, 2023.

ZYF+20. Jiang Zhang, Yu Yu, Shuqin Fan, Zhenfeng Zhang, and Kang Yang. Tweaking the asymmetry of asymmetric-key cryptography on lattices: Kems and signatures of smaller sizes. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *Public-Key Cryptography - PKC 2020*, volume 12111 of *Lecture Notes in Computer Science*, pages 37–65. Springer, 2020. Full version available at Cryptology ePrint Archive, Report 2019/510, https://eprint.iacr.org/2019/510.

A Summary of the Hardness Estimation for the Asymmetric Problems

This appendix specifies the attack models and numerical procedure used for the parameter comparison in Section 3.2. We follow the approach to asymmetric lattice problems of [ZYF+20], accounting for the secret distribution, module dimensions, and verification bound of Fig. 4. The analysis proceeds from the lattice instances induced by the scheme to their attack costs, and then to the choice of Gaussian widths and compression parameters.

A.1 Attack Instances and Cost Model

Write $r = k - m - 1$ and identify each element of $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$ with its coefficient vector. Before compression and rejection, the key-generation relation $\mathbf{b} - \mathbf{a} = \mathbf{A}_1\mathbf{s}_1 + \mathbf{s}_2 \bmod q$ gives an AMLWE instance with coefficient secret dimension $N_s = nr$ and at most $M_{\max} = nm$ samples. Forgery extraction yields a normal-form MSIS instance with nm equations and nk unknown coefficients. These coefficient representations retain the ring structure of the public matrices. We estimate attacks on the resulting lattices using BKZ-shape and isotropic-projection heuristics.

The numerical model comprises an augmented primal attack and a dual attack for AMLWE, together with a short-vector attack for MSIS. We use the BKZ-shape formulas implemented by the G+G reference estimator [DPS26b], fixed to commit `f145c4f21e0f2746c98487bbca2d09250d692ba2`. A BKZ block size b is assigned sieving

cost $2^{c_{\mathrm{svp}}b}$, where

$$c_{\mathrm{svp}} = \begin{cases} \log_2(3/2)/2, & \text{classical cost,} \\ 0.257, & \text{quantum lattice cost.} \end{cases}$$

All reported costs are base-two logarithms of the resulting work factors. The search minimizes over the implemented attack families. Hybrid attacks, TSM, and other distribution-sensitive methods are outside this model. The AMLWE inputs below describe the unconditioned coefficient distributions. The accepted-key decisional assumption used by the security reduction is specified separately in Appendix B.3.

A.2 Primal and Dual Attacks on AMLWE

The coefficients of $\mathbf{s}_1$ and $\mathbf{s}_2$ are independently uniform in $[-\zeta_1, \zeta_1]$ and $[-\zeta_2, \zeta_2]$, respectively. Their variances before public-key decomposition are

$$\tau_s^2 = \frac{\zeta_1(\zeta_1+1)}{3}, \qquad \tau_e^2 = \frac{\zeta_2(\zeta_2+1)}{3}.$$

Thus (τ_s^2, τ_e^2) equals $(2/3, 2)$ for Set-I, $(2/3, 4)$ for Set-II, and $(2/3, 2/3)$ for the symmetric G+G controls. The corresponding coefficient standard deviations are the inputs to the attack model. The response widths $\sigma_1, \sigma_2, \sigma_3$ instead determine the signing distribution and the verification bound.

To account for the asymmetric variances, we scale the error coordinates by $\xi > 0$. Let $\xi_{\mathrm{bal}} = \tau_s/\tau_e$. For each $1 \le M \le M_{\max}$, the two asymmetric sets use the scaling grid

$$\{\xi_{\mathrm{bal}} 2^{j/12-1} : 0 \le j \le 24\} \cup \{1\},$$

where $\xi_{\mathrm{bal}} = 1/\sqrt{3}$ and $1/\sqrt{6}$ for Set-I and Set-II. The symmetric controls use $\xi = 1$. For a fixed (M, ξ), write $q_{\mathrm{eff}} = q\xi$ for the determinant parameter of the normalized lattice. The search considers all legal block sizes beginning at $b = 50$, up to the dimension of the corresponding lattice.

Primal attack. The embedding contains the coefficient secret $\mathbf{s}_1$, the M selected error coefficients $\mathbf{e}$, and one embedding coordinate. After scaling, the target is $(\mathbf{s}_1, \xi\mathbf{e}, 1)$, and the lattice has dimension $N_s + M + 1$ and determinant q_{eff}^M. Under the isotropic-projection heuristic, the squared length of its projection onto the final b coordinates is modeled by

$$\ell_{\mathrm{p}}(b)^2 = b\,\frac{N_s\tau_s^2 + M\xi^2\tau_e^2 + 1}{N_s + M + 1}.$$

The embedding coordinate is included in both the numerator and the dimension. The success criterion compares $\ell_{\mathrm{p}}(b)$ with the Gram–Schmidt length predicted at the beginning of the final BKZ block. The smallest block size satisfying this criterion gives cost $c_{\mathrm{svp}}b$ for the selected (M, ξ).

Dual attack. The normalized dual lattice has dimension $N_s + M$ and determinant $q_{\mathrm{eff}}^{N_s}$. The same variance model gives

$$\tau_{\mathrm{eff}}^2 = \frac{N_s\tau_s^2 + M\xi^2\tau_e^2}{N_s + M}.$$

If L_b is the predicted short-vector length after BKZ-b, set $u = L_b\tau_{\mathrm{eff}}/q_{\mathrm{eff}}$. The Gaussian estimate for the distinguishing bias is $\exp(-2\pi^2u^2)$. Combining lattice reduction with the repetitions needed for constant distinguishing advantage gives logarithmic cost

$$c_{\mathrm{svp}}b + \max\left\{0, \frac{4\pi^2u^2}{\ln 2} - \frac{\log_2(4/3)}{2}b\right\}.$$

The subtracted term accounts for the number of vectors produced by a sieve in the reference model. Minimizing these primal and dual costs over (M, ξ, b) gives the AMLWE entries in Table 2.

A.3 MSIS Bounds and Parameter Selection

The verifier in Fig. 4 checks one Euclidean norm bound on the reconstructed response. Accordingly, the short-vector estimator uses the global MSIS radius. The identity block of the normal-form matrix makes its coefficient map onto $\mathbb{Z}_q^{nm}$ surjective, so its kernel lattice has dimension nk and determinant q^{nm}. For each radius $\beta \in \{\gamma, 2\gamma\}$, we find the smallest $50 \le b \le nk$ for which the predicted first-vector length is at most β, and report cost $c_{\mathrm{svp}} b$. Both radii are smaller than q for the parameter sets considered here. The radius γ supplies the direct short-response comparison in Table 2. The radius 2γ is the extraction bound in Theorem 1 and is reported separately in Table 3, together with the challenge-space limits.

The choice of response widths links these attack costs to signature size. For each candidate, we impose the weighted spectral condition of Lemma 2 on the actual secret $\mathbf{s} = (1, \mathbf{s}_1^\top, (\mathbf{s}_2 - \mathbf{b}_0)^\top)^\top$. Writing $\mathbf{s}_2' = \mathbf{s}_2 - \mathbf{b}_0$, negacyclic Fourier evaluation gives

$$\|\mathbf{D}^{-1}\,\mathrm{rot}(\eta\mathbf{s})\|_{\mathrm{op}}^2 = \nu^2 \max_{\omega^n=-1} \left(\frac{1}{\sigma_1^2} + \frac{\sum_{i=1}^{r} |s_{1,i}(\omega)|^2}{\sigma_2^2} + \frac{\sum_{j=1}^{m} |s_{2,j}'(\omega)|^2}{\sigma_3^2} \right),$$

with $\nu^2 = 2$ for $\eta = 1 + x^{n/2}$ and $\nu^2 = 4$ for $\eta = 2$. The uniform public vector $\mathbf{a}$ makes $\mathbf{b}$ independent of the uncorrected short secret. Acceptance experiments therefore sample its uniform marginal and apply the same Decomp routine as KeyGen before evaluating the displayed expression. The symmetric controls retain their scalar spectral thresholds 72.98 and 123.93.

Width selection uses 10000 preliminary samples per family at $q_{\mathrm{train}} = 65537$. Each selected parameter set is then validated on 100000 independent trials at its actual modulus. The recorded seeds, acceptance counts, and confidence intervals specify these experiments. Validation uses floating-point Fourier arithmetic. An implementation of KeyGen must enforce the spectral inequality with certified bounds or an outward numerical error margin. For the selected widths, Lemma 4 determines the verification bound, including the compression allowance in Equation (4).

We evaluate signature size using the ideal entropy expression of Appendix C. The hint contribution is computed from Equation (10) for every phase modulo q, including the exceptional quantization bin of size $a+1$, where $a = \alpha/2$. The one-dimensional Gaussian sums extend to 14 corresponding continuous standard deviations, leaving omitted mass below 2^{-140} per coefficient for the listed widths. The same tail bound and cyclic-hint accounting are applied to our sets and the G+G controls, giving a common entropy-based size metric.

The source archive records the selected parameters, the finite search, and the numerical outputs in the `reproducibility` directory. The command

```
python3 reproduce.py --all --samples 100000
```

reproduces the attack costs, size estimates, and acceptance experiments. The options `--deterministic`, `--lwe`, and `--acceptance` select the corresponding stages, while `check_compression.py` verifies quantizer boundaries and the recovery identity. The archived reference output also records the G+G estimator under its original parameter and size conventions.

B Supplementary Security Analysis

We prove correctness and unforgeability of the signature scheme in Fig. 4. The proof first establishes a public simulation of the three-block response distribution and a bound on its norm. We then specify the public-key distributions induced by spectral rejection and combine key indistinguishability with the commitment-recovery properties of Appendix C to obtain the ROM reduction.

B.1 Three-Block Gaussian Convolution

We use the notation of Section 3. Let $\mathcal{R} = \mathbb{Z}[x]/(x^n+1)$, where $n > 1$ is a power of two, and put $r = k - m - 1$. The accepted secret is $\mathbf{s} = (1, \mathbf{s}_1^\top, (\mathbf{s}_2')^\top)^\top \in \mathcal{R}^k$, where $\mathbf{s}_2' = \mathbf{s}_2 - \mathbf{b}_0$ includes the public-key rounding correction. All norms and matrix inequalities below refer to coefficient embeddings. For $\eta\eta^* = 2$, define

$$\mathbf{D} = \mathrm{diag}(\sigma_1 \mathbf{I}_n, \sigma_2 \mathbf{I}_{rn}, \sigma_3 \mathbf{I}_{mn}), \qquad \mathbf{W} = \mathbf{D}^2, \qquad \mathbf{T} = \mathrm{rot}(\eta\mathbf{s}) \in \mathbb{Z}^{kn \times n}.$$

The key generator accepts exactly when the perturbation shape $\mathbf{V} = \Sigma(\mathbf{s})$ satisfies

$$\mathbf{V} = \mathbf{W} - s^2\mathbf{T}\mathbf{T}^\top, \qquad \mathbf{W} \succeq 2s^2\mathbf{T}\mathbf{T}^\top. \tag{1}$$

Throughout, s and the σ_i are Gaussian widths under the convention $\rho_\sigma(x) = \exp(-\pi x^2/\sigma^2)$. Thus $\mathbf{V}$ is the shape matrix of the perturbation, and its continuous Gaussian counterpart has covariance $\mathbf{V}/(2\pi)$.

The following identity relates the acceptance test to the rescaled secret matrix used in Section 3.1. It also provides the spectral quantity whose distribution is evaluated when selecting the Gaussian widths.

Lemma 2 (Weighted spectral condition). *Let $\mathbf{U}_1 = \mathrm{rot}(\mathbf{s}_1)$, $\mathbf{U}_2 = \mathrm{rot}(\mathbf{s}_2')$, and $t = \sigma_2/\sigma_3$, and set $\mathbf{U} = (\mathbf{U}_1^\top, t\mathbf{U}_2^\top)^\top \in \mathbb{R}^{(k-1)n\times n}$. For $\eta = 1 + x^{n/2}$ set $\nu = \sqrt{2}$, and for $\eta = 2$ set $\nu = 2$. Then*

$$\|\mathbf{D}^{-1}\mathbf{T}\|_{\mathrm{op}}^2 = \nu^2\big(\sigma_1^{-2} + \sigma_2^{-2}\|\mathbf{U}\|_{\mathrm{op}}^2\big). \tag{2}$$

In particular, (1) holds if and only if the right-hand side of (2) is at most $1/(2s^2)$.

Proof. For $\mathbf{G} = \mathrm{rot}(\eta)$, both choices of η satisfy $\mathbf{G}^\top\mathbf{G} = \nu^2\mathbf{I}_n$. Since $\mathbf{T} = (\mathbf{I}_n, \mathbf{U}_1^\top, \mathbf{U}_2^\top)^\top\mathbf{G}$, we have

$$\mathbf{T}^\top\mathbf{D}^{-2}\mathbf{T} = \mathbf{G}^\top\big(\sigma_1^{-2}\mathbf{I}_n + \sigma_2^{-2}\mathbf{U}^\top\mathbf{U}\big)\mathbf{G}.$$

The matrix $\mathbf{G}/\nu$ is orthogonal, so taking the largest eigenvalue gives (2). Congruence by $\mathbf{D}^{-1}$ proves the equivalence with (1).

For every accepted key, (1) implies $\mathbf{V} \succeq \mathbf{W}/2 \succ 0$. To sample the perturbation efficiently, we also require

$$\frac{\min_i \sigma_i}{\sqrt{2}} \geq \sqrt{\frac{\ln(2kn+4)}{\pi}}.$$

This is sufficient for the nonspherical Gaussian sampler of Lemma 4 in [DPS26a], applied to the standard basis of $\mathbb{Z}^{kn}$. All parameter sets in Table 1 satisfy this condition. With the perturbation distribution fixed, the next lemma establishes the response distribution for each accepted key and each fixed challenge.

Lemma 3 (Convolution with three widths). *Let $\epsilon \in (0, \min(1/2, 0.086n))$ and suppose that*

$$s \geq \sqrt{\frac{2\ln(n-1+2n/\epsilon)}{\pi}}$$

and (1) holds. Fix a representative c of an element of $\mathcal{R}/\eta\mathcal{R}$, and independently sample

$$\mathbf{y} \leftarrow \mathcal{D}_{\mathbb{Z}^{kn},\sqrt{\mathbf{V}}}, \qquad u \leftarrow \mathcal{D}_{\mathbb{Z}^n,s,-\eta^* c/2}.$$

Set $\mathbf{z} = \mathbf{y} + (\eta u + c)\mathbf{s}$ and $R_\epsilon = (1+\epsilon)/(1-\epsilon)$. If P_c denotes the distribution of $\mathbf{z}$ and $P_{\mathrm{id}} = \mathcal{D}_{\mathbb{Z}^{kn},\mathbf{D}}$, then

$$R_\epsilon^{-1} \leq \frac{P_c(\mathbf{z})}{P_{\mathrm{id}}(\mathbf{z})} \leq R_\epsilon \quad \text{for every } \mathbf{z} \in \mathbb{Z}^{kn}. \tag{3}$$

Consequently, the infinite-order Rényi divergence in either direction, in the multiplicative convention, is at most R_ϵ, and $\Delta(P_c, P_{\mathrm{id}}) \leq 2\epsilon/(1-\epsilon)$.

Proof. Let $\mathbf{d} \in \mathbb{R}^n$ and $\mathbf{u} \in \mathbb{Z}^n$ be the coefficient vectors of $\eta^* c/2$ and u. The identity $\eta\eta^* = 2$ gives $\mathbf{T}\mathbf{d} = c\mathbf{s}$ in coefficient embedding, and hence $\mathbf{z} = \mathbf{y} + \mathbf{T}(\mathbf{u} + \mathbf{d})$. For $\mathbf{v} \in \mathbb{Z}^n + \mathbf{d}$, completing the square yields

$$\begin{aligned}&(\mathbf{z} - \mathbf{T}\mathbf{v})^\top\mathbf{V}^{-1}(\mathbf{z} - \mathbf{T}\mathbf{v}) + s^{-2}\|\mathbf{v}\|^2 \\ &\quad = \mathbf{z}^\top\mathbf{W}^{-1}\mathbf{z} + (\mathbf{v} - s^2\mathbf{T}^\top\mathbf{W}^{-1}\mathbf{z})^\top\mathbf{K}^{-1}(\mathbf{v} - s^2\mathbf{T}^\top\mathbf{W}^{-1}\mathbf{z}),\end{aligned}$$

where

$$\mathbf{K} = (s^{-2}\mathbf{I}_n + \mathbf{T}^\top\mathbf{V}^{-1}\mathbf{T})^{-1} = s^2\mathbf{I}_n - s^4\mathbf{T}^\top\mathbf{W}^{-1}\mathbf{T} \succeq (s^2/2)\mathbf{I}_n.$$

The lower bound on s ensures $s/\sqrt{2} \geq \eta_\epsilon(\mathbb{Z}^n)$. Summing the second Gaussian factor over $\mathbb{Z}^n + \mathbf{d}$ therefore gives a value in $\sqrt{\det\mathbf{K}}\,[1-\epsilon, 1+\epsilon]$, uniformly in $\mathbf{z}$. Thus $P_c(\mathbf{z})$ is proportional to $\rho_{\mathbf{D}}(\mathbf{z})$ times a factor in this interval. Normalization proves (3), from which the divergence and statistical-distance bounds follow.

The lemma applies to the bounded-uniform secret distributions in the construction because its hypothesis concerns the accepted matrix $\mathbf{T}$. It already includes the correction $\mathbf{s}_2 - \mathbf{b}_0$. For the integer-matrix construction in Section 3.1, substitute k for kn, ℓ for n, $2\mathbf{S}$ for $\mathbf{T}$, and $\mathbf{c}/2$ for $\mathbf{d}$. The sufficient conditions become $s/\sqrt{2} \geq \eta_\epsilon(\mathbb{Z}^\ell)$ and $\mathbf{D}^2 \succeq 8s^2\mathbf{S}\mathbf{S}^\top$, as stated in the main text.

B.2 Public Simulation and Correctness

The public-key relation $\mathbf{A}\mathbf{s} = q\mathbf{j} \bmod 2q$ and the identity $\eta\mathbf{j} = (2, 0, \ldots, 0)^\top$ imply that every honest response satisfies

$$\mathbf{A}\mathbf{z} - qc\mathbf{j} = \mathbf{A}\mathbf{y} \pmod{2q}.$$

Accordingly, for a fixed challenge c, a public simulator samples $\mathbf{z} \leftarrow P_{\mathrm{id}}$, computes $\mathbf{w} = \mathbf{A}\mathbf{z} - qc\mathbf{j} \bmod 2q$, and applies PackCommit and Compress from Appendix C. These deterministic maps preserve the pointwise bound of Lemma 3, giving a simulation of the complete transcript without the secret key. To set the verification radius, we bound the norm of the same response distribution and then account for the reconstruction error.

Lemma 4 (Response tail bound). *Assume the hypotheses of Lemma 3. Let $(d_1, d_2, d_3) = (n, rn, mn)$ and $v_i = \sigma_i^2/(2\pi)$, and define*

$$\tau_1 = \sum_{i=1}^{3} d_i v_i, \qquad \tau_2 = \sum_{i=1}^{3} d_i v_i^2, \qquad v_{\max} = \max_i v_i.$$

For $t > 0$, put

$$B(t) = \sqrt{\tau_1 + 2\sqrt{\tau_2 t} + 2v_{\max} t}.$$

A response in an honest-verifier transcript satisfies $\Pr[\|\mathbf{z}\| > B(t)] \leq R_\epsilon e^{-t}$.

Proof. For $X \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}$ and $0 \leq \lambda < \pi/\sigma^2$, Poisson summation gives

$$\mathbb{E}[e^{\lambda X^2}] = \frac{\rho_{\sigma/\sqrt{1-\lambda\sigma^2/\pi}}(\mathbb{Z})}{\rho_\sigma(\mathbb{Z})} \leq (1 - \lambda\sigma^2/\pi)^{-1/2}.$$

Indeed, $\rho_a(\mathbb{Z}) = a\rho_{1/a}(\mathbb{Z})$ and $\rho_{1/a}(\mathbb{Z})$ decreases with a. Multiplying over the independent coordinates of P_{id} and applying the weighted chi-square Chernoff bound yields

$$\Pr_{P_{\mathrm{id}}}[\|\mathbf{z}\|^2 > \tau_1 + 2\sqrt{\tau_2 t} + 2v_{\max} t] \leq e^{-t}.$$

Equation (3) transfers the bound to the real response.

By Lemma 5, compression increases the response norm by at most $\sqrt{mn}(\alpha/4 + 1)$. For a desired error bound $\epsilon_c \in (0, 1)$, it suffices to take

$$\gamma = B\big(\ln(R_\epsilon/\epsilon_c)\big) + \sqrt{mn}(\alpha/4 + 1). \tag{4}$$

This gives honest-verifier correctness error at most ϵ_c. For Fiat–Shamir signing, fix an accepted key and a preceding oracle history, and let Fresh denote that the new signing input is absent from the H-table. On this event the challenge is sampled uniformly, so restricting the honest-verifier experiment to fresh inputs gives $\Pr[\mathsf{Reject} \wedge \mathsf{Fresh}] \leq \epsilon_c$. If the history contains L distinct hash inputs, Lemma 6 further gives

$$\Pr[\mathsf{Reject}] \leq \epsilon_c + L2^{-h},$$

where h is the commitment min-entropy bound in that lemma. In particular, correctness error is at most ϵ_c in the usual experiment with an initially empty H-table. These probabilities average over the random oracle and the randomness of the single signing attempt specified in Fig. 4.

We use exact Gaussian samples in the analysis. If the total sampling error per signing attempt is at most δ_{samp}, uniformly over accepted keys and prior oracle histories, coupling adds at most δ_{samp} to the correctness error and $Q_S\delta_{\mathrm{samp}}$ to the security bound for Q_S signing queries.

B.3 Key Indistinguishability under Spectral Rejection

The key-indistinguishability experiment follows the rejection-aware formulation of G+G [DPS26a, Theorem 6] and incorporates the public-key decomposition used in Fig. 4. In an unseeded generation attempt, sample independent uniform $\mathbf{a} \in \mathcal{R}_q^m$ and $\mathbf{A}_1 \in \mathcal{R}_q^{m\times r}$, independent short vectors $\mathbf{s}_1 \leftarrow \chi_{\zeta_1}^r$ and $\mathbf{s}_2 \leftarrow \chi_{\zeta_2}^m$, and set

$$\mathbf{b} = \mathbf{a} + \mathbf{A}_1\mathbf{s}_1 + \mathbf{s}_2 \bmod q, \qquad (\mathbf{b}_0, \mathbf{b}_1) = \mathsf{Decomp}(\mathbf{b}).$$

Let Acc be the event (1) for $\mathbf{s} = (1, \mathbf{s}_1^\top, (\mathbf{s}_2 - \mathbf{b}_0)^\top)^\top$. Since this event depends on $\mathbf{b}_0$, it also affects the marginal law of the encoded public part $\mathbf{b}_1/2$. We retain that marginal in the comparison experiment by means of the following public sampler.

The sampler $\mathcal{M}_{\mathsf{Acc}}$ independently draws $\mathbf{b} \leftarrow \mathcal{U}(\mathcal{R}_q^m)$, $\mathbf{s}_1 \leftarrow \chi_{\zeta_1}^r$, and $\mathbf{s}_2 \leftarrow \chi_{\zeta_2}^m$, applies Decomp, and repeats until the same spectral test accepts. It outputs $\mathbf{b}_1/2$ and discards the auxiliary short vectors. The uniform offset $\mathbf{a}$ makes $\mathbf{b}$ independent of $(\mathbf{A}_1, \mathbf{s}_1, \mathbf{s}_2)$ before rejection. Consequently, $\mathcal{M}_{\mathsf{Acc}}$ samples exactly the accepted public-part marginal under fresh uniform expansions, using only public parameters and $1/p$ trials in expectation, where $p = \Pr[\mathsf{Acc}]$.

The unseeded assumption $\mathsf{AMLWE}^{\mathsf{Acc}}$ is computational indistinguishability of

$$(\mathbf{a}, \mathbf{A}_1, \mathbf{b}_1/2) \mid \mathsf{Acc} \quad \text{and} \quad (\mathbf{a}_0, \mathbf{A}_0, \mathbf{v}),$$

where $\mathbf{a}_0 \leftarrow \mathcal{U}(\mathcal{R}_q^m)$, $\mathbf{A}_0 \leftarrow \mathcal{U}(\mathcal{R}_q^{m\times r})$, and $\mathbf{v} \leftarrow \mathcal{M}_{\mathsf{Acc}}$ are independent. The unrounded vector $\mathbf{b}$ is internal to key generation. This formulation isolates the pseudorandomness of the matrix relation while preserving the public marginal produced by rounding and rejection.

For the seeded public keys in Fig. 4, let

$$G : \{0,1\}^\kappa \longrightarrow \mathcal{R}_q^{m\times(r+1)}, \qquad \kappa = 256,$$

be a random oracle independent of H. We use the corresponding oracle-aware assumption, defined by the following two experiments:

- In the real experiment, run KeyGen with oracle G, including every rejection attempt, and output its verification key $(\tau, \mathbf{b}_1/2)$.
- In the comparison experiment, sample a uniform $\tau \in \{0,1\}^\kappa$ and $\mathbf{v} \leftarrow \mathcal{M}_{\mathsf{Acc}}$ independently of G, and output $(\tau, \mathbf{v})$.

The distinguisher has access to G in both experiments, whereas $\mathcal{M}_{\mathsf{Acc}}$ uses fresh internal randomness and makes no queries to G. We require KeyGen and the auxiliary sampler to run in expected polynomial time. This seeded experiment includes the correlation between the accepted seed and the expansion oracle and is the key assumption used for the concrete representation.

Both comparison experiments admit an exact embedding of a uniform normal-form matrix. Given independent uniform $\mathbf{u} \in \mathcal{R}_q^m$ and $\mathbf{A}_1 \in \mathcal{R}_q^{m\times r}$, sample $\mathbf{v} \leftarrow \mathcal{M}_{\mathsf{Acc}}$, set $\mathbf{b}_1 = 2\mathbf{v}$, and choose $\mathbf{a} = \mathbf{u} + \mathbf{b}_1 \bmod q$. Then $(\mathbf{a}, \mathbf{A}_1)$ is uniform and independent of $\mathbf{v}$, while the effective matrix satisfies

$$2^{-1}\mathbf{A} \bmod q = (\mathbf{u} \mid \mathbf{A}_1 \mid \mathbf{I}_m).$$

For a seeded key, choose a uniform seed τ and initialize $G(\tau) = (\mathbf{u} + \mathbf{b}_1 \mid \mathbf{A}_1)$ before starting the adversary. All other values of G are sampled lazily. The programmed value is a uniform expansion even conditioned on $\mathbf{v}$, which gives the exact seeded comparison experiment. This initialization takes place at setup, before the adversary receives its public key.

The distinction between fresh expansions and repeated seeds can be quantified directly. Couple actual key generation to a process that uses an independent uniform expansion on every attempt until a seed repeats. In the fresh process the number N of attempts is geometric with parameter p, independently of the seed labels. Hence

$$\Pr[\text{a seed repeats before acceptance}] \le 2^{-\kappa}\mathbb{E}\binom{N}{2} = \frac{1-p}{p^2}\,2^{-\kappa}.$$

This bounds the distance between the encoded public-part marginals of the two generation processes. The real experiment above uses the actual key generator, so its distinguishing advantage already includes seed reuse.

These distribution-specific decisional assumptions are separate from the unconditioned AMLWE estimates in Appendix A. For the key-recovery sensitivity calculation in Section 3.2, conditioning yields

$$\Pr[\text{key recovery succeeds} \mid \mathsf{Acc}] \le \frac{\Pr[\text{key recovery succeeds}]}{p} \tag{5}$$

for any fixed recovery algorithm and $p > 0$. The inequality concerns recovery probabilities. The key hybrid below uses the decisional experiment just defined, including the acceptance event and its public marginal.

B.4 Unforgeability in the Random-Oracle Model

We now combine the public simulation with key indistinguishability and special soundness. The extraction problem is normal-form $\mathsf{MSIS}_{n,m,k,q,\beta}$: given $\mathbf{M} = (\mathbf{M}_0 \mid \mathbf{I}_m)$ with $\mathbf{M}_0 \leftarrow \mathcal{U}(\mathcal{R}_q^{m\times(k-m)})$, find a nonzero $\mathbf{x} \in \mathcal{R}^k$ such that $\mathbf{M}\mathbf{x} = 0 \bmod q$ and the Euclidean norm of its coefficient vector is at most β. This matrix has the identity block used by Recover, while the Euclidean bound matches the condition enforced by Verify.

Theorem 1 (ROM security with compressed commitments). *Assume the conditions of Lemma 3, Appendix C, and the sampling condition above. Suppose that key generation and $\mathcal{M}_{\mathsf{Acc}}$ run in expected polynomial time, and that $q > 2\gamma$. Let h be a lower bound on the min-entropy of the compressed commitment for every accepted key. For every classical EUF-CMA adversary $\mathcal{A}$ making at most Q_S signing queries and $Q_H \ge 1$ hash queries, there exist adversaries $\mathcal{B}_L$ and $\mathcal{B}_S$ against $\mathsf{AMLWE}^{\mathsf{Acc}}$ and normal-form $\mathsf{MSIS}_{n,m,k,q,2\gamma}$, respectively, such that*

$$\begin{aligned}\mathrm{Adv}^{\mathsf{EUF-CMA}}(\mathcal{A}) &\le Q_S(Q_H+Q_S)2^{-h} \\ &\quad + R_\epsilon^{Q_S}\left[\mathrm{Adv}^{\mathsf{AMLWE}^{\mathsf{Acc}}}(\mathcal{B}_L) + Q_H\left(\sqrt{\mathrm{Adv}^{\mathsf{MSIS}}(\mathcal{B}_S)} + \frac{2}{|\mathcal{C}|}\right)\right].\end{aligned} \tag{6}$$

Here $\mathcal{C} = \mathcal{R}/\eta\mathcal{R}$, and the hash input contains the complete encoding of the compressed commitment and the message. The key term uses the unseeded experiment for explicitly represented matrices and the seeded experiment of Appendix B.3 for the representation in Fig. 4. In the latter case G and H are independent random oracles. The adversary starts after key generation.

Proof. We first simulate signing, then replace the public key, and finally extract an MSIS solution from two accepting responses.

Signing simulation. On a message μ, the simulator samples a uniform challenge c and $\mathbf{z} \leftarrow P_{\mathrm{id}}$, computes $\mathbf{w} = \mathbf{A}\mathbf{z} - qc\mathbf{j} \bmod 2q$, and obtains the compressed commitment and signature through the public compression routines. It programs $H(\overline{\mathbf{w}}, \mu) = c$, aborting if that input is already in the oracle table.

Fix an accepted secret key and the entire history before a real signing query. There are at most $Q_H + Q_S$ fixed hash inputs, and the next commitment has min-entropy at least h. The probability that any of the Q_S real signing attempts hits such an input is therefore at most $Q_S(Q_H+Q_S)2^{-h}$. For a fixed table $\mathcal{T}$, let $F_{\mathcal{T}}(c,\mathbf{z})$ indicate that $(\mathsf{PackCommit}(\mathbf{A}\mathbf{z} - qc\mathbf{j} \bmod 2q), \mu)$ is absent from that table. The fresh parts of the real and simulated transcript laws are respectively

$$\frac{1}{|\mathcal{C}|}P_c(\mathbf{z})F_{\mathcal{T}}(c,\mathbf{z}) \quad \text{and} \quad \frac{1}{|\mathcal{C}|}P_{\mathrm{id}}(\mathbf{z})F_{\mathcal{T}}(c,\mathbf{z}).$$

Lemma 3 bounds the first by R_ϵ times the second. These are subprobability distributions, so no normalization by the probability of freshness is introduced. Multiplying over adaptive signing queries shows that real forgery probability is at most the collision term plus $R_\epsilon^{Q_S}$ times the forgery probability with simulated signing.

Public-key hybrid and embedding. The simulated signer uses only the public key. Replacing the accepted key by its comparison key therefore changes the forgery probability by at most $\mathrm{Adv}^{\mathsf{AMLWE}^{\mathsf{Acc}}}(\mathcal{B}_L)$. Given an MSIS instance $\mathbf{M} = (\mathbf{u} \mid \mathbf{A}_1 \mid \mathbf{I}_m)$, use the embedding from Appendix B.3 to generate the comparison public key. It gives $2^{-1}\mathbf{A} \bmod q = \mathbf{M}$ and the prescribed public-key distribution, including oracle access in the seeded experiment.

Forking and extraction. Let δ be the forgery probability in the comparison experiment. A successful forgery whose hash input was never queried has probability at most $1/|\mathcal{C}|$. For the remaining executions, guess the index of the first adversarial query to the eventual forgery input. At that query, regard the compressed commitment as an interactive prover commitment, answer with the verifier's challenge, and continue the adversary with the locally simulated signing oracle. Duplicate queries retain their previous answers. Because the forged message was never signed, this input was not programmed by the signing oracle. The interactive prover succeeds with probability $p_{\mathrm{f}} \geq (\delta - 1/|\mathcal{C}|)/Q_H$.

Reset at this commitment and give two independent uniform challenges with independent continuation randomness. Averaging the squared conditional acceptance probability gives joint acceptance probability at least p_{f}^2. Joint acceptance with coinciding challenges has probability at most $p_{\mathrm{f}}/|\mathcal{C}|$. Consequently, the probability of obtaining two accepting transcripts with distinct challenges is at least $p_{\mathrm{f}}(p_{\mathrm{f}} - 1/|\mathcal{C}|)$.

For such transcripts, Lemma 5 gives full responses $\widehat{\mathbf{z}}^{(0)}, \widehat{\mathbf{z}}^{(1)}$ with norms at most γ satisfying (8) with the same reconstructed commitment modulo $2q$. Their difference $\mathbf{x}$ satisfies

$$\mathbf{M}\mathbf{x} = 0 \pmod{q}, \qquad \|\mathbf{x}\| \leq 2\gamma.$$

If $\mathbf{x} = 0$, subtracting the two recovery identities and reducing modulo 2 gives $\eta^*(c_0 - c_1) = 0 \bmod 2\mathcal{R}$. Since $2 = \eta\eta^*$ and $\mathcal{R}$ is an integral domain, cancellation implies $c_0 - c_1 \in \eta\mathcal{R}$, contradicting the distinct challenge classes. Moreover, $q > 2\gamma$ excludes a nonzero vector all of whose coefficients are divisible by q. The extracted vector is therefore a valid solution to the embedded MSIS instance.

It follows that $\mathrm{Adv}^{\mathsf{MSIS}}(\mathcal{B}_S) \geq p_{\mathrm{f}}(p_{\mathrm{f}} - 1/|\mathcal{C}|)$. Solving this inequality and using $Q_H \geq 1$ gives

$$\delta \leq Q_H \left(\sqrt{\mathrm{Adv}^{\mathsf{MSIS}}(\mathcal{B}_S)} + \frac{2}{|\mathcal{C}|} \right).$$

Combining the simulation and public-key hybrid bounds proves (6).

The key distinguisher runs the adversary once. The MSIS reduction uses at most two executions with one reset and one auxiliary public-part sample, which takes an expected $1/p$ trials. The remaining work consists of Gaussian sampling, oracle simulation, and polynomial-time matrix and compression operations.

The reduction has expected running time. Stopping $\mathcal{M}_{\mathsf{Acc}}$ after T fresh trials incurs abort probability $(1-p)^T$. For the truncated reduction, replace the MSIS advantage inside the square root in (6) by its truncated advantage plus $(1-p)^T$. The theorem uses the usual post-setup EUF-CMA experiment. If Q_G^{pre} expansion-oracle queries are allowed before setup, the MSIS embedding additionally fails with probability at most $Q_G^{\mathrm{pre}} 2^{-\kappa}$, which is included at that hybrid.

The reduction identifies the normal-form MSIS instance at radius 2γ used in Table 3. Together with its explicit query losses, it gives the classical ROM security statement for the compressed scheme in Fig. 4.

C Compression and Commitment Recovery

We specify the encoding and recovery procedures used by Fig. 4. Public-key compression removes one bit from each coefficient of the public vector. Response compression omits the last m polynomials and replaces them with a hint from which the verifier recovers the commitment and a nearby response. We first give these procedures, then establish the recovery and entropy properties used in Theorem 1 and Table 2. Throughout, $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$ and $r = k - m - 1$, as in Section 3. All coefficient encodings are canonical.

C.1 Public-Key Encoding

For $q \equiv 1 \pmod 4$ and $b \in \{0, \ldots, q-1\}$, define

$$b_0 = \begin{cases} 0, & b \equiv 0 \pmod 2, \\ 1, & b \equiv 1 \pmod 4, \\ -1, & b \equiv 3 \pmod 4, \end{cases} \qquad b_1 = b - b_0.$$

The procedure Decomp applies this map coefficientwise and returns $(\mathbf{b}_0, \mathbf{b}_1)$. Each coefficient of $\mathbf{b}_1$ is an even integer in $[0, q-1]$, so the verification key stores $\mathbf{b}_1/2$ together with the public seed τ. A coefficient of $\mathbf{b}_1/2$ requires $\lceil \log_2((q+1)/2) \rceil$ bits. Including the 256-bit seed, the public-key encoding therefore occupies

$$256 + mn\lceil \log_2((q+1)/2) \rceil$$

bits for the parameter sets in Table 1.

The low-bit correction is incorporated into the secret before the spectral test. Indeed, for $\mathbf{b} = \mathbf{a} + \mathbf{A}_1\mathbf{s}_1 + \mathbf{s}_2 \bmod q$, put

$$\mathbf{s} = (1, \mathbf{s}_1^\top, (\mathbf{s}_2 - \mathbf{b}_0)^\top)^\top, \qquad \mathbf{A} = (2(\mathbf{a} - \mathbf{b}_1) + q\mathbf{j} \mid 2\mathbf{A}_1 \mid 2\mathbf{I}_m).$$

Then $\mathbf{A}\mathbf{s} = q\mathbf{j} \bmod 2q$, which is the relation required by the signing algorithm. Thus KeyGen applies its weighted spectral test to $\mathbf{s}_2 - \mathbf{b}_0$. The verifier expands τ to obtain $(\mathbf{a}, \mathbf{A}_1)$ and reconstructs $\mathbf{A}$ from $\mathbf{b}_1/2$.

C.2 Commitment and Response Compression

Let $a = \alpha/2$ be a positive even integer and write $q = Ka + 1$ with $K \geq 1$. Define the cyclic quantizer

$$Q(v) = \left\lfloor \frac{v + a/2}{a} \right\rfloor \bmod K, \qquad v \in \{0, \ldots, q-1\},$$

with output in $\{0, \ldots, K-1\}$. We apply Q coefficientwise and write cent_q for the centered representative in $[-(q-1)/2, (q-1)/2]$. The quantizer has the two properties

$$|\mathrm{cent}_q(v - aQ(v))| \leq a/2 + 1, \qquad \max_j |Q^{-1}(j)| = a + 1. \tag{7}$$

To see this, for $j \neq 0$ the preimage is $[ja - a/2, ja + a/2 - 1] \cap \mathbb{Z}$, of size a. The zero bin consists of $[0, a/2 - 1] \cap \mathbb{Z}$ and $[Ka - a/2, Ka] \cap \mathbb{Z}$, of total size $a + 1$. This bin also determines the additional unit in the reconstruction bound.

Write $\mathbf{A} = (\mathbf{A}_0 \mid 2\mathbf{I}_m)$ and split the response as $\mathbf{z} = (\mathbf{z}_0^\top, \mathbf{z}_3^\top)^\top$, where $\mathbf{z}_0 \in \mathcal{R}^{k-m}$ and $\mathbf{z}_3 \in \mathcal{R}^m$. The identity block allows the verifier to reconstruct the contribution of $\mathbf{z}_3$ using the following three procedures.

PackCommit. Given $\mathbf{w} \in \mathcal{R}_{2q}^m$, compute

$$\mathbf{p} = \mathbf{w} \bmod 2, \qquad \mathbf{v} = (\mathbf{w} - \mathbf{p})/2 \bmod q, \qquad \overline{\mathbf{w}} = (\mathbf{p}, Q(\mathbf{v})).$$

The output $\overline{\mathbf{w}}$ is the commitment hashed by the signer. In Fig. 4, $\mathbf{w} = \mathbf{A}\mathbf{y} \bmod 2q$ and $c = H(\overline{\mathbf{w}}, \mu)$. The response generated thereafter satisfies $\mathbf{A}\mathbf{z} - qc\mathbf{j} = \mathbf{w} \bmod 2q$.

Compress. Given $(\mathbf{A}, \mathbf{z}, c, \overline{\mathbf{w}})$, with $\overline{\mathbf{w}} = (\mathbf{p}, Q(\mathbf{v}))$, compute

$$\mathbf{w}' = \mathbf{A}_0\mathbf{z}_0 - qc\mathbf{j} \bmod 2q, \qquad \mathbf{v}' = (\mathbf{w}' - \mathbf{p})/2 \bmod q, \qquad \mathbf{h} = Q(\mathbf{v}) - Q(\mathbf{v}') \bmod K.$$

Return the compressed signature $(\mathbf{z}_0, c, \mathbf{h})$. For a valid input, $\mathbf{w} - \mathbf{w}' = 2\mathbf{z}_3 \bmod 2q$, so $\mathbf{w}'$ and $\mathbf{w}$ have the same parity. All divisions by 2 in these procedures are integer divisions after subtracting that parity.

Recover. Given $(\mathbf{A}, \mathbf{z}_0, c, \mathbf{h})$, compute

$$\mathbf{w}' = \mathbf{A}_0\mathbf{z}_0 - qc\mathbf{j} \bmod 2q, \qquad \mathbf{p} = \mathbf{w}' \bmod 2, \qquad \mathbf{v}' = (\mathbf{w}' - \mathbf{p})/2 \bmod q,$$

and set

$$\mathbf{t} = Q(\mathbf{v}') + \mathbf{h} \bmod K, \qquad \widehat{\mathbf{z}}_3 = \mathrm{cent}_q(a\mathbf{t} - \mathbf{v}').$$

Return $((\mathbf{p}, \mathbf{t}), \widehat{\mathbf{z}})$, where $\widehat{\mathbf{z}} = (\mathbf{z}_0^\top, \widehat{\mathbf{z}}_3^\top)^\top$. After checking the encodings, the verifier accepts precisely when $c = H((\mathbf{p}, \mathbf{t}), \mu)$ and $\|\widehat{\mathbf{z}}\| \leq \gamma$. The parity is therefore part of the hashed commitment and is recovered from the transmitted response and challenge.

C.3 Recovery and Correctness

The next lemma provides both the norm bound for honest signatures and the exact relation required for extraction from adversarial signatures.

Lemma 5 (Commitment recovery). *Suppose* $\mathbf{Az}-qc\mathbf{j} = \mathbf{w} \bmod 2q$ *and set* $\overline{\mathbf{w}} = \mathsf{PackCommit}(\mathbf{w})$. *Apply* Compress *to* $(\mathbf{A}, \mathbf{z}, c, \overline{\mathbf{w}})$. *Then* Recover *returns the original compressed commitment and* $\widehat{\mathbf{z}}$ *satisfying*

$$\|\widehat{\mathbf{z}}\| \le \|\mathbf{z}\| + \sqrt{mn}(a/2+1).$$

Moreover, for every canonically encoded $(\mathbf{z}_0, c, \mathbf{h})$, *the output* $((\mathbf{p}, \mathbf{t}), \widehat{\mathbf{z}})$ *of* Recover *satisfies*

$$\mathbf{A}\widehat{\mathbf{z}} - qc\mathbf{j} = \mathbf{p} + 2a\mathbf{t} \pmod{2q}. \tag{8}$$

Proof. For an honest compression, $\mathbf{w} = \mathbf{w}' + 2\mathbf{z}_3 \bmod 2q$. The two procedures therefore obtain the same parity and satisfy $\mathbf{v} = \mathbf{v}' + \mathbf{z}_3 \bmod q$. The definition of $\mathbf{h}$ gives $\mathbf{t} = Q(\mathbf{v})$, proving recovery of the original commitment. Put $\mathbf{d} = \mathrm{cent}_q(\mathbf{v} - aQ(\mathbf{v}))$. Then

$$\widehat{\mathbf{z}}_3 = \mathrm{cent}_q(\mathbf{z}_3 - \mathbf{d}), \qquad \|\mathbf{d}\| \le \sqrt{mn}(a/2+1)$$

by (7). Taking a centered representative cannot increase the absolute value of an integer coefficient. The triangle inequality hence yields

$$\|\widehat{\mathbf{z}}\| \le \|(\mathbf{z}_0^\top, (\mathbf{z}_3 - \mathbf{d})^\top)^\top\| \le \|\mathbf{z}\| + \sqrt{mn}(a/2+1).$$

For an arbitrary encoded input, the recovery procedure gives $\mathbf{w}' = \mathbf{p} + 2\mathbf{v}' \bmod 2q$ and $\widehat{\mathbf{z}}_3 = a\mathbf{t} - \mathbf{v}' \bmod q$. Substituting these identities into $\mathbf{A}\widehat{\mathbf{z}} - qc\mathbf{j} = \mathbf{w}' + 2\widehat{\mathbf{z}}_3$ proves (8), independently of how the hint was chosen.

Since $a/2+1 = \alpha/4+1$, the norm bound gives exactly the compression allowance in (4). Equation (8) also shows that two accepting signatures with the same compressed commitment satisfy the same public linear relation. This is the commitment-recovery property used in the MSIS extraction in Theorem 1.

C.4 Commitment Min-Entropy

The ROM proof bounds the probability that a signing commitment coincides with a previously queried hash input. The identity block in $\mathbf{A}$ lets us obtain the required entropy from the last perturbation block, even after conditioning on all other coefficients.

Lemma 6 (Compressed commitment min-entropy). *Suppose* (1) *holds. Choose* $\epsilon_h \in (0,1)$ *such that*

$$\sigma_3/\sqrt{2} \ge \eta_{\epsilon_h}(\mathbb{Z}^{mn}), \qquad q > \sigma_3\eta_{\epsilon_h}(\mathbb{Z}^{mn}).$$

For every accepted key and $\mathbf{y} \leftarrow \mathcal{D}_{\mathbb{Z}^{kn},\sqrt{\mathbf{V}}}$, *the min-entropy of* $\mathsf{PackCommit}(\mathbf{Ay} \bmod 2q)$ *is at least* $\max(0,h)$, *where*

$$h = mn\log_2\left(\frac{\sigma_3}{\sqrt{2}(a+1)}\right) - \log_2\left(\frac{1+\epsilon_h}{1-\epsilon_h}\right). \tag{9}$$

Proof. Use the perturbation shape $\mathbf{V}$ and target shape $\mathbf{W}$ from Appendix B, and split $\mathbf{y} = (\mathbf{y}_0^\top, \mathbf{y}_3^\top)^\top$ according to $\mathbf{A} = (\mathbf{A}_0 \mid 2\mathbf{I}_m)$. Conditioned on any value of $\mathbf{y}_0$, the vector $\mathbf{y}_3$ is a discrete Gaussian with some center $\mathbf{d} \in \mathbb{R}^{mn}$ and shape matrix $\mathbf{C}$, the Schur complement of the first block of $\mathbf{V}$. From $\mathbf{W}/2 \preceq \mathbf{V} \preceq \mathbf{W}$, inversion followed by restriction to the lower principal block gives

$$\sigma_3^2\mathbf{I}_{mn}/2 \preceq \mathbf{C} \preceq \sigma_3^2\mathbf{I}_{mn}.$$

For $\mathbf{b} \in \mathbb{Z}_q^{mn}$, its conditional residue probability is

$$\Pr[\mathbf{y}_3 = \mathbf{b} \bmod q \mid \mathbf{y}_0] = \frac{\rho_{\sqrt{\mathbf{C}},\mathbf{d}}(\mathbf{b} + q\mathbb{Z}^{mn})}{\rho_{\sqrt{\mathbf{C}},\mathbf{d}}(\mathbb{Z}^{mn})} \le \frac{1+\epsilon_h}{1-\epsilon_h}\left(\frac{\sqrt{2}}{\sigma_3}\right)^{mn}.$$

Here $\rho_{\sqrt{\mathbf{C}},\mathbf{d}}(T) = \rho_{\sqrt{\mathbf{C}}}(T - \mathbf{d})$. The smoothing condition bounds the denominator below by $(1-\epsilon_h)\sqrt{\det \mathbf{C}}$, with $\sqrt{\det \mathbf{C}} \geq (\sigma_3/\sqrt{2})^{mn}$. For the numerator, Poisson summation bounds the Gaussian mass of a translate of $q\mathbb{Z}^{mn}$ by the mass of $q\mathbb{Z}^{mn}$ itself. The upper bound $\mathbf{C} \preceq \sigma_3^2 \mathbf{I}_{mn}$ and the condition on q then bound this mass by $1+\epsilon_h$.

Conditioned on $\mathbf{y}_0$, the parity $\mathbf{p}$ is fixed and $\mathbf{v}$ is a fixed translate of $\mathbf{y}_3 \bmod q$. By (7), a quantized vector has at most $(a+1)^{mn}$ preimages. Thus each compressed commitment has conditional probability at most

$$\frac{1+\epsilon_h}{1-\epsilon_h}\left(\frac{\sqrt{2}(a+1)}{\sigma_3}\right)^{mn}.$$

This bound holds for every $\mathbf{y}_0$, so it also holds after averaging over $\mathbf{y}_0$. Taking the negative base-two logarithm and using nonnegativity of min-entropy proves the claim.

C.5 Signature-Size Calculation

We calculate signature entropy in the ideal transcript distribution of Appendix B: sample c uniformly and $\mathbf{z} \leftarrow \mathcal{D}_{\mathbb{Z}^{kn},\mathbf{D}}$ independently, then apply the compression procedure before the oracle-input collision check. The transmitted signature consists of the challenge, the first $k-m$ response polynomials, and the complete hint. Its entropy can be evaluated by decoding the challenge and retained response first, since they determine every coefficient of $\mathbf{v}'$.

Let $Z \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_3}$ and let H denote Shannon entropy in bits. Conditional on a coefficient $v' = t$, independence of the ideal response blocks gives the hint distribution

$$P_t(j) = \sum_{z \in \mathbb{Z}} \Pr[Z = z]\, \mathbf{1}\{Q(t+z \bmod q) - Q(t) = j \bmod K\}. \tag{10}$$

Setting $\mathsf{H}_{\max} = \max_{t \in \mathbb{Z}_q} \mathsf{H}(P_t)$ accounts for every possible public phase, including the exceptional zero bin of the quantizer. The chain rule and the product form of the ideal response yield the signature-entropy bound

$$\log_2 |\mathcal{C}| + n\mathsf{H}(\mathcal{D}_{\mathbb{Z},\sigma_1}) + rn\mathsf{H}(\mathcal{D}_{\mathbb{Z},\sigma_2}) + mn\mathsf{H}_{\max}.$$

Table 2 reports this quantity divided by eight. It includes all transmitted signature fields under the ideal entropy model. Finite-precision coding, framing, and coder state are implementation costs outside this entropy calculation.