HAETAE: Shorter Lattice-Based Fiat-Shamir Signatures

Jung Hee Cheon[1,2], Hyeongmin Choe[1], Julien Devevey[3], Tim Güneysu[4,5], Dongyeon Hong[6], Markus Krausz[4], Georg Land[4], Marc Möller[4], Junbum Shin[2], Damien Stehlé[7] and MinJune Yi[1]

[1] Seoul National University, Seoul, Republic of Korea
jhcheon@snu.ac.kr,sixtail528@snu.ac.kr,yiminjune@snu.ac.kr
[2] CryptoLab Inc., Seoul, Republic of Korea junbum.shin@cryptolab.co.kr
[3] ANSSI, Paris, France julien.devevey@ens-lyon.fr
[4] Ruhr University Bochum, Bochum, Germany firstname.lastname@rub.de,mail@georg.land
[5] DFKI, Bremen, Germany
[6] The Affiliated Institute of ETRI, Daejeon, Republic of Korea jjoker041@gmail.com
[7] CryptoLab Inc., Lyon, France damien.stehle@cryptolab.co.kr

Version 3.0
(July 4, 2024)

The most up-to-date versions of HAETAE can be accessed at:
www.kpqc.cryptolab.co.kr/haetae.

Abstract. We present HAETAE (Hyperball bimodAl modulE rejecTion signAture schemE), a new lattice-based signature scheme. Like the NIST-selected Dilithium signature scheme, HAETAE is based on the Fiat-Shamir with Aborts paradigm, but our design choices target an improved complexity/compactness compromise that is highly relevant for many space-limited application scenarios. We primarily focus on reducing signature and verification key sizes so that signatures fit into one TCP or UDP datagram while preserving a high level of security against a variety of attacks. As a result, our scheme has signature and verification key sizes up to 39% and 25% smaller, respectively, compared than Dilithium. We provide a portable, constant-time reference implementation together with an optimized implementation using AVX2 instructions and an implementation with reduced stack size for the Cortex-M4. Moreover, we describe how to efficiently protect HAETAE against implementation attacks such as side-channel analysis, making it an attractive candidate for use in IoT and other embedded systems.

Keywords: Signature, Fiat-Shamir, Lattice-based Cryptography, Bimodal Distribution

This work is submitted to 'Korean Post-Quantum Cryptography Competition' (www.kpqc.or.kr) and 'NIST Post-Quantum Cryptography Standardization of Additional Digital Signature Schemes.'

Changelog

July 04, 2024 (version 3.0)

We corrected the private key sizes in the specification, which were missing 32 bytes. B' has been reduced by 0.01 to satisfy Lemma 5, the effect on the implementation is negligible. We significantly improved the key generation procedure and achieved 40-60% reduced clock cycles in AVX2 implementation by replacing the 512-point Fast Fourier Transform (FFT) from key generation by a 256-point FFT.

February 23, 2024 (version 2.1)

We newly provide a full reasoning about HVZK property of the compressed HAETAE in Section 4.3 based on the reduction from [BBD$^+$23]. We also add a precise security bound in Section 4.5, explicitly stating how the security of each problem is related to the overall security of HAETAE and satisfies NIST security levels. This includes "refined" security estimations for the MLWE problem based on the analysis of [DDGR20] in Table 1, since our security estimations rely on the Core-SVP method, which underestimates the security of the scheme.

Reflecting the requirements of KpqC round 2, we provide the list of reported attacks and fixes in Section 8, which were mainly due to the implementation errors. For the masking strategies against side-channel attacks, we also remark Appendix B.

Some changes that were missed in the previous changelog are now added.

November 29, 2023 (version 2.0)

To clearly specify the design of HAETAE to both theorists and the implementers, we significantly update the specification of HAETAE. The building blocks that are specific to HAETAE are clearly identified in Section 3, which can be directly linked to the parameters. The name of the challenge sampler is changed from SampleInBall to SampleBinaryChallenge. Some parameters are renamed to match the reference implementation. Also, an implementation-oriented specification is available in Section 5.

The Hyperball sampler for HAETAE-260 and the rANS encoding are modified to match the implementation with our specification, and to avoid implementation-specific attacks [Saa23, Tea23], respectively. Note, these attacks were specifically due to an implementation error in the rANS encoding, which are now fixed. The parameters for HAETAE-120 and HAETAE-180 are changed to correctly reflect the verification key compression, which ends up with a small increase in the signature sizes.

In addition, we provide how HAETAE can be securely and efficiently masked against side-channel attacks in Appendix B, which will naturally prevent the well-known power/electromagnetic vulnerabilities.

Lastly, the AVX2 optimized implementation and the Cortex M4 embedded implementation of HAETAE are now available in Section 6 and Section 7, respectively.

May 2, 2023 (version 1.0)

First, all the missing parts in the KpqC first-round submission are included in the reference code, which can be found on the team HAETAE website: https://kpqc.cryptolab.co.kr/haetae. We implement the rANS encoding for $\mathbf{z}_1$ and additionally apply it to the hint vector ($\mathbf{h}$) compression along with slightly modified HighBits and LowBits algorithms. We implement the rejection algorithm for the secret key rejection but with a new rejection condition. The new condition reduces the magnitude of the secret key when multiplied by the challenge, hence the SIS bound. As a result, we update the signature sizes to the actual sizes with the encodings.

Second, we remove all the algorithms using floating-point arithmetic. The continuous Gaussian sampling, required for the continuous hyperball uniform sampling, is implemented in fixed-point arithmetic by discretizing the hyperball with an appropriate scale factor. Based on this, we further make the reference implementation to be constant-time; that is, the execution time for the signing is independent of the secret key.

Third, we introduce a verification key truncation algorithm which is adopted from Dilithium. This is applied to the first two parameter sets HAETAE-120 and HAETAE-180.

We change the parameter set by considering all the above changes, which give a trade-off between the sizes and speed.

Lastly, we change our security proof to rely on the analysis of [DFPS23] since the analysis of [KLS18] is flawed, as pointed out in [DFPS23] and [BBD+23].

1 Introduction

We introduce HAETAE[1], a new post-quantum digital signature scheme, whose security is based on the hardness of the module versions of the lattice problems LWE and SIS. The scheme design follows the "Fiat-Shamir with Aborts" paradigm [Lyu09, Lyu12], which relies on rejection sampling: rejection sampling is used to transform a signature trial whose distribution depends on sensitive information, into a signature whose distribution can be publicly simulated. Our scheme is in part inspired from CRYSTALS-Dilithium [DKL+18], a post-quantum "Fiat-Shamir with Aborts" signature scheme which was selected for standardization by the American National Institute of Standards and Technology (NIST). HAETAE differs from Dilithium in two major aspects: (i) we use a bimodal distribution for the rejection sampling, like in the BLISS signature scheme [DDLL13], instead of a "unimodal" distribution like Dilithium, (ii) we sample from and reject to hyperball uniform distributions, instead of discrete hypercube uniform distributions. This last aspect also departs from BLISS, which relies on discrete Gaussian distributions, and follows a suggestion from [DFPS22], which studied rejection sampling in lattice-based signatures following the "Fiat-Shamir with Aborts" paradigm.

1.1 Design rationale

A brief recap on Fiat-Shamir with Aborts. The Fiat-Shamir with Aborts paradigm was introduced in lattice-based cryptography in [Lyu09, Lyu12]. The verification key is a pair of matrices $(\mathbf{A}, \mathbf{T} = \mathbf{AS} \bmod q)$, where $\mathbf{A}$ is a uniform matrix modulo some integer q and $\mathbf{S}$ is a small-magnitude matrix that makes up the secret key. A signature for a message M is comprised of an integer vector $\mathbf{z}$ of the form $\mathbf{y} + \mathbf{Sc}$, for some random small-magnitude $\mathbf{y}$ and some small-magnitude challenge $\mathbf{c} = H(\mathbf{Ay} \bmod q, M)$. Rejection sampling is then used to ensure that the distribution of the signature becomes independent from the secret key. Finally, the verification algorithm checks that the vector $\mathbf{z}$ is short and that $\mathbf{c} = H(\mathbf{Az} - \mathbf{Tc} \bmod q, M)$.

Improving compactness. As analyzed in [DFPS22], The choice of the distributions to sample from and reject to has a major impact on the signature size. Dilithium relies on discrete uniform distributions in hypercubes, which makes the scheme easier to implement. However, such distributions are far from optimal in terms of resulting signature sizes. We choose a different trade-off: by losing a little on ease of implementation, we obtain more compact signatures.

[1] The haetae is a mythical Korean lion-like creature with the innate ability to distinguish right from wrong.

Uniform distributions in hyperballs. A possibility would be to consider Gaussian distributions, which are superior to uniform distributions in hypercubes, in terms of resulting signature compactness (see, e.g., [DFPS22]). However, this choice has two downsides. First, the rejection step involves the computation of a transcendental function on an input that depends on the secret key. This is cumbersome to implement and sensitive to side-channel attacks [EFGT17]. Second, since the final signature follows a Gaussian distribution there is a nonzero probability that the final signature is too large and does not pass the verification. The signer must realise that and reject the signature, making the expected number of rejects slightly grow in practice. Uniform distributions over hyperballs have been put forward in [DFPS22] as an alternative choice of distributions leading to signatures with compactness between those obtained with Gaussians and those obtained with hypercube uniforms. Compared to Gaussians, they do not suffer from the afore-mentioned downsides: the rejection step is simply checking whether Euclidean norms are sufficiently small; and as there is no tail, there is no need for an extra rejection step to ensure that verification will pass. HAETAE showcases that this provides an interesting simplicity/compactness compromise.

Bimodal distributions. A modification of Lyubashevsky's signatures was introduced in [DDLL13]. It allows for the use of bimodal distributions in the signature generation. The signature is now of the form $\mathbf{y} + (-1)^b\mathbf{Sc}$, where $\mathbf{y}$ is sampled from a fixed distribution and $b \in \{0, 1\}$ is sampled uniformly. The signature is then rejected to a given secret-independent target distribution. To make sure that the verification test passes, computations are performed modulo $2q$ and key generation forces the equality $\mathbf{AS} = q\mathbf{Id}$. It turns out that this modification can lead to more compact signatures than the unimodal setup. In [DDLL13], the authors relied on discrete Gaussian distributions. We instead use uniform distributions over hyperballs: like for Gaussians, switching from unimodal to bimodal for hyperball-uniforms leads to more compact signatures.

Flexible design by working with modules. The original design for BLISS [DDLL13] relies on Ring-LWE and Ring-SIS, and a variant of the key generation algorithm relied on ratios of polynomials, *à la* NTRU. This setup forces to choose a working polynomial ring for any desired security level. In order to offer more flexibility without losing in terms of implementation efficiency, we choose to rely on module lattices, like Dilithium, with a fixed working polynomial ring $\mathcal{R} = \mathbb{Z}[x]/(x^{256} + 1)$ across all security levels. In our instantiations, we target the NIST PQC security levels 2, 3 and 5. Varying the security and updating the parameters is easily achievable and we provide a security estimator that is able to help one reach a given target security.

Secret key rejection sampling. We introduce a new rejection procedure in the key generation algorithm to minimize the magnitude of the secret key when multiplied by the challenge. This facilitates rejection sampling in the signing algorithm and leads to smaller signatures. The key generation rejection is also designed to be efficient and simple to implement. It significantly improves over a procedure with a similar objective in the key generation of BLISS.

A compact verification key. The flexibility provided by modules allows us to reduce the verification key size. Instead of taking the challenge $\mathbf{c}$ as a vector over $\mathcal{R}$, we choose it in $\mathcal{R}$: the main condition on the challenge is that it has high min-entropy, which is already the case for binary vectors over $\mathcal{R}$. As a result, the secret $\mathbf{S}$ can be chosen as a vector over $\mathcal{R}$ rather than a matrix. The key-pair equation $\mathbf{AS} = q\mathbf{Id}$ then becomes $\mathbf{As} = q\mathbf{j}$, where $\mathbf{j}$ is the vector starting with 1 and then continuing with 0's. To further compress the verification key, we use verification key truncation adopted from Dilithium by taking into account the residue modulo 2. Our key generation algorithm just creates an MLWE sample $(\mathbf{A}_{\mathsf{gen}}, \mathbf{b} - \mathbf{a} = \mathbf{A}_{\mathsf{gen}}\mathbf{s}_{\mathsf{gen}} + \mathbf{e}_{\mathsf{gen}})$ modulo q, where $\mathbf{a}$ is uniform random over $\mathcal{R}_q^k$. By truncating $\mathbf{b}$ as $\mathbf{b} = \mathbf{b}_1 + \mathbf{b}_0$, we define a $k \times (k + \ell)$ matrix $\mathbf{A}$ as $\mathbf{A} = (-2(\mathbf{a} - \mathbf{b}_1) + q\mathbf{j} | \ 2\mathbf{A}_{\mathsf{gen}} | \ 2\mathbf{Id}_k) \bmod 2q$. The key-pair equation is satisfied for $\mathbf{s} = (1 | \ \mathbf{s}_{\mathsf{gen}} | \ \mathbf{e}_{\mathsf{gen}} - \mathbf{b}_0)$. The verification key consists of $(\mathbf{A}_{\mathsf{gen}}, \mathbf{a}, \mathbf{b}_1)$. As $(\mathbf{a} | \ \mathbf{A}_{\mathsf{gen}})$ is uniformly distributed, we can generate it from a seed using an extendable output function, and the verification key is reduced to the seed and the vector $\mathbf{b}_1$. If we had kept the original key-pair equation $\mathbf{AS} = q\mathbf{Id}$, then the appropriately modified variant of our key-generation algorithm would have led to a verification key that is a matrix (with a seed) rather than a vector (with a seed).

Compression techniques to lower the signature size. We use two techniques to compress the signatures. First, as the verification key $\mathbf{A}$ is in (almost)-HNF, we can use the Bai-Galbraith technique [BG14]. Namely, the second part of the signature, which is multiplied by $2\mathbf{Id}$ in the challenge computation and verification algorithm, can be aggressively compressed by cutting its low bits. This requires in turn modifying the computation of the challenge $\mathbf{c}$ and the verification algorithm, in order to account for this precision loss. Usually, this is done by keeping only the high bits of $\mathbf{Ay}$ in the computation of the challenge. However, as we multiply everything by 2, we do not keep the lowest bit of those high bits and keep the (overall) least significant bit instead. As in Dilithium, our decomposition of bits technique is a Euclidean division with a centered remainder, and we choose a representative range for modular integers that starts slightly below zero to further reduce the support of the high bits. The second compression technique, suggested in [ETWY22] in the context of lattice-based hash-and-sign signatures, concerns the choice of the binary representation of the signature. As the largest part of it consists in a vector that is far from being uniform, we can choose some entropic coding to obtain a signature size close to its entropy. In particular, as in [ETWY22], we choose the efficient range Asymmetric Numeral System to encode our signature, as it allows us to encode the whole signature and not lose a fraction of a bit per vector coordinate, like with Huffman coding. We can further apply the two techniques to the hint vector $\mathbf{h}$, which is also a part of the signature, to reduce the signature sizes.

Efficient choice of modulus. We choose the prime q to be a good prime in the sense that the ring operations can be implemented efficiently and that the decomposition of bits algorithms, are correctly operated. For ring operations, we use the Number Theoretic Transform (NTT) with a fully splitting polynomial ring. The polynomial ring $\mathcal{R}$ fully splits modulo q when the multiplicative group $\mathbb{Z}_q^\times$ has an element of order 512, or equivalently when $q = 1 \bmod 512$. We choose $q = 64513$, which indeed satisfies this property. Interestingly, it fits in 16 bits, which allows dense storing on embedded devices. Furthermore, it is close to the next power of two, which is convenient for the sampling of uniform integers modulo q.

Fixed-point algorithm for hyperball sampling. Unlike uniform Gaussian sampling or uniform hypercube sampling, uniform hyperball sampling has not been considered in the cryptographic protocols before the suggestion of [DFPS22]. To narrow the gap between the hyperball uniforms sampled in the real and the ideal world, we discretize the hyperball and bound the numerical error and their effect by analyzing their propagation. This leads to a fixed-point hyperball sampling algorithm and, therefore, the fixed-point implementation of the whole signing process.

Deterministic and randomized version. HAETAE can be set in a deterministic or randomized mode. We focus on the deterministic version, but we also give the randomized version. Note that in the randomized version, a significant part of the signing algorithm can be executed off-line as it does not depend on the message.

1.2 Advantages and limitations

1.2.1 Advantages

- Our scheme relies on the difficulty of hard lattice problems, which have been well-studied for a long time.
- Signature sizes are 29% to 39% smaller than those of Dilithium at comparable security levels, and verification keys are 20% to 25% smaller.
- Implementation-wise, while our design rationale departs from Dilithium's, the scheme remains implementation-friendly. In particular,
  - the rejection step only involves computations of Euclidean norms,
  - the whole signing process can be implemented with fixed-point arithmetic,
  - a significant message-independent part of signing can be performed "off-line", for the randomized version of the scheme.

Comparison with hash-and-sign lattice signatures. In terms of ease of implementation, our scheme favorably compares to lattice signatures based on the hash and sign paradigm such as Falcon [FHK+18]. The efforts for making it masking-friendly, namely Mitaka [EFG+22], were recently broken [Pre23]. HAETAE, Falcon and Mitaka all three rely on some form of Gaussian sampling, which are typically difficult to implement and protect against side-channel attacks. Falcon makes sequential calls to a Gaussian sampler over $\mathbb{Z}$ with arbitrary centers. Mitaka also relies on an integer Gaussian sampler with arbitrary centers, but the calls to it can be massively parallelized. It also uses a continuous Gaussian sampler, which is arguably simpler. HAETAE, however, only relies on a (zero-centered) continuous Gaussian sampler, used to sample uniformly in hyperballs. The calls to it can also be massively parallelized. This difference makes HAETAE possible to have a fixed-point signing algorithm and easier masking. Further, in the randomized version of the signature scheme, these samples can be computed offline as they are independent of the message to be signed. The online tasks are far simpler than those of Falcon and Mitaka. Finally, we note that key generation is much simpler for HAETAE than in Falcon and Mitaka.

1.2.2 Limitations

- The key generation algorithm restarts if the secret key does not satisfy the key rejection condition. This makes the key generation algorithm of HAETAE slower than Dilithium's.
- While HAETAE is simpler from an implementation perspective, its verification key and signature sizes are larger than Falcon's.

2 Preliminaries

Before introducing specific results adapted to the setting in HAETAE in Section 3 and the HAETAE scheme itself in Section 4, we start by defining notations used throughout this paper and recapitulate relevant fundamental works.

2.1 Notations

Matrices are denoted in bold font and upper case letters (e.g., $\mathbf{A}$), while vectors are denoted in bold font and lowercase letters (e.g., $\mathbf{y}$ or $\mathbf{z}_1$). The i-th component of a vector is denoted with subscript i (e.g., y_i for the i-th component of $\mathbf{y}$).

Every vector is a column vector. We denote concatenation between vectors by putting the rows below as $(\mathbf{u}, \mathbf{v})$ and the columns on the right as $(\mathbf{u}|\mathbf{v})$. We naturally extend the latter notation to concatenations between matrices and vectors (e.g., $(\mathbf{A}|\mathbf{b})$ or $(\mathbf{A}|\mathbf{B})$).

We let $\lfloor y \rceil$ be a rounding of $y \in \mathbb{R}$ to the nearest integer. We naturally extend the rounding notation to vectors and polynomials by applying it component-wise.

We let $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$ be a polynomial ring where n is a power of 2 integer and for any positive integer q the quotient ring $\mathcal{R}_q = \mathbb{Z}[x]/(q, x^n + 1) = \mathbb{Z}_q[x]/(x^n + 1)$. We abuse notations and identify $\mathcal{R}_2$ with the set of elements in $\mathcal{R}$ with binary coefficients. We also let $\mathcal{R}_{\mathbb{R}} = \mathbb{R}[x]/(x^n + 1)$ be a polynomial ring over real numbers. For an integer η, we let S_η denote the set of polynomials of degree less than n with coefficients in $[-\eta, \eta] \cap \mathbb{Z}$. Given $\mathbf{y} = (\sum_{0 \le i < n} y_i \ x^i, \cdots, \sum_{0 \le i < n} y_{nk-n+i} \ x^i)^\top \in \mathcal{R}^k$ (or $\mathcal{R}_{\mathbb{R}}^k$), we define its ℓ_2-norm as the ℓ_2-norm of the corresponding "flattened" vector $\|\mathbf{y}\|_2 = \|(y_0, \cdots, y_{nk-1})^\top\|_2$.

Let $\mathcal{B}_{\mathcal{R},m}(r, \mathbf{c}) = \{\mathbf{x} \in \mathcal{R}_{\mathbb{R}}^m | \|\mathbf{x} - \mathbf{c}\|_2 \le r\}$ denote the continuous hyperball with center $\mathbf{c} \in \mathcal{R}^m$ and radius $r > 0$ in dimension $m > 0$. When $\mathbf{c} = \mathbf{0}$, we omit it. Let $\mathcal{B}_{(1/N)\mathcal{R},m}(r, \mathbf{c}) = (1/N)\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r, \mathbf{c})$ denote the discretized hyperball with radius $r > 0$ and center $\mathbf{c} \in \mathcal{R}^m$ in dimension $m > 0$ with respect to a positive integer N. When $\mathbf{c} = \mathbf{0}$, we omit it. Given a measurable set $X \subseteq \mathcal{R}^m$ of finite volume, we let $U(X)$ denote the continuous uniform distribution over X. It admits $\mathbf{x} \mapsto \chi_X(\mathbf{x})/\mathsf{Vol}(X)$ as a probability density, where χ_X is the indicator function of X and $\mathsf{Vol}(X)$ is the volume of the set X. For the normal distribution over $\mathbb{R}$ centered at μ with standard deviation σ, we use the notation $\mathcal{N}(\mu, \sigma)$.

For a positive integer α, we define $r \bmod^{\pm} \alpha$ as the unique integer r' in the range $[-\alpha/2, \alpha/2)$ satisfying the relation $r = r' \bmod \alpha$. We also define $r \bmod^{+} \alpha$ as the unique integer r' in the range $[0, \alpha)$ that satisfies $r = r' \bmod \alpha$. We denote the least significant bit of an integer r with $\mathsf{LSB}(r)$. We naturally extend this to integer polynomials and vectors of integer polynomials, by applying it component-wise.

For a sequence of real numbers $a_0, \ldots, a_n$, we denote the i-th maximum as $\overset{i\text{-th}}{\max_j} a_j$.

2.2 Signatures

We briefly recall the formalism of digital signatures.

Definition 1 (Digital Signature)**.** A signature scheme is a tuple of PPT algorithms $(\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ with the following specifications:

- $\mathsf{KeyGen} : 1^\lambda \to (\mathsf{vk}, \mathsf{sk})$ outputs a verification key vk and a signing key sk;
- $\mathsf{Sign} : (\mathsf{sk}, \mu) \to \sigma$ takes as inputs a signing key sk and a message μ and outputs a signature σ;
- $\mathsf{Verify} : (\mathsf{vk}, \mu, \sigma) \to b \in \{0, 1\}$ is a deterministic algorithm that takes as inputs a verification key vk, a message μ, and a signature σ and outputs a bit $b \in \{0, 1\}$.

Let $\gamma > 0$. We say that it is γ-correct if for any pair $(\mathsf{vk}, \mathsf{sk})$ in the range of KeyGen and μ,

$$\Pr[\mathsf{Verify}(\mathsf{vk}, \mu, \mathsf{Sign}(\mathsf{sk}, \mu)) = 1] \ge \gamma,$$

where the probability is taken over the random coins of the signing algorithm. We say that it is correct in the (Q)ROM if the above holds when the probability is also taken over the randomness of the random oracle modeling the hash function used in the scheme.

We also give two security notions, namely the existential unforgeability under chosen message attacks, and under no-message attacks.

Definition 2 (Security)**.** Let $T, \delta \geq 0$. A signature scheme $\mathsf{sig} = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ is said to be (T, δ)-UF-CMA secure in the QROM if for any quantum adversary $\mathcal{A}$ with runtime $\leq T$ given (classical) access to the signing oracle and (quantum) access to a random oracle H, it holds that

$$\Pr_{(\mathsf{vk},\mathsf{sk})}[\mathsf{Verify}(\mathsf{vk}, \mu^*, \sigma^*) = 1 | (\mu^*, \sigma^*) \leftarrow \mathcal{A}^{H,\mathsf{Sign}}(\mathsf{vk})] \leq \delta,$$

where the randomness is taken over the random coins of $\mathcal{A}$ and $(\mathsf{vk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$. The adversary should also not have issued a sign query for μ^*. The above probability of forging a signature is called the advantage of $\mathcal{A}$ and denoted by $\mathsf{Adv}^{\mathsf{UF\text{-}CMA}}_{\mathsf{sig}}(\mathcal{A})$. If $\mathcal{A}$ does not output anything, then it automatically fails.

Existential unforgeability against no-message attack, denoted by UF-NMA is defined similarly except that the adversary is not allowed to query any signature per message. Strong existential unforgeability, denoted by sUF-CMA, allows an adversary to query signatures for its target message, as long as it does not output a queried signature.

2.3 Lattice Assumptions

We first recall the well-known lattice assumptions MLWE and MSIS on algebraic lattices.

Definition 3 (Decision-$\mathsf{MLWE}_{n,q,k,\ell,\eta}$)**.** For positive integers q, k, ℓ, η and the dimension n of $\mathcal{R}$, we say that the advantage of an adversary $\mathcal{A}_1$ solving the decision-$\mathsf{MLWE}_{n,q,k,\ell,\eta}$ problem is

$$\mathsf{Adv}^{\mathsf{MLWE}}_{n,q,k,\ell,\eta}(\mathcal{A}_1) = \left| \begin{array}{l} \Pr\left[b = 1 \mid \mathbf{A} \leftarrow \mathcal{R}_q^{k\times\ell}; \mathbf{b} \leftarrow \mathcal{R}_q^k; b \leftarrow \mathcal{A}_1(\mathbf{A}, \mathbf{b})\right] \\ \quad - \Pr\left[b = 1 \,\middle|\, \begin{array}{c} \mathbf{A} \leftarrow \mathcal{R}_q^{k\times\ell}; (\mathbf{s}_1, \mathbf{s}_2) \leftarrow S_\eta^\ell \times S_\eta^k; \\ b \leftarrow \mathcal{A}_1(\mathbf{A}, \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2) \end{array}\right] \end{array} \right|.$$

Definition 4 (Search-$\mathsf{MSIS}_{n,q,k,\ell,\beta}$)**.** For positive integers q, k, ℓ, a positive real number β and the dimension n of $\mathcal{R}$, we say that the advantage of an adversary $\mathcal{A}_2$ solving the search-$\mathsf{MSIS}_{n,q,k,\ell,\beta}$ problem is

$$\mathsf{Adv}^{\mathsf{MSIS}}_{n,q,k,\ell,\beta}(\mathcal{A}_2) = \Pr\left[\begin{array}{c} 0 < \|\mathbf{y}\|_2 < \beta \,\wedge \\ (\mathbf{A}|\ \mathbf{Id}_k) \cdot \mathbf{y} = 0 \mod q \end{array} \,\middle|\, \mathbf{A} \leftarrow \mathcal{R}_q^{k\times\ell}; \mathbf{y} \in \mathcal{R}_q^{k+\ell} \leftarrow \mathcal{A}_2(\mathbf{A})\right].$$

Moreover, we finally introduce a variant of the SelfTargetMSIS problem introduced in Dilithium [DKL+18], which corresponds to our setting.

Definition 5 ($\mathsf{BimodalSelfTargetMSIS}_{H,n,q,k,\ell,\beta}$)**.** Let $H : \{0,1\}^* \times \mathcal{M} \to \mathcal{R}_2$ be a cryptographic hash function, where $\mathcal{M} \subseteq \{0,1\}^*$ is a message space. Let $q, k, \ell > 0$, $\beta \geq 0$ and the dimension n of $\mathcal{R}$. An adversary $\mathcal{A}_3$ solving the search-$\mathsf{BimodalSelfTargetMSIS}_{H,n,q,k,\ell,\beta}$ problem with respect to $\mathbf{j} \in \mathcal{R}_2^k \setminus \{0\}$ has an advantage of

$$\mathsf{Adv}^{\mathsf{BimodalSelfTargetMSIS}}_{H,n,q,k,\ell,\beta}(\mathcal{A}_3) = \Pr\left[\begin{array}{c} 0 < \|\mathbf{y}\|_2 < \beta \,\wedge \\ H(\mathbf{A}\mathbf{y} - qc\mathbf{j} \mod 2q, \mu) = c \\ \hline (\mathbf{A}_0|\mathbf{b}) \leftarrow \mathcal{R}_q^{k\times\ell}; \\ \mathbf{A} = (2\mathbf{b} + q\mathbf{j}|\ 2\mathbf{A}_0|\ 2\mathbf{Id}_k) \mod 2q; \\ (\mathbf{y}, c, \mu) \in \mathcal{R}_q^{k+\ell} \times \mathcal{R}_2 \times \mathcal{M} \leftarrow \mathcal{A}_3^{|H(\cdot)\rangle}(\mathbf{A}) \end{array}\right].$$

In the ROM (resp. QROM), the adversary is given classical (resp. quantum) access to H.

The following classical reduction from MSIS to BimodalSelfTargetMSIS is very similar to the reduction from MSIS to SelfTargetMSIS introduced in [DKL+18] and is similarly non-tight. As this latter reduction, it cannot be straightforwardly extended to a reduction in the QROM, since it relies on the forking lemma.

Theorem 1 (Classical Reduction from MSIS to $\mathsf{BimodalSelfTargetMSIS}$)**.** *Let $q > 0$ be an odd modulus, $H : \{0,1\}^* \times \mathcal{M} \to \mathcal{R}_2$ be a cryptographic hash function modeled as a random oracle and that every polynomial-time classical algorithm has a negligible advantage against $\mathsf{MSIS}_{n,q,k,\ell,\beta}$. Then every polynomial-time classical algorithm has negligible advantage against $\mathsf{BimodalSelfTargetMSIS}_{n,q,k,\ell,\beta/2}$.*

Proof sketch. Consider a $\mathsf{BimodalSelfTargetMSIS}_{n,q,k,\ell,\beta/2}$ classical algorithm $\mathcal{A}$ that is polynomial-time and has classical access to H. If $\mathcal{A}^{H(\cdot)}(\mathbf{A})$ makes Q hash queries $H(\mathbf{w}_i, \mu_i)$ for $i = 1, \cdots, Q$ and outputs a solution $(\mathbf{y}, c, \mu_j)$ for some $j \in [Q]$, then we can construct an adversary $\mathcal{A}'$ for $\mathsf{MSIS}_{n,q,k,\ell,\beta}$ as follows.

The adversary $\mathcal{A}'$ can first rewind $\mathcal{A}$ to the point at which the j-th query was made and reprogram the hash as $H(\mathbf{w}_j, \mu_j) = c'(\neq c)$. Then, with probability approximately $1/Q$, algorithm $\mathcal{A}$ will produce another solution $(\mathbf{y}', c', \mu_j)$. We then have

$$\begin{cases} \mathbf{A}\mathbf{y} - qc\mathbf{j} = \mathbf{z}_j = \mathbf{A}\mathbf{y}' - qc'\mathbf{j} \bmod 2q, \\ \|\mathbf{y}\|_2, \|\mathbf{y}'\|_2 < \beta/2. \end{cases}$$

As q is odd, we have $\mathbf{A}(\mathbf{y} - \mathbf{y}') = (c - c')\mathbf{j} \bmod 2$. The fact that $c' \neq c$ implies that the latter is non-zero modulo 2, and hence so is $\mathbf{y} - \mathbf{y}'$ over the integers. As it also satisfies $(\mathbf{b}|\ \mathbf{A}_0|\ \mathbf{Id}_k) \cdot (\mathbf{y} - \mathbf{y}') = 0 \bmod q$ and $\|\mathbf{y} - \mathbf{y}'\| < \beta$, it provides a $\mathsf{MSIS}_{n,q,k,\ell,\beta}$ solution for the matrix $(\mathbf{b}|\ \mathbf{A}_0|\ \mathbf{Id}_k)$, where the submatrix $(-\mathbf{b}|\ \mathbf{A}_0) \in \mathcal{R}_q^{k\times\ell}$ is uniform. □

2.4 Sampling from the Continuous Hyperball-uniform

In order to sample in practice from hyperball uniform, we rely on the following result.

```
$\mathbf{y} \leftarrow U(\mathcal{B}_{\mathcal{R},k}(r))$:
  $y_i \leftarrow \mathcal{N}(0,1)$ for $i = 0, \cdots, nk+1$
  $L \leftarrow \|(y_0, \cdots, y_{nk+1})^\top\|_2$
  $\mathbf{y} \leftarrow r/L \cdot (\sum_{i=0}^{n-1} y_i\ x^i, \cdots, \sum_{i=nk-n}^{nk-1} y_i\ x^i)^\top$
  return $\mathbf{y}$                                        $\triangleright\ \mathbf{y} \in \mathcal{R}_{\mathbb{R}}^k$
```

Figure 1: Hyperball uniform sampling

Lemma 1 ([VGS17])**.** *The distribution of the output of the algorithm in Figure 1 is $U(\mathcal{B}_{\mathcal{R},k}(r))$.*

Sampling from continuous hyperball-uniform can be done using the algorithm in Figure 1 due to Lemma 1. However, to allow side-channel secure implementations of HAETAE, we sample from discrete hyperball-uniform. We delay to Section 3.2 the analysis of a discretized version which turns discrete Gaussian samples to discrete hyperball-uniform distribution.

2.5 Signature Encoding via Range Asymmetric Numeral System

A HAETAE signature is essentially a vector $\mathbf{z}$, that is compressed into $\mathbf{z}_2$ with smaller dimension and a hint $\mathbf{h}$, that are then encoded. While Huffman coding would be applied on each coordinate at a time, an arithmetic coding encodes the entire vector coordinates in a single number. In contrast to Huffman coding, arithmetic coding gets close to entropy also for alphabets, where the probabilities of the symbols are not powers of two. We recall a recent type of entropy coding, named range Asymmetric Numeral systems (rANS) [Dud14], that encodes the state in a natural number and thus allows faster implementations. The rANS encoding technique was recently used in [ETWY22] and we adapt it to hyperball uniform distributions. As a stream variant, rANS can be implemented with finite precision integer arithmetic by using renormalization.

Definition 6 (Range Asymmetric Numeral System (rANS) Coding)**.** Let $t > 0$ and $S \subseteq [0, 2^t - 1]$. Let $g : [0, 2^t - 1] \to \mathbb{Z} \cap (0, 2^t]$ such that $\sum_{x \in S} g(x) \leq 2^t$ and $g(x) = 0$ for all $x \notin S$. We define the following:

- $\mathsf{CDF} : S \to \mathbb{Z}$, defined as $\mathsf{CDF}(s) = \sum_{y=0}^{s-1} g(y)$.
- $\mathsf{symbol} : \mathbb{Z} \to S$, where $\mathsf{symbol}(y)$ is defined as $s \in S$ satisfying $\mathsf{CDF}(s) \leq y < \mathsf{CDF}(s+1)$.
- $C : \mathbb{Z} \times S \to \mathbb{Z}$, defined as

$$C(x, s) = \left\lfloor \frac{x}{g(s)} \right\rfloor \cdot 2^t + (x \bmod^+ g(s)) + \mathsf{CDF}(s).$$

Then, we define the rANS encoding/decoding for the set S and frequency $g/2^t$ as in Figure 2.

```
Encode((s_1, ..., s_m) ∈ S^m):
$x_0 = 0$
for $i = 0, \cdots, m-1$ do
    $x_{i+1} = C(x_i, s_{i+1})$
return $x_m$

Decode(x ∈ Z):
$y_0 = x$
$i = 0$
while $y_i > 0$ do
    $s'_{i+1} = \mathsf{symbol}(y_i \bmod^+ 2^t)$
    $y_{i+1} = \lfloor y_i / 2^t \rfloor \cdot g(s'_{i+1}) + (y_i \bmod^+ 2^t) - \mathsf{CDF}(s'_{i+1})$
    $i$++
$m = i - 1$
return $(s'_m, \cdots, s'_1) \in S^m$
```

Figure 2: rANS encoding and decoding procedures

Lemma 2 (Adapted from [Dud14])**.** *The rANS coding is correct, and the size of the rANS code is asymptotically equal to Shannon entropy of the symbols. That is, for any choice of* $\mathbf{s} = (s_1, \cdots, s_m) \in S^m$, $\mathsf{Decode}(\mathsf{Encode}(\mathbf{s})) = \mathbf{s}$*. Moreover, for any positive* x *and any probability distribution* p *over* S*, it holds that*

$$\sum_{s \in S} p(s) \log_2(C(x, s)) \leq \log_2(x) + \sum_{s \in S} p(s) \log_2 \left(\frac{g(s)}{2^t} \right) + \frac{2^t}{x}.$$

Finally, the cost in bitlength of encoding the first symbol is $\leq t$*, i.e., for any* $s \in S$*, we have* $\log_2(C(0, s)) \leq t$*.*

We determine the frequency of the symbols experimentally, by executing the signature computation and collecting several million samples. Finally, we apply some rounding strategy in order to heuristically minimize the empirical entropy $\sum_{s \in S} p(s) \log(g(s)/2^n)$.

3 HAETAE-specific Results

While our scheme is reminiscent of Dilithium, the bimodal setting hinders the use of some of its base components. In this section, we describe parts that are specifically adapted to HAETAE. First, the key generation algorithm departs from known key generation algorithms for BLISS, as we work in the module setting. Second, we study the precision needed when discretizing the hyperball sampler from Section 2.4 to enable fixed-point arithmetic. Then, we explain how challenges are computed in HAETAE. Next, we describe the rejection sampling procedure and estimate its expected number of iterations depending on the fixed-point arithmetic precision. Finally, we explain how to split the coordinates of a signature vector into high and low bits, allowing for signature compression via low bits drop. This order is consistent with the order in which those results are used during signing.

3.1 Key Generation

When using bimodal rejection sampling, the verification step relies on a specific key pair $(\mathbf{A}, \mathbf{s}) \in \mathcal{R}_p^{k\times(k+\ell)} \times \mathcal{R}_p^{k+\ell}$ such that the bimodal centers $(-1)^b\mathbf{As} \mod p$ $(b = 0, 1)$ are the same, regardless of the bit b. To generate such a pair, following [DDLL13], we choose $p = 2q$ and aim at $\mathbf{As} = q\mathbf{j} \bmod 2q$ for $\mathbf{j} = (1, 0, \dots, 0)^\top$.

3.1.1 Key Generation and Encoding

To build such a key pair $(\mathbf{A}, \mathbf{s})$, we do as follows. We first generate an MLWE sample $\mathbf{b} = \mathbf{A}_{\mathsf{gen}}\mathbf{s}_{\mathsf{gen}} + \mathbf{e}_{\mathsf{gen}} \bmod q$, where $\mathbf{A}_{\mathsf{gen}} \hookleftarrow U(\mathcal{R}_q^{k\times(\ell-1)})$ and $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) \hookleftarrow U(S_\eta^{\ell-1} \times S_\eta^k)$. We then define $\mathbf{A} = (-2\mathbf{b} + q\mathbf{j}|\ 2\mathbf{A}_{\mathsf{gen}}|\ 2\mathbf{Id}_k) \mod 2q$ as well as $\mathbf{s}^\top = (1|\mathbf{s}_{\mathsf{gen}}^\top|\mathbf{e}_{\mathsf{gen}}^\top)$. This is a valid verification key pair for HAETAE, but the choice of even modulus $2q$ makes it hard to truncate the least significant bits of $\mathbf{b}$ as in Dilithium.

To enable the verification key truncation, we modify the key generation algorithm, as follows. We use an extra randomness $\mathbf{a}_{\mathsf{gen}} \hookleftarrow U(\mathcal{R}_q^k)$ and let $\mathbf{b}-\mathbf{a}_{\mathsf{gen}} = \mathbf{A}_{\mathsf{gen}}\mathbf{s}_{\mathsf{gen}}+\mathbf{e}_{\mathsf{gen}} \bmod q$. For any decomposition $\mathbf{b} = \mathbf{b}_1 + \mathbf{b}_0$, we then define $\mathbf{A} = (2(\mathbf{a}_{\mathsf{gen}} - \mathbf{b}_1) + q\mathbf{j}|2\mathbf{A}_{\mathsf{gen}}|2\mathbf{I}_k)$ as well as $\mathbf{s}^\top = (1|\mathbf{s}_{\mathsf{gen}}^\top|(\mathbf{e}_{\mathsf{gen}} - \mathbf{b}_0)^\top)$. One sees that $\mathbf{As} = q\mathbf{j} \bmod 2q$. In practice, the verification key is then comprised of $\mathbf{b}_1$ and the seed that allows generating $\mathbf{A}_{\mathsf{gen}}$ and $\mathbf{a}_{\mathsf{gen}}$. The secret key is the seed used to generate $\mathbf{s}$ and $(\mathbf{A}_{\mathsf{gen}}, \mathbf{a}_{\mathsf{gen}})$.

It remains to choose the decomposition of $\mathbf{b}$, that we see as an nk-dimensional vector with coordinates in $[0, q-1]$. We set the coordinates of $\mathbf{b}_1$ as follows. If some coordinate of $\mathbf{b}$ is even, then we take the same value for the corresponding coordinate of $\mathbf{b}_1$. Else, we take the rounding of this coordinate to the nearest multiple of 4 as value for $\mathbf{b}_1$. Next we set $\mathbf{b}_0 = \mathbf{b} - \mathbf{b}_1$ and we note that coordinates of $\mathbf{b}_0$ lie in $[-1, 1]$, i.e., $\mathbf{b}_0 \in S_1^k$. We can then write $\mathbf{b} = \mathbf{b}_0 + 2\mathbf{b}_1'$, where $\mathbf{b}_1'$ is encoded using $\lceil \log_2(q) - 1\rceil$ bits per coordinate, i.e. one less bit than $\mathbf{b}$. This is computed coordinate-wise with $\mathbf{b}_0 = (-1)^{\lfloor \mathbf{b}/2 \rfloor \bmod 2}\mathbf{b} \bmod 2$. In all of the following, we let $(\mathsf{LowBits}^{\mathsf{vk}}(\mathbf{b}), \mathsf{HighBits}^{\mathsf{vk}}(\mathbf{b}))$ denote $(\mathbf{b}_0, \mathbf{b}_1)$.

When $\mathbf{b}$ is uniform, we notice that the coordinates of $\mathbf{b}_0$ roughly follow a (centered) binomial law with parameters $(2, 1/2)$, which experimentally leads to smaller choices for γ, which we discuss and introduce below.

Note that the truncation reduces each coefficient of $\mathbf{b}$ by 1 bit. So the verification key becomes shorter, but not significantly. Thus, we use the truncation for lower security levels and keep the no-truncation version for the highest level. In the following, we refer to the truncated version as $d = 1$ and the non-truncated version as $d = 0$, where d is the vk truncation bit.

3.1.2 Rejection Sampling on the Key

A critical step of our scheme is bounding $\|c\mathbf{s}\|_2$, where $\mathbf{s}$ is generated as before and $c \in \mathcal{R}$ is a polynomial with coefficients in $\{0,1\}$ and has less than or equal to τ nonzero coefficients. The lower this bound is, the smaller the signature is, which in turn leads to the harder forging. In the key generation algorithm, we apply the following rejection condition for some heuristic value γ, bounding $\|c\mathbf{s}\|_2 \le \gamma\sqrt{\tau}$:

$$\mathcal{N}(\mathbf{s}) := \tau \cdot \sum_{i=1}^{m} \overset{i\text{-th}}{\max_{0\le j<2n}} \|\mathbf{s}(\omega_j)\|_2^2 + r \cdot \overset{(m+1)\text{-th}}{\max_{0\le j<2n}} \|\mathbf{s}(\omega_j)\|_2^2 \le \gamma^2 n,$$

where $m = \lfloor n/\tau \rfloor$, $r = n \bmod \tau$, ω_j's are the primitive $2n$-th roots of unity ($1 \le j \le n$). Note that $\mathbf{s}(\omega_j)$ is defined as $(s_1(\omega_j), \cdots, s_{k+\ell}(\omega_j)) \in \mathbb{C}^{k+\ell}$ given the secret key $\mathbf{s} = (s_1, \cdots, s_{k+\ell}) \in \mathcal{R}^{k+\ell}$. Below, we prove that the left hand side is a bound on $\frac{n}{\tau} \cdot \|c\mathbf{s}\|_2^2$ and that this condition leads to asserting $\|c\mathbf{s}\|_2 \le \gamma\sqrt{\tau}$.

Lemma 3. *For any challenge $c \in \{0,1\}^n$ with Hamming weight τ and a secret $\mathbf{s} \in S_\eta^{k+\ell}$, the value $\|c\mathbf{s}\|_2^2$ is upper bounded by*

$$\frac{\tau}{n}\left(\tau \cdot \sum_{i=1}^{m} \overset{i\text{-}th}{\max_{0\le j<2n}} \|\mathbf{s}(\omega_j)\|_2^2 + r \cdot \overset{(m+1)\text{-}th}{\max_{0\le j<2n}} \|\mathbf{s}(\omega_j)\|_2^2\right),$$

where $m = \lfloor n/\tau \rfloor$, $r = n \bmod \tau$, and ω_j's are the primitive $2n$-th roots of unity.

Proof. We first rewrite $\|c\mathbf{s}\|_2^2$ as:

$$\|c\mathbf{s}\|_2^2 = \frac{\sum_j |c(\omega_j)|^2 \cdot \|\mathbf{s}(\omega_j)\|_2^2}{n},$$

where $\mathbf{s}(\omega_j) = (s_1(\omega_j), \cdots, s_{k+\ell}(\omega_j))$. We have that $\sum_{j=1}^n |c(\omega_j)|^2 = n\tau$ and $|c(\omega_j)|^2 = |\omega_{j,1} + \cdots + \omega_{j,\tau}|^2 \le \tau^2$. We can bound $\sum_{j=1}^n |c(\omega_j)|^2 \cdot \|\mathbf{s}(\omega_j)\|_2^2$ by rearranging the order. Let $m = \lfloor n/\tau \rfloor$ and $r = n \bmod \tau$. Then m is the maximum number of $|c(\omega_j)|^2$'s that can be τ^2. By sorting $\|\mathbf{s}(\omega_j)\|_2$ in a decreasing order,

$$\|\mathbf{s}(\omega_{\sigma(1)})\|_2 \ge \|\mathbf{s}(\omega_{\sigma(2)})\|_2 \ge \cdots \ge \|\mathbf{s}(\omega_{\sigma(n)})\|_2,$$

where σ is a permutation for the indices, we have

$$\sum_{j=1}^{n} |c(\omega_j)|^2 \cdot \|\mathbf{s}(\omega_j)\|_2^2 \le \sum_{j=1}^{m} |c(\omega_{\sigma(j)})|^2 \cdot \|\mathbf{s}(\omega_{\sigma(j)})\|_2^2 + \sum_{j=m+1}^{n} |c(\omega_{\sigma(j)})|^2 \cdot \|\mathbf{s}(\omega_{\sigma(m+1)})\|_2^2.$$

Then it reaches the maximum when the m largest $\|\mathbf{s}(\omega_j)\|_2^2$'s are multiplied with τ^2's, i.e.,

$$\begin{aligned}\sum_{j=1}^{n} |c(\omega_j)|^2 \cdot \|\mathbf{s}(\omega_j)\|_2^2 &\le \sum_{j=1}^{m} \tau^2 \cdot \|\mathbf{s}(\omega_{\sigma(j)})\|_2^2 + \Big(\sum_{j=1}^{n} |c(\omega_j)|^2 - m\tau^2\Big) \cdot \|\mathbf{s}(\omega_{\sigma(m+1)})\|_2^2 \\ &= \tau^2 \cdot \sum_{j=1}^{m} \|\mathbf{s}(\omega_{\sigma(j)})\|_2^2 + r \cdot \tau \cdot \|\mathbf{s}(\omega_{\sigma(m+1)})\|_2^2.\end{aligned}$$

This concludes the proof. □

3.2 Sampling in a Discrete Hyperball

In order to generate a hyperball uniform sample $\mathbf{y}$, we apply a rounding-and-reject strategy to the discretization of the continuous hyperball uniform sampling from Figure 1, which allows to generate rightly distributed samples. Our approach in sampling is to avoid the use of floating point arithmetic for two reasons: First, many microarchitectures do not provide floating-point units and even if so, the execution time of floating-point instructions may be data-dependent and thus unsuitable [AKM+15] for a constant-time implementation. Floating-point computation would also prohibit a masked implementation, that is protected against power side-channel attacks, because known masking techniques are only applicable to integers. And second, the required precision is higher than achievable even in IEEE double. In order to do so, we replace the continuous Gaussian sampler from Lemma 1 and instead use discrete Gaussian distributions, as we know that they approximate well continuous Gaussian distribution for large standard deviation.

Discretizing the Output. Once we obtain an "hyperball" sample, we choose to round it. Then, if the resulting sample lies too close to the border of the hyperball, we reject it. This ensures that for any possible sample, they have the same amount of pre-rounding predecessors. This also decreases the precision but the output is now discrete in a hyperball with a somewhat-smaller radius. We simply increase the starting radius to compensate.

```
$\mathbf{y} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},m}(B))$:
$\mathbf{y} \leftarrow U(\mathcal{B}_{\mathcal{R},m}(NB + \sqrt{mn}/2))$    ▷ continuous sampling in Figure 1
if $\|\lfloor \mathbf{y} \rceil\|_2 \leq NB$ then
    return $\lfloor \mathbf{y} \rceil / N$
else, restart    ▷ $\mathbf{y} \in \mathcal{B}_{(1/N)\mathcal{R},m}(B) \subset (1/N)\mathcal{R}^m$
```

Figure 3: Discrete hyperball uniform sampling

We study in the following lemma the rejection probability of this step.

Lemma 4. *Let n be the degree of $\mathcal{R}$, $M_0 \geq 1$, $B, m, N > 0$. At each iteration, the algorithm from Figure 3 succeeds with probability $\geq 1/M_0$ and the distribution of the output is $U(\mathcal{B}_{(1/N)\mathcal{R},m}(B))$ if we set*

$$N \geq \frac{\sqrt{mn}}{2B} \cdot \frac{M_0^{1/(mn)} + 1}{M_0^{1/(mn)} - 1}.$$

The proof of this lemma can be found in Appendix C.

3.3 Challenge Sampling

Challenges in HAETAE are polynomials $c \in \mathcal{R}$ with binary coefficients and exactly τ of them are nonzero. As simply hashing the message and the commitment only yields a binary string ρ, we now explain how to format it to get such a challenge. Since $n = 256$ across all three parameter sets, the challenge space has size $\binom{n}{\tau}$ exceeding the required entropy 2^{192} and 2^{225} for HAETAE-120 and HAETAE-180, respectively. To sample such challenges we rely on the (binary version of) SampleInBall algorithm from Dilithium, which we specify in the first half of Figure 4.

For HAETAE-260, however, we require 255 bits of entropy for the challenge space, which cannot be reached with the fixed Hamming weights for $n = 256$. To achieve it, we replace the challenge space by a set containing exactly half of the bitstrings of length 256.

```
SampleBinaryChallenge$_\tau(\rho)$
// for HAETAE-120 or HAETAE-180
Initialize $\mathbf{c} = c_0c_1 \dots c_{255} = 00\dots0$
for $i = 256 - \tau$ to $255$ do
    $j \hookleftarrow \{0, \dots, i\}$
    $c_i = c_j$
    $c_j = 1$
Return $\mathbf{c}$
// for HAETAE-260
Initialize $\mathbf{c} = c_0c_1 \dots c_{255} = H(\rho)$
if $\mathrm{wt}(c) > 128$ then
    $c = c \otimes 11 \cdots 1$
else if $\mathrm{wt}(c) = 128$ then
    $c = c \otimes c_0c_0 \cdots c_0$
Return $\mathbf{c}$
```

Figure 4: Challenge sampling algorithm

Specifically, we choose a set containing all elements of Hamming weight strictly less than 128 and half of the elements of Hamming weight 128, using the following algorithm. Given a 256-bits hash with Hamming weight w, do the following. If $w < 128$, we do nothing, and if $w > 128$, we flip all the bits. If $w = 128$, we decide whether to flip or not, depending on the first bit. Exactly half of all binary polynomials are reachable this way, which means that the challenge set has size 2^{255} as desired. The algorithm is specified in the second half of Figure 4.

As a side note, this means that the hash function with which we instantiate the Fiat-Shamir transform is the composition of these two steps, hashing and formatting. Looking ahead, this corresponds to steps 6 and 7 of Figure 8. Contrary to Dilithium, we do not stray away from the Fiat-Shamir transform and include the challenge c in the signature as it is no bigger than ρ when encoded.

3.4 Bimodal Hyperball Rejection Sampling

Recently, Devevey et al. [DFPS22] conducted a study of rejection sampling in the context of lattice-based Fiat-Shamir with Aborts signatures. They observe that (continuous) uniform distributions over hyperballs can be used to obtain compact signatures, with a relatively simple rejection procedure. To make masking easier, HAETAE uses (discretized) uniform distributions over hyperballs, in the bimodal context. The proof of the following lemma is available in Appendix C.

Lemma 5 (Bimodal Hyperball Rejection Sampling)**.** *Let n be the degree of $\mathcal{R}$, $c > 1$, $r, t, m > 0$, and $B \geq \sqrt{B'^2 + t^2}$. Define $M = 2(B/B')^{mn}$ and set*

$$N \geq \frac{1}{c^{1/(mn)} - 1} \frac{\sqrt{mn}}{2} \left(\frac{c^{1/(mn)}}{B'} + \frac{1}{B} \right).$$

Let $\mathbf{v} \in \mathcal{R}^m \cap \mathcal{B}_{(1/N)\mathcal{R},m}(t)$. Let $p : \mathbb{R}^m \to \{0, 1/2, 1\}$ be defined as follows

$$p(\mathbf{z}) = \begin{cases} 0 & \text{if } \|\mathbf{z}\| \geq B', \\ 1/2 & \text{else if } \|\mathbf{z} - \mathbf{v}\| < B \wedge \|\mathbf{z} + \mathbf{v}\| < B, \\ 1 & \text{otherwise.} \end{cases}$$

Then there exists $M' \leq cM$ such that the output distributions of the two algorithms from Figure 6 are identical.

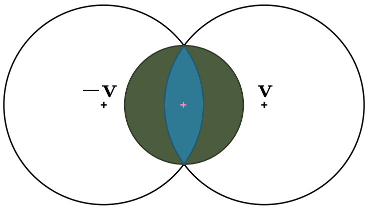

Figure 5: The HAETAE eyes

Figure 5 illustrates (the continuous version) of the rejection sampling that we consider. The black empty circles have radii equal to B and the green circle has radius B'. We sample a vector $\mathbf{z}$ uniformly inside one of the black circles (with probability 1/2 for each) and keep $\mathbf{z}$ with $p(\mathbf{z}) = 1/2$ if $\mathbf{z}$ lies in the blue zone, with probability $p(\mathbf{z}) = 1$ if it lies in the green zone, and with probability $p(\mathbf{z}) = 0$ everywhere else.

We now have all necessary ingredients in Figures 1, 3, 5, and 6 to make sure the resulting distribution of $\mathbf{z}$ is indeed uniform over the discretized hyperball. Thanks to Lemma 4 and Lemma 5, we already know the level of precision required for $\mathbf{y}$ to maintain the provable security of HAETAE.

3.5 High and Low Bits

Recall that a HAETAE signature is principally a vector $\mathbf{z}$, whose lower part is replaced with a (smaller) hint. HAETAE makes use of two different high and low bits decompositions: one helps encoding a signature while the other is used when computing a hint. Following [ETWY22], the first is helpful in the sense that if we correctly choose the number of low bits, they will be distributed almost uniformly and can then be excluded from the encoding step. The high bits on the other hand, will then follow a distribution with a very small variance and we apply the rANS encoding on them only, making it much more efficient as the size of the alphabet greatly shrunk.

The second decomposition allows to reduce the alphabet size of the resulting hint, and thus to reduce the size of its encoding.

We use the following base method of decomposing an element in high and low bits. We first recall the Euclidean division with a centered remainder.

Lemma 6. *Let $a \geq 0$ and $b > 0$. It holds that*

$$a = \left\lfloor \frac{a + b/2}{b} \right\rfloor \cdot b + (a \bmod^{\pm} b),$$

and this writing as $a = bq + r$ with $r \in [-b/2, b/2)$ is unique.

```
A(v) :
    $\mathbf{y} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},m}(B))$
    $\mathbf{b} \leftarrow U(\{0,1\})$
    $\mathbf{z} \leftarrow \mathbf{y} + (-1)^b \mathbf{v}$
    return $\mathbf{z}$ with probability $p(\mathbf{z})$, else $\perp$

B :
    $\mathbf{z} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},m}(B'))$
    return $\mathbf{z}$ with probability $1/M'$, else $\perp$
```

Figure 6: Bimodal hyperball rejection sampling

We define our decomposition for compressing the upper part of the signature.

Definition 7 (High and low bits)**.** Let $r \in \mathbb{Z}$ and α be a power of two integer. Define $r_1 = \lfloor (r + \alpha/2)/\alpha \rfloor$ and $r_0 = r \bmod^{\pm} \alpha$. Finally, define the tuple:

$$(\mathsf{LowBits}(r, \alpha), \mathsf{HighBits}(r, \alpha)) \;=\; (r_0, r_1).$$

We extend these definitions to vectors by applying them component-wise. We state that this decomposition lets us recover the original element and bound the components of the decomposition in Lemma 7. The proof is available in Appendix C.

Lemma 7. *Let α be a power of two. Let $q > 2$ be a prime with $\alpha | 2(q-1)$ and $r \in \mathbb{Z}$. Then it holds that*

$$\begin{aligned} & r = \alpha \cdot \mathsf{HighBits}(r, \alpha) + \mathsf{LowBits}(r, \alpha), \\ & \mathsf{LowBits}(r, \alpha) \in [-\alpha/2, \alpha/2), \\ & r \in [0, 2q-1] \implies \mathsf{HighBits}(r, \alpha) \in [0, (2q-1)/\alpha]\,. \end{aligned}$$

We define $\mathsf{HighBits}^{z1}(r) = \mathsf{HighBits}(r, 256)$ and $\mathsf{LowBits}^{z1}(r) = \mathsf{LowBits}(r, 256)$.

3.5.1 High and Low Bits for h

In order to produce the hint that we send instead of the lower part of $\mathbf{z}$, we could use the previous bit decomposition. However, as noted in [DKL+18, Appendix B] in a preliminary version, a slight modification allows to further reduce the entropy of the hint.

The idea is to pack the high bits in the range $[0, 2(q-1)/\alpha_{\mathsf{h}})$. This is possible if we use the range $[-\alpha_{\mathsf{h}}/2 - 2, 0)$ to represent the integers that are close to $2q-1$.

Definition 8 (High and low bits for h)**.** Let $r \in \mathbb{Z}$. Let q be a prime and $\alpha_{\mathsf{h}} | 2(q-1)$ be a power of two. Let $m = 2(q-1)/\alpha_{\mathsf{h}}$, $r_1 = \mathsf{HighBits}(r \bmod^{+} 2q, \alpha_{\mathsf{h}})$, and $r_0 = \mathsf{LowBits}(r \bmod^{+} 2q, \alpha_{\mathsf{h}})$. If $r_1 = m$, let $(r_0', r_1') = (r_0 - 2, 0)$. Else, $(r_0', r_1') = (r_0, r_1)$. We define:

$$(\mathsf{LowBits}^{\mathsf{h}}(r), \mathsf{HighBits}^{\mathsf{h}}(r)) \;=\; (r_0', r_1').$$

As before, we extend these definitions to vectors by applying them component-wise. We state that this decomposition lets us recover the original element and bound the decomposition components.

Lemma 8. *Let $r \in \mathbb{Z}$. Let q be a prime, $\alpha_{\mathsf{h}} | 2(q-1)$ be a power of two and define $m = 2(q-1)/\alpha_{\mathsf{h}}$. It holds that*

$$\begin{aligned} & r = \alpha_{\mathsf{h}} \cdot \mathsf{HighBits}^{\mathsf{h}}(r) + \mathsf{LowBits}^{\mathsf{h}}(r) \mod 2q, \\ & \mathsf{LowBits}^{\mathsf{h}}(r) \in [-\alpha/2 - 2, \alpha/2), \\ & \mathsf{HighBits}^{\mathsf{h}}(r) \in [0, m-1]\,. \end{aligned}$$

The proof of Lemma 8 is available in Appendix C.

4 The HAETAE Signature Scheme

In this section, we describe three different versions of HAETAE. As a warm-up, we give an uncompressed, un-truncated version of HAETAE, implementing the Fiat-Shamir with aborts paradigm in the bimodal hyperball-uniform setting. We then give the full description of optimized and deterministic HAETAE as we implemented it. Finally, we discuss the parts of the signing algorithm which can be pre-computed.

4.1 Uncompressed Description

As a first approach, we give a high-level, uncompressed, description of our signature scheme in Figure 7. In all of the following sections, we let $\mathbf{j} = (1, 0, \ldots, 0) \in \mathcal{R}^k$, as well as k, ℓ be two dimensions, $N > 0$ the fix-point precision and $\tau > 0$ the challenge min-entropy parameter. The parameters B, B', and B'' refer to the radii of hyperballs. Let q be an odd prime and $\alpha_{\mathsf{h}} | 2(q-1)$ is a power of two. We recall the key rejection function based on Lemma 3:

$$\mathcal{N} : \mathbf{s} \mapsto \tau \cdot \sum_{i=1}^{m} \overset{i\text{-th}}{\max_{j}} \|\mathbf{s}(\omega_j)\|_2^2 + r \cdot \overset{(m+1)\text{-th}}{\max_{j}} \|\mathbf{s}(\omega_j)\|_2^2.$$

With the parameter γ, we bound $\mathcal{N}(\mathbf{s}) \leq \gamma^2 n$, which ensures that $\|c\mathbf{s}\|_2 \leq \gamma\sqrt{\tau}$ for all $c \in \mathcal{R}_2$ satisfying $\mathrm{wt}(c) \leq \tau$. The key generation algorithm is a simplified version from Section 3.1, which removes the verification key truncation, for conceptual simplicity.

Effectively, this requires computing a complex (fix-point) 512-point FFT for each polynomial in $\mathbf{s}$, where the input is padded with zeros. The absolute value of the FFT outputs is accumulated point-wise, which yields the vector that is the input to the above function f. In fact, this can be optimized this by replacing the 512-point FFT by a 256-point one, where the j-th input coefficient x_j is multiplied by $e^{-\frac{i\pi}{256}j}$, which is functionally equivalent.

```
$\mathsf{KeyGen}(1^\lambda)$:
  $(\mathbf{A}_{\mathsf{gen}}) \leftarrow \mathcal{R}_q^{k\times(\ell-1)}$ and $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) \leftarrow S_\eta^{\ell-1} \times S_\eta^k$
  $\mathbf{b} = \mathbf{A}_{\mathsf{gen}} \cdot \mathbf{s}_{\mathsf{gen}} + \mathbf{e}_{\mathsf{gen}} \in \mathcal{R}_q^k$
  $\mathbf{A} = (-2\mathbf{b} + q\mathbf{j} |\ 2\mathbf{A}_{\mathsf{gen}} |\ 2\mathbf{Id}_k) \bmod 2q$
  $\mathbf{s} = (1, \mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}})$
  if $\mathcal{N}(\mathbf{s}) > \gamma^2 n$ then restart
  return $\mathsf{sk} = (\mathbf{A}, \mathbf{s})$, $\mathsf{vk} = \mathbf{A}$

$\mathsf{Sign}(\mathsf{sk}, M)$:
  $\mathbf{y} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},(k+\ell)}(B))$
  $\mathbf{w} \leftarrow \mathbf{A}\lfloor\mathbf{y}\rceil$
  $c = H(\mathbf{w}, M) \in \mathcal{R}_2$
  $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2) = \mathbf{y} + (-1)^b c\mathbf{s}$ for $b \leftarrow U(\{0,1\})$
  if $\|\mathbf{z}\|_2 \geq B'$ then restart
  else if $\|2\mathbf{z} - \mathbf{y}\|_2 < B$ then restart with probability 1/2
  return $\sigma = (\lfloor\mathbf{z}\rceil, c)$

$\mathsf{Verify}(\mathsf{vk}, M, \sigma = (\mathbf{z}, c))$:
  $\tilde{\mathbf{w}} = \mathbf{A}\mathbf{z} - qc\mathbf{j} \bmod 2q$
  return $(c = H(\tilde{\mathbf{w}}, M)) \wedge \left(\|\mathbf{z}\| < B + \frac{\sqrt{n(k+\ell)}}{2}\right)$
```

Figure 7: Uncompressed description of HAETAE .

4.2 Specification of HAETAE

We now give the full description of the signature scheme HAETAE in Figure 8 with the following building blocks:

- Hash function H_{gen} for generating the seeds and hashing the messages,

- Hash function H for signing, returning a seed ρ for sampling a challenge,
- Extendable output function $\mathsf{expandA}$ for deriving $\mathbf{a}_{\mathsf{gen}}$ and $\mathbf{A}_{\mathsf{gen}}$ from $\mathsf{seed}_{\mathbf{A}}$,
- Extendable output function $\mathsf{expandS}$ for deriving $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) \in S_\eta^{\ell-1} \times S_\eta^k$ from $\mathsf{seed}_{\mathsf{sk}}$ and $\mathsf{counter}_{\mathsf{sk}}$,
- Extendable output function $\mathsf{expandYbb}$ for deriving $\mathbf{y}$, b and b' from seed_{ybb} and $\mathsf{counter}$,

The above building blocks can be implemented with symmetric primitives.

Note that at Step 6 of the Verify algorithm, the division by 2 is well-defined as the operand is even.

4.3 Theoretical Analysis

In this section, we prove the theoretical correctness and security of HAETAE.

4.3.1 Correctness and Runtime

Correctness. Before showing the correctness of the HAETAE scheme, we prove the following intermediary result.

Lemma 9. *We borrow the notations from Figure 8. If we run* $\mathsf{Verify}(\mathsf{vk}, M, \sigma)$ *on the signature* σ *returned by* $\mathsf{Sign}(\mathsf{sk}, M)$ *for an arbitrary message* M *and an arbitrary key-pair* $(\mathsf{sk}, \mathsf{vk})$ *returned by* $\mathsf{KeyGen}(1^\lambda)$*, then the following relations hold:*

1) $\mathbf{w}_1 = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w})$*,*

2) $w'\mathbf{j} = \mathsf{LSB}(\lfloor y_0 \rceil) \cdot \mathbf{j} = \mathsf{LSB}(\mathbf{w}) = \mathsf{LSB}(\mathbf{w} - 2\lfloor \mathbf{z}_2 \rceil)$*,*

3) $2\lfloor \mathbf{z}_2 \rceil - 2\tilde{\mathbf{z}}_2 = \mathsf{LowBits}^{\mathsf{h}}(\mathbf{w}) - \mathsf{LSB}(\mathbf{w})$ *assuming* $B' + \alpha_{\mathsf{h}}/4 + 1 \leq B'' < q/2$*.*

Proof. Let $m = 2(q-1)/\alpha_{\mathsf{h}}$. Let us prove the first statement. By definition of $\mathbf{h}$, it holds that $\mathbf{w}_1 = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}) \mod m$. However, the latter part of the equality already lies in $[0, m-1]$ by Lemma 8. The first part lies in the same range as we reduce $\mod^+ m$. Hence, the equality stands over $\mathbb{Z}$ too.

We move on to the second statement. By considering only the first component of $\mathbf{z} = \mathbf{y} + (-1)^b c\mathbf{s}$, we obtain, modulo 2:

$$\tilde{z}_0 = \lfloor z_0 \rceil = \lfloor y_0 \rceil + (-1)^b c = \lfloor y_0 \rceil + c.$$

Moreover, considering everywhere a 2 appears in the definition of $\mathbf{A}$, we obtain that

$$\mathbf{w} = \mathbf{A}_1 \lfloor \mathbf{z}_1 \rceil - qc\mathbf{j} = (\lfloor z_0 \rceil - c)\mathbf{j} \bmod 2.$$

For the last statement, let us use the two preceding results. In particular, we note

$$\mathbf{w}_1 \cdot \alpha_{\mathsf{h}} + w'\mathbf{j} = \mathbf{w} - \mathsf{LowBits}^{\mathsf{h}}(\mathbf{w}) + \mathsf{LSB}(\mathbf{w}).$$

We note that the last two elements have same parity, as the former one has the same parity as $\mathsf{LowBits}(\mathbf{w}, \alpha_{\mathsf{h}})$. By Lemma 8 their sum has infinite norm $\leq \alpha_{\mathsf{h}}/2 + 2$. Hence from its definition, it holds that

$$2\tilde{\mathbf{z}}_2 = 2\lfloor \mathbf{z}_2 \rceil - \mathsf{LowBits}^{\mathsf{h}}(\mathbf{w}) + \mathsf{LSB}(\mathbf{w}) \mod^{\pm} 2q.$$

Finally, this holds over the integers as the right-hand side has infinite norm at most $2B' + \alpha_{\mathsf{h}}/2 + 2 < q$. □

```
KeyGen($1^\lambda$):
 1: $\mathsf{seed} \leftarrow \{0,1\}^{\rho_0}$                                   ▷ KeyGen for $d = 1$
 2: $(\mathsf{seed}_{\mathbf{A}}, \mathsf{seed}_{\mathsf{sk}}, K) = H_{\mathsf{gen}}(\mathsf{seed})$
 3: $(\mathbf{a}_{\mathsf{gen}} |\ \mathbf{A}_{\mathsf{gen}}) := \mathsf{expandA}(\mathsf{seed}_{\mathbf{A}}) \in \mathcal{R}_q^{k \times \ell}$
 4: $\mathsf{counter}_{\mathsf{sk}} = 0$
 5: $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) := \mathsf{expandS}(\mathsf{seed}_{\mathsf{sk}}, \mathsf{counter}_{\mathsf{sk}})$
 6: $\mathbf{b} = \mathbf{a}_{\mathsf{gen}} + \mathbf{A}_{\mathsf{gen}} \cdot \mathbf{s}_{\mathsf{gen}} + \mathbf{e}_{\mathsf{gen}} \in \mathcal{R}_q^k$
 7: $(\mathbf{b}_0, \mathbf{b}_1) = (\mathsf{LowBits}^{\mathsf{vk}}(\mathbf{b}), \mathsf{HighBits}^{\mathsf{vk}}(\mathbf{b}))$
 8: $\mathbf{A} = (2(\mathbf{a}_{\mathsf{gen}} - 2\mathbf{b}_1) + q\mathbf{j} |\ 2\mathbf{A}_{\mathsf{gen}} |\ 2\mathbf{Id}_k) \bmod 2q$
 9: $\mathbf{s} = (1, \mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}} - \mathbf{b}_0)$
10: if $\mathcal{N}(\mathbf{s}) > \gamma^2 n$ then $\mathsf{counter}_{\mathsf{sk}}{+}{+}$ and Go to 5
11: return $\mathsf{sk} = (\mathbf{s}, K)$, $\mathsf{vk} = (\mathsf{seed}_{\mathbf{A}}, \mathbf{b}_1)$

Sign($\mathsf{sk}, M$):
 1: $\mu = H_{\mathsf{gen}}(\mathsf{seed}_{\mathbf{A}}, \mathbf{b}_1, M)$
 2: $\mathsf{seed}_{ybb} = H_{\mathsf{gen}}(K, \mu)$
 3: $\mathsf{counter} = 0$
 4: $(\mathbf{y}, b, b') := \mathsf{expandYbb}(\mathsf{seed}_{ybb}, \mathsf{counter})$
 5: $\mathbf{w} \leftarrow \mathbf{A} \lfloor \mathbf{y} \rceil$
 6: $\rho = H(\mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}), \mathsf{LSB}(\lfloor y_0 \rceil), \mu)$
 7: $c = \mathsf{SampleBinaryChallenge}_\tau(\rho)$
 8: $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2) = \mathbf{y} + (-1)^b c\mathbf{s}$
 9: $\mathbf{h} = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}) - \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w} - 2\lfloor \mathbf{z}_2 \rceil) \bmod^{+} \frac{2(q-1)}{\alpha_{\mathsf{h}}}$
10: if $\|\mathbf{z}\|_2 \geq B'$ then
11:     $\mathsf{counter}{+}{+}$ and Go to 4
12: else if $\|2\mathbf{z} - \mathbf{y}\|_2 < B \ \wedge\ b' = 0$ then
13:     $\mathsf{counter}{+}{+}$ and Go to 4
14: else
15:     $x = \mathsf{Encode}(\mathsf{HighBits}^{z1}(\lfloor \mathbf{z}_1 \rceil))$
16:     $\mathbf{v} = \mathsf{LowBits}^{z1}(\lfloor \mathbf{z}_1 \rceil)$
17:     return $\sigma = (x, \mathbf{v}, \mathsf{Encode}(\mathbf{h}), c)$

Verify($\mathsf{vk}, M, \sigma = (x, \mathbf{v}, h, c)$):
 1: $\tilde{\mathbf{z}}_1 = \mathsf{Decode}(x) \cdot 256 + \mathbf{v}$ and $\tilde{\mathbf{h}} = \mathsf{Decode}(h)$
 2: $(\mathbf{a}_{\mathsf{gen}} |\ \mathbf{A}_{\mathsf{gen}}) = \mathsf{expandA}(\mathsf{seed}_{\mathbf{A}})$
 3: $\mathbf{A}_1 = (2(\mathbf{a}_{\mathsf{gen}} - 2\mathbf{b}_1) + q\mathbf{j} |\ 2\mathbf{A}_{\mathsf{gen}}) \bmod 2q$
 4: $\mathbf{w}_1 = \tilde{\mathbf{h}} + \mathsf{HighBits}^{\mathsf{h}}(\mathbf{A}_1 \tilde{\mathbf{z}}_1 - qc\mathbf{j}) \bmod^{+} \frac{2(q-1)}{\alpha_{\mathsf{h}}}$
 5: $w' = \mathsf{LSB}(\tilde{z}_0 - c)$
 6: $\tilde{\mathbf{z}}_2 = (\mathbf{w}_1 \cdot \alpha_{\mathsf{h}} + w'\mathbf{j} - (\mathbf{A}_1 \tilde{\mathbf{z}}_1 - qc\mathbf{j}))/2 \bmod^{\pm} q$
 7: $\tilde{\mathbf{z}} = (\tilde{\mathbf{z}}_1, \tilde{\mathbf{z}}_2)$
 8: $\tilde{\mu} = H_{\mathsf{gen}}(\mathsf{seed}_{\mathbf{A}}, \mathbf{b}_1, M)$
 9: $\tilde{\rho} = H(\mathbf{w}_1, w', \tilde{\mu})$
10: return $(c = \mathsf{SampleBinaryChallenge}_\tau(\tilde{\rho})) \wedge (\|\tilde{\mathbf{z}}\| < B'')$
```

Figure 8: Full description of deterministic HAETAE. The KeyGen algorithm is slightly different for $d = 0$ (HAETAE-260), which do not truncate $\mathbf{b}$. See Section 3.1.1 for details.

Theorem 2 (Completeness)**.** *Assume that* $B'' = B' + \sqrt{n(k+\ell)}/2 + \sqrt{nk} \cdot (\alpha_{\mathsf{h}}/4 + 1) < q/2$. *Then the signature schemes of Figure 8 is complete, i.e., for every message M and every key-pair* $(\mathsf{sk}, \mathsf{vk})$ *returned by* $\mathsf{KeyGen}(1^\lambda)$*, we have:*

$$\mathsf{Verify}(\mathsf{vk}, M, \mathsf{Sign}(\mathsf{sk}, M)) = 1.$$

Proof. We use the notations of the algorithms. The first and second equations from Lemma 9 state that $\rho = \tilde{\rho}$ and thus

$$c = \mathsf{SampleBinaryChallenge}_{\tau}(\tilde{\rho}).$$

On the other hand, we use the last equation from the same lemma to bound the size of $\tilde{\mathbf{z}}$. We have:

$$\begin{aligned}
\|\tilde{\mathbf{z}}\| &\leq \|\mathbf{z}\| + \|\mathbf{z} - \lfloor \mathbf{z} \rceil\| + \|\lfloor \mathbf{z} \rceil - \tilde{\mathbf{z}}\| \\
&\leq B' + \sqrt{n(k+\ell)} \cdot \|\mathbf{z} - \lfloor \mathbf{z} \rceil\|_\infty + \|\lfloor \mathbf{z}_2 \rceil - \tilde{\mathbf{z}}_2\| \\
&\leq B' + \frac{\sqrt{n(k+\ell)}}{2} + \sqrt{nk} \cdot \|\mathsf{LowBits}^{\mathsf{h}}(\mathbf{w})\|_\infty \\
&\leq B' + \frac{\sqrt{n(k+\ell)}}{2} + \sqrt{nk} \cdot \left(\frac{\alpha_{\mathsf{h}}}{4} + 1\right).
\end{aligned}$$

The definition of B'' implies that the scheme is correct.

□

Runtime. A non-trivial task is analyzing the runtime of Fiat-Shamir with aborts schemes. As shown in [DFPS23, Section 6.1], degenerate cases exist where the signing algorithm does not terminate. In particular, it is impossible, even in the ROM, to show that the generic signing algorithm has polynomial expected runtime. However, in the case of HAETAE, all starting points $\mathbf{y}$ lying in the Euclidean ball of radius $B' - \gamma\sqrt{\tau}$ have probability at least $1/2$ of being accepted in the end, whatever shift $\pm\mathbf{s}c$ is added, whatever instance of the hash function H is chosen. This gives rise to the following two points.

- With M' defined as in Lemma 5, the probability of doing more than i iterations is bounded from above by $(1 - 1/M')^i + 2^{-\alpha} M'^3$, where α is the commitment min-entropy.
- The signing algorithm has a finite expected runtime.

These are insufficient to conclude that the number of iterations is M' on average, but this is nonetheless what happens for our choice of parameters.

4.3.2 Theoretical Security

```
Compress($\mathbf{z}, c$):
if $\mathbf{z} = \perp$ then
    return $(\perp, \perp, \perp)$
$\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2) = \mathbf{y} + (-1)^b c\mathbf{s}$
$\mathbf{h} = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w})$
    $-\mathsf{HighBits}^{\mathsf{h}}(\mathbf{w} - 2\lfloor \mathbf{z}_2 \rceil) \bmod^+ \frac{2(q-1)}{\alpha_{\mathsf{h}}}$
$x = \mathsf{Encode}(\mathsf{HighBits}^{z1}(\lfloor \mathbf{z}_1 \rceil))$
$\mathbf{v} = \mathsf{LowBits}^{z1}(\lfloor \mathbf{z}_1 \rceil)$
return $(x, \mathbf{v}, \mathsf{Encode}(\mathbf{h}))$
```

```
Decompress($\mathbf{A}, (x, \mathbf{v}, h), c$):
if $(x, \mathbf{v}, h) = (\perp, \perp, \perp)$ then
    return $\perp$
$\tilde{\mathbf{z}}_1 = \mathsf{Decode}(x) \cdot 256 + \mathbf{v}$
$\tilde{\mathbf{h}} = \mathsf{Decode}(h)$
$\mathbf{A} = (\mathbf{A}_1 | \; 2\mathbf{I}) \bmod 2q$
$\mathbf{w}_1 = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{A}_1 \tilde{\mathbf{z}}_1 - qc\mathbf{j})$
    $+\tilde{\mathbf{h}} \bmod^+ \frac{2(q-1)}{\alpha_{\mathsf{h}}}$
$w' = \mathsf{LSB}(\tilde{z}_0 - c)$
$\tilde{\mathbf{w}} = \mathbf{A}_1 \tilde{\mathbf{z}}_1 - qc\mathbf{j}$
$\tilde{\mathbf{z}}_2 = (\mathbf{w}_1 \cdot \alpha_{\mathsf{h}} + w'\mathbf{j} - \tilde{\mathbf{w}}) / 2 \bmod^{\pm} q$
return $\tilde{\mathbf{z}} = (\tilde{\mathbf{z}}_1, \tilde{\mathbf{z}}_2)$
```

Figure 9: Compression and decompression algorithms.

We analyze the security of HAETAE by following standard arguments and proof techniques. We start by detailing the underlying canonical identification scheme that makes up HAETAE and give their properties in Lemma 10. We then show that HAETAE is UF-NMA secure in Lemma 11. Finally, we put everything together and give the security bound in Theorem 3 based on the reduction from [BBD+23]. Reminders on canonical identification schemes and the Fiat-Shamir with aborts transform can be found in Appendix F.

In Figure 10, we first describe in Figure 10a an "uncompressed" version of the identification scheme, as well as the HAETAE canonical identfication scheme in Figure 10b. The latter uses the Compress and Decompress functions, as described in Figure 9. We omit the instance generator for both CID, as it is exactly the key generation algorithm from HAETAE, except that do not do truncation for simplicity. Adding it can be done by noting that $\mathbf{a}_{\mathsf{gen}}$ gives uniform low bits that would otherwise be missing.

$P(\mathbf{A}, \mathbf{s})$		$V(\mathbf{A})$
$\mathbf{y} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},(k+\ell)}(B))$ $\mathbf{w} = \mathbf{A}\lfloor \mathbf{y} \rceil \bmod 2q$		
	$\xrightarrow{\mathbf{w}}$	$c \leftarrow U(\mathcal{C})$
	$\xleftarrow{c}$	
$b \leftarrow U(\{0,1\})$ $\mathbf{z} = \mathbf{y} + (-1)^b c\mathbf{s}$ If $p(\mathbf{z})$, restart		
Else	$\xrightarrow{\lfloor \mathbf{z} \rceil}$	Accept if $\|\lfloor \mathbf{z} \rceil\| \leq B''$ and $\mathbf{A}\mathbf{z} - qc\mathbf{j} = \mathbf{w} \bmod 2q$

(a) Uncompressed HAETAE

$P(\mathbf{A}, \mathbf{s})$		$V(\mathbf{A})$
$\mathbf{y} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},(k+\ell)}(B))$ $\mathbf{w} = \mathbf{A}\lfloor \mathbf{y} \rceil \bmod 2q$ $\mathbf{w}_1 = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w})$		
$w_0 = \mathsf{LSB}(\lfloor y_0 \rceil)$	$\xrightarrow{w_0, \mathbf{w}_1}$	$c \leftarrow U(\mathcal{C})$
	$\xleftarrow{c}$	
$b \leftarrow U(\{0,1\})$ $\mathbf{z} = \mathbf{y} + (-1)^b c\mathbf{s}$ If $p(\mathbf{z})$, restart		
Else $\sigma = \mathsf{Compress}(\lfloor \mathbf{z} \rceil, c)$	$\xrightarrow{\sigma}$	$\tilde{\mathbf{z}} = \mathsf{Decompress}(\mathbf{A}, \sigma, c)$ Accept if $\|\tilde{\mathbf{z}}\| \leq B''$ and $\mathbf{A}\tilde{\mathbf{z}} - qc\mathbf{j} = \mathbf{w}_1 \alpha_{\mathsf{h}} + w_0 \mathbf{j} \bmod 2q$

(b) Compressed HAETAE

Figure 10: Underlying canonical identification schemes.

We sum up the properties of the CID from Figure 10b.

Lemma 10. *The HAETAE canonical identification scheme from Figure 10b satisfies the following properties.*

Commitment Min-entropy. *The commitment $(w_0, \mathbf{w}_1)$ has min-entropy $\geq n$.*

paHVZK. *Assuming $B^2 \geq B'^2 + \gamma^2 \tau$, the HAETAE CID satisfies the paHVZK property with the simulator described in Figure 11.*

```
$\mathsf{Sim}_{\not{\mathcal{L}}}(\mathsf{vk}, c)$ :
$\mathbf{z} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},k+\ell}(B'))$
$\mathbf{w} \leftarrow \mathbf{A}\lfloor \mathbf{z} \rceil - qc\mathbf{j}$
return $(\mathbf{w}, c, \mathbf{z})$
```

Figure 11: Uncompressed HAETAE simulator.

Computational Unique Response. *Let $\mathcal{A}$ be an adversary against the CUR property of the scheme. There exist two adversaries $\mathcal{B}$ and $\mathcal{B}'$ with roughly the same runtime such that:*

$$\mathsf{Adv}^{\mathsf{CUR}}_{\mathsf{HAETAE}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{MLWE}}_{n,q,k,\ell-1,\eta}(\mathcal{B}) + \mathsf{Adv}^{\mathsf{MSIS}}_{n,q,k,\ell,2B''}(\mathcal{B}').$$

Proof. Let us prove each point.

Commitment Min-entropy. Recall that $w_0 = LSB(\lfloor y_0 \rceil)$ is a binary polynomial of degree $\leq n$. As the coefficients of y_0 are large and distributed following the distribution of a subset of a uniform-hyperball, we expect the coefficients of w_0 to be roughly uniform, hence the lower bound on the commitment min-entropy.

paHVZK. We first prove that the CID from Figure 10a satisfies paHVZK. Consider the simulator described in Figure 11.

The distribution of real non-aborting transcripts and the output distribution of $\mathsf{Sim}_{\not{\mathcal{L}}}(\mathsf{vk}, c)$ are identical under the condition on the radii. Indeed, by conditioning the distribution of the output of algorithms $\mathcal{A}$ and $\mathcal{B}$ from Lemma 5 on not being $\perp$, the conditional (on $\mathbf{A}$ and c) distributions of $\mathbf{z}$ in the two cases are identical, and we note that $(\mathbf{w}, c, \mathbf{z})$ is a deterministic function of $(\mathbf{A}, c, \mathbf{z})$.

Finally, given an uncompressed transcript $(\mathbf{w}, c, \mathbf{z})$ for the Figure 10a CID, we note that $((\mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}), \mathsf{LSB}(w_0)), c, \mathsf{Compress}(\mathbf{z}, c))$ is a valid transcript for the HAETAE CID. Indeed, note that $w_0 = \lfloor y_0 \rceil \bmod 2$ by definition of $\mathbf{A}$. Moreover, applying this deterministic transform turns the distribution of real transcripts for the uncompressed CID to the distribution of real transcripts for the compressed CID. By an immediate reduction, the HAETAE CID also satisfies paHVZK.

CUR. Given $((\mathbf{w}_1, w_0), c, \sigma, \sigma')$ such that $\mathbf{A}\mathbf{z} - qc\mathbf{j} = \mathbf{w}_1\alpha_{\mathsf{h}} + w_0\mathbf{j} = \mathbf{A}\mathbf{z}' - qc\mathbf{j} \bmod 2q$, we get $\mathbf{A}(\mathbf{z} - \mathbf{z}') = 0 \bmod 2q$, which we can reduce mod$q$:

$$(-\mathbf{b}|\mathbf{A}_{\mathsf{gen}}|\mathbf{Id}_k)(\mathbf{z} - \mathbf{z}') = 0 \bmod q.$$

We prove that $\mathbf{z} \neq \mathbf{z}' \bmod q$. Let $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_1)$ and $\mathbf{z}' = (\mathbf{z}'_1, \mathbf{z}'_2)$. Let $\sigma = (x, \mathbf{v}, h)$ and $\sigma' = (x', \mathbf{v}', h)$. If $\mathbf{z}_1 = \mathbf{z}'_1$ then $x = x'$ and $\mathbf{v} = \mathbf{v}'$ due to the encoding being unique. This means that $h \neq h'$, which in turn will give two different values for $\mathbf{z}_2$ and $\mathbf{z}'_2$ as the other values composing $\mathbf{z}_2$ and $\mathbf{z}'_2$ are identical.

However, $(\mathbf{A}_{\mathsf{gen}}, \mathbf{b})$ is a $\mathsf{MLWE}_{n,q,k,\ell-1,\eta}$ instance instead of being uniform. Both adversaries $\mathcal{B}$ and $\mathcal{B}'$ can run $\mathcal{A}$ by simulating a random oracle in the case of $\mathcal{B}$ and by forwarding the queries to its own oraclein the case of $\mathcal{B}'$. Then, if $\mathcal{A}$ was successful, adversary $\mathcal{B}$ outputs "MLWE" while $\mathcal{B}'$ computes the MSIS solution as described above, and if $\mathcal{A}$ is unsuccessful, adversary $\mathcal{B}$ outputs "unif" while $\mathcal{B}'$ aborts. By the triangular inequality, we get the desired inequality.

□

We reduce the UF-NMA security of the signature scheme to the BimodalSelfTargetMSIS problem.

Lemma 11 (UF-NMA security)**.** *Let $\mathcal{A}$ be an adversary against the* UF-NMA *security of* HAETAE*. There exist two adversaries $\mathcal{B}$ and $\mathcal{B}'$ with essentially the same runtime as $\mathcal{A}$*

such that:

$$\mathsf{Adv}^{\mathsf{UF\text{-}NMA}}_{\mathsf{HAETAE}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{MLWE}}_{n,q,k,\ell-1,\eta}(\mathcal{B}) + \mathsf{Adv}^{\mathsf{BimodalSelfTargetMSIS}}_{H,n,q,k,\ell,B''}(\mathcal{B}').$$

Proof. Let us first describe $\mathcal{B}'$. On input the public matrix $\mathbf{A}$, it calls $\mathcal{A}$ on it. The adversary $\mathcal{B}'$ defines the random oracle $H' : (\mathbf{w}_1, w_0, \mu) \mapsto H(\mathbf{w}_1\alpha_{\mathsf{h}} + w_0\mathbf{j} \bmod 2q, \mu)$, to take care of the compression that does not appear in the BimodalSelfTargetMSIS problem. This can be simulated by $\mathcal{B}$ by applying the right function on the queries of $\mathcal{A}$ and forwarding it to its own random oracle H.

When $\mathcal{A}$ outputs a valid forgery $\sigma^* = ((x, \mathbf{v}, h), c)$ for some message μ^*, adversary $\mathcal{B}'$ runs $\mathsf{Decompress}(\mathbf{A}, (x, \mathbf{v}, h), c)$ to get a vector $\mathbf{z}$ such that $\|\mathbf{z}\| \leq B''$ and $H'(\mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}^*), \mathsf{LSB}(w_0^*), \mu) = c$, where we let $\mathbf{w}^* = \mathbf{Az} - qc\mathbf{j} \bmod 2q$. By construction $\mathbf{w}^* = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}^*)\alpha_{\mathsf{h}} + \mathsf{LSB}(w_0^*)\mathbf{j} \bmod 2q$, meaning that $c = H(\mathbf{w}^*, \mu)$. Then $(\mathbf{z}, c, \mu^*)$ is exactly a solution to the $\mathsf{BimodalSelfTargetMSIS}_{H,n,q,k,\ell,B''}$ problem.

Note that $\mathcal{A}$ is called on a "lossy" instance of HAETAE, where the verification key $\mathbf{b}$ does not have a signing key associated and is actually uniform. This difference in setting exactly corresponds to the $\mathsf{MLWE}_{n,q,k,\ell-1,\eta}$ problem. As such the design of $\mathcal{B}$ is as follows: on input $(\mathbf{A}_{\mathsf{gen}}, \mathbf{b})$, it puts together the verification key $\mathbf{A}$ and calls $\mathcal{A}$ on it. It can simulate a random oracle for it, and outputs "MLWE" if $\mathcal{A}$ is successful at forging, "unif" otherwise. □

We combine the previous results with Theorem 4 to get the following security bound for HAETAE.

Theorem 3. *Let $\mathcal{A}$ be an adversary against the* UF-CMA *security of HAETAE making Q_s signature queries and Q_h hash queries. Let α be the commitment min-entropy of HAETAE. Let $B^2 \geq B'^2 + \gamma^2\tau$. There exist two adversaries $\mathcal{B}$ and $\mathcal{B}'$ such that*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{UF\text{-}CMA}}_{\mathsf{HAETAE}}(\mathcal{A}) \leq{}& \mathsf{Adv}^{\mathsf{MLWE}}_{n,q,k,\ell-1,\eta}(\mathcal{B}) + \mathsf{Adv}^{\mathsf{BimodalSelfTargetMSIS}}_{H,n,q,k,\ell,B''}(\mathcal{B}') \\ &+ \frac{2^{-\alpha/2+1}Q_s}{1-\beta}\sqrt{Q_h + 1 + \frac{Q_s}{1-\beta}} + 2^{-\alpha/2+1}(Q_h+1)\sqrt{\frac{Q_s}{1-\beta}}.\end{aligned} \tag{4.1}$$

If $\mathcal{A}$ is an adversary against the sUF-CMA *security of HAETAE, then there exist two more adversaries $\mathcal{B}''$ and $\mathcal{B}'''$ such that the previous bound holds by adding the extra term $\mathsf{Adv}^{\mathsf{MLWE}}_{n,q,k,\ell-1,\eta}(\mathcal{B}'') + \mathsf{Adv}^{\mathsf{MSIS}}_{n,q,k,\ell,2B''}(\mathcal{B}''')$.*

4.4 HAETAE with Pre-computation

We observe that in the randomized signing process of HAETAE, many operations do not depend on the message M, and some do not even depend on the signing key. This enables efficient "offline" procedures, i.e., precomputations that speed up the "online" phase.

Specifically, there are two levels of offline signing that can be applied to randomized HAETAE:

1. **Generic.** If neither the message M nor the signing key is chosen in advance, it is still possible to perform hyperball sampling. This removes the most time-consuming operation from the online phase.

2. **Designated signing key.** Here, only the message M is unknown during offline signing, while the signing key is fixed. This allows us to perform even more precomputations by using only the verification key, as shown in Figure 12. Most notably, there is no longer a matrix-vector multiplication in the online phase.

We showcase the offline and online parts of the (randomized) version of HAETAE in Figure 12.

```
$\mathsf{Sign}^{\mathsf{offline}}(\mathsf{vk})$:
1: $(\mathbf{a}_{\mathsf{gen}} |\ \mathbf{A}_{\mathsf{gen}}) = \mathsf{expandA}(\mathsf{seed}_{\mathbf{A}})$
2: $\mathbf{A}_1 = (2(\mathbf{a}_{\mathsf{gen}} - 2\mathbf{b}_1) + q\mathbf{j} |\ 2\mathbf{A}_{\mathsf{gen}}) \bmod 2q$
3: $\mathsf{List} = ()$
4: for $\mathsf{iter}$ in $[L]$ do
5:     $\mathbf{y} \leftarrow U(\mathcal{B}_{(1/N)\mathcal{R},(k+\ell)}(B))$
6:     $\mathbf{w} = \mathbf{A}\lfloor \mathbf{y} \rceil$
7:     $\mathbf{w}_1 = \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w})$
8:     $\mathsf{List.append}(\mathbf{y}, \mathbf{w}, \mathbf{w}_1, \mathsf{LSB}(\lfloor y_0 \rceil))$
9: return $\mathsf{List}$

$\mathsf{Sign}^{\mathsf{online}}(\mathsf{sk}, \mathsf{List}, M)$:
1: $\mu = H_{\mathsf{gen}}(\mathsf{seed}_{\mathbf{A}}, \mathbf{b}_1, M)$
2: $\mathsf{tuple} = (\mathbf{y}, \mathbf{w}, \mathsf{tuple}_3, \mathsf{tuple}_4) \leftarrow \mathsf{List}$
3: $\mathsf{List.delete}(\mathsf{tuple})$
4: $c = \mathsf{SampleBinaryChallenge}_{\tau}(H(\mathsf{tuple}_3, \mathsf{tuple}_4, \mu))$
5: $(b, b') \leftarrow \{0,1\}^2$
6: $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2) = \mathbf{y} + (-1)^b c\mathbf{s}$
7: $\mathbf{h} = \mathsf{tuple}_3 - \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w} - 2\lfloor \mathbf{z}_2 \rceil) \bmod^+ \frac{2(q-1)}{\alpha_{\mathsf{h}}}$
8: if $\|\mathbf{z}\|_2 \geq B'$ then Go to 2
9: else if $\|2\mathbf{z} - \mathbf{y}\|_2 < B \ \wedge\ b' = 0$ then Go to 2
10: else
11:     $x = \mathsf{Encode}(\mathsf{HighBits}^{z1}(\lfloor \mathbf{z}_1 \rceil))$
12:     $\mathbf{v} = \mathsf{LowBits}^{z1}(\lfloor \mathbf{z}_1 \rceil)$
13:     return $\sigma = (x, \mathbf{v}, \mathsf{Encode}(\mathbf{h}), c)$
```

Figure 12: Randomized, on/off-line signing.

4.5 Parameter Sets

To choose parameters reaching the desired NIST security levels, we estimated the costs of practical attacks, as in Dilithium, Falcon, and many other NIST-submitted schemes. In particular, our methodology is directly inspired from the one used in Dilithium, and we sum it up in the following.

Looking at Equation 4.1, we note that we have highly underestimated the commitment min-entropy of the scheme in Lemma 10 as we ignored its biggest part $\mathbf{w}_1$, and we choose to ignore terms on the second line as we are confident that they are small enough. We first evaluate the cost of attacks on MLWE, i.e. key-recovery attacks. Second, to evalute the cost of forgery attacks, i.e. attacks on BimodalSelfTargetMSIS, we use the fact that the only known way to solve BimodalSelfTargetMSIS is to solve MSIS. Heuristically, the hash function is not aware of the algebraic structure of its input, and the random oracle assumption that c is uniform and independent from the input is sound. Thus, an adversary has no choice but to choose some $\mathbf{w}$, hash its high and low bits along with some message, and try to compute a short preimage of $\mathbf{w} - qc\mathbf{j} \bmod 2q$. If the adversary succeeds, the preimage is in particular an $\mathsf{MSIS}_{n,q,k,\ell+1,\sqrt{B''^2+1}}$ solution for the matrix $[\mathbf{w}|\mathbf{A}]$. Finally, we evaluate the cost of MSIS with a bound twice as loose to evaluate the strong unforgeability of the HAETAE scheme. Our cryptanalysis of these problems follows the approach taken in Dilithium. We provide a modification of their security estimation script[1], where we also updated the cost of quantum attacks following recent works [CL21]. These estimations follow the CoreSVP

[1] Available in the submission package.

Table 1: HAETAE parameters sets. Hardness is measured with the Core-SVP methodology and a refined analysis is given for LWE. The numbers in parenthesis for SIS are for the strong unforgeability property.

Security		120	180	260
n	Degree of $\mathcal{R}$ (2.1)	256	256	256
(k, ℓ)	Dimensions of $\mathbf{z}_2$, $\mathbf{z}_1$ (4.2)	(2,4)	(3,6)	(4,7)
q	Modulus for MLWE & MSIS (2.3)	64513	64513	64513
η	Range of sk coefficients (2.1)	1	1	1
τ	Weight of c (3.3)	58	80	128
γ	sk rejection parameter (3.1)	48.858	57.707	55.13
	Resulting key acceptance rate (3.1)	0.1	0.1	0.1
d	Truncated bits of vk (3.1)	1	1	0
M	Expected # of repetitions (3.4)	6.0	5.0	6.0
B	$\mathbf{y}$ radius (3.4)	9846.02	18314.98	22343.66
B'	Rejection radius (3.4)	9838.98	18307.70	22334.95
B''	Verify radius (4.2)	12777.52	21906.65	24441.49
α	$\mathbf{z}_1$ compression factor (3.5)	256	256	256
α_h	$\mathbf{h}$ compression factor (3.5)	512	512	256
Forgery: SIS Hardness (Core-SVP)				
BKZ block-size b (GSA)		409 (333)	617 (512)	878 (735)
Classical Core-SVP		119 (97)	180 (149)	256 (214)
Quantum Core-SVP		105 (85)	158 (131)	225 (188)
Key-recovery: LWE Hardness (Core-SVP and refined)				
BKZ block-size b (GSA)		428	810	988
Classical Core-SVP		125	236	288
Quantum Core-SVP		109	208	253
BKZ block-size b (simulation)		439	834	1019
$\log_2$ Classical Gates		159	270	322
$\log_2$ Classical Memory		99	177	214

approach. In the case of MLWE, we also give refined estimates, computed in a branch[1] of the leaky LWE estimator from [DDGR20].

We propose three different parameter sets with varying security levels, where we prioritize low signature and verification key sizes over faster execution time. The parameter choices are versatile, adaptable and allow size vs. speed trade-offs at consistent security levels. For example at cost of larger signatures, a smaller repetition rate M is possible and thus a faster signing process. This versatility is a notable advantage over Falcon and Mitaka.

Like in Dilithium, our modulus q is constant over the parameter sets and allows an optimized NTT implementation shared for all sets. With only 16-bit in size, our modulus also allows storing coefficients memory-efficiently without compression.

5 Implementation Details

In this section, we detail how to efficiently implement HAETAE. We start this section with an implementation-oriented specification. Specifically, Figure 13 demonstrates how to implement the key generation, Figure 14 the signature generation and Figure 15 the signature verification. These illustrate the use of the Chinese Remainder Theorem (CRT) and Number-Theoretic Transform (NTT) for efficient polynomial arithmetic. The most

[1] https://github.com/jdevevey/refined-haetae

notable remark from an implementation point of view is that $\mathbf{b}$ can be transmitted in NTT domain if no rounding is applied, and most arithmetic is carried out modulo q, and recovering the values modulo $2q$ is only required for computing the low and high bits. Relevant for signing and verification is the unpacking routine for $\widehat{\mathbf{A}}$, depicted in Algorithm 1.

```
KeyGen(1^λ) for d > 0
seed ← $\{0,1\}^{\rho_0}$
$(\mathsf{seed}_{\mathbf{A}}, \mathsf{seed}_{\mathsf{sk}}, K) := H_{\mathsf{gen}}(\mathsf{seed})$
$(\mathbf{a}_{\mathsf{gen}}, \widehat{\mathbf{A}}_{\mathsf{gen}}) := \mathsf{expandA}_d(\mathsf{seed}_{\mathbf{A}})$    ▷ $(\mathbf{a}_{\mathsf{gen}}, \widehat{\mathbf{A}}_{\mathsf{gen}}) \in \mathcal{R}_q^k \times \mathcal{R}_q^{k\times\ell-1}$
$(\mathsf{counter}_{\mathsf{sk}}, \mathsf{flag}) := (0, \mathsf{true})$
while flag do
    $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) := \mathsf{expandS}(\mathsf{seed}_{\mathsf{sk}}, \mathsf{counter}_{\mathsf{sk}})$    ▷ $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) \in \mathcal{S}_\eta^{\ell-1} \times \mathcal{S}_\eta^k$
    $\mathbf{b} := \mathbf{a}_{\mathsf{gen}} + \mathsf{NTT}^{-1}(\widehat{\mathbf{A}}_{\mathsf{gen}} \circ \mathsf{NTT}(\mathbf{s}_{\mathsf{gen}})) + \mathbf{e}_{\mathsf{gen}} \bmod q$    ▷ $\mathbf{b} \in \mathcal{R}_q^k$
    $(\mathbf{b}_0, \mathbf{b}_1) := (\mathsf{LowBits}^{\mathsf{vk}}(\mathbf{b}), \mathsf{HighBits}^{\mathsf{vk}}(\mathbf{b}))$
    $(\mathbf{s}_1, \mathbf{s}_2) := (\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}} - \mathbf{b}_0)$
    $\mathsf{counter}_{\mathsf{sk}} := \mathsf{counter}_{\mathsf{sk}} + 1$
    if $\mathcal{N}(\mathbf{s}_1, \mathbf{s}_2) \le \gamma^2 n$ then
        flag := false
$tr := H(\mathsf{seed}_{\mathbf{A}}, \mathbf{b}_1)$
return $\mathsf{sk} = (\mathbf{s}_1, \mathbf{s}_2, K, tr, \mathsf{seed}_{\mathbf{A}}, \mathbf{b}_1)$, $\mathsf{vk} = (\mathsf{seed}_{\mathbf{A}}, \mathbf{b}_1)$
```

```
KeyGen(1^λ) for d = 0
seed ← $\{0,1\}^{\rho_0}$
$(\mathsf{seed}_{\mathbf{A}}, \mathsf{seed}_{\mathsf{sk}}, K) := H_{\mathsf{gen}}(\mathsf{seed})$
$(\mathsf{counter}_{\mathsf{sk}}, \mathsf{flag}) := (0, \mathsf{true})$
while flag do
    $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) := \mathsf{expandS}(\mathsf{seed}_{\mathsf{sk}}, \mathsf{counter}_{\mathsf{sk}})$    ▷ $(\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}}) \in \mathcal{S}_\eta^{\ell-1} \times \mathcal{S}_\eta^k$
    $(\mathbf{s}_1, \mathbf{s}_2) := (\mathbf{s}_{\mathsf{gen}}, \mathbf{e}_{\mathsf{gen}})$
    $\mathsf{counter}_{\mathsf{sk}} := \mathsf{counter}_{\mathsf{sk}} + 1$
    if $\mathcal{N}(\mathbf{s}_1, \mathbf{s}_2) \le \gamma^2 n$ then
        flag := false
$\widehat{\mathbf{A}}_{\mathsf{gen}} := \mathsf{expandA}_d(\mathsf{seed}_{\mathbf{A}})$    ▷ $\widehat{\mathbf{A}}_{\mathsf{gen}} \in \mathcal{R}_q^{k\times\ell-1}$
$\widehat{\mathbf{b}} := -2\left(\widehat{\mathbf{A}}_{\mathsf{gen}} \circ \mathsf{NTT}(\mathbf{s}_{\mathsf{gen}}) + \mathsf{NTT}(\mathbf{e}_{\mathsf{gen}})\right) \bmod q$    ▷ $\mathbf{b} \in \mathcal{R}_q^k$
$tr := H(\mathsf{seed}_{\mathbf{A}}, \widehat{\mathbf{b}})$
return $\mathsf{sk} = (\mathbf{s}_1, \mathbf{s}_2, K, tr, \mathsf{seed}_{\mathbf{A}}, \widehat{\mathbf{b}})$, $\mathsf{vk} = (\mathsf{seed}_{\mathbf{A}}, \widehat{\mathbf{b}})$
```

Figure 13: Implementation specification (deterministic version) of HAETAE key generation

In the following, we detail about the hyperball sampler, packing, and the performance of our reference implementation.

5.1 Hyperball Sampler

Essentially, the hyperball sampling procedure consists of four steps:

1. Sample $n(k+\ell)+2$ discrete Gaussians with $\sigma = 2^{76}$, sum up their squares, and drop two samples eventually.

2. Compute the inverse of the square root of the sum of squares, multiply the result by $B_0 + \sqrt{nm}/(2N)$.

```
$\mathsf{Sign}(\mathsf{sk}, M)$
$(\mathbf{s}_1, \mathbf{s}_2, K, tr, \mathsf{seed}_{\mathbf{A}}, \psi) := \mathsf{sk}$
$\widehat{\mathbf{A}} := \mathsf{unpackA}_d(\mathsf{seed}_{\mathbf{A}}, \psi)$    ▷ Algorithm 1, $\mathbf{A} \in \mathcal{R}_q^{k \times \ell}$
$\mu := H_{\mathsf{gen}}(tr, M)$
$\mathsf{seed}_{ybb} := H_{\mathsf{gen}}(K, \mu)$
$(\kappa, \sigma) := (0, \perp)$
while $\sigma = \perp$ do    ▷ pre-compute $(\widehat{\mathbf{s}}_1, \widehat{\mathbf{s}}_2) := (\mathsf{NTT}(\mathbf{s}_1), \mathsf{NTT}(\mathbf{s}_2))$
    $(\mathbf{y}_1, \mathbf{y}_2, b, b', \kappa) := \mathsf{expandYbb}(\mathsf{seed}_{ybb}, \kappa)$    ▷ $(\mathbf{y}_1, \mathbf{y}_2) \in (1/N)\mathcal{R}^\ell \times (1/N)\mathcal{R}^k$
    $\mathbf{w} := \mathsf{NTT}^{-1}(\widehat{\mathbf{A}} \circ \mathsf{NTT}(\lfloor \mathbf{y}_1 \rceil)) + 2 \cdot \lfloor \mathbf{y}_2 \rceil \bmod q$    ▷ $\mathbf{w} \in \mathcal{R}_q^k$
    $\mathbf{w}' := \mathsf{fromCRT}(\mathbf{w}, \lfloor y_{1,1} \rceil)$    ▷ Algorithm 2, $\mathbf{w}' \in \mathcal{R}_{2q}^k$
    $\mathbf{w}'_1 := \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}')$
    $\rho = H(\mathbf{w}'_1, \mathsf{LSB}(\lfloor y_{1,1} \rceil), \mu)$
    $c = \mathsf{SampleBinaryChallenge}_\tau(\rho)$
    $\widehat{c} := \mathsf{NTT}(c)$
    $z_{1,1} := y_{1,1} + (-1)^b \cdot c$    ▷ $(\mathbf{z}_1, \mathbf{z}_2) \in (1/N)\mathcal{R}^\ell \times (1/N)\mathcal{R}^k$
    $(\mathbf{z}_1)_{2..\ell} := (\mathbf{y}_1)_{2..\ell} + (-1)^b \mathsf{NTT}^{-1}(\widehat{c} \circ \widehat{\mathbf{s}}_1)$
    $\mathbf{z}_2 := \mathbf{y}_2 + (-1)^b \mathsf{NTT}^{-1}(\widehat{c} \circ \widehat{\mathbf{s}}_2)$
    if $\|(\mathbf{z}_1, \mathbf{z}_2)\|_2 < B'$ and $(\|2(\mathbf{z}_1, \mathbf{z}_2) - (\mathbf{y}_1, \mathbf{y}_2)\|_2 > B$ or $b' = 1)$ then
                                                    ▷ Check this condition in constant time.
        $\mathbf{h} := \mathbf{w}'_1 - \mathsf{HighBits}^{\mathsf{h}}(\mathbf{w}' - 2\lfloor \mathbf{z}_2 \rceil) \bmod^+ \frac{2(q-1)}{\alpha_{\mathsf{h}}}$
        $\sigma := \mathsf{packSig}(\mathsf{HighBits}^{z1}(\lfloor \mathbf{z}_1 \rceil), \mathsf{LowBits}^{z1}(\lfloor \mathbf{z}_1 \rceil), \mathbf{h}, c)$
                                                    ▷ Section 5.2 (can fail and return $\perp$)
```

Figure 14: Implementation specification (deterministic version) of HAETAE signing.

```
$\mathsf{Verify}(\mathsf{vk}, M, \sigma)$
$(\mathsf{seed}_{\mathbf{A}}, \psi) := \mathsf{vk}$
$\widehat{\mathbf{A}} := \mathsf{unpackA}_d(\mathsf{seed}_{\mathbf{A}}, \psi)$    ▷ Algorithm 1, $\mathbf{A} \in \mathcal{R}_q^{k \times \ell}$
$(\mathsf{HighBits}^{z1}(\lfloor \mathbf{z}_1 \rceil), \mathsf{LowBits}^{z1}(\lfloor \mathbf{z}_1 \rceil), \mathbf{h}, c) := \mathsf{unpackSig}(\sigma)$
                                    ▷ Section 5.2 (can fail and cause a rejection)
$\tilde{\mathbf{z}}_1 := \mathsf{HighBits}^{z1}(\lfloor \mathbf{z}_1 \rceil) \cdot 256 + \mathsf{LowBits}^{z1}(\lfloor \mathbf{z}_1 \rceil)$
$w' := \mathsf{LSB}(\tilde{z}_{1,1} - c)$
$\tilde{\mathbf{w}} := \widehat{\mathbf{A}} \circ \mathsf{NTT}(\tilde{\mathbf{z}}_1) \bmod q$
$\tilde{\mathbf{w}}' := \mathsf{fromCRT}(\tilde{\mathbf{w}}, w')$    ▷ Algorithm 2
$\tilde{\mathbf{w}}'_1 := \tilde{\mathbf{h}} + \mathsf{HighBits}^{\mathsf{h}}(\tilde{\mathbf{w}}') \bmod^+ \frac{2(q-1)}{\alpha_{\mathsf{h}}}$
$\tilde{\mathbf{z}}_2 := [\tilde{\mathbf{w}}'_1 \cdot \alpha_{\mathsf{h}} + w'\mathbf{j} - \tilde{\mathbf{w}}' \mod 2q]/2$ ▷ addition with $w'$ only for first vector element
$\tilde{\mu} = H_{\mathsf{gen}}(\mathsf{seed}_{\mathbf{A}}, \psi, M)$
Return $(c = \mathsf{SampleBinaryChallenge}_\tau(H(\tilde{\mathbf{w}}'_1, w', \tilde{\mu}))) \wedge (\|(\tilde{\mathbf{z}}_1, \tilde{\mathbf{z}}_2)\| < B'')$
```

Figure 15: Implementation specification (deterministic version) of HAETAE verification

3. Multiply every sample from Step 1 by the result of Step 2.

4. Check the ℓ_2 norm of the resulting vector, start from Step 1 if this is bigger than $B_0 N$.

In the following, we explain how the Gaussian sampling and the approximation of the inverse of the square root can be implemented efficiently. Besides, we choose to generate each of the $k + \ell$ polynomials independently, which helps parallelizing the randomness generation for implementations that use vectorization and hardware implementations.

```
$\mathsf{expandYbb}(\mathsf{seed}_{ybb}, \kappa)$:
$(\mathbf{y}_1, \mathbf{y}_2) := \perp$    ▷ $(\mathbf{y}_1, \mathbf{y}_2) \in (1/N)\mathcal{R}^\ell \times (1/N)\mathcal{R}^k$
while $(\mathbf{y}_1, \mathbf{y}_2) = \perp$ do
    $\mathbf{t}_1 := \mathsf{sampleGauss}(\mathsf{SHAKE256}(\mathsf{seed}_{ybb}, \kappa), n+1)$    ▷ Section 5.1.1
    $\mathbf{t}_2 := \mathsf{sampleGauss}(\mathsf{SHAKE256}(\mathsf{seed}_{ybb}, \kappa+1), n+1)$    ▷ Section 5.1.1
    for $i := 3$ to $k+\ell$ do
        $\mathbf{t}_i := \mathsf{sampleGauss}(\mathsf{SHAKE256}(\mathsf{seed}_{ybb}, \kappa+i), n)$    ▷ Section 5.1.1
    $s := \sum_{i=1}^{n+1} t_{1,i}^2 + \sum_{i=1}^{n+1} t_{2,i}^2 + \sum_{i=3}^{k+\ell} \sum_{j=1}^{n} t_{i,j}^2$
    drop $t_{1,n}$ and $t_{2,n}$
    approximate $1/\sqrt{s}$    ▷ Section 5.1.2
    for $i := 0$ to $k+\ell-1$ do
        for $j := 0$ to $n-1$ do
            $t_{i,j} := \left\lfloor t_{i,j} \cdot \left(BN + \frac{\sqrt{nm}}{2}\right) \cdot \frac{1}{\sqrt{s}} \right\rceil$    ▷ round to $\log_2 N$ fix-point bits
    $\kappa := \kappa + k + \ell$
    if $\sum_{i=0}^{k+\ell-1} \sum_{j=0}^{n-1} t_{i,j}^2 \leq (BN)^2$ then
        arrange $\mathbf{t}_1, \ldots, \mathbf{t}_\ell$ as $\mathbf{y}_1$
        arrange $\mathbf{t}_{\ell+1}, \ldots, \mathbf{t}_{k+\ell}$ as $\mathbf{y}_2$
sample $b, b'$ as the first two bits from the output of $\mathsf{SHAKE256}(\mathsf{seed}_{ybb}, \kappa)$
$\kappa := \kappa + 1$
return $(\mathbf{y}_1, \mathbf{y}_2, b, b', \kappa)$
```

Figure 16: Deterministic hyperball sampling.

Then, for the first two polynomials, we generate one more Gaussian sample each, which is never stored but included in the sum of squared samples.

5.1.1 Discrete Gaussian Sampling

As we will lose precision when computing the inverse square root of a Gaussian sample, we require a Gaussian sampler with high fix-point precision. This is achieved by sampling over $\mathbb{Z}$ with a large standard deviation and then scaling the resulting sample to our convenience. We use [Ros20, Algorithm 12] to sample from a discrete Gaussian distribution with $\sigma = 2^{76}, k = 2^{72}$.

In essence, we start by sampling a discrete Gaussian x with $\sigma = 16$ using a Cumulative Distribution Table (CDT) and a uniform $y \in \{0, \ldots, 2^{72}-1\}$ and set the Gaussian sample candidate as $r = x2^{72} + y$. Subsequently, this candidate is accepted with probability $\exp(-y(y + x2^{73})/2^{153})$. Fortunately, we achieve a very low rejection rate of less than 5 %.

Specifically, the CDT we use has 64 entries and uses a precision of 16 bit. Then, to compute the sample candidate's square and the input to the exponential, we first compute r^2 and round the result to 76-bit precision, which is accumulated later if the sample is accepted. Subsequently, $r^2 - 2^{76}x^2$ yields the input to the exponential.

Approximating the Exponential. For this, we need to approximate the exponential function e^{-x} by a polynomial $f(x)$ on the closed interval $[c - \frac{w}{2}, c + \frac{w}{2}]$, with center c and width w. We first determine an upper bound for the polynomial order required to approximate e^{-x}, given an upper bound for the absolute error. We obtain $f(x)$ by truncating the expansion of e^{-x} into a series of Chebyshev polynomials of the first kind $T_n(x)$ with linearly transformed input, as this is known to yield small absolute approximation errors for a given polynomial order. We find:

$$e^{-x} = -e^{-c} + 2e^{-c}\sum_{n=0}^{\infty}(-1)^n I_n\left(\tfrac{w}{2}\right)T_n\left(\frac{x-c}{w/2}\right) \qquad\qquad x \in [c - \tfrac{w}{2}, c + \tfrac{w}{2}]$$

where $I_n(z)$ are modified Bessel functions of the first kind, which rapidly converge to zero for growing n. We recall $\|T_n(x)\| \leq 1$ for $\|x\| \leq 1$. For intervals $[0, w]$ with not too large widths we find $2e^{-c}I_{m+1}(\frac{w}{2})$ to be a useful estimate of the maximum absolute error, when truncating the series at order $m > 1$. This relation allows us to directly limit m according to the interval to cover and the maximum permissible error.

We then determine the polynomial $f(x)$ of at most order m by using the Chebyshev approximation formula, which has been shown to result in a nearly optimal approximation polynomial in the case of the exponential function [Li04]. The number of fraction bits is chosen to match the error. The numerical evaluation is performed in fixed-point arithmetic using the Horner's scheme and multiplying with shifts to retain significant bits. When shifting right, we round half up, which retains about one additional bit of accuracy when compared to truncation.

Barthe et al. [BBE+19] introduced the GALACTICS toolbox to derive suitable polynomials approximating e^{-x}. They numerically evaluate and modify trial polynomials, minimizing the *relative* error, until an acceptable level is reached. The polynomials are evaluated using a Horner's scheme, similar to this work, but rely on truncation. When comparing to polynomials derived using the GALACTICS toolbox, our approximation has a slightly smaller absolute error for intervals of interest in this work, while maintaining the same polynomial order and constant time properties. This holds even when introducing rounding to the GALACTICS evaluation of polynomials. Moreover, our approach is somewhat less heuristic than the GALACTICS method. Practically, as can be seen in Listing 1, the approximation consists of six signed 48-bit multiplications with subsequent rounding (`smulh48`), several constant shifts with rounding and constant additions.

Listing 1: Fix-point approximation of the exponential function with 48 bit of precision.

```
static uint64_t approx_exp(const uint64_t x) {
 int64_t result;
 result = -0x0000B6C6340925AELL;
 result = ((smulh48(result, x) + (1LL << 2)) >> 3) + 0x0000B4BD4DF85227LL;
 result = ((smulh48(result, x) + (1LL << 2)) >> 3) - 0x0000887F727491E2LL;
 result = ((smulh48(result, x) + (1LL << 1)) >> 2) + 0x0000AAAA643C7E8DLL;
 result = ((smulh48(result, x) + (1LL << 1)) >> 2) - 0x0000AAAAA98179E6LL;
 result = ((smulh48(result, x) + 1LL) >> 1) + 0x0000FFFFFFFB2E7ALL;
 result = ((smulh48(result, x) + 1LL) >> 1) - 0x0000FFFFFFFFF85FLL;
 result = ((smulh48(result, x))) + 0x0000FFFFFFFFFFFCLL;
 return result;
}
```

Finalization. If the sample is accepted eventually, it is (implicitly) scaled by the factor 2^{-76} to obtain a continuous sample from the standard normal distribution. Moreover, we only need to store the upper 64 bits of the sample and round off the rest.

In summary, each Gaussian sample candidate requires **72 bit** randomness for the lower part of the candidate (y), **16 bit** randomness for the CDT sampling, and **48 bit** randomness for rejecting the candidate conditionally according to the output of the exponential. This results in a vast randomness demand per hyperball sample, and explains the dominance of hashing in the cycle count performance.

Table 2: Symbol mapping and encoding size parameters.

Scheme	cut_h	$offset_h$	$\lvert\{S_h(n)\}\rvert$	cut_{z_1}	$\lvert\{S_{z_1}(n)\}\rvert$	$base_h$	$base_{z_1}$
HAETAE-120	6	239	13	6	13	7	132
HAETAE-180	8	235	17	8	17	127	376
HAETAE-260	16	471	33	9	19	358	501

5.1.2 Approximating the Inverse of the Square Root

To turn the vector of standard normal distributed variates into a hyperball sample candidate, we must compute its norm. For this, we accumulate all squared samples and approximate the inverse of the square root of the accumulated value. The approximation result is then multiplied by the constant $r' + \sqrt{nm}/(2N)$, which yields the scaling factor that is multiplied to each Gaussian sample. For the inverse square root, we deploy Newton's method, which is a well-known technique for that purpose. However, Newton's method requires a starting approximation that is, with each iteration, turned into a better approximation. Fortunately, we know that the sum of $nm + 2$ independent squared standard normal variables follows a χ^2 distribution with expected value $nm + 2$. Hence, the starting approximation can be fixed and precomputed as $1/\sqrt{nm + 2}$. The number of iterations for a targeted precision can be determined experimentally. Therefore, we performed the approximation for the first input values that have negligible probabilities either for the cumulative distribution function of $\chi^2(nm + 2)$ or its survival function, and checked how many iterations are required to still reach reasonable precision.

5.2 Signature Packing and Sizes

The last step of the signature generation is to compress and pack the elements of the signature. A packed HAETAE signature consists of the challenge c, the low bits of $\mathbf{z_1}$ (LN coefficients), the high bits of $\mathbf{z_1}$ and $\mathbf{h}$ (KN coefficients). Because the distributions of the values in the high bits of $\mathbf{z_1}$ and the coefficients in $\mathbf{h}$ are both very dense, we can compress both polynomial vectors with encoding. Figure 17 displays the frequencies for the possible values for both vectors in HAETAE-120. Before compressing the values, we map them to a smaller symbol space and thereby reject the very unlikely values and the corresponding signatures. For $\mathbf{h}$ we cut out most of the values in the middle of the range, for HAETAE-120 this reduces the size of the symbol space from 252 to 13.

$$S_h(n) = \left\{ \begin{array}{lr} n, & \text{for } 0 \leq n \leq cut_h \\ \perp, & \text{for } cut_h < n \leq cut_h + offset_h \\ n - offset_h, & \text{for } cut_h + offset_h < n \end{array} \right\}$$

For the high bits of $\mathbf{z_1}$ we tail-cut the distribution left and right of the center at 0, and then shift the remaining values to the non-negative range beginning at 0. For HAETAE-120 this reduces the size of the symbol space from 37 to 13.

$$S_{z_1}(n) = \left\{ \begin{array}{lr} \perp, & \text{for } n < -cut_{z_1} \\ n + cut_{z_1}, & \text{for } -cut_{z_1} \leq n \leq cut_{z_1} \\ \perp, & \text{for } cut_{z_1} < n \end{array} \right\}$$

The parameters for these mappings are defined in Table 2. At the signature verification, the mapping must be reverted after decoding the compressed symbols.

The reason for these mappings is mainly to get significantly smaller precomputation tables for the rANS encoding and decoding. Also, all symbols can now be represented with 8-bits, which simplifies the rANS implementation. Furthermore, for the high bits of $\mathbf{z_1}$, a mapping to non-negative values is necessary to be able to use rANS encoding. The effect on the resulting signature size is insignificant.

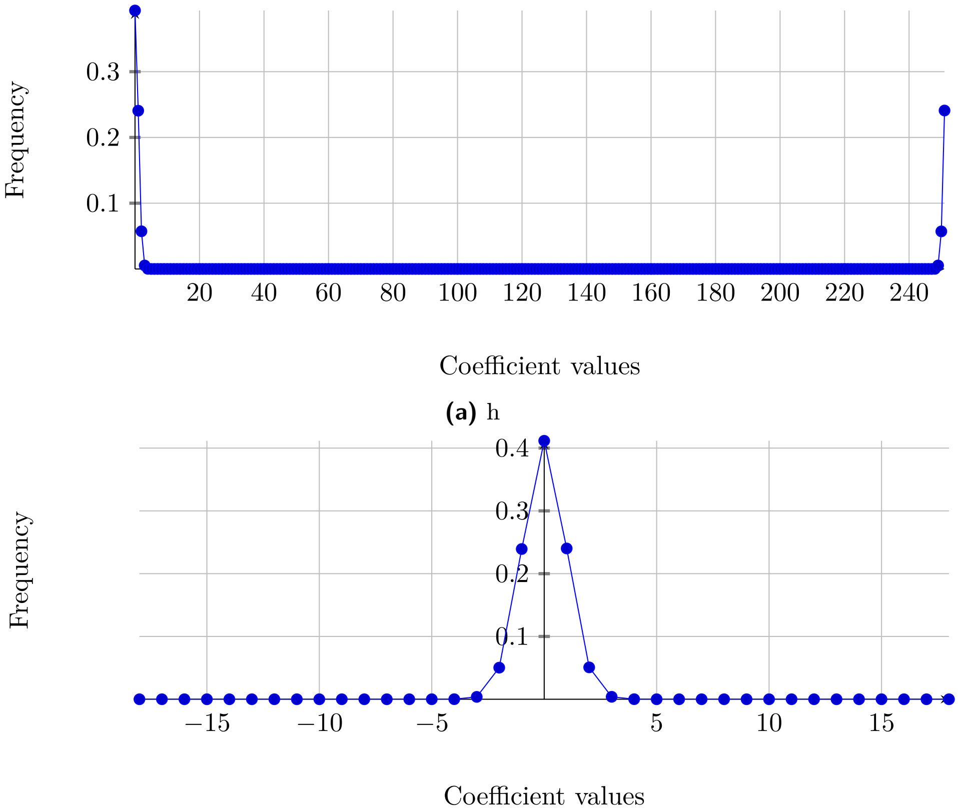

(a) h

(b) HB($\mathbf{z_1}$)

Figure 17: Distribution of the coefficients of h and HB($\mathbf{z_1}$) in HAETAE-120.

The size of the compressed high bits of $\mathbf{z_1}$ and $\mathbf{h}$ varies and must be included in the signature, to allow a correct unpacking and decoding. The size of one compressed polynomial vector is often more than 255 bytes, and can thus not be expressed by one byte. Its variance however, is limited, and thus we encode the size the compressed high bits of $\mathbf{z_1}$ and $\mathbf{h}$ as positive offset to a fixed base value. This unsigned offset value fits into one byte in most of the cases, if not, the signature gets rejected. The base values can be found in Table 2.

The final signature is then built as following: The first 32 bytes contain the seed for the challenge polynomial c. Following, we have LN bytes for the low bits of $\mathbf{z_1}$. The next first byte consists of the offset to the base size for the encoding of the high bits of $\mathbf{z_1}$ and the second byte is the offset for $\mathbf{h}$. Then we have the encoding of the high bits of $\mathbf{z_1}$ and directly afterwards the encoding of $\mathbf{h}$, both with varying sizes, which are indicated beforehand. Lastly, the signature is padded with zero bytes to reach the fixed signature size, if any bytes remain. Signatures that would exceed the fixed limit get rejected.

To prevent signature forgeries, during signature unpacking and decoding multiple sanity checks have to be performed: the zero padding must be correct, the decoding must not fail and decode the expected number of coefficients while using exactly the amount of bytes indicated with the offset. Furthermore, rANS decoding must end with the fixed predefined start value to be unique. Our rANS encoding is based on an implementation by Fabian Giesen [Gie14].

To set the fixed signature size as reported in Table 3, we evaluated the distribution empirically and determined a threshold that requires a rejection in less than 0.1% of the

Table 3: NIST security level, signature and key sizes (bytes) of HAETAE, Dilithium, and Falcon.

Scheme	Lvl.	vk	Signature	Sum	Secret key
HAETAE-120	2	992	1,474	2,466	1,408
HAETAE-180	3	1,472	2,349	3,821	2,112
HAETAE-260	5	2,080	2,948	5,028	2,752
Dilithium-2	2	1,312	2,420	3,732	2,528
Dilithium-3	3	1,952	3,293	5,245	4,000
Dilithium-5	5	2,592	4,595	7,187	4,864
Falcon-512	1	897	666	1,563	1,281
Falcon-1024	5	1,792	1,280	3,072	2,305

cases. Figure 18 displays the raw signature size distribution of 20000 executions (without the size-based rejection sampling).

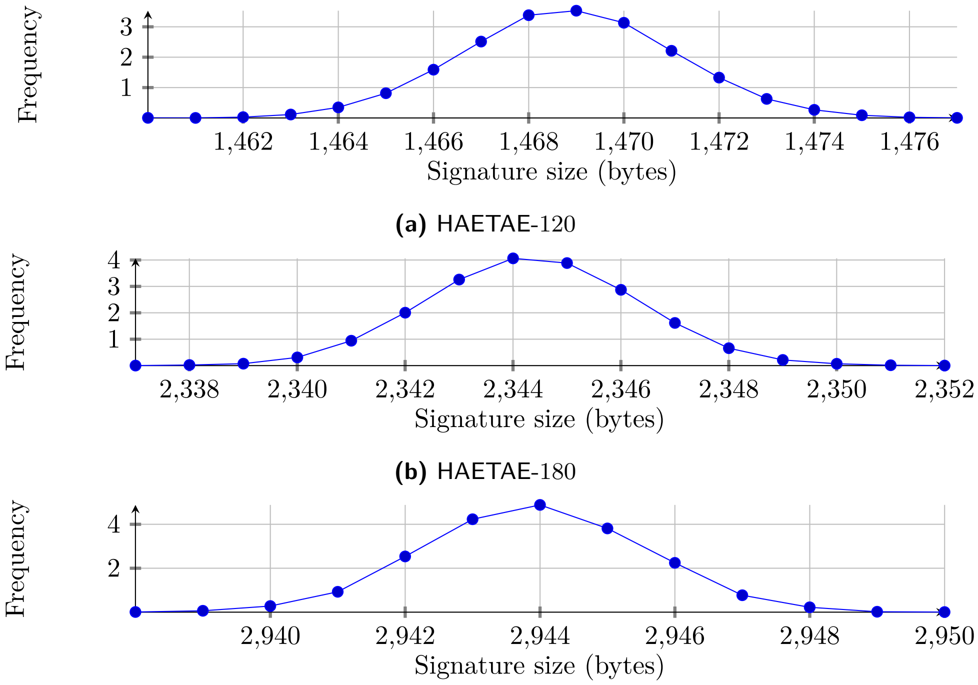

(a) HAETAE-120

(b) HAETAE-180

(c) HAETAE-260

Figure 18: Raw signature size distribution over 20000 executions. We set the bound for the size-based rejection to result in a rejection rate of less than 0.1%.

In Table 3 we compare the signature and key sizes of HAETAE, Dilithium, and Falcon. The verification keys in HAETAE are 20% (HAETAE-260) to 25% (HAETAE-120 and HAETAE-180) smaller, than their counterparts in Dilithium. The advantage of the hyperball sampling manifests itself in the signature sizes, HAETAE has 29% to 39% smaller signatures than Dilithium. Less relevant are the secret key sizes, that are almost half the size in HAETAE compared to Dilithium. A direct comparison to Falcon for the same claimed security level is only possible for the highest parameter set, Falcon-1024 has a signature of less than half the size compared to HAETAE-260, and its verification key is about 14% smaller.

5.3 Performance Reference Implementation

We developed an unoptimized, portable and constant-time implementation in C for HAETAE and report median and average cycle counts of one thousand executions for each parameter set in Table 4. Due to the key and signature rejection steps, the median and average values for key generation and signing respectively differ clearly, whereas the two values are much closer for the verification.

Table 4: Reference implementation speeds. Median and average cycle counts of 1000 executions for HAETAE, Dilithium, and Falcon. Cycle counts were obtained on one core of an Intel Core i7-10700k, with TurboBoost and hyperthreading disabled.

Scheme		KeyGen	Sign	Verify
HAETAE-120	*med*	1,403,402	6,039,674	376,486
	ave	1,827,567	9,458,682	376,631
HAETAE-180	*med*	2,368,038	9,161,312	691,652
	ave	3,448,185	11,611,868	692,014
HAETAE-260	*med*	3,101,280	11,444,678	895,098
	ave	4,088,383	17,229,712	896,622
Dilithium-2	*med*	343,222	1,191,218	376,008
	ave	343,639	1,527,406	376,543
Dilithium-3	*med*	630,170	2,061,816	612,538
	ave	630,607	2,603,237	612,852
Dilithium-5	*med*	945,776	2,522,834	987,154
	ave	949,662	3,080,734	988,250
Falcon-512	*med*	53,778,476	17,332,716	103,056
	ave	60,301,272	17,335,484	103,184
Falcon-1024	*med*	154,298,384	38,014,050	224,378
	ave	178,516,059	38,009,559	224,840

For a fair comparison, we also performed measurements on the same system with identical settings of the reference implementation of Dilithium[1] and the implementation with emulated floating-point operations, and thus also fully portable, of Falcon[2], as given in Table 4. The performance of the signature verification for HAETAE is very close to Dilithium throughout the parameter sets. HAETAE-180 verification is 13% slower than its counter-part, HAETAE-260 on the other hand, is 9% faster than the respective Dilithium parameter set. For key generation and signature computation, our current implementation of HAETAE is clearly slower than Dilithium. We measure a slowdown of factors three to five. In comparison to Falcon, however, HAETAE has 38-50 times faster key generation and around three times faster signing speed. For the verification, Falcon outperforms both Dilithium and HAETAE by roughly a factor of four.

A closer look at the key generation reveals that the complex Fast Fourier Transformation, that is required for the rejection step, is with 53% by far the most expensive operation and a sensible target for optimized implementations.

Profiling the signature computation reveals that the slowdown compared to Dilithium is mainly caused by the sampling from a hyperball, where about 80% of the computation time is spent. The hyperball sampling itself is dominated by the generation of randomness, which we derive from the extendable output function SHAKE256 [Dwo15], which is also used in the Dilithium implementation. Almost 60% of the signature computation time is spent in SHAKE256.

Based on the profiling and benchmarking of subcomponents, we estimate the performance of a randomized HAETAE implementation with pre-computation. The generic

[1] https://github.com/pq-crystals/dilithium/tree/master/ref

[2] https://falcon-sign.info/falcon-round3.zip

version, which is independent of the key, would already achieve a speedup of a factor five for its online signing, because the expensive hyperball sampling can be done offline. For the pre-computation variant with a designated signing key, additionally, a lot of matrix-vector multiplications and therefore most of the transformations from and to the NTT domain, can be precomputed. We estimate about 12% of the full deterministic signing running time, for the online signing in this case.

6 Optimized Implementation for AVX2

Advanced Vector Extensions 2 (AVX2) is an extension to the x86 instruction set architecture and available in processors since 2011. It provides Single Instruction Multiple Data (SIMD) operations on 256-bit registers, and thus allows to e.g. do an operation on eight 32-bit values in parallel. In this section we explain, how to exploit this parallelization.

The three major components, that significantly determine the computation time of HAETAE are Keccak, the NTT and the hyperball sampling. For the first two components we can fall back to existing optimized code. For the NTT in particular, we can reuse the implementation in Dilithium with only slight adaptions with regard to constants. In the following we demonstrate how to implement the third component, the hyperball sampling efficiently using AVX2 instructions.

6.1 Vectorized Hyperball Sampling

After the parallel generation of the randomness, generally, we have two options to parallelize the hyperball sampling. First, we can sample four different polynomials in parallel, and second, we can generate the Gaussian samples within a polynomial in parallel, since they are generated independently. We opt for the first approach.

As the sampling process is relatively complex, we cannot load input vectors, generate samples from them, and eventually store the samples. Instead, we pass several times over the internal memory state, dividing the procedure into seven separate steps:

1. parsing of the input randomness: separating the three parts for each sample candidate into separate memory locations such that later steps can process them quickly,
2. CDT sampling,
3. constructing the sample candidate, its square, and the input to the exponential approximation,
4. approximate the exponential,
5. generate masks which candidates to reject,
6. accumulate the squares of the non-rejected samples, and
7. storing only the accepted samples into the correct final memory positions.

In the following, we detail Steps 2, 3, and 4.

6.1.1 Parallel CDT Sampling

Although we perform 16-bit CDT sampling, we cannot use the 16-fold parallel `vpcmpgtw` comparison, since it is a *signed* comparison, and use `vpcmpgtd` instead, which operates on eight signed 32-bit integers. As we want to sample four samples in parallel, we store the CDT and the input randomness redundantly, such that we can perform the comparison with all 64 entries with 32 comparisons. However, since AVX2 only offers 16 vector registers, we have to perform this comparison in three chunks.

`vpcmpgtd c, a, b` writes -1 to c if the respective vector element in `a` is in fact greater than its counterpart in `b`, and else 0. Thus, we can use `vpaddd` to accumulate these results, but they will be negative. For the final chunk, we use `vpsrlq` to line up the two intermediate results, then add them and negate them.

6.1.2 Multi-precision Squaring

Since the sample candidate is 72 bits long and AVX2 only supports 32-bit multiplication, we perform vectorized multi-precision arithmetic. Therefore, we split the candidate into its low 32 bits, the middle 16 bits, and the upper 31 bits. For the schoolbook multiplication, we perform the six partial multiplications with `vpmuludq` consecutively, such that they can be executed pipelined and in parallel. The subsequent recombination and rounding can be performed with a sequence of 16 instructions of shifts (`vpsrlq`, `vpsllq`), additions (`vpaddq`), and ANDs (`vpand`).

6.1.3 Vectorized Approximation of the Exponential

The exponential approximation as explained in Section 5.1 consists of six signed 48-bit multiplications, which is not supported natively by AVX2. Consequently, we implement this operation with a vectorized multi-precision approach.

More specifically, we know that the first operand of this multiplication is signed, and the second is not. Thus, splitting the second operand into a low and a high half is trivial, but for the signed operand, this requires a slightly more sophisticated approach: Here, the upper half is obtained by an arithmetic right-shift by 24. By shifting this result left again by 24 (shifting in zeros), and subtracting the result from the original value, we obtain the lower half.

Since AVX2 does not offer a signed right shift over 64-bit entries, we generate a mask of sign bits and simulate a signed right shift by performing a bitwise OR. Unfortunately, we require this operation three times during a single signed multiplication operation. Notably, AVX512 offers a signed right shift, which will speed up this operation considerably.

Eventually, we perform a vectorized signed 48-bit multiplication using 32 instructions, out of which 17 are used for the emulation of a signed right shift. Moreover, we use seven variable and three constant vector registers (out of which one is the second input, cf. Listing 1), which leaves six registers for other constants. Apart from the multiplication, we only make use of addition and shifts (`vpaddq`, `vpsrlq`).

6.2 Performance and Comparison AVX2

The impact of the parallelized Keccak can be observed by looking at the cycle costs for unpacking the matrix $\mathbf{A}$, which is between five to seven times faster compared to the reference implementation. For e.g. HAETAE-120, the costs went down from around 132k cycles to 24k cycles. The picture is similar for the hyperball sampling, where we measure a speed-up of factor six to eight. The cycle counts for one function call in HAETAE-120 are around 1640k in the reference and 270k in the optimized implementation. The highly optimized NTT taken from Dilithium, is almost 19 times faster than the one in our portable reference code.

Table 5 provides cycle numbers for the AVX2 optimized implementations of HAETAE, Dilithium [1] and Falcon [2]. Compared to our reference implementation, the signature

[1] https://github.com/pq-crystals/dilithium/tree/master/avx2

[2] https://falcon-sign.info/falcon-round3.zip

Table 5: AVX2 optimized implementation speeds. Median and average cycle counts of 1000 executions for HAETAE, Dilithium, and Falcon. Cycle counts were obtained on one core of an Intel Core i7-10700k, with TurboBoost and hyperthreading disabled.

Scheme		KeyGen	Sign	Verify
HAETAE-120	*med*	535,210	1,306,112	108,840
	ave	733,350	1,916,063	108,890
HAETAE-180	*med*	820,570	1,793,342	167,722
	ave	1,122,867	2,198,920	167,787
HAETAE-260	*med*	872,280	2,049,008	207,146
	ave	1,202,911	3,013,239	207,229
Dilithium-2	*med*	87,020	200,242	92,148
	ave	86,937	252,905	92,190
Dilithium-3	*med*	146,560	334,898	148,810
	ave	146,688	402,012	148,883
Dilithium-5	*med*	233,976	415,228	232,146
	ave	233,895	484,119	232,241
Falcon-512	*med*	24,663,306	863,076	100,540
	ave	26,637,878	863,420	100,709
Falcon-1024	*med*	71,013,520	1,740,188	228,086
	ave	78,797,658	1,740,520	228,326

generation is around five times faster in the optimized implementation. For the signature verification we observe an acceleration of around three to four.

The comparison with Dilithium does not change distinctly with respect to the reference implementations. Key generation and signature generation are about six times faster with Dilithium and the speed differences in the verification are up to 24% with the optimized implementations. Falcon on the other hand is considerably faster at the signature generation with its AVX2 implementation, compared to its portable reference code. The performance of the reference code in Table 4 showed around three times faster signature generation for HAETAE compared to Falcon. Optimized for AVX2 and also using the floating-point unit, Falcon becomes faster than HAETAE for signature generation. Key generation is still at least forty times faster with HAETAE compared to Falcon, and the verification speed is very similar for both schemes.

We note, however, that both Dilithium and Falcon went through a multi-year process of incrementally optimized implementations, whereas this process has just started for HAETAE. Moreover, when we apply the heuristic that sending one byte via internet costs *at least* 1000 cycles [BBC$^+$20, Sec. 5.4], remarkably, HAETAE is already nearly as performant as Dilithium in terms of signing plus sending the signature.

7 Embedded Implementation on Cortex-M4

To evaluate the suitability of HAETAE for embedded environments we developed an implementation for the `STM32f4-Discovery` board, featuring 128 KiB RAM and a Cortex-M4F processor which implements the ARMv7E-M ISA and operates on 32-bit words. We use the PQM4 framework [KRSS19] for development and evaluation, as it is the de facto standard for comparison of post-quantum cryptography schemes on Cortex-M4 processors. The Cortex-M4F, in contrast to the Cortex-M4, features a floating-point unit. Its floating point registers can be used to store and load intermediate values within a single cycle to reduce the pressure on the 13 general purpose registers.

The profiling of the reference implementation already indicates that replacing the portable Keccak implementation with one optimized for the Cortex-M4 is an important and straightforward step towards fast execution time. The two other major components that are highly relevant are the polynomial arithmetic and the Gaussian sampler, both will be discussed in the following.

7.1 Polynomial Arithmetic

In this section, we will address the issue of how to implement the required arithmetic operations on these rings and mappings between them on a Cortex-M4 platform.

Modular Reductions. In modular arithmetic, the Barret reduction [Bar87], the Montgomery modular multiplication [Mon85], and related techniques are indispensable for efficient computation, the first for reducing given numbers, the second for yielding the reduced result of a multiplication with a constant mod$q$. The algorithms avoid computationally expensive divisions by q and replace them with a multiplication by a suitably chosen number and division by powers of two, which can be realized with shift operations. Both methods initially reduce the result to the interval $[0, 2q]$ and perform the full reduction with a conditional subtraction, which can be done in constant time. In many cases one can forgo the final reduction for intermediate results, an approach dubbed lazy reduction.

The prime chosen in the HAETAE scheme is $q = 64513 =$ `0xFC01`, the largest unsigned 16-bit prime with a 512^{th} root of unity. Fully reduced elements of $\mathbb{Z}_q$ can be stored efficiently in 16-bit in the bottom or top half of a 32-bit register.

Unfortunately, this does not carry over to arithmetic operations. A lazy reduction or an addition already requires 17 bits to store the result, a combination of a lazily reduced multiplication followed by an addition requires 18 bits. The recent advance of the Plantard multiplication [Pla21] is not useful within this work, as the prime in HAETAE is not compatible. Plantard multiplication requires $q < \frac{2R}{1+\sqrt{5}} \approxeq 0.618\,R$, i.e., $q \leq 40503$ for $R = 2^{16}$. So the prime of HAETAE is too large for this use-case with Cortex-M4 16-bit DSP-instructions. The same goes for Seiler's variant [Sei18] of signed Montgomery modular multiplication, which is only well-defined for $q < \frac{R}{2}$.

16-bit vs 32-bit. Quite a few post-quantum schemes use primes that are 13-bit values or smaller. In this case, one can both store and manipulate the coefficients graciously and efficiently as two signed 16-bit values packed into one 32-bit register, as the Cortex-M4 offers a wide range of instructions intended for Digital Signal Processor (DSP) applications, like mixed multiplication of upper and lower halves of two registers that can be used for this purpose. In the case of HAETAE, trade-offs need to be found between the compactness of 16-bit storage and doubled speed of access for consecutive coefficients on the one hand, and the required overhead to fully reduce the coefficients before writing them to memory.

If coefficients are written once and afterwards are read repeatedly without alterations, the 16-bit representation can be worthwhile. When polynomials of the public key are expanded, the coefficients are sampled in fully reduced state, we therefore store them in halfwords.

While it is feasible to use a modified Montgomery reduction with unsigned 16-bit integers as input and 17 bits output (or a 16-bit value with overflow flag), there are no corresponding instructions available to exploit this. In contrast, the Montgomery reduction in the Dilithium implementation for the Cortex-M4 uses $R = 2^{32}$ and takes only three instructions. We determined the overhead associated with full reductions required to store coefficients as 16-bit values to be too large to outperform the 32-bit variant for the NTT. The same applies to other polynomial arithmetic operations in HAETAE, besides the expanded polynomials of the public key, we therefore operate on 32-bit coefficients.

NTT. HAETAE, as other lattice-based schemes, extensively employs polynomial multiplication. The NTT is a generalization of the FFT and is the state-of-the-art technique to perform this operation, speeding up the computation considerably, as compared to, e.g., schoolbook multiplication. The addition and multiplication of polynomials transformed into the NTT domain are carried out coefficient-wise, greatly reducing the cost of the latter operation. The overhead to perform the transform and inverse transform, where required, is usually outweighed by the performance gain in the multiplication. HAETAE is specified such that large parts of the public key are expanded directly to the NTT domain.

Fortunately, the closeness to Dilithium, which also uses polynomials of degree $n = 256$, allows us the reuse its highly optimized assembly code for the NTT developed for the Cortex-M4 by Abdulrahman et al. [AHKS22], which improves over previous work [GKOS18]. The code adjustments required for operation in HAETAE are limited to adjusting constants like the prime being used, the root of unity, which is chosen as 426 (the primitive 512^{th} root of unity of $q = 64513$) and the twiddle factors.

Replacing the portable C code from the reference implementation with optimized assembly derived from the Dilithium implementation reduces the cycle count per invocation from 37506 cycles to 8047 cycles, a speed-up by a factor of **4.6**. For the inverse NTT, the cycle count dropped from 43116 to 8369, a speed-up by a factor of **5.1**.

7.2 Gaussian Sampler

The major numerical components of the Gaussian sampler are the CDT sampler for sampling the most significant bits and the fixed-point exponential function used in the rejection step. As both components are called repeatedly, both have been implemented in assembly code.

The CDT sampler accumulates ones and zeros, depending on comparisons of a uniformly sampled 16-bit random value to tabled threshold values. We use the `uadd16`, `usub16` and `sel` SIMD instructions from the Cortex-M4 instruction set to carry out two comparisons and accumulations in parallel. We furthermore optimize the memory access and unroll the loop. By doing so we reduce the instruction count from 800 cyles for the reference implementation to 206 cyles, a speed-up by a factor of **3.9**.

The exponential is approximated by a polynomial, which is evaluated using Horner's scheme. The reference implementation of the exponential function uses fixed-point arithmetic with 48 fraction bits. Values are embedded in 64-bit integers and 64x64 to 128-bit multiplication is used. The latter operation has no native support on the Cortex-M4. To circumvent this limitation, the Cortex-M4 implementation splits the value into two signed components at the start of each multiplication, namely the most significant bits `ah = a ≫ 24` and the least significant bits as `al = a - ah`. For the accumulation of the results a 64-bit integer is used again, taking advantage of the `smlal` instruction. This repeated switch between representations allows for efficient computation of the individual Horner's scheme iterations. Whereas the reference implementation of the exponential function takes 1658 cycles to execute, this is reduced to 563 instructions in the optimized Cortex-M4 code, a speed-up by a factor of **2.9**.

7.3 Stack Optimization

Besides execution time, also the memory footprint is an important metric for constrained devices. The target device in this work has 128 KiB of RAM available as stack memory. In this context, data structures typically encountered in lattice based cryptography need to be considered as rather large. A single polynomial in HAETAE takes 512 B or 1 KiB of memory to store, depending on whether the data is represented as 16-bit or 32-bit values. So vectors or matrices of polynomials can occupy a considerable share of the available

RAM, if stored in their entirety. In this section we explore how to minimize memory use; in some cases significant trade-offs between memory usage and execution speed can be made.

The reference implementation of HAETAE is designed with an emphasis on readability and close similarity to the mathematical specification. This results in top level functions, which consist of long monolithic blocks with many large data structures, which do not necessarily have overlapping lifetimes, but nevertheless occupy stack space for the entire function lifecycle. Most stack memory is required for the signature generation. Due to the unnecessarily high stack usage of the reference implementation, HAETAE-180 and HAETAE-260 do not run on the `STM32f4-Discovery` board without optimizations.

To reduce the memory footprint we executed two strategies. First, we carefully analyzed the liveness of each relevant variable and refactored the monolithic code into subroutines to reduce the scope of variables and thereby the total stack usage. This slightly impacts the readability of the code, but does not affect the speed.

In a second variant we additionally opted for a more aggressive memory reduction, by recomputing polynomials on demand, this obviously comes with performance costs. We adapt the data structures to be primarily polynomial oriented instead of vector and matrix oriented representations. Recomputations are done during public key usage, where we generate each polynomial on demand, and during hyperball sampling, where we sample each polynomial twice, once for the evaluation of the normalization factor and a second time to sample the actual $\mathbf{y}$ values. Since the hyperball sampling is computationally very expensive, this leads to a severe overhead in runtime.

7.4 Performance and Comparison Cortex-M4

Table 6 shows the maximum stack size of our two Cortex-M4 implementations of HAETAE and values reported in the PQM4 framework about Dilithium and Falcon. With *speed-opt*, we refer to our implementation, that is optimized for the Cortex-M4 and includes multiple stack-size optimizations, but does not trade speed for better memory requirements. *stack-opt* refers to the version, where we additionally exploit speed vs memory trade-offs.

First, we can observe that the memory requirements of HAETAE are small enough to run on the `STM32f4-Discovery` board for all parameter sets, even in the *speed-opt* version. Second, the stack sizes for HAETAE are in the same order of magnitude as Dilithium and Falcon. Compared to *speed-opt* HAETAE, Dilithium requires around two to three times more memory during key generation, and a similar overhead for signature verification. The difference is at most 20% for signature generation, for this operation Dilithium requires less memory than HAETAE for the first two parameter sets.

Our *stack-opt* version reduces the stack-size up to 34% during signature generation and key generation, but does not differ for the verification. However, this comes with higher costs in terms of computation time.

Falcon stands out for its stack-size below 10 KiB for both parameter sets during verification.

Table 6 shows the cycles spend by our two Cortex-M4 implementations of HAETAE and values reported in the PQM4 framework[1] for Dilithium and Falcon. Our version using aggressive stack reduction techniques based on recomputations does not impact the signature verification time, but almost doubles the computation time for the signature generation. The time overhead for key generation is up to 30%. Similar to our AVX2 optimized implementation, the relative performance comparison of HAETAE to Dilithium and Falcon does not change drastically.

[1] https://github.com/mupq/pqm4/blob/master/benchmarks.md

Table 6: Maximum stack size in bytes and average kilo cycle counts of 1000 (100 for Falcon) executions for Cortex-M4 implementations of HAETAE, Dilithium, and Falcon.

Scheme		*Stack size*			*Cycle count*		
		KeyGen	Sign	Verify	KeyGen	Sign	Verify
HAETAE-120	*speed-opt*	19,796	54,564	22,532	8,905	27,312	1,054
	stack-opt	17,364	40,732	22,532	10,819	51,017	1,054
HAETAE-180	*speed-opt*	29,612	69,631	31,020	17,666	34,466	2,026
	stack-opt	22,444	57,116	31,020	22,860	65,855	2,026
HAETAE-260	*speed-opt*	34,108	102,964	36,428	22,851	50,175	2,733
	stack-opt	22,356	68,380	36,428	23,213	99,472	2,733
Dilithium-2		38,408	49,380	36,212	1,598	4,112	1,572
Dilithium-3		60,836	68,836	57,724	2,830	6,588	2,691
Dilithium-5		97,692	115,932	92,788	4,826	8,779	4,706
Falcon-512		18,416	42,508	4,724	155,758	38,979	481
Falcon-1024		36,296	82,532	8,820	480,072	85,125	995

8 Corrections Against Reported Attacks

Since the beginning of KpqC round 1 then until now, there were several reported attacks aside from their impacts.

Buffer overflows. Saarinen [Saa23] reported an attack based on the buffer overflow of HAETAE version 1.0, which basically due to the non-deterministic packing and no error checking unpacking functions. From HAETAE version 2.0, this implementation error was fixed.

Bit-flipping attack. The HUFU team [Tea23] reported an attack based on an implementation error of HAETAE version 1.0, in the rANS encoding. As the HUFU team pointed out, the fix is simple, by sanity check before decoding. From HAETAE version 2.0, this implementation error was fixed, and is indeed replaced to simpler rANS with reduced symbol space, avoiding the excessive RAM usage.

Hyperball sampler. As noticed by HAETAE team during the 7th KpqC Workshop, there was a bug in the hyperball sampling algorithm, in `polyfix.c`. The variable square-sum was not initialized to zero in each loop iteration, but only in the beginning of the function. If a hyperball sample is rejected, the next candidate is automatically very small, which may be exploitable cryptanalytically. From HAETAE version 2.0, this error was fixed.

Side-channel analysis. During the 6th KpqC Workshop, Lee et al. presented a side-channel vulnerability of non-masked implementation of HAETAE, version 1.0. We give a justification on how HAETAE can be securely and efficiently masked against side-channel attacks in Appendix B, which will naturally prevent their, and other possible well-known power/electromagnetic vulnerabilities.

9 Conclusion

With HAETAE, we close an important gap between the two state-of-the-art digital signature schemes Dilithium and Falcon. Novel contributions in key generation and rejection sampling allow us to reach *significantly* smaller signature and verification key sizes, while still allowing physical side-channel protected implementations for Internet of Things (IoT) use-cases. Moreover, our first set of optimized implementations exhibits that our proposed algorithms run in feasible time both on embedded and high-end platforms and compete with existing schemes when considering sending latency.

A possible direction to achieve a similar signature sizes with HAETAE could be following the BLISS approach, using the bimodal Gaussians in the module LWE setting with the compression techniques, which we did not go through to avoid computing a secret-dependent transcendental function during signature rejection to enhance the implementation security.

Acknowledgements

Part of this work was done while Damien Stehlé was in École Normale Supérieure de Lyon and Institut Universitaire de France, and Julien Devevey was in École Normale Supérieure de Lyon, and Dongyeon Hong and MinJune Yi were in CryptoLab Inc(during KpqC round 1). Julien Devevey and Damien Stehlé were supported by the AMIRAL ANR grant (ANR-21-ASTR-0016), the PEPR quantique France 2030 programme (ANR-22-PETQ-0008) and the PEPR Cyber France 2030 programme (ANR-22-PECY-0003). Tim Güneysu, Markus Krausz, Georg Land and Marc Möller have been supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972, by the German Federal Ministry of Education and Research BMBF through the projects 6GEM (16KISK038) and FlexKI (01IS22086I), and by the European Commission under the grant agreement number 101070374 (CONVOLVE). Jung Hee Cheon and Hyeongmin Choe were supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2022R1A5A6000840).

References

[ABC$^+$23] Melissa Azouaoui, Olivier Bronchain, Gaëtan Cassiers, Clément Hoffmann, Yulia Kuzovkova, Joost Renes, Tobias Schneider, Markus Schönauer, François-Xavier Standaert, and Christine van Vredendaal. Protecting dilithium against leakage revisited sensitivity analysis and improved implementations. *IACR TCHES*, 2023(4):58–79, 2023.

[AHKS22] Amin Abdulrahman, Vincent Hwang, Matthias J. Kannwischer, and Amber Sprenkels. Faster kyber and dilithium on the cortex-M4. In Giuseppe Ateniese and Daniele Venturi, editors, *ACNS 22*, volume 13269 of *LNCS*, pages 853–871. Springer, Heidelberg, June 2022.

[AKM$^+$15] Marc Andrysco, David Kohlbrenner, Keaton Mowery, Ranjit Jhala, Sorin Lerner, and Hovav Shacham. On subnormal floating point and abnormal timing. In *2015 IEEE Symposium on Security and Privacy*, pages 623–639. IEEE, 2015.

[Bar87] Paul Barrett. Implementing the Rivest Shamir and Adleman public key encryption algorithm on a standard digital signal processor. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 311–323. Springer, Heidelberg, August 1987.

[BBC$^+$20] Daniel J. Bernstein, Billy Bob Brumley, Ming-Shing Chen, Chitchanok Chuengsatiansup, Tanja Lange, Adrian Marotzke, Bo-Yuan Peng, Nicola Tuveri, Christine van Vredendaal, and Bo-Yin Yang. NTRU Prime: round 3. https://ntruprime.cr.yp.to/nist/ntruprime-20201007.pdf, 2020.

[BBD$^+$23] Manuel Barbosa, Gilles Barthe, Christian Doczkal, Jelle Don, Serge Fehr, Benjamin Grégoire, Yu-Hsuan Huang, Andreas Hülsing, Yi Lee, and Xiaodi Wu. Fixing and mechanizing the security proof of Fiat-Shamir with aborts and Dilithium. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 358–389. Springer, Heidelberg, August 2023.

[BBE$^+$19] Gilles Barthe, Sonia Belaïd, Thomas Espitau, Pierre-Alain Fouque, Mélissa Rossi, and Mehdi Tibouchi. GALACTICS: Gaussian sampling for lattice-based constant- time implementation of cryptographic signatures, revisited. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2147–2164. ACM Press, November 2019.

[BG14] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *CT-RSA 2014*, volume 8366 of *LNCS*, pages 28–47. Springer, Heidelberg, February 2014.

[BP18] Leon Groot Bruinderink and Peter Pessl. Differential fault attacks on deterministic lattice signatures. *IACR TCHES*, 2018(3):21–43, 2018. https://tches.iacr.org/index.php/TCHES/article/view/7267.

[CL21] André Chailloux and Johanna Loyer. Lattice sieving via quantum random walks. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 63–91. Springer, Heidelberg, December 2021.

[DDGR20] Dana Dachman-Soled, Léo Ducas, Huijing Gong, and Mélissa Rossi. LWE with side information: Attacks and concrete security estimation. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 329–358. Springer, Heidelberg, August 2020.

[DDLL13] Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 40–56. Springer, Heidelberg, August 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, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 34–64. Springer, Heidelberg, December 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, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 327–357. Springer, Heidelberg, August 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 TCHES*, 2018(1):238–268, 2018. https://tches.iacr.org/index.php/TCHES/article/view/839.

[dPEK$^+$] Rafaël del Pino, Thomas Espitau, Shuichi Katsumata, Mary Maller, Fabrice Mouhartem, Thomas Prest, Mélissa Rossi, and Markku-Juhani Saarinen. Raccoon: A side-channel secure signature scheme. Submission to the NIST Additional Post-Quantum Digital Signature Schemes standardization process. Available at https://raccoonfamily.org/wp-content/uploads/2023/07/raccoon.pdf.

[Dud14] Jarek Duda. Asymmetric numeral systems: entropy coding combining speed of Huffman coding with compression rate of arithmetic coding, 2014. ArXiv preprint, available at https://arxiv.org/abs/1311.2540.

[Dwo15] Morris J Dworkin. SHA-3 standard: Permutation-based hash and extendable-output functions. *FIPS 202*, 2015.

[EFG+22] Thomas Espitau, Pierre-Alain Fouque, François Gérard, Mélissa Rossi, Akira Takahashi, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. Mitaka: A simpler, parallelizable, maskable variant of falcon. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 222–253. Springer, Heidelberg, May / June 2022.

[EFGT17] Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Side-channel attacks on BLISS lattice-based signatures: Exploiting branch tracing against strongSwan and electromagnetic emanations in microcontrollers. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1857–1874. ACM Press, October / November 2017.

[ENST23] Thomas Espitau, Guilhem Niot, Chao Sun, and Mehdi Tibouchi. Square unstructured integer euclidean lattice signature. *Submission to the NIST's post-quantum cryptography standardization process*, 2023.

[ETWY22] Thomas Espitau, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. Shorter hash-and-sign lattice-based signatures. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 245–275. Springer, Heidelberg, August 2022.

[FHK+18] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, Zhenfei Zhang, et al. Falcon: Fast-fourier lattice-based compact signatures over NTRU. *Submission to the NIST's post-quantum cryptography standardization process*, 36(5), 2018.

[Gie14] Fabian Giesen. Interleaved entropy coders. *arXiv preprint arXiv:1402.3392*, 2014.

[GKOS18] Tim Güneysu, Markus Krausz, Tobias Oder, and Julian Speith. Evaluation of lattice-based signature schemes in embedded systems. In *25th IEEE International Conference on Electronics, Circuits and Systems, ICECS 2018, Bordeaux, France, December 9-12, 2018*, pages 385–388. IEEE, 2018.

[HDS23] Abiodoun Clement Hounkpevi, Sidoine Djimnaibeye, and Michel Seck. EagleSign: A new post-quantum ElGamal-like signature over lattices. *Submission to the NIST's post-quantum cryptography standardization process*, 2023.

[HPP+23] Thomas Pornin Huang, Eamonn W Postlethwaite, Thomas Prest, Ludo N Pulles, and Wessel van Woerden. HAWK., 2023. Version 1.0.1 (July 19, 2023).

[KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, Heidelberg, April / May 2018.

[KLS$^+$23] Markus Krausz, Georg Land, Florian Stolz, Dennis Naujoks, Jan Richter-Brockmann, Tim Güneysu, and Lucie Johanna Kogelheide. Generic accelerators for costly-to-mask PQC components. *IACR Cryptol. ePrint Arch.*, page 1287, 2023.

[KRSS19] Matthias J. Kannwischer, Joost Rijneveld, Peter Schwabe, and Ko Stoffelen. pqm4: Testing and benchmarking NIST PQC on ARM cortex-M4. Cryptology ePrint Archive, Report 2019/844, 2019. https://eprint.iacr.org/2019/844.

[Li04] Ren-Cang Li. Near optimality of chebyshev interpolation for elementary function computations. *IEEE Trans. Computers*, 53(6):678–687, 2004.

[Lyu09] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 598–616. Springer, Heidelberg, December 2009.

[Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Heidelberg, April 2012.

[MGTF19] Vincent Migliore, Benoît Gérard, Mehdi Tibouchi, and Pierre-Alain Fouque. Masking Dilithium - efficient implementation and side-channel evaluation. In Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung, editors, *ACNS 19*, volume 11464 of *LNCS*, pages 344–362. Springer, Heidelberg, June 2019.

[Mon85] Peter L Montgomery. Modular multiplication without trial division. *Mathematics of computation*, 44(170):519–521, 1985.

[MUTS22] Soundes Marzougui, Vincent Ulitzsch, Mehdi Tibouchi, and Jean-Pierre Seifert. Profiling side-channel attacks on Dilithium: A small bit-fiddling leak breaks it all. Cryptology ePrint Archive, Report 2022/106, 2022. https://eprint.iacr.org/2022/106.

[Pla21] Thomas Plantard. Efficient word size modular arithmetic. *IEEE Transactions on Emerging Topics in Computing*, 9(3):1506–1518, 2021.

[Pre17] Thomas Prest. Sharper bounds in lattice-based cryptography using the Rényi divergence. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 347–374. Springer, Heidelberg, December 2017.

[Pre23] Thomas Prest. A key-recovery attack against mitaka in the t-probing model. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 205–220. Springer, Heidelberg, May 2023.

[Ros20] Melissa Rossi. *Extended Security of Lattice-Based Cryptography.* PhD thesis, École Normale Supérieure de Paris, 2020.

[RS23] Keegan Ryan and Adam Suhl. Round 1 (additional signatures) official comment: EHT, 2023.

[Saa23] Markku-Juhani O. Saarinen. Buffer overflows in HAETAE / On crypto vs implementation errors, 2023. Available at https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/ImcSqGLFdoo/m/fnL4HxZ6BQAJ.

[Sei18] Gregor Seiler. Faster AVX2 optimized NTT multiplication for ring-LWE lattice cryptography. Cryptology ePrint Archive, Report 2018/039, 2018. https://eprint.iacr.org/2018/039.

[SF23] Igor Semaev, , and Martin Feussner. Digital signature algorithms EHTV3 and EHTV4 submission to NIST PQC. *Submission to the NIST's post-quantum cryptography standardization process*, 2023.

[Tea23] HuFu Team. rans signature compression done right, 2023. Available at http://123.56.244.4/rANS.pdf.

[Tib23] Mehdi Tibouchi. Round 1 (additional signatures) official comment: EagleSign, 2023.

[VGS17] Aaron R Voelker, Jan Gosmann, and Terrence C Stewart. Efficiently sampling vectors and coordinates from the n-sphere and n-ball. *Centre for Theoretical Neuroscience-Technical Report*, 01 2017.

[YJL+23] Yang Yu, Huiwen Jia, Leibo Li, Delong Ran, Zhiyuan Qiu, Shiduo Zhang, Xiuhan Lin, and Xiaoyun Wang. HuFu: Hash-and-sign signatures from powerful gadgets. Algorithm specifications and supporting documentation., 2023. Version 1.1 (September 2, 2023).

A Recent lattice-based signatures

In this section, we give a size comparison of the recent lattice-based signature schemes. In Table 7, we compare the verification key and signature sizes (in bytes) of some selected lattice signature schemes, including the NIST's on-ramp candidates. Note that we additionally showcase some disadvantages, e.g., some recent analyses reported at pqc-forum (some are verified by the submitters, and others are not) or the use of new family of assumptions, which may degrade the claimed security.

Table 7: Size comparison of recent lattice-based signatures.

Scheme	Levels 1 & 2			Level 3			Level 5			Ref.
	vk	sig.	sum ↑	vk	sig.	sum	vk	sig.	sum	
EHTv4[1]	1,110	369	1,479	—	—	—	1,110	369	1,479	[SF23]
Falcon [2]	897	666	1,563	—	—	—	897	666	1,563	[FHK+18]
HAWK[3]	1,024	555	1,579	—	—	—	1,024	555	1,579	[HPP+23]
Mitaka [4]	896	713	1,609	—	—	—	896	713	1,609	[EFG+22]
HAETAE	992	1,474	2,466	1,472	2,349	3,821	2,080	2,948	5,028	**ours**
Dilithium	1,312	2,420	3,732	1,952	3,293	5,245	2,592	4,595	7,187	[DKL+18]
EagleSign[5]	1,824	2,144	3,968	2,842	2,336	5,160	3,616	3,488	7,104	[HDS23]
Raccoon	2,256	11,524	13,780	3,160	14,544	17,704	4,064	20,330	24,394	[dPEK+]
SQUIRRELS	682K	1,019	683K	1,600K	1,554	1,602K	—	—	—	[ENST23]
HUFU[6]	1,059K	2,450	1,061K	2,177K	3,540	2,181K	3,573K	4,520	3,578K	[YJL+23]

[1] reported attack [RS23] [2] infeasible to mask, uses floating-point arithmetic [3] new assumptions smLIP, omSVP [4] hard to mask, proposed masked Gaussian sampler recently broken [Pre23] [5] reported attack [Tib23] [6] uses floating-point arithmetic

B Security against Physical Attacks

Implementation security is a crucial aspect of making cryptosystems feasible in real-world applications. A significant advantage of HAETAE is that it can be protected against power side-channel attacks efficiently and with reasonable overhead. In this context, we emphasize the similarity of HAETAE to Dilithium. Hence, past works analyzing concrete attacks [BP18, MUTS22], but also countermeasures [MGTF19, ABC+23], mainly apply to HAETAE as well.

While there is no known method to efficiently mask Falcon, Mitaka [EFG+22] was designed to be easy to protect against implementation attacks, while still having the advantage of similarly small signatures as Falcon. For Mitaka, the crux regarding side-channel security is sampling Gaussian-distributed values. Together with Mitaka, an efficient, masked algorithm for discrete Gaussian sampling was presented. However, Prest broke its security proof recently [Pre23]. In this respect, HAETAE has the strong advantage that Gaussian sampling only needs to be secured against the much stronger Simple Power Analysis (SPA) attacker model, which allows for simpler countermeasures, while Mitaka's side-channel security will always depend on a masked sampler.

While a fully protected implementation of HAETAE is out of scope for this paper, we briefly sketch its feasibility.

Protecting the Arithmetic. Most notably, HAETAE does not deploy floating-point arithmetic at any point, and only few secrets require fix-point arithmetic. Remarkably, addition of fix-point arithmetic can be masked relatively easy, and HAETAE never requires a multiplication of fix-point values.

During signing, the most critical operation is multiplying the (public) challenge polynomial c with $\mathbf{s}$ and subsequently adding the result to $\mathbf{y}$. Since this operation may leak information about the secret key statistically over many executions, implementers must protect it accordingly. As countermeasures against these so-called Differential Power Analysis (DPA) attacks, masking has been proven effective.

This operation is straightforward to mask at arbitrary order by splitting the secret key polynomials into multiple additive shares in $\mathcal{R}_q$. A masked implementation then stores the NTT of each share of $\mathbf{s}$ and multiplies them to c, obtaining a shared $c\mathbf{s}$. Following this, the inverse NTT is applied share-wise. Since $\mathbf{y}$ is a polynomial vector in $(1/N)\mathcal{R}$, it is not trivially possible to add our shares of $c\mathbf{s} \in \mathcal{R}_q^{k+\ell}$. On the other hand, $\mathbf{y}$ is not a secret-key-dependent value. Therefore, it is not required to be protected against DPA but only against the much stronger attacker model of a SPA. In fact, coefficient-wise shuffling of the addition might be sufficient at this point.

Independent of whether the addition is shuffled or masked, this involves a masking conversion from $\mathbb{Z}_q$ to $\mathbb{Z}_{2^{32}}$. Subsequently, the computation of $2\mathbf{z} - \mathbf{y}$ and the bound checks can be shuffled without applying costly masking.

Protecting the Hyperball Sampler. The same idea applies to the whole hyperball sampling procedure. Since the order of the Gaussian samples is, in principle, irrelevant, they can be generated in random order. This is particularly an advantage for randomized HAETAE.

For the deterministic version, a masked CDT sampler, and a masked approximation of the exponential function are required. The former was shown to be feasible recently by Krausz et *al.* [KLS+23], while the latter is a sequence of multiplications, shifting by a constant amount, and addition by constants, which is expected to be costly but feasible.

It is noteworthy that the random oracle hash (which outputs the challenge) is only required to be protected against SPA as well. Since the input order into the hash function cannot be randomized, the preceding values must still be protected by masking. Therefore, if no masked hyperball sampling has been performed, we propose to perform a shuffled point-wise multiplication of $\mathbf{A}$ and $\mathbf{y}$, directly followed by freshly masking the resulting coefficients. Then, a share-wise inverse NTT and a masking conversion to the Boolean domain will be performed, which enables a secure HighBits operation. For the LSBs of

y_0, generating a fresh Boolean masking during the shuffled generation of the hyperball sample's coefficients is sufficient.

C Additional Proofs

C.1 Useful Lemma

We will rely on the following claim.

Lemma 12. *Let n be the degree of $\mathcal{R}$. Let $m, N, r > 0$ and $\mathbf{v} \in \mathcal{R}^m$. Then the following statements hold:*

1. $|(1/N)\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r)| = |\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(Nr)|$,
2. $|\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r, \mathbf{v})| = |\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r)|$,
3. $\mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(r - \frac{\sqrt{mn}}{2})) \leq |\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r)| \leq \mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(r + \frac{\sqrt{mn}}{2}))$.

Proof. For the first statement, note that we only scaled $(1/N)\mathcal{R}^m$ and $\mathcal{B}_{\mathcal{R},m}(r)$ by a factor N. For the second statement, note that the translation $\mathbf{x} \mapsto \mathbf{x} - \mathbf{v}$ maps $\mathcal{R}^m$ to $\mathcal{R}^m$.

We now prove the third statement. For $x \in \mathcal{R}^m$, we define $T_{\mathbf{x}}$ as the hypercube of $\mathcal{R}^m_{\mathbb{R}}$ centered in $\mathbf{x}$ with side-length 1. Observe that the $T_{\mathbf{x}}$'s tile the whole space when $\mathbf{x}$ ranges over $\mathcal{R}^m$ (the way boundaries are handled does not matter for the proof). Also, each of those tiles has volume 1. As any element in $T_{\mathbf{x}}$ is at Euclidean distance at most $\sqrt{mn}/2$ from $\mathbf{x}$, the following inclusions hold:

$$\mathcal{B}_{\mathcal{R},m}\Big(r - \frac{\sqrt{mn}}{2}\Big) \subseteq \bigcup_{\mathbf{x} \in \mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r)} T_{\mathbf{x}} \subseteq \mathcal{B}_{\mathcal{R},m}\Big(r + \frac{\sqrt{mn}}{2}\Big).$$

Taking the volumes gives the result. □

C.2 Proof of Lemma 4

Proof. To ease the notation, let us use $B = r'$. Let $\mathbf{y} \in \mathcal{B}_{\mathcal{R},m}(Nr' + \sqrt{mn}/2)$ and set $\mathbf{z} = \lfloor \mathbf{y} \rceil$. Note that $\mathbf{z}$ is sampled (before the rejection step) with probability

$$\frac{\mathsf{Vol}(T_{\mathbf{z}} \cap \mathcal{B}_{\mathcal{R},m}(Nr' + \sqrt{mn}/2))}{\mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(Nr'))},$$

where $T_{\mathbf{z}}$ is the hypercube of $\mathcal{R}^m_{\mathbb{R}}$ centered in $\mathbf{z}$ with side-length 1. By the triangle inequality, this probability is equal to $1/\mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(Nr' + \sqrt{mn}/2)$ when $\mathbf{z} \in \mathcal{B}_{\mathcal{R},m}(Nr')$. Hence the distribution of the output is exactly $U(\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(Nr'))$, as each element is sampled with equal probability and as the algorithm almost surely terminates (its runtime follows a geometric law of parameter the rejection probability).

It remains to consider the acceptance probability.

$$\frac{\sum_{\mathbf{y} \in \mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(Nr')} \mathsf{Vol}(T_{\mathbf{y}} \cap \mathcal{B}_{\mathcal{R},m}(Nr' + \sqrt{mn}/2))}{\mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(Nr' + \sqrt{mn}/2))}.$$

By the triangle inequality and Lemma 12, it is

$$\frac{|\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(Nr')|}{\mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(Nr' + \sqrt{mn}/2))} \geq \left(\frac{Nr' - \sqrt{mn}/2}{Nr' + \sqrt{mn}/2}\right)^{mn}.$$

Note that by our choice of N, this is $\geq 1/M_0$. □

C.3 Proof of Lemma 5

Proof. Figure 6 is the bimodal rejection sampling algorithm applied to the source distribution $U((1/N)\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r'))$ and target distribution $U((1/N)\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r))$ (see, e.g., [DFPS22]). It then suffices that the support of the bimodal shift of the source distribution by $\mathbf{v}$ contains the support of the target distribution. It is implied by $r' \geq \sqrt{r^2+t^2}$.

We now consider the number of expected iterations, i.e., the maximum ratio between the two distributions. To guide the intuition, note that if we were to use continuous distributions, the acceptance probability $1/M'$ would be bounded by $1/M$. In our case, the acceptance probability can be bounded as follows (using Lemma 12):

$$\begin{aligned}\frac{1}{M'} =& \frac{|(1/N)\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r)|}{2|(1/N)\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(r')|} = \frac{|\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(Nr)|}{2|\mathcal{R}^m \cap \mathcal{B}_{\mathcal{R},m}(Nr')|}\\ &\geq \frac{\mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(Nr-\sqrt{mn}/2))}{2\mathsf{Vol}(\mathcal{B}_{\mathcal{R},m}(Nr'+\sqrt{mn}/2))}\\ &= \frac{1}{2}\left(\frac{Nr-\sqrt{mn}/2}{Nr'+\sqrt{mn}/2}\right)^{mn}.\end{aligned}$$

It now suffices to bound the latter term from below by $1/(cM) = 1/(2c(r'/r)^{mn})$. This inequality is equivalent to:

$$c \geq \frac{1}{2}\cdot\left(\frac{r}{r-\sqrt{mn}/(2N)}\right)^{mn}\cdot\left(\frac{r'+\sqrt{mn}/(2N)}{r'}\right)^{mn},$$

and to:

$$N \geq \frac{1}{c^{1/(mn)}-1}\cdot\frac{\sqrt{mn}}{2}\left(\frac{c^{1/(mn)}}{r}+\frac{1}{r'}\right),$$

which allows to complete the proof. □

C.4 Proof of Lemma 7

Proof. By Lemma 6, there exists a unique representation

$$r = \lfloor (r+\alpha/2)/\alpha \rfloor \alpha + (r \bmod^{\pm} \alpha).$$

By identifying $\mathsf{HighBits}(r,\alpha)$ and $\mathsf{LowBits}(r,\alpha)$ in the above equation, we obtain the first result.

By definition of $\bmod^{\pm} \alpha$, we have the second range.

Finally, since $r \mapsto \lfloor (r+\alpha/2)/\alpha \rfloor$ is a non-decreasing function, it is sufficient to show that $\lfloor (2q-1+\alpha/2)/\alpha \rfloor \leq \lfloor (2q-1)/\alpha \rfloor$. We have $(2q-1+\alpha/2) \leq \lfloor (2q-1)/\alpha \rfloor \alpha + \alpha - 1$ by assumption on q. Dividing by α and taking the floor yields the result.

□

C.5 Proof of Lemma 8

Proof. Let $r \in [0, 2q-1]$. Let r_0, r_1, r_0', and r_1' defined as in Definition 8. If $r_0' = r_0$ and $r_1' = r_1$, the equality $r_0' + r_1' \cdot \alpha_{\mathsf{h}} = r_0 + r_1 \cdot \alpha_{\mathsf{h}} \bmod 2q$ holds vacuously.

If not, then $r_0' = r_0 - 2$ and $r_1' = r_1 - 2(q-1)/\alpha_{\mathsf{h}}$ and $r_0' + r_1'\alpha_{\mathsf{h}} = r_0 + r_1\alpha_{\mathsf{h}} - 2q$. By Lemma 7, we get the first equality.

The second property stems from the second property in Lemma 7. The modifications to r_0 make r_0' lie in the range $[-\alpha_{\mathsf{h}}/2 - 2, \alpha_{\mathsf{h}}/2)$.

The last property stems from the third property in Lemma 7 and the fact that if $r_1 = m$, then we have $r_1' = 0$.

□

D Additional Implementation Specification

Algorithm 1 describes how to implement the unpacking of $\widehat{\mathbf{A}}$, and in Algorithm 2 we demonstrate how to apply the CRT.

Algorithm 1 Unpacking routine for $\widehat{\mathbf{A}}$.

```
$\underline{\mathsf{unpackA}_d(\mathsf{seed}_{\mathbf{A}}, \psi)}$
if $d = 0$ then
    $\widehat{\mathbf{A}}_{\mathsf{gen}} := \mathsf{expandA}_d(\mathsf{seed}_{\mathbf{A}})$
    $\widehat{\mathbf{b}} := \psi$
else
    $(\mathbf{a}_{\mathsf{gen}}, \widehat{\mathbf{A}}_{\mathsf{gen}}) := \mathsf{expandA}_d(\mathsf{seed}_{\mathbf{A}})$
    $\widehat{\mathbf{b}} := 2 \cdot \mathsf{NTT}(\mathbf{a}_{\mathsf{gen}} - \psi \cdot 2^d) \bmod q$
return $\widehat{\mathbf{A}} \in \mathcal{R}_q^{k \times \ell} := (\widehat{\mathbf{b}} \mid 2 \cdot \widehat{\mathbf{A}}_{\mathsf{gen}}) \bmod q$
```

E Notes Regarding Hardware Implementations

Hashing and generation of randomness are the most time-consuming operations of HAETAE. Therefore, we assume that hardware implementations will bring significant speedup and can be competitive to Dilithium, particularly through efficient Keccak cores. Furthermore, hardware implementations will benefit significantly from applying the offline approach. Naturally, a module generating hyperball samples can be instantiated and run parallel to the online phase, thus, hiding its latency behind the online phase. Moreover, high-speed applications could adopt the offline approach with designated signing key, including the multiplication of $\mathbf{A}$ and $\mathbf{y}$, to further reduce the latency of the online phase.

Algorithm 2 Mapping from $(\mathcal{R}_q^k, \mathcal{R}_q)$ to $\mathcal{R}_{2q}^k$

```
$\underline{\mathsf{fromCRT}(\mathbf{w}, x)}$
parse $\mathbf{w}$ as vector of integers $\overline{\mathbf{w}}$ of size $kn$
parse $x$ as vector of integers $\overline{x}$ of size $n$
for $i := 0$ to $n - 1$ do
    if $\mathsf{LSB}(\overline{x}_i) = \mathsf{LSB}(\overline{\mathbf{w}}_i)$ then          ▷ Implement in constant time.
        $\overline{\mathbf{w}}'_i := \overline{\mathbf{w}}_i$
    else
        $\overline{\mathbf{w}}'_i := \overline{\mathbf{w}}_i + q$
for $j := 1$ to $k - 1$ do
    for $i := 0$ to $n - 1$ do
        if $\mathsf{LSB}(\overline{\mathbf{w}}_{nj+i}) = 0$ then          ▷ Implement in constant time.
            $\overline{\mathbf{w}}'_{nj+i} := \overline{\mathbf{w}}_{nj+i}$
        else
            $\overline{\mathbf{w}}'_{nj+i} := \overline{\mathbf{w}}_{nj+i} + q$
arrange $\overline{\mathbf{w}}'$ to $\mathbf{w}'$, an element in $\mathcal{R}_{2q}^k$
return $\mathbf{w}'$
```

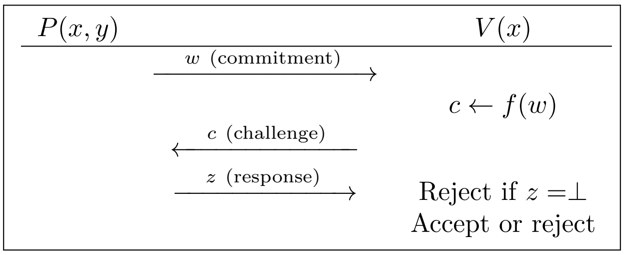

Figure 19: Interaction between P and V. We let $\langle P(x,y) \leftrightarrow V(x)\rangle_{f(\cdot)}$ denote the transcript (w, c, z).

F Cryptographic Reminders

In this section, we recall the definition of a canonical identification scheme, which is the main building brick of HAETAE and the Fiat-Shamir with aborts transform, which allows one to turn a canonical identification scheme into a signature scheme.

F.1 Canonical Identification Schemes

We start by defining a canonical identification scheme.

Definition 9 (Canonical Identification Scheme with Aborts (CID))**.** A canonical identification scheme ID for a relation R is a three rounds interactive protocol between a prover P and a verifier V, as defined in Figure 19. The challenge is generated from $f(w)$ set as the distribution $U(\mathcal{C})$ for some challenge set $\mathcal{C}$. The prover holds a pair $(x, y) \in R$ while the verifier only has x, where the pair (x, y) was generated by a PPT algorithm Gen, called an instance generator. The event "$z = \perp$" is called an abort, and its probability β is called the probability of aborting.

The canonical identification schemes considered in this work are such that given c and z, there is only one w such that V accepts, which can be efficiently computed from c and z.

We need the following flavor of zero knowledge, denoted perfect accepting honest verifier zero-knowledge.

Definition 10 (Perfect Accepting Honest Verifier Zero-Knowledge)**.** Let ID be a canonical identification scheme. It satisfies paHVZK if there exists a PPT simulator Sim such that the output distribution of (w, c, z) resulting from the interaction $\langle P(x,y) \leftrightarrow V(x)\rangle_{U(\mathcal{C})}$ conditioned on $z \neq \perp$ and (w', c', z') generated by $\mathsf{Sim}(x)$ for any (x, y) generated by $\mathsf{Gen}(1^\lambda)$ are identical.

In the case of the HAETAE CID, the challenge c can be sampled uniformly from the challenge space $\mathcal{C}$ and passed over as input to the simulator Sim.

To avoid using the same commitment twice, we require the min-entropy of the commitment to be high.

Definition 11 (Commitment Min-Entropy)**.** For $\alpha \geq 0$, we say that an identification scheme ID with instance generator Gen has commitment min-entropy α if $H_\infty[w|(w, c, z) \leftarrow \langle P(x,y) \leftrightarrow V(x)\rangle_{U(\mathcal{C})}] \geq \alpha$, for all $(x, y) \leftarrow \mathsf{Gen}(1^\lambda)$.

Finally we need the notion of computational unique response to enable the strong unforgeability property of the final signature scheme.

Definition 12 (Computational Unique Response)**.** Let ID be a canonical identification scheme with instance generator Gen. The advantage $\mathsf{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathcal{A})$ of an adversary $\mathcal{A}$ against

```
KeyGen($1^\lambda$):
    $(x, y) \leftarrow \mathsf{Gen}(1^\lambda)$
    $(\mathsf{vk}, \mathsf{sk}) = (x, (x, y))$
    return $(\mathsf{vk}, \mathsf{sk})$

Sign($\mathsf{sk}, \mu$):
    While $z = \perp$
        $(w, c, z) \leftarrow \langle P(x, y) \leftrightarrow V(x)\rangle_{H(\cdot,\mu)}$
    return $\sigma = (c, z)$

Verify($\mathsf{vk}, \mu, \sigma$):
    Parse $\sigma = (c, z)$
    Recover $w$ from $\sigma$
    return $c = H(w, \mu)$
```

Figure 20: Fiat-Shamir with Aborts tranform $\mathsf{FS}[\mathsf{ID}, H]$.

the CUR property of ID is its probability of computing (w, c, z, z') such that $V(x)$ accepts both (w, c, z) and (w, c, z') and $z \neq z'$ on input x generated from $(x, y) \leftarrow \mathsf{Gen}(1^\lambda)$.

F.2 The Fiat-Shamir Transform

We now briefly recall how to turn a canonical identification scheme into a full-fledge signature scheme. To do so, the signing scheme runs the interactive protocol and outputs the transcript minus the commitment. The challenge is however generated by hashing the message along with the commitment to prevent an adversary from tampering with it.

We recall the assumptions necessary to show the unforgeability of the resulting signature in the following theorem. The following theorem works in the case of quantum adversaries working in the QROM, but tighter statements can be obtained when restricting to classical adversaries in the ROM.

Theorem 4 (Adapted from [BBD+23, Theorem 2])**.** *Let* ID *be a canonical identification scheme with* α *commitment min-entropy and that satisfies* paHVZK *with probability of aborting* β*. For any quantum adversary* $\mathcal{A}$ *against the* $\mathsf{UF\text{-}CMA}$ *security of* $\mathsf{FS}[\mathsf{ID}, H]$ *in the QROM making* Q_s *signature queries and* Q_h *hash queries, there exists an adversary* $\mathcal{B}$ *against the* $\mathsf{UF\text{-}NMA}$ *security of* $\mathsf{FS}[\mathsf{ID}, H]$ *in the QROM such that:*

$$\mathsf{Adv}^{\mathsf{UF\text{-}CMA}}_{\mathsf{FS}[\mathsf{ID},H]}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{UF\text{-}NMA}}_{\mathsf{FS}[\mathsf{ID},H]}(\mathcal{B}) + \frac{2^{-\alpha/2+1}Q_s}{1-\beta}\sqrt{Q_h + 1 + \frac{Q_s}{1-\beta}} + 2^{-\alpha/2+1}(Q_h+1)\sqrt{\frac{Q_s}{1-\beta}}.$$

Furthermore, if $\mathcal{A}$ *is an adversary against the* $\mathsf{sUF\text{-}CMA}$ *security, there exists an adversary* $\mathcal{B}'$ *against the CUR property of* ID *such that the previous bound holds by adding* $\mathsf{Adv}^{\mathsf{CUR}}_{\mathsf{ID}}(\mathcal{B}')$ *on the right-hand side.*

The strong unforgeability statement comes from the fact that in the last game of their proof, the reduction fails if and only if the forgery uses the same commitment (and thus reprogrammed challenge) than the signature query. As this requires a different answer, this corresponds to an attack against the computational unique response property of the canonical identification scheme.

G Fixed-point Sampling Analysis

In this appendix, we quantify the precision needed to sample from the uniform-hyperball distribution by sampling from the discrete Gaussian distribution.

We first describe the representation of numbers and operations. A fixed-point number in precision p will consist in a p-bit signed integer $k \in \mathbb{Z} \cap [-2^{p-1}, 2^{p-1})$ along with an implicit scaling exponent e: the represented number is $x = k \cdot 2^{e-p} \in [-2^{e-1}, 2^{e-1})$. The data can for example be stored in a p-bit integer in two's complement representation. The scaling exponent e is not stored, it only exists on paper. For convenience, a precision p fixed-point number x with implicit exponent e will be referred to as a (p, e)-number.

When performing arithmetic operations on fixed-point numbers, particular care must be taken with overflows: in the analysis, we make sure that during the algorithm execution, any (p, e)-number x will satisfy $|x| < 2^{e-1}$. The following assumes no overflow occurs. We can add, subtract and negate (p, e)-numbers exactly (note that we only consider the situation where the operands of those operations share the same exponent). We assume that we can multiply (p, e_0)-number x_0 with a (p, e_1)-number x_1 into a $(p, e_\times)$-number $x_\times$ as if the multiplication was exact and then rounded to a nearest representable number. Finally, we assume that we can compute an inverse square-root of a (p, e)-number x into a (p, e')-number y with possibly slightly more error than that. This is summarized as follows:

$$x_0 \oplus x_1 = x_0 + x_1; \quad x_0 \ominus x_1 = x_0 - x_1; \quad \ominus x = -x;$$
$$|(x_0 \otimes_{e_0,e_1}^{e_\times} x_1) - (x_0 \cdot x_1)| \le 2^{e_\times - p - 1}; \quad |(1/\sqrt{\cdot})_e^{e_t} - 1/\sqrt{x}| \le 2^{e_t - p}.$$

G.1 Discrete Gaussian samples as rounded Gaussian samples

The hyperball-uniform sampler relies on an algorithm that samples from the continuous Gaussian distribution. In our fixed-point sampling algorithm, we make do with fixed-point approximations to sample from the discrete Gaussian distribution, as described in Section 5.1.1. We analyze how this changes the final distribution in the following lemma, which could be rephrased using the smooth Rényi divergence introduced in [DFPS22].

Lemma 13. *Let $\sigma > 0$. Let $D_{\mathbb{Z},\sigma}$ (resp. D_σ) be the distribution D over $\mathbb{Z}$ (resp. $\mathbb{R}$) such that $D(k) \sim \exp(-k^2/(2\sigma^2))$ for all $k \in \mathbb{Z}$ (resp. $k \in \mathbb{R}$). Then we have:*

$$\Pr_{k \leftarrow D_{\mathbb{Z},\sigma}}[|k| \ge 14 \cdot \sigma] \le 2^{-140} \quad \textit{and} \quad \max_{|k| \le 14 \cdot \sigma} \frac{D_{\mathbb{Z},\sigma}(k)}{\lfloor D_\sigma \rceil(k)} \le \frac{1}{1 - 8/\sigma}.$$

Proof. Using the discrete Gaussian tail bound from [Lyu12, Lemma 4.4], the weight of $D_{\mathbb{Z},\sigma}$ out of the interval $[-14 \cdot \sigma, 14 \cdot \sigma]$ is $\le 2^{-140}$. Using the Poisson Summation Formula, we have that:

$$\forall k \in \mathbb{Z},\ D_{\mathbb{Z},\sigma}(k) \le \frac{\exp(-k^2/(2\sigma^2))}{\sigma\sqrt{2\pi}}.$$

Further, for $k \in \mathbb{Z} \cap [-14 \cdot \sigma, 14 \cdot \sigma]$, the following inequalities hold:

$$\begin{aligned}
\lfloor D_\sigma \rceil(k) &= \frac{1}{\sigma\sqrt{2\pi}} \int_{k-1/2}^{k+1/2} \exp(-x^2/(2\sigma^2)) \mathrm{d}x \\
&\ge \frac{\exp(-k^2/(2\sigma^2))}{\sigma\sqrt{2\pi}} \cdot \exp(-(|k| + 1/4)/(2\sigma^2)) \\
&\ge \frac{\exp(-k^2/(2\sigma^2))}{\sigma\sqrt{2\pi}} \cdot \big(1 - \frac{|k| + 1/4}{2\sigma^2}\big) \\
&\ge \frac{\exp(-k^2/(2\sigma^2))}{\sigma\sqrt{2\pi}} \cdot \big(1 - \frac{8}{\sigma}\big).
\end{aligned}$$

This completes the proof. □

Lemma 13 implies that a signature forger for the imperfect Gaussian sampler succeeds with essentially the same probability with the ideal Gaussian sampler.

G.2 Normalization of Gaussian samples

In the following, we assume that we have access to arbitrarily many statistically independent (p, e)-numbers $\overline{y_i}$ that approximate (perfect) samples y_i from $D_1 = \mathcal{N}(0, 1)$. We first consider the algorithm of Figure 1 with radius 1, using fixed-point arithmetic. We show that the vector $\overline{\mathbf{y}}$ output by the approximate algorithm is close to the vector $\mathbf{y}$ output by the exact algorithm. As $\mathbf{y}$ is uniformly distributed in a hyperball, the computed vector $\overline{\mathbf{y}}$ is an approximation to such a sample.

We first bound the quantities involved during the computations. These bounds are for the exact quantities. To avoid overflows, we actually need them for the corresponding computed quantities. We will see later that as the numerical errors are low, the bounds still essentially hold. The bounds are probabilistic, and hold with probability extremely close to 1.

Lemma 14. *Let* $d_{\min} = 6 \cdot 256 + 2$ *and* $d_{\max} = 11 \cdot 256 + 2$*. The following bounds hold for all* $d \in [d_{\min}, d_{\max}]$*:*

$$\Pr_{y \leftarrow D_1}[|y| \geq 2^4] < 2^{-188},$$

$$\Pr_{\substack{y_i \leftarrow D_1 \\ \forall i \in [d]}}[\|\mathbf{y}\|^2 \geq 2^{12}] < 2^{-144}, \qquad \Pr_{\substack{y_i \leftarrow D_1 \\ \forall i \in [d]}}[\|\mathbf{y}\|^2 \leq 2^{9}] < 2^{-144}.$$

$$\Pr_{\mathbf{z} \leftarrow U(\mathcal{B}_{d-2}(1))}[|z_1| \geq 2^{-2}] < 2^{-150}.$$

Proof. The first probability is $1 - \mathrm{erf}(2^4/\sqrt{2})$. The two others are bounded with the Laurent-Massart bounds for the chi-squared distribution, i.e., for all d, t:

$$\Pr_{\substack{y_i \leftarrow D_1 \\ \forall i \in [d]}}[\|\mathbf{y}\|^2 \geq d + 2\sqrt{dt} + 2t] \quad \leq \exp(-t) \text{ and } \Pr_{\substack{y_i \leftarrow D_1 \\ \forall i \in [d]}}[\|\mathbf{y}\|^2 \leq d - 2\sqrt{dt}] \quad \leq \exp(-t).$$

For the last bound, we use [DFPS22, Lemma A.13]. The probability is exactly $I_{1-1/\eta^2}((d+1)/2, 1/2)$ where I refers to the regularized incomplete Beta function and $1/\eta$ is probabilistic magnitude upper bound. We then use numerical computations. □

Throughout the execution of the approximate version of the algorithm of Figure 1, we fix the precision to $p \geq 64$. The implicit exponents vary depending on the algorithm step: the y_i's are represented by $(p, 5)$-numbers, their squares by $(p, 13)$-numbers, the squared-norm $\|\mathbf{y}\|^2$ by a $(p, 13)$-number, the inverse-norm $1/\|\mathbf{y}\|$ by a $(p, -3)$-number and the output coordinates on $(p, -1)$-numbers.

Assume that we have $|\overline{y_i} - y_i| \leq \epsilon_0$ for all i, for some $\epsilon_0 \geq 2^{-p+5}/2 = 2^{-p+4}$. To avoid overflows of $\overline{y_i}$'s, it suffices that $|y_i| \leq 2^4 - 2^{-p+5} - \epsilon_0$. The first bound from Lemma 14 still holds for any $\epsilon_0 \leq 2^{-5}$.

We now consider the computations of the approximations $\overline{y_i^2}$'s to the y_i^2's. We have:

$$\begin{aligned} \left|\overline{y_i^2} - y_i^2\right| &\leq |(\overline{y_i} \otimes_{5,5}^{13} \overline{y_i}) - \overline{y_i}^2| + |\overline{y_i} - y_i| \cdot |\overline{y_i} + y_i| \\ &\leq 2^{-p+12} + |\overline{y_i} - y_i| \cdot [|\overline{y_i} - y_i| + 2|y_i|] \\ &\leq 2^{-p+12} + 2^6 \cdot \epsilon_0. \end{aligned}$$

As addition is exact, we obtain:

$$\left|\overline{\|\mathbf{y}\|^2} - \|\mathbf{y}\|^2\right| \leq d_{\max} \cdot (2^{-p+12} + 2^6 \cdot \epsilon_0) =: \epsilon_1.$$

To avoid overflow of $\overline{\|\mathbf{y}\|^2}$ and hence of the $\overline{y_i^2}$'s, it suffices that $\|\mathbf{y}\|^2 \leq 2^{12} - 2^{-p-13} - \epsilon_1$. The second bound from Lemma 14 still holds for any $\epsilon_0 \leq 2^{-5}$.

We continue with the inverse square root computation. The following holds:

$$\begin{aligned}
\left|\overline{1/\|\mathbf{y}\|} - \frac{1}{\|\mathbf{y}\|}\right| &\leq \left|(1/\sqrt{\cdot})^{-3}_{13}(\overline{\|\mathbf{y}\|^2}) - \frac{1}{\overline{\|\mathbf{y}\|}}\right| + \left|\frac{1}{\|\mathbf{y}\|} - \frac{1}{\overline{\|\mathbf{y}\|}}\right| \\
&\leq 2^{-p-3} + \frac{|\overline{\|\mathbf{y}\|^2} - \|\mathbf{y}\|^2|}{2[\|\mathbf{y}\|^2 - |\overline{\|\mathbf{y}\|^2} - \|\mathbf{y}\|^2|]^{3/2}} \\
&\leq 2^{-p-3} + \frac{\epsilon_1}{2[2^9 - \epsilon_1]^{3/2}} \\
&\leq 2^{-p-3} + 2^{-15}(1 + 2^{-1})\epsilon_1 =: \epsilon_2,
\end{aligned}$$

where the last inequality holds for any $\epsilon_0 \leq 2^{-15}$. To avoid overflow of $\overline{1/\|\mathbf{y}\|}$, it suffices that $1/\|\mathbf{y}\| \leq 2^{-4} - 2^{-p-3} - \epsilon_2$. The third bound from Lemma 14 holds for any $\epsilon_0 \leq 2^{-15}$.

We finally evaluate the accuracy of the output vector $\bar{\mathbf{z}}$ with respect to $\mathbf{z} := (y_1, \ldots, y_d)^\top / \|\mathbf{y}\|^2$. We have, for all i:

$$\begin{aligned}
\left|\overline{z_i} - z_i\right| &\leq \left|\overline{y_i} \otimes^{-1}_{5,-3} \overline{1/\|\mathbf{y}\|} - \overline{y_i} \cdot \overline{1/\|\mathbf{y}\|}\right| + \left|\overline{y_i} \cdot \overline{1/\|\mathbf{y}\|} - y_i/\|\mathbf{y}\|\right| \\
&\leq 2^{-p-2} + \left|\overline{1/\|\mathbf{y}\|} - 1/\|\mathbf{y}\|\right| \cdot |\overline{y_i}| + |\overline{y_i} - y_i|/\|\mathbf{y}\| \\
&\leq 2^{-p-2} + 2^5 \cdot \epsilon_2 + 2^{-4} \cdot \epsilon_0 =: \epsilon_3.
\end{aligned}$$

To avoid overflow of $\overline{z_i}$, it suffices that $|z_i| \leq 2^{-2} - 2^{-p-3} - \epsilon_3$. The fourth bound from Lemma 14 still holds for any $\epsilon_0 \leq 2^{-20}$.

Note that ϵ_3 is of the order of $2^{14}2^{-p}$. This is a crude upper bound, as it assumes that errors are always in the same direction.

G.3 Rounding Continuous Hyperball-uniform

Let us now consider the algorithm from Figure 3. We consider a uniform $\mathbf{y}_0 \leftarrow U(\mathcal{B}_{\mathcal{R},m}(1))$, its representation $\bar{\mathbf{y}}_0$ such that $\|\mathbf{y}_0 - \bar{\mathbf{y}}_0\|_\infty \leq \epsilon_4$, the scaled sample $\mathbf{y} = r'' \cdot \mathbf{y}_0$ and its representation $\bar{\mathbf{y}} = r'' \cdot \bar{\mathbf{y}}_0$, computed exactly, where we let $r'' = NB + \sqrt{mn}/2$. In particular $\|\mathbf{y} - \bar{\mathbf{y}}\|_\infty \leq r''\epsilon_4$.

Now given any $\mathbf{t}$ with $\|\mathbf{t}\| \leq r' = NB$, the probability that the random variable $\mathbf{y}$ gets rounded to it is exactly the inverse of the volume of the hyperball $V_{mn}(r'')$. However, when considering $\bar{\mathbf{y}}$, this probability lies in $[(1 - 2r''\epsilon_4)^{mn}/V_{mn}(r''), (1 + 2r''\epsilon_4)^{mn}/V_{mn}(r'')]$. The ratio of the probability densities thus always lies in $[(1 - 2r''\epsilon_4)^{mn}, (1 + 2r''\epsilon_4)^{mn}]$.

G.4 Rejection Sampling with Approximate Source Distribution

Here, we study the evolution of the rejection sampling steps from Figure 3 and Figure 6 when used with real, approximate distributions.

Theorem 5. *Let $M > 1$. Let $\delta_M \geq 1 \geq \delta_m > 0$. Let Q_r, Q_i and P_i such that for any $x \in \mathsf{Supp}(P_i)$, $Q_r(x)/Q_i(x) \in [\delta_m, \delta_M]$ and $R_\infty(P_i \| Q_i) \leq M$. Consider P_r, the probability distribution whose mass function is proportional to $x \mapsto P_i(x)/(MQ_i(x))Q_r(x)$. Then $P_r(x)/P_i(x) \in [\delta_m/\delta_M, \delta_M/\delta_m]$ for any $x \in \mathsf{Supp}(P_i) = \mathsf{Supp}(P_r)$.*

Note that P_r is the resulting distribution when using rejection sampling to go from Q_i to P_i with source distribution Q_r.

Proof. The normalization factor is

$$\sum_{z \in \mathsf{Supp}(P_i)} \frac{P_i(z)}{M} \cdot \frac{Q_r(z)}{Q_i(z)} \in \left[\frac{\delta_m}{M}, \frac{\delta_M}{M}\right] .$$

As we consider positive values, its inverse then lies in $[M/\delta_M, M/\delta_m]$. We conclude by considering that $Q_r(x)/(MQ_i(x))$ lies in $[\delta_m/M, \delta_M/M]$ for any $x \in \mathsf{Supp}(P_i)$. □

This lemma is applied twice in a row, first to get the final discrete hyperball sample $\mathbf{y}$ and second to get the final signature $\mathbf{z}$. In the second case, note that instead of bounding from above and below the distribution of $\mathbf{z}$ (before rejection), we can use the bound we get for $\mathbf{y}$, due to data processing arguments.

G.5 Putting Everything together

The final interval in which lies the ratio of real and ideal distributions of $\mathbf{z}$ is

$$\left[\left(\frac{1-2r''\epsilon_4}{1+2r''\epsilon_4}\right)^{2mn}, \left(\frac{1+2r''\epsilon_4}{1-2r''\epsilon_4}\right)^{2mn}\right] ,$$

where $\epsilon_4 \approx 2^{14-p}$, $r'' \leq 2^{28}$ and $mn \leq 2^{11}$. As $\max(1-1/x, x-1) = x-1$ for positive $x \leq 1$, we apply [Pre17, Lemma 3] and follow their subequently described strategy to derive the necessary precision. In order to get about 2^{-37} precision in the bounds after rejection sampling, it is necessary to choose $p \approx 14 + 28 + 37 + 1$, as the radius $r'' \leq 2^{28}$. Going back to the σ of our discrete Gaussian distribution, as elements sampled from it will have absolute value at most $14\sigma \approx p$, we choose $\sigma = 2^{76}$, which is also sufficiently big to use discrete Gaussians instead of rounded ones.