

Quantum-safe HIBE: does it cost a LATTE?

Raymond K. Zhao

Faculty of Information Technology,
Monash University
Clayton, Australia
raymond.zhao@monash.edu

Sarah McCarthy*

Institute for Quantum Computing,
University of Waterloo
Waterloo, Canada
sarah.mccarthy@uwaterloo.ca

Ron Steinfeld

Faculty of Information Technology,
Monash University
Clayton, Australia
ron.steinfeld@monash.edu

Amin Sakzad

Faculty of Information Technology,
Monash University
Clayton, Australia
amin.sakzad@monash.edu

Máire O'Neill

Centre for Secure Information
Technologies, Queen's University
Belfast
Belfast, United Kingdom
m.oneill@ecit.qub.ac.uk

ABSTRACT

The UK government is considering advanced primitives such as identity-based encryption (IBE) for adoption as they transition their public-safety communications network from TETRA to an LTE-based service. However, the current LTE standard relies on elliptic curve based IBE, which will be vulnerable to quantum computing attacks, expected within the next 20–30 years. Lattices can provide quantum-safe alternatives for IBE. These schemes have shown promising results in terms of practicality. To date, several IBE schemes over lattices have been proposed but there has been little in the way of practical evaluation.

This paper provides the first complete C implementation and benchmarking of LATTE, a promising Hierarchical IBE scheme proposed by the United Kingdom (UK) National Cyber Security Centre (NCSC) in 2017 and endorsed by European Telecommunications Standards Institute (ETSI). We propose optimisations for the KeyGen, Delegate, Extract and Gaussian sampling components of LATTE thereby increasing attack costs, reducing decryption key lengths by $2x-3x$, ciphertext sizes by up to 33% and improving speed. In addition, we conduct a precision analysis, bounding the Rényi divergence of the Gaussian sampling procedures from the ideal distribution, in corroboration of our claimed security levels. Our resulting implementation of the Delegate function takes 0.4 seconds at 80-bit security level on a desktop machine at 4.2GHz, significantly faster than the order of minutes estimated in the ETSI technical report. Furthermore, our optimised LATTE Encrypt/Decrypt implementation reaches speeds up to 5.8x faster than the ETSI implementation.

KEYWORDS

lattice-based cryptography, hierarchical identity-based encryption, advanced primitives, software design, post-quantum

1 INTRODUCTION

The UK Government anticipates the migration of its mission-critical communications network from Airwave TETRA to LTE-based Emergency Services Network (ESN) [46] will be complete by 2026 [18].

However, the current standard [2] relies on Elliptic Curve (ECC)-based Identity Based Encryption (IBE) scheme MIKEY-SAKKE for securing messages. An IBE scheme removes the need for a certificate repository by deriving a user's public key from their already established public identity. This provides a low latency setup with instantaneous communication capabilities, hence is ideal for this use-case. However, ECC will be rendered insecure under quantum computing attacks, as acknowledged by current post-quantum cryptography standardization efforts by National Institute of Standards and Technology (NIST) [44].

One of the advantages of lattice-based cryptography (LBC), a contender for quantum-secure cryptographic solutions, is the ability to build advanced primitives such as IBE. Furthermore, a hierarchy can be built into an IBE scheme to provide a more distributed workload and allow for finer grained control over private key distribution. Hierarchical identity-based encryption (HIBE) schemes extend the concept of using a personal identity as a public key to a multi-levelled scenario, such as one would find within a functioning company. HIBE has further applications such as forward-secure encryption [11] and public key broadcast encryption [21]. However, it is still new territory within the post-quantum field. Additionally, with the growth of the Internet of Things (IoT), which brings with it complex interconnected systems of constrained devices, there is a greater requirement for lightweight, advanced primitives unlike ever before. The long-term security considerations indicate that these should be made quantum-secure today. The aim of this paper is to assess the practicality and optimise the implementation/integration of a quantum-safe HIBE scheme.

The DLP-IBE scheme [22] was in 2017 combined with the Bonsai tree HIBE scheme introduced in [12] to create LATTE by Campbell and Groves [10]. This research was carried out by the National Cyber Security Centre (NCSC), with a view to utilising the scheme in UK public-safety communications. They are currently working with the European Telecommunications Standards Institute (ETSI) in a move towards standardising the scheme [1]. However, the proposed specification [1] only provides the Encrypt and Decrypt performance results, and it is unclear if LATTE KeyGen, Delegate, and Extract are practical at all. There remains substantial analysis to be performed to determine if and how this scheme will work in the real world. This is the gap our research endeavours to bridge.

*Also with Centre for Secure Information Technologies, Queen's University Belfast.

This paper provides the first performance benchmarking of a quantum-safe HIBE scheme, *LATTE*, written in C.¹ We also identify bottlenecks, propose optimisations, and provide further statistical and security analysis for *LATTE* and consider its suitability for such applications. In more detail, the contributions of this paper are:

- **Precision Analysis of *LATTE*:** We develop a statistical model for floating-point arithmetic errors in our efficient *Latte* implementation, verified by experimental analysis. This allows us to quantify the security impact on *Latte* of the arithmetic precision. In particular, we bound the Rényi divergence from ideal, as recommended in [49], of the Gaussian lattice sampler (with its underlying fast ff-Sampling algorithm), and deduce that 53 bits of precision retains our claimed security levels for *LATTE*-1 and 2 with up to 2^{44} key Extract/Delegate queries. For *LATTE*-3 and 4, our analysis shows that about 90 bits precision should be sufficient.
- **Optimised *LATTE* (Sub-)Algorithms:** We first reduce the module dimension of the extracted user keys by 1 compared to [1], by extending a similar approach used in the DLP IBE [22]. This leads to faster performance and reduces user private key sizes by $2x-3x$ and ciphertext length by up to 33%. In addition, we also show a faster fFLDL algorithm for (Mod)NTRU basis in Sec. 5.1. We then adapt the NTRU-Solve function from *FALCON* [51] in order to efficiently solve the NTRU equation in our optimised *LATTE* KeyGen algorithm. The NTRUSolve is both faster and more compact [48] than the resultant method [22] used in [1]. In addition, we adapt the technique from *MODFALCON* [17] and the length reduction technique by using Cramer’s rule [1] in order to efficiently solve the NTRU equation for higher lattice dimensions in our optimised *LATTE* Delegate algorithm. We further adapt the FFT sampling procedures from *FALCON* [51], which is faster than the Klein-GPV sampler [26] used in [1]. In addition, the proposed *LATTE* specification [1] did not discuss the integer discrete Gaussian sampling techniques suitable for the needed standard deviations. We integrate efficient sampling techniques including FACCT [56] and the variant [54] of COSAC [55] in our optimised *LATTE* implementation.
- **New Parameter Sets for *LATTE*:** We provide slightly revised parameter sets for *LATTE*, fixing a bug in the computation of a lattice smoothing parameter in the ETSI technical report [1], and also modify a Gaussian sampling standard deviation parameter to accommodate the more efficient FACCT [56] sampler for the Key Generation algorithm. Security estimates for these revised parameters is also presented, as we discover that our redesign reduces decryption failure rate and increases the cost of recovering the user key.
- **First Full Implementation of *LATTE*:** We give the first complete performance results of a lattice-based HIBE scheme, including the KeyGen, Delegate, and Extract algorithms, for which the implementation results were unclear in [1]. The proposed specification [1] estimated that Delegate would

have run-time in the order of minutes on a desktop machine. In contrast, we show that our efficient implementation can perform the Delegate function in 0.4s (resp. 1.3s) for 80-bit (resp. 160-bit) security level on a desktop machine. In addition, for the same ring dimension, our optimised *LATTE* implementation is up to 8.1x faster than the DLP IBE implementation result from [1] for the corresponding algorithms, and our *LATTE* Extract run-time overhead is less than 3.7x over the *FALCON* Sign algorithm run-time with the same lattice dimension.

The structure of the paper is as follows. Sec. 2 gives the background to HIBE and the lattice-based concepts used in HIBE schemes. Sec. 3 describes our improved *LATTE* HIBE scheme. Sec. 4 provides the precision and security analyses. Sec. 5 discusses our implementation techniques in making the scheme practical for real world applications. Performance results for the scheme are given in Sec. 6.

2 PRELIMINARIES

A lattice can be expressed as a collection of integer linear combinations of a set of basis vectors. Popular underlying hard lattice problems believed to be secure against quantum computing attacks include the Shortest Vector Problem (SVP) and Learning With Errors (LWE) alongside its ring variant (over ideal lattices), Ring-LWE. These are all concerned with finding short vectors in the lattice which can be attempted to be solved by lattice reduction algorithms such as LLL [36] and BKZ [14, 53]. Another common lattice problem is the NTRU assumption [29]; that is, given a polynomial \mathbf{h} , one must find \mathbf{f}, \mathbf{g} such that $\mathbf{h} = \mathbf{g} \cdot \mathbf{f}^{-1}$.

In this paper, vectors or, interchangeably through the canonical embedding, polynomials will be denoted by bold small letters like \mathbf{f} , matrices \mathbf{M} , polynomial ring mod q as \mathcal{R}_q , and lattices as Λ . The field of integers mod q is denoted as \mathbb{Z}_q . Discrete Gaussian distributions with centre t and standard deviation σ are denoted as $\mathcal{D}_{\sigma,t}$, and we omit the center if it is zero i.e. \mathcal{D}_σ if $t = 0$. A distribution is B -bounded for some $B \in \mathbb{R}^+$, if its support is in the interval $[-B, B]$. The smoothing parameter of \mathbb{Z} is denoted as $\eta_\varepsilon(\mathbb{Z}) = (1/\pi)\sqrt{\ln(2+2/\varepsilon)}/2$. The Euclidean norm of a vector/polynomial \mathbf{f} is denoted $\|\mathbf{f}\|$. The transpose \mathbf{f}^* of polynomial $\mathbf{f} = f_0 + f_1x + \dots + f_{N-1}x^{N-1}$ is defined as $\mathbf{f}^* = f_0 - f_{N-1}x - \dots - f_1x^{N-1}$. We denote \mathbf{M}^* as the transpose of matrix \mathbf{M} where $M_{i,j}^* = (M_{j,i})^*$. The Hermitian product of vectors \mathbf{a}, \mathbf{b} is denoted as $\langle \mathbf{a}, \mathbf{b} \rangle$. The concatenation of several vectors $\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_N$ will be written as $(\mathbf{f}_1|\mathbf{f}_2|\dots|\mathbf{f}_N)$. In HIBE schemes, user identities at level ℓ are denoted by ID_ℓ . A hash function from an arbitrary length input to a vector of integers of length N is written as $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^N$. An arrow \leftarrow_s is used to show the uniform random sampling of an element from a set e.g. $\mathbf{f} \leftarrow_s \mathbb{Z}_q^N$. The operator \oplus means XOR. A Gram-Schmidt orthogonalised basis of \mathbf{B} is denoted as $\tilde{\mathbf{B}} = \{\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_N\}$. The notation $\mathcal{A}(\mathbf{f})$ refers to the anti-circulant matrix associated with polynomial \mathbf{f} . The notation $\lfloor k \rfloor$ indicates the real number k is to be rounded to the nearest integer. The rounding $\lfloor \mathbf{f} \rfloor$ of a polynomial \mathbf{f} is taken to be coefficient-wise rounding. The Fast Fourier Transform (FFT) and Number Theoretic Transform (NTT) of polynomial \mathbf{f} are the evaluations $\mathbf{f}(\zeta^i)$ for $i \in \{0, \dots, N-1\}$, where ζ is the $2N$ -th complex root of unity in the FFT, and ζ is the $2N$ -th root of unity mod q in the NTT. Let \odot be the point-wise multiplication.

¹The source code is available at <https://gitlab.com/raykzhao/latte>.

DEFINITION 2.1 (RÉNYI DIVERGENCE [5]). *For two discrete distributions P and Q such that $\text{Supp}(P) \subseteq \text{Supp}(Q)$, the Rényi divergence (RD) of order $a \in (1, +\infty)$ is defined as:*

$$R_a(P||Q) = \left(\sum_{x \in \text{Supp}(P)} \frac{P(x)^a}{Q(x)^{a-1}} \right)^{\frac{1}{a-1}}.$$

In addition, for $a = +\infty$, we have: $R_\infty(P||Q) = \max_{x \in \text{Supp}(P)} \frac{P(x)}{Q(x)}$.

LEMMA 2.1. *Let $a \in [1, +\infty]$. Let P and Q denote distributions with $\text{Supp}(P) \subseteq \text{Supp}(Q)$. Then the following properties hold:*

- **Data Processing Inequality:** $R_a(P^f||Q^f) \leq R_a(P||Q)$ for any function f , where P^f (resp. Q^f) denotes the distribution of $f(y)$ induced by sampling $y \leftarrow P$ (resp. $y \leftarrow Q$).
- **Probability Preservation:** Let $A \subseteq \text{Supp}(Q)$ be an arbitrary event. If $a \in (1, +\infty)$, then $Q(A) \geq P(A)^{\frac{a}{a-1}} / R_a(P||Q)$. Further, we have $Q(A) \geq P(A) / R_\infty(P||Q)$.

We use the notation \lesssim, \sim as in [42], in order to “absorb” all higher-order terms of negligible elements: for example, if $\delta = o(1)$, then $\delta + \delta^2 \sim \delta$. The following remark bounds $R_\infty(D_2; D_1)$.

REMARK 1. *Let $\tau \in \mathbb{Z}$ be the tailcut bound as above, and let $Q = 2^k$ for some $k \in \mathbb{Z}$. If $\tau \geq \sqrt{2 \ln(2Q)}$, then:*

$$R_\infty(D_2; D_1) \leq 1/(1 - Q^{-1}) \lesssim 1 + 1/Q. \quad (1)$$

This can be verified by using classical tailcut bounds [38, Lemma 4.4].

In our analysis, we will apply the following proposition, adapted from Proposition 4 of [31].

PROPOSITION 2.2. [31] *Let P and Q denote two distributions of a N -tuple of random variables $(x_i)_{i < N}$. For $0 \leq i < N$, assume P_i (resp. Q_i) is the marginal distribution of x_i , and let $P_{i < i}(\cdot | x < i)$ denote the conditional distribution of x_i given that $(x_0, \dots, x_{i-1}) =: x < i$. Let $a > 1$. Suppose that for all $0 \leq i < N$, there exists $B_i \geq 1$ such that for all i -tuples $x < i$ in the support of Q restricted to its first i variables, $R_a(Q_i | x < i, P_i | x < i) \leq B_i$. Then $R_a(Q, P) \leq \prod_{i < N} B_i$.*

THEOREM 2.3 (TAIL-CUT BOUND, ADAPTED FROM [5], THM. 2.11). *Let \mathcal{D}'_σ be the B -bounded distribution of \mathcal{D}_σ by cutting its tail. For M independent samples, we have $R_\infty((\mathcal{D}'_\sigma)^M || (\mathcal{D}_\sigma)^M) \leq \exp(1)$ if $B \geq \sigma \cdot \sqrt{2 \ln(2M)}$.*

2.1 Hierarchical Identity-based Encryption

Hierarchical identity-based encryption (HIBE) schemes were introduced by Horwitz and Lynn [30] and can be considered a generalisation of an IBE scheme to multiple levels. An HIBE scheme consists of five components: Keygen, Delegate, Extract, Encrypt, and Decrypt. Appendix A explains the role of each component and discusses advantages of HIBE schemes over 2-level IBE. Gentry and Silverberg proposed the first secure HIBE scheme in the random oracle model (ROM) in 2002 [27], which was an extension of the Boneh-Franklin IBE scheme [8], a Weil-pairing based scheme, the security of which relies on the bilinear Diffie-Hellman problem. This was shown to be secure against adaptive identity and chosen ciphertext attacks, by use of the Fujisaki-Okamoto (FO) transformation [25], although the security degrades exponentially with the number of levels.

Boneh-Boyer built on this in 2004 to create a scheme which was secure without random oracles [6]. However, as both the ciphertext and private keys grew linearly with the number of levels of the hierarchy, in 2005 a scheme [7] was proposed which fixed the ciphertext size to three group elements, and curtailed private key growth to within level ℓ group elements. In 2018, an isogeny-based version of the Decisional Bilinear Diffie-Hellman-based scheme was proposed [33]. Despite isogenies possessing quantum-safe properties, this variant only serves to strengthen the existing classical security, by proving it secure under the assumption of *either* the classical version or the isogeny-based version of the problem and therefore is not necessarily quantum-safe. To the best of the authors’ knowledge, the only quantum-safe HIBE schemes so far proposed are based on lattices. We now introduce the schemes upon which LATTE is built.

2.2 The Ingredients of LATTE

DLP IBE Scheme. In 2014, Ducas, Lyubashevsky and Prest proposed the first efficient lattice-based identity-based encryption (IBE) scheme [22]. They based their construction on the IBE scheme by Gentry, Peikert and Vaikuntanathan [26], using a variant of NTRU lattices. The underlying security problems are the NTRU problem for key generation and R-LWE for the encryption. The ciphertexts therefore have more practical sizes than previous constructions, for example 30kb (kilobits) for 192-bit classical security. The use of structured lattices also allowed for implementation optimisations such as the Number Theoretic Transform (NTT), as demonstrated by [40], whose software performance of the DLP-IBE outperformed that of the elliptic-curve based Boneh-Franklin IBE scheme.

Bonsai Trees HIBE. Cash [12] proposed the use of Bonsai trees to create a hierarchical structure for IBE. They model the hierarchical network of users as a tree, whereby arborists, or sub-key-managers, have control over the sub-trees, and have the authority to delegate user private keys. Delegation requires the knowledge of a trapdoor basis of the lattice at that level. During the process whereby keys are delegated down the tree, the lattice is extended, and therefore its dimension and hence the key and ciphertext sizes increase. The public key size is of $\mathcal{O}(d^3 kn^2)$ and ciphertext size is of $\mathcal{O}(d^3 kn)$ at depth d , for security parameter n and hash output length k . The root authority has control of the whole tree by knowing the short trapdoor basis for the master root lattice. The security of this HIBE scheme is based on LWE over standard lattices.

3 IMPROVED LATTE HIBE SCHEME

LATTE was proposed in 2017 [10] and can be considered as a combination of the DLP-IBE scheme [22] and Bonsai Tree HIBE scheme [12] to create a lattice-based hierarchical IBE (HIBE) scheme. It can be shown to be ID-IND-CCA-secure, the proof for which is given in [1], based on the NTRU and R-LWE hardness assumptions.

3.1 Proposed Design Optimisations

For the optimised LATTE scheme presented in this section and used in our software design and implementation, features of the FALCON [51] and the MODFALCON [17] signature schemes were utilised. This is the first time these features have been considered in LATTE and so the rationale for this is expanded on in Sec. 5. For now, it suffices

Table 1: LATTE Algorithm Summary.

Alg.	Inputs	Outputs
KeyGen	$N, q \in \mathbb{Z}, \sigma_0 \in \mathbb{R}$	$\mathbf{h} \in \mathcal{R}_q, \mathbf{B}_0, \mathbf{S}_0 \in \mathcal{R}_q^{2 \times 2}$
Delegate	$S_{\ell-1} \in \mathcal{R}_q^{(\ell+1) \times (\ell+1)},$ $\sigma_\ell \in \mathbb{R}, \text{ID}_\ell,$ $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^N$	$S_\ell \in \mathcal{R}_q^{(\ell+2) \times (\ell+2)}$
Extract	$S_{\ell-1} \in \mathcal{R}_q^{(\ell+1) \times (\ell+1)},$ $\sigma_\ell \in \mathbb{R}, \text{ID}_\ell,$ $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^N$	$\mathbf{t}_0, \dots, \mathbf{t}_\ell \in \mathcal{R}_q$
Encrypt	$\sigma_e \in \mathbb{R}, \mathbf{h} \in \mathcal{R}_q,$ KDF, $\text{ID}_\ell, \mu \in \{0, 1\}^{256},$ $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^N$	$Z \in \{0, 1\}^{256},$ $C_1, \dots, C_\ell, C_h \in \mathcal{R}_q$
Decrypt	$Z \in \{0, 1\}^{256},$ $C_1, \dots, C_\ell, C_h \in \mathcal{R}_q,$ $\mathbf{t}_0, \dots, \mathbf{t}_\ell \in \mathcal{R}_q$	$\mu' \in \{0, 1\}^{256}$

to acknowledge that the sub-algorithms NTRUSolve, ffSampling and Tree are taken from FALCON and are displayed in Appendix B.1. The currently presented LATTE in this Section also improves on the efficiency of the original proposal [10] by reducing the module dimension of the extracted secret keys by 1, by extending a similar approach used in the DLP IBE [22]. More concretely, we eliminate public key polynomial \mathbf{b} by modifying the equation satisfied by the decryption key at level ℓ from the original rank $\ell + 2$ module relation over \mathcal{R}_q :

$$\mathbf{t}_0 + \mathbf{t}_1 \cdot \mathbf{h} + \mathbf{t}_2 \cdot \mathbf{A}_1 + \dots + \mathbf{t}_\ell \cdot \mathbf{A}_{\ell-1} + \mathbf{t}_{\ell+1} \cdot \mathbf{A}_\ell = \mathbf{b}, \quad (2)$$

where $\mathbf{A}_i = H(\text{ID}_1 | \dots | \text{ID}_i)$ for $1 \leq i \leq \ell$, to the following rank $\ell + 1$ relation over \mathcal{R}_q :

$$\mathbf{t}_0 + \mathbf{t}_1 \cdot \mathbf{h} + \mathbf{t}_2 \cdot \mathbf{A}_1 + \dots + \mathbf{t}_\ell \cdot \mathbf{A}_{\ell-1} = \mathbf{A}'_\ell, \quad (3)$$

where $\mathbf{A}'_\ell = H_E(\text{ID}_1 | \dots | \text{ID}_\ell) := H("E" | \text{ID}_1 | \dots | \text{ID}_\ell)$. Furthermore, we remove the need for Extract algorithm to be stateful. This is achieved by deriving randomness deterministically from the ID (see Sec. 4.4 for security discussion).

3.2 Scheme Description

The full pseudocode for LATTE KeyGen, Delegate, Extract, Encryption, and Decryption are presented in Alg. 2–6 in Appendix B, respectively. Table 1 further summarises the inputs and outputs of the LATTE algorithms. The KeyGen algorithm, given in Alg. 2 in Appendix B, generates an NTRU-type basis. This is performed by sampling the short basis polynomials \mathbf{f}, \mathbf{g} from a Gaussian distribution. Operations are over the polynomial ring $\mathcal{R}_q = \mathbb{Z}_q[x] / \langle x^N + 1 \rangle$, a variant of the NTRU ring. For the purposes of optimisation in the implementation, variables are stored in NTT representation where appropriate. The Gram-Schmidt norm of the associated basis is computed to ensure smallness allowing for short private keys to be delegated to the next level. If not, the polynomials are re-sampled. The rest of the basis, polynomials \mathbf{F}, \mathbf{G} , are computed so that they satisfy the NTRU equation, $\mathbf{f}\mathbf{G} - \mathbf{g}\mathbf{F} = q \pmod{x^N + 1}$. This sub-algorithm is referred to as NTRUSolve, and its implementation will be discussed in Sec. 5. The solution to this is not unique, but any solution suffices provided it is short enough. This is taken care of by

reduction of the coefficients. The public key consists of polynomial $\mathbf{h} = \mathbf{g} \cdot \mathbf{f}^{-1}$. The master public basis \mathbf{B}_0 and private basis \mathbf{S}_0 at level 0 are implicit in the polynomial master keys, as follows:

$$\mathbf{B}_0 = \begin{bmatrix} -\mathcal{A}(\mathbf{h}) & \mathbf{I}_N \\ q\mathbf{I}_N & \mathbf{0}_N \end{bmatrix}, \quad \mathbf{S}_0 = \begin{bmatrix} \mathbf{g} & -\mathbf{f} \\ \mathbf{G} & -\mathbf{F} \end{bmatrix}.$$

The Delegate process, given in Alg. 3 in Appendix B, creates a public/secret key pair for the next level in the tree, allowing it to become a sub key management service (sub-KMS). Suppose the KMS wishes to delegate a key from level $\ell - 1$ to level ℓ . Then it can extend the public basis of the user at level ℓ , denoted by \mathbf{B}_ℓ by placing $\mathbf{A}_\ell = H(\text{ID}_1 | \dots | \text{ID}_\ell)$, where H is a hash function, to the beginning of the first column and filling the extra row with \mathbf{I}_N and $\mathbf{0}_N$, as shown below. The dimension of the new matrix becomes $(\ell + 2)N \times (\ell + 2)N$. The corresponding private basis, \mathbf{S}_ℓ , can then be generated. The i^{th} row $(s_{i,0}, s_{i,1}, \dots, s_{i,\ell+1})$ of the private basis is a short solution to the equation:

$$s_{i,0} + s_{i,1} \cdot \mathbf{h} + s_{i,2} \cdot \mathbf{A}_1 + \dots + s_{i,\ell+1} \cdot \mathbf{A}_\ell = \mathbf{0} \pmod{q}.$$

This can be found by sampling short vectors (using the Klein-GPV sampler [26] or its variant from FALCON [51]) from the $(\ell - 1)$ -level lattice using its secret basis, with centre vector $(-s_{i,\ell+1} \cdot \mathbf{A}_\ell, \mathbf{0}, \dots, \mathbf{0})$, where $s_{i,\ell+1}$ is sampled from a discrete Gaussian distribution $\mathcal{D}_{\sigma_\ell}$ over \mathcal{R} . A check is made to ensure the GS-norm of the sampled lattice vector is within the bound $\sigma_\ell \cdot \sqrt{(\ell + 2)N}$ to ensure the delegated basis will be of sufficient quality.

The remainder of the Delegate algorithm, in which the bottom row $(s_{\ell+1,0}, s_{\ell+1,1}, \dots, s_{\ell+1,\ell+1})$ is generated, is a higher-dimensional analogue of LATTE KeyGen. The resulting matrix has a determinant of size q . The final row is then reduced similarly to as in the KeyGen component to ensure the basis is of the required quality for further delegation. Cramer's rule is utilised here to find the reduction coefficients, the details of which are given in Appendix H. Generalising to level ℓ , the public basis \mathbf{B}_ℓ and the private basis \mathbf{S}_ℓ , respectively become:

$$\mathbf{B}_\ell = \begin{bmatrix} -\mathcal{A}(\mathbf{A}_\ell) & \mathbf{0}_N & \dots & \mathbf{I}_N \\ \vdots & \vdots & \ddots & \vdots \\ -\mathcal{A}(\mathbf{h}) & \mathbf{I}_N & \dots & \mathbf{0}_N \\ q\mathbf{I}_N & \mathbf{0}_N & \dots & \mathbf{0}_N \end{bmatrix}, \quad \mathbf{S}_\ell = \begin{bmatrix} s_{0,\ell+1} & \dots & s_{0,\ell+1} \\ s_{1,\ell+1} & \dots & s_{1,\ell+1} \\ \vdots & \ddots & \vdots \\ s_{\ell+1,\ell+1} & \dots & s_{\ell+1,\ell+1} \end{bmatrix}.$$

In the LATTE Extract algorithm (Alg. 4 in Appendix B), the user private key is a short solution $(\mathbf{t}_0, \mathbf{t}_1, \dots, \mathbf{t}_\ell)$ to:

$$\mathbf{t}_0 + \mathbf{t}_1 \cdot \mathbf{h} + \mathbf{t}_2 \cdot \mathbf{A}_1 + \dots + \mathbf{t}_\ell \cdot \mathbf{A}_{\ell-1} = \mathbf{A}_\ell \pmod{q}, \quad (4)$$

where $\mathbf{A}_i = H(\text{ID}_1 | \dots | \text{ID}_i)$ for $i = 1, \dots, \ell$. Again, this is found using the Klein-GPV style sampler over the short basis from the previous level. An extended version of traditional R-LWE encryption/decryption [39] is used for ciphering messages as given in Alg. 5 and 6 in Appendix B, respectively. A random *seed* is sampled and used together with a Key Derivation Function (KDF) to one-time-pad the message. The *seed* is encoded² and then encrypted using Ring-LWE and sent. The ciphertext consists of the encrypted message Z and deterministically sampled ephemeral public keys C_1, \dots, C_ℓ, C_h . This is a variant of the FO transform [25] to protect against invalid ciphertexts. The Decrypt process takes the user

²The Encode/Decode are the same as described in [1].

private key to decrypt the *seed* and reconstruct the message. This works as follows by operations over \mathcal{R}_q :

$$\begin{aligned} V &= C_\ell - t_1 \cdot C_h - t_2 \cdot C_1 - \dots - t_\ell \cdot C_{\ell-1} \\ &= (A_\ell \cdot \mathbf{e} + \mathbf{e}_\ell + \mathbf{m}) - t_1(\mathbf{h} \cdot \mathbf{e} + \mathbf{e}_h) - \dots - t_\ell(A_{\ell-1} \cdot \mathbf{e} + \mathbf{e}_{\ell-1}) \\ &= \mathbf{e}_\ell + \mathbf{m} - t_1 \cdot \mathbf{e}_h - t_2 \cdot \mathbf{e}_1 - \dots - t_\ell \cdot \mathbf{e}_{\ell-1} + t_0 \cdot \mathbf{e}, \end{aligned}$$

where the first equality is derived by substituting C_h, C_i with their definitions, and the second equality holds based on Eq. (4). By construction, the error and private key terms are small enough so that \mathbf{m} is decoded successfully to recover the *seed*. From the *seed*, the message is straightforwardly recovered from Z , which is sent as part of the ciphertext.

4 SECURITY ANALYSIS

A recurring concern around LBC is the precision requirements of the implementation, in particular of the discrete Gaussian sampler. As noted in [52], the precision used is often excessive, leading to slow and impractical implementations. Traditional measures of statistical distance have recently been substituted for Rényi divergence or Kullback-Leibler divergence to reduce memory and computational resources, whilst maintaining security. In this section, we make use of the Rényi divergence argument initially proposed in [49] to answer the question of how low we can allow the precision of our implementation to be, without allowing an adversary to detect any distinction between the actual distribution and the ideal distribution of a true Gaussian sample *over the lattice*, hence maintaining our claimed security levels. In particular, we analyze the security impact on LATTE of finite precision errors in the floating-point arithmetic and in the \mathbb{Z} -samplers used in our implementation of Extract and Delegate algorithm based on the the ffSampling lattice Gaussian algorithm. For this, we follow the following steps:

- (1) **Rényi divergence security reduction:** We give a security reduction (Sec. 4.1) based Rényi divergence analysis to relate security of finite precision LATTE to the security of its ideal (infinite precision) implementation, and bounds on the errors in the center and standard deviation parameters of \mathbb{Z} Gaussian samples used in lattice Gaussian ffSampling algorithm.
- (2) **Rényi divergence between \mathbb{Z} -Gaussians with errors in parameters:** To support above security reduction, we give a tight Lemma (in Sec. 4.2) giving a bound on Rényi divergence between the output distribution of \mathbb{Z} Gaussian samplers with errors in the center and standard deviation parameters, extending the sharp Rényi divergence results of [49].
- (3) **Statistical Model for ffSampling Precision Errors:** For use with the above security reduction, we introduce and empirically verify (in Sec. 4.3) a heuristic statistical model to compute upper bounds on the finite precision errors in the lattice Gaussian ffSampling algorithm. We give empirical evidence for the validity of our model, and use it to compute estimated error bounds for the LATTE parameter sets, and apply the results with the above security reductions to evaluate the security impact of precision errors on LATTE.

Then, in Sec. 4.5, we evaluate the security against best known lattice attacks of our Improved LATTE for the four recommended LATTE parameter sets.

4.1 Rényi Divergence Security Reduction

The security reduction to establish the (full chosen identity) indistinguishability against chosen-ciphertext attack security (ID-IND-CCA) of the *IDEAL* (infinite precision) LATTE HIBE encryption scheme is summarised in Annex C of the LATTE specification [1] and proceeds in two steps. Here, we show how to obtain a security reduction that takes into account and quantifies the security impact of the *REAL* (finite precision) implementation of LATTE. To do so, we introduce an additional middle step (step 2 below) in the security reduction steps for LATTE, and so we end up with the following three security reduction steps:

Step 1 - FO Transform: This generic reduction transforms any ID-IND-CCA attack (chosen ID indistinguishability against chosen ciphertext attacks) against *REAL* (finite precision) LATTE to an ID-OW-CPA (chosen ID one-wayness against chosen plaintext attacks) of *REAL* (finite precision) LATTE', assuming the random oracle model for the LATTE KDF hash function. Here, LATTE' denotes the ID-OW-CPA encryption scheme underlying LATTE: the encryption/decryption algorithms of LATTE can be obtained by applying the tag-based Fujisaki-Okamoto KEM-DEM transform of [3] to the LATTE' scheme. As pointed out in Annex C of [1], this reduction step follows directly from an (ID-based variant) of the composition of Theorem 3.1 and Theorem D.1 in [3].

Step 2 - Rényi Divergence - REAL to IDEAL: This presented reduction (in Lemma 4.1 below) transforms any ID-OW-CPA attack against *REAL* (finite precision) LATTE' to an ID-OW-CPA attack against *IDEAL* (infinite precision) LATTE', by using Rényi Divergence (RD) analysis techniques [5, 49]. For the latter RD reduction to apply, we exploit the fact that the one-wayness notion ID-OW-CPA is a *search* problem, rather than a *decision* problem.

Step 3 - ID-OW-CPA to NTRU/RLWE: This reduction transforms any attack against *IDEAL* LATTE' ID-OW-CPA security into attacks against the NTRU or RLWE problems, assuming the random oracle model for the ID hash function H . As pointed out in Annex C of [1], this reduction is a variant of the Bonsai tree reduction presented in Theorem 5.2 of [13], with a minor modification for our improved LATTE construction (see Sec. 4.4 for more details).

The following result fills the missing Step 2 above (where we apply Lemma 4.1 with the ID-OW-CPA attack game and the event E being the winning of this game by the adversary), and quantifies the security impact of finite precision in the discrete integer Gaussian samplers and the floating point arithmetic used in the FFT lattice Gaussian sampler used within the LATTE' Delegate algorithm (Alg. 3 in Appendix B) and Extract algorithm (Alg. 4 in Appendix B). The latter security impact is expressed as a function of upper bounds $\delta_{\sigma_i}^U$ and $\Delta_{t_i}^U$ on the relative (resp. absolute) finite precision errors in the integer discrete Gaussian standard deviation parameters $\sigma^{(i)}$ (resp. Gaussian center parameters $t^{(i)}$) used inside the Delegate and Extract algorithms, and an upper bound Δ_z^U on the absolute error in the final output value of the FFT sampling algorithm. We also allow for a negligible probability p_U (over the randomness of the key generation and discrete Gaussian samplers)

that the above error upper bounds fail to hold. The next subsection explains our statistical model and results for estimating the latter error upper bounds and the probability p_U for the chosen implementation finite precision.

Consider an attack game REAL against LATTE' with depth parameter d where the attack algorithm \mathcal{A} is run input a LATTE' master public key h (where $(S_0, h) \leftarrow \text{KeyGen}(N, q, \sigma_0)$), makes at most Q_D total number of queries to the Delegate algorithm (Alg. 3 in Appendix B) and Q_E queries to the Extract algorithm (Alg. 4 in Appendix B) implemented with:

- A finite precision p_D 1-dimensional Discrete Gaussian \mathbb{Z} -sampling algorithm in lines 3–4 of Alg. 10 in Appendix B.1 outputting samples from a distribution $\tilde{\mathcal{D}}_{\sigma,t}$ within Rényi divergence of order a at most B from the ideal Discrete Gaussian distribution $\mathcal{D}_{\sigma,t}$ i.e. $R_a(\tilde{\mathcal{D}}_{\sigma,t}, \mathcal{D}_{\sigma,t}) \leq B$.
- A finite precision p_{fp} floating-point arithmetic for Alg. 3, Alg. 4 in Appendix B, and lines 10–13 of Alg. 10 in Appendix B.1.

Let IDEAL denote the attack game against the ideal implementation of LATTE' where both p_D and p_{fp} are infinite precision. Let $(t^{(i)}, \sigma^{(i)})$ denote the center and std dev. parameter (resp.) for the i 'th query to the 1-dim. \mathbb{Z} Gaussian sampler (i.e. at line 1 of Alg. 2 in Appendix B, line 3 or 4 of Alg. 10 in Appendix B.1, or line 4 of Alg. 3 in Appendix B) in the game IDEAL, and let $\bar{z}^{(j)}$ denote the value of \bar{z} in the output of the j 'th query to FFT^{-1} (ffSampling) in the game (i.e. at line 6 of Alg. 3 or line 5 of Alg. 4 in Appendix B). Suppose that, except for an event B_U , the absolute errors $\Delta_{t^{(i)}}$ in centers $t^{(i)}$ relative to $\sigma^{(i)}$ (i.e. $\Delta_{t^{(i)}}/\sigma^{(i)}$) are upper bounded by $\Delta_{t^{(i)}}^U$ and relative errors $\delta_{\sigma^{(i)}}$ in standard deviations $\sigma^{(i)}$ are upper bounded by $\delta_{\sigma^{(i)}}^U$ for all $1 \leq i \leq M_{\mathbb{Z}}$, and the infinity-norm absolute errors $\Delta_{\bar{z}^{(j)}}$ in $\bar{z}^{(j)}$ is upper bounded by $\Delta_{\bar{z}^{(j)}}^U < 1/2$ for all $1 \leq j \leq M_f$. The above errors are computed with respect to the same game with finite precision floating-point arithmetic. Here, $M_{\mathbb{Z}} \leq K \cdot (Q_E + (d+1) \cdot Q_D) + 2$ denotes the total number of queries to the 1-dim. \mathbb{Z} Gaussian sampler in the game, K denotes the number of \mathbb{Z} sampler calls of each call of Alg. 10 in Appendix B.1, and $M_f \leq Q_E + (d+1) \cdot Q_D$ denotes the number of calls of Alg. 10 in the game. Let $Q, \epsilon > 0$ with $\tau \geq \sqrt{2 \ln(2Q)}$ and assume that $\sigma^{(i)} \geq \eta_\epsilon(\mathbb{Z})$ for $1 \leq i \leq m_{\mathbb{Z}}$.

LEMMA 4.1. *Let p_U denote the probability of event B_U in game IDEAL. Let E denote any event defined over the view of \mathcal{A} , $B_T := B^{M_{\mathbb{Z}}}$, $C_T := \prod_{i < M_{\mathbb{Z}}} C^{(i)}$, where $C^{(i)}$ is given by the right hand side of (6) in Lemma 4.2 with $\delta_\sigma := \delta_{\sigma^{(i)}}^U$, $\Delta'_t := \Delta_{t^{(i)}}^U$ for $1 \leq i \leq M_{\mathbb{Z}}$. Then,*

$$\Pr[E_{\text{IDEAL}}] \geq \frac{1}{C_T} \cdot \left(\Pr[E_{\text{REAL}}]^{a/(a-1)} / B_T - \eta \right)^{a/(a-1)} - p_U,$$

where $\eta := C_T(p_U + 1/Q)^{(a-1)/a}$.

4.2 Rényi divergence between \mathbb{Z} -Gaussians with errors in parameters

This step builds upon unpublished work by Prest [50]. We will consider the following Gaussians:

- D_1 is an ideal Gaussian of standard deviation σ and center t .

Table 2: Empirical Results of the Error Estimation from Our Statistical Model.

Set	Prcn.	$\ell = 1$			$\ell = 2$		
		δ_σ	Δ'_t	$\Delta_{\bar{z}}$	δ_σ	Δ'_t	$\Delta_{\bar{z}}$
LATTE-1	53	2^{-51}	2^{-35}	2^{-23}	-	-	-
LATTE-2	53	2^{-51}	2^{-34}	2^{-21}	-	-	-
LATTE-3	113	2^{-111}	2^{-89}	2^{-71}	2^{-102}	2^{-59}	2^{-35}
LATTE-4	113	2^{-111}	2^{-87}	2^{-68}	2^{-102}	2^{-55}	2^{-30}

Table 3: Empirical Results of the Actual Arithmetic Errors.

Set	Prcn.	$\ell = 1$			$\ell = 2$		
		δ_σ	Δ'_t	$\Delta_{\bar{z}}$	δ_σ	Δ'_t	$\Delta_{\bar{z}}$
LATTE-1	53	2^{-50}	2^{-34}	2^{-22}	-	-	-
LATTE-2	53	2^{-50}	2^{-32}	2^{-20}	-	-	-
LATTE-3	113	2^{-110}	2^{-88}	2^{-70}	2^{-108}	2^{-58}	2^{-34}
LATTE-4	113	2^{-110}	2^{-85}	2^{-67}	2^{-105}	2^{-53}	2^{-29}

- D_2 is a Gaussian of standard deviation σ and center t , restricted to the interval $I = [t - \tau \cdot \sigma, t + \tau \cdot \sigma]$.
- D_3 is a Gaussian of standard deviation $\bar{\sigma}$ and center \bar{t} , restricted to the interval I .

Now, we present Lemma 4.2 showing for adequate values of τ , $|\bar{t} - t|/\sigma$, $|\frac{\bar{\sigma}}{\sigma} - 1|$, D_1 and D_3 are close in the Rényi divergence sense. The proof appears in Appendix C.

LEMMA 4.2. *Consider D_1, D_2, D_3, τ, Q as defined above. Suppose that there exist $\delta_\sigma, \epsilon, \delta > 0$ such that:*

- (1) $\max(\delta_\sigma, \epsilon) \leq \delta = o(1)$;
- (2) $|\bar{t} - t|/\sigma \leq \Delta'_t$ (bounded absolute error);
- (3) $|\frac{\bar{\sigma}}{\sigma} - 1| = \delta_\sigma$ (bounded relative error);
- (4) $\sigma \geq \eta_\epsilon(\mathbb{Z})$;
- (5) $\tau \geq \sqrt{2 \ln(2Q)}$ (motivated by remark 1).
- (6) $\text{num}(a, b, c) := |a^2 + 2a \frac{\sqrt{2\pi}b}{1-b} + (2c + c^2)(1 + \frac{2\pi b}{1-b})|/2(1 - c^2)$;
- (7) $\text{ub} := 2/Q + \text{num}(\Delta'_t, \epsilon, \delta_\sigma) + \frac{1}{1-\delta_\sigma} \cdot (\tau \Delta'_t + \tau^2 \delta_\sigma)$.

Then the Rényi divergences of D_3 and D_2 (resp. D_1) is:

$$R_a(D_3; D_2) \leq 1 + \frac{a \cdot \text{ub}^2}{2}. \quad (5)$$

$$R_a(D_3; D_1) \leq 1 + \frac{1}{Q} + \frac{a \cdot \text{ub}^2}{2}. \quad (6)$$

4.3 Statistical Model for ffSampling Precision Errors

In this Section, we present a statistical model to estimate bounds for the floating point arithmetic errors in the LATTE ffSampling algorithm using our chosen implementation floating point precision for the LATTE parameter sets, and we use those bounds to analyse the security impact of those errors on our LATTE implementation by applying Lemma 4.1.

Our statistical model makes the heuristic but natural assumption the floating point error introduced in each arithmetic operation in the ffSampling algorithm can be modelled as independent

zero-center continuous Gaussian random variable, and the model estimates the maximum standard deviations $\delta_\sigma, \Delta'_t, \Delta_{\bar{z}}$ of the errors $\delta_{\sigma^{(i)}}, \Delta'_{t^{(i)}}, \Delta_{\bar{z}^{(j)}}$ over all \mathbb{Z} -sampler query indices $1 \leq i \leq M_{\mathbb{Z}}$ and ffSampler query indices $1 \leq j \leq M_f$ in the IDEAL game of Lemma 4.1 by propagating the standard deviations of the independent errors through the ffSampling algorithm arithmetic steps, assuming uniformly random input matrices $A_i \in R_q$ at the input to the Extract and Delegate algorithms. We explain at the end of this section how we apply the standard deviations in the \mathbb{Z} sampler queries to derive the security impact of floating point errors on LATTE. We remark that the use of the random oracles H and G to hash the attacker's choice of identities queried to Extract or Delegate algorithms to derive the ffSampling input ring elements A_i uniformly at random in R_q and seed for Extract uniformly random supports our statistical (rather than adversarial) model of floating point errors, since the attacker cannot control the randomness of H and G and the Delegate, Extract and Key Generation algorithms. A similar heuristic statistical model is commonly used in the context of evaluating the propagation of LWE errors via a circuit computed homomorphically with Fully Homomorphic Encryption schemes [16].

We now present the details our statistical model for estimating the standard deviations of errors, i.e. δ_σ, Δ'_t , and $\Delta_{\bar{z}}$. For a complex number $a = \mu_R + i\mu_I$, with $\mu_R, \mu_I \in \mathbb{R}$, let denote the absolute error of the real part μ_R as σ_R and the absolute error of the imaginary part μ_I as σ_I , respectively. We assume the real and imaginary parts of every complex number in our statistical model are independent Gaussian variables e.g. for complex number $a = \mu_R + i\mu_I$, $\text{Re}(a)$ follows the normal Gaussian distribution $\mathcal{N}(\mu_R, \sigma_R^2)$ and $\text{Im}(a)$ follows the normal distribution $\mathcal{N}(\mu_I, \sigma_I^2)$, respectively, no matter whether $\text{Re}(a), \text{Im}(a)$ are linear combinations of one or more independent normal variables. Therefore, we use the tuple $(\mu_R, \sigma_R^2, \mu_I, \sigma_I^2)$ to represent a complex number with errors.

DEFINITION 4.1 (ADDBOUND, SUBBOUND, AND MULTBOUND). For independent $a = (\mu_{a,R}, \sigma_{a,R}^2, \mu_{a,I}, \sigma_{a,I}^2)$ and $b = (\mu_{b,R}, \sigma_{b,R}^2, \mu_{b,I}, \sigma_{b,I}^2)$, let define addition $\text{AddBound}(a, b)$, subtraction $\text{SubBound}(a, b)$ as

$$\text{AddBound}(a, b) := (\mu_{a,R} + \mu_{b,R}, \sigma_{a,R}^2 + \sigma_{b,R}^2, \mu_{a,I} + \mu_{b,I}, \sigma_{a,I}^2 + \sigma_{b,I}^2),$$

$$\text{SubBound}(a, b) := (\mu_{a,R} - \mu_{b,R}, \sigma_{a,R}^2 + \sigma_{b,R}^2, \mu_{a,I} - \mu_{b,I}, \sigma_{a,I}^2 + \sigma_{b,I}^2),$$

and multiplication $\text{MultBound}(a, b)$ as:

$$\begin{aligned} &(\mu_{a,R}\mu_{b,R} - \mu_{a,I}\mu_{b,I}, \mu_{a,R}^2\sigma_{b,R}^2 + \mu_{b,R}^2\sigma_{a,R}^2 + \sigma_{a,R}^2\sigma_{b,R}^2 + \mu_{a,I}^2\sigma_{b,I}^2 + \\ &\mu_{b,I}^2\sigma_{a,I}^2 + \sigma_{a,I}^2\sigma_{b,I}^2, \mu_{a,R}\mu_{b,I} + \mu_{a,I}\mu_{b,R}, \mu_{a,R}^2\sigma_{b,I}^2 + \mu_{b,I}^2\sigma_{a,R}^2 + \\ &\sigma_{a,R}^2\sigma_{b,I}^2 + \mu_{a,I}^2\sigma_{b,R}^2 + \mu_{b,R}^2\sigma_{a,I}^2 + \sigma_{a,I}^2\sigma_{b,R}^2). \end{aligned}$$

DEFINITION 4.2 (DIVBOUND, ADAPTED FROM [20]). For $a = (\mu_{a,R}, \sigma_{a,R}^2, \mu_{a,I}, \sigma_{a,I}^2)$ and real number $b = (\mu_{b,R}, \sigma_{b,R}^2, 0, 0)$, assuming $\text{Re}(a), \text{Im}(a)$, and b are independent normal variables such that $\sqrt{\frac{\sigma_{a,R}^2}{\mu_{a,R}^2} + \frac{\sigma_{b,R}^2}{\mu_{b,R}^2}} < 1$, $\sqrt{\frac{\sigma_{a,I}^2}{\mu_{a,I}^2} + \frac{\sigma_{b,R}^2}{\mu_{b,R}^2}} < 1$, let define the division $\text{DivBound}(a, b)$ between a and b as:

$$\left(\frac{\mu_{a,R}}{\mu_{b,R}}, \frac{\mu_{a,R}^2}{\mu_{b,R}^2} \left(\frac{\sigma_{a,R}^2}{\mu_{a,R}^2} + \frac{\sigma_{b,R}^2}{\mu_{b,R}^2} \right), \frac{\mu_{a,I}}{\mu_{b,R}}, \frac{\mu_{a,I}^2}{\mu_{b,R}^2} \left(\frac{\sigma_{a,I}^2}{\mu_{a,I}^2} + \frac{\sigma_{b,R}^2}{\mu_{b,R}^2} \right) \right).$$

Table 4: LATTE Security Impact of Finite Precision.

Set	Pfp	PD	$\ell = 1$		$\ell = 2$	
			Q_{\max}^C	Q_{\max}^B	Q_{\max}^C	Q_{\max}^B
LATTE-1	53	48	2^{46}	2^{76}	-	-
LATTE-2	53	48	2^{44}	2^{75}	-	-
LATTE-3	113	96	2^{73}	2^{173}	2^{72}	2^{77}
LATTE-4	113	96	2^{149}	2^{171}	2^{85}	2^{68}

Table 5: LATTE σ_ℓ and Decryption Fail. Prob.

Set	σ_ℓ			Fail. Prob.	
	$\ell = 0$	$\ell = 1$	$\ell = 2$	$\ell = 1$	$\ell = 2$
LATTE-1	106.2	5513.3	-	2^{-191}	-
LATTE-2	106.2	7900.2	-	2^{-380}	-
LATTE-3	6777.6	351968.4	22559988.0	$2^{-\text{inf}}$	2^{-126}
LATTE-4	9583.7	713167.	64997288.2	$2^{-\text{inf}}$	2^{-246}

DEFINITION 4.3 (ABSQRBOUND). For $a = (\mu_{a,R}, \sigma_{a,R}^2, \mu_{a,I}, \sigma_{a,I}^2)$, assuming $\text{Re}(a)$ and $\text{Im}(a)$ are independent normal variables, let define the squared absolute value $\text{AbsSqrBound}(a)$ of a i.e. $|a|^2$ as $\text{AddBound}((\text{Re}(a))^2, (\text{Im}(a))^2)$, that is:

$$(\mu_{a,R}^2 + \sigma_{a,R}^2 + \mu_{a,I}^2 + \sigma_{a,I}^2, 2\sigma_{a,R}^2 + 4\mu_{a,R}^2\sigma_{a,R}^2 + 2\sigma_{a,I}^2 + 4\mu_{a,I}^2\sigma_{a,I}^2, 0, 0).$$

We can use the above absolute arithmetic error bound approximations to rewrite our optimised fLDL in Alg. 1 in Sec. 5.1 and ffSampling in Alg. 10 in Appendix B.1, in order to estimate δ_σ and Δ'_t , respectively. For δ_σ , we first use the fLDLBounds in Alg. 11 in Appendix E to estimate the absolute errors of the leaf values (real numbers) in fLDL tree T for a given Gram matrix G i.e. the standard deviation $\sigma_{\text{leaf},R}$. The splitfftBounds used by Alg. 11 is shown in Alg. 13 in Appendix E. Since the σ for the 1-D integer Gaussian sampler is computed by $\sigma_\ell / \sqrt{\mu_{\text{leaf},R}}$ in the Tree computation Alg. 9 in Appendix B.1, assuming the relative error of the floating-point arithmetic is u , we have the following arithmetic error bound:

$$\begin{aligned} \delta_\sigma &\leq \max_{\text{all leaves}} \left[\frac{(1+u) \frac{(1+u)\sigma_\ell}{(1-u)\sqrt{\mu_{\text{leaf},R} - \sigma_{\text{leaf},R}}} - 1}{\frac{\sigma_\ell}{\sqrt{\mu_{\text{leaf},R}}}} \right] \\ &= \frac{(1+u)^2}{1-u} \sqrt{\max_{\text{all leaves}} \frac{\mu_{\text{leaf},R}}{\mu_{\text{leaf},R} - \sigma_{\text{leaf},R}} - 1}. \end{aligned}$$

Similarly, we can use the ffSamplingBounds in Alg. 12 in Appendix E to output Δ'_t for a given vector t and fLDL tree T . The splitfftBounds and mergefftBounds in Alg. 12 are shown in Alg. 13 and Alg. 14 in Appendix E, respectively. In addition, we can compute the rounding errors $\Delta_{\bar{z}}$ i.e. Line 8 in Alg. 3 and Line 6 in Alg. 4 in Appendix B, by combining the FFT/FFT⁻¹ errors of the input and the errors $\sigma_{z,R}, \sigma_{z,I}$ of z computed by ffSamplingBounds in Alg. 12 in Appendix E. We can also use similar statistical modelling approach to estimate the errors of FFT/FFT⁻¹ for a given vector a , as shown in Alg. 15 and Alg. 16 in Appendix E, respectively.

Here we show the empirical results of the errors δ_σ, Δ'_t , and $\Delta_{\bar{z}}$ estimated by our statistical model. For the target floating-point precisions used by our implementation of the LATTE scheme (see

Sec. 6 for the rationale behind the chosen precision), we compute the errors for 100 random (S, \mathbf{t}) pairs, where S is the basis and \mathbf{t} is the input of the `ffSampling` in Alg. 10 in Appendix B.1. The maximal δ_σ , Δ'_t , and $\Delta_{\bar{z}}$ among these 100 iterations are shown in Table 2. To provide empirical evidence for supporting the accuracy of our statistical model, for the same 100 pairs of (S, \mathbf{t}) , we also give the actual arithmetic errors between the values computed by using a very high precision (1024 bits) and the values computed by using the target precisions. The actual arithmetic errors computed by this approach are shown in Table 3. By comparing the results in Table 2 and Table 3, the actual arithmetic errors is larger than the estimated errors from our statistical model by at most 2 bits in this empirical experiment. We will leave modelling the distributions of (S, \mathbf{t}) to make our statistical model fully deterministic as future works.

Security Impact of Finite Precision errors. In order to use the results in Table 2 with Lemma 4.1 to derive the security impact of floating point errors, we first derive corresponding upper bounds $\delta_\sigma^U := \tau_U \cdot \delta_\sigma$, $\Delta'_t{}^U := \tau_U \cdot \Delta'_t$, and $\Delta_{\bar{z}}^U := \tau_U \cdot \Delta_{\bar{z}}$ on the absolute value of the errors, where τ_U is chosen so that each individual Gaussian error's absolute value exceeds its bound with probability $\leq 2^{-\lambda}$, which by the standard Gaussian tail bound is satisfied by setting $2 \exp(-\tau^2/2) \leq 2^{-\lambda}$. Therefore by a union bound, all bounds hold except with a negligible probability $p_U \leq (2M_{\bar{z}} + M_f)2^{-\lambda}$, with λ denoting the target security level. Applying Lemma 4.1 with $a := 2\lambda$ we conclude using $a/(a-1) \approx 1$ and p_U is negligible, that $\Pr[E]_{IDEAL} \approx \frac{1}{B_T C_T} \Pr[E]_{REAL}$ so that finite precision causes a bit security loss $L \approx \log_2(B_T) + \log_2(C_T)$ bits. We use the above floating point arithmetic upper bounds to compute an estimate for the maximum number of delegate/extract queries Q_{\max}^C (resp. Q_{\max}^B) that ensures $\log_2(C_T) \leq 1$ (resp. $\log_2(B_T) \leq 1$) so that if $\max(Q_D, Q_E) \leq \min(Q_{\max}^C, Q_{\max}^B)$, then $L \leq 2$ bits of security are lost overall for our finite arithmetic precision p_{fp} LATTE implementation versus the infinite precision implementation. To compute $B_T \leq B^{M_{\bar{z}}}$, we use the RD bound B on the COSAC \mathbb{Z} sampler Rényi divergence from the ideal \mathbb{Z} sampler distribution derived in [55] corresponding to the COSAC sampler precision $p_{\mathcal{D}}$ used in our COSAC implementation (see Sec. 5.3 for the discussions). The finite precision security impact results are summarised in Table 4. The results show that for LATTE-1 and LATTE-2, 2^{44} Extract and/or Delegate queries can be supported with at most 2 bits of security loss with our 53-bit double precision floating-point precision implementation. This should suffice for most practical applications. For LATTE-3 and LATTE-4, the main bottleneck in precision is the $\Delta_{\bar{z}}$ bound, but the results indicate that even reducing the precision by about 25 bits from our chosen 113 bit arithmetic precision to ≈ 90 bits precision would suffice for security.

4.4 Provable Security for ID-OW-CPA of Improved LATTE

Recall from Sec. 3 that our improved LATTE scheme achieves improved efficiency and shorter decryption keys output by the Extract algorithm relative to the original LATTE scheme. This change to Extract necessitates a different strategy for simulating the Extract oracle at level ℓ in the ID-OW-CPA security proof (step 3 in the overview of Sec. 4.1), compared with the strategy outlined in [1]

based on Theorem 5.2 in [13]. In particular, the Extract oracle simulation at level ℓ must simulate the decryption key \mathbf{t} at that level without knowing the delegation secret key $S_{\ell-1}$ at level $\ell-1$. In the original LATTE using Eq. (2), this can be done by programming $A_\ell = H(|ID_1| \dots |ID_\ell|)$ to be a matrix with an embedded NTRU trapdoor and using the basis extension method used in the Delegate oracle and its simulation at level ℓ . But with our improved LATTE Extraction using Eq. (3), we cannot use a trapdoor for A'_ℓ to simulate multiple such decryption key vectors \mathbf{t} ; indeed, if this were possible then subtracting two such distinct short vectors would reveal a short vector \mathbf{s} in the (secret) level $\ell-1$ delegation module lattice $\mathbf{s} : \mathbf{s}_0 + \mathbf{s}_1 \mathbf{h} + \dots + \mathbf{t}_\ell A_{\ell-1} = \mathbf{0}$.

Instead, our Extract simulator generates a *single* such short vector \mathbf{t} using the GPV signature simulation strategy [26], i.e. programming $A'_\ell = H_E(|ID_1| \dots |ID_\ell|) := \mathbf{t}_0 + \mathbf{t}_1 \mathbf{h} + \dots + \mathbf{t}_\ell A_{\ell-1}$ for short discrete Gaussian \mathbf{t}_i 's sampled by the Extract simulator. To avoid a contradiction with the different programming strategy for A_ℓ , our modified LATTE uses a different hash function H_E modelled as a random oracle (obtained from the random oracle H by using the prefix "E") for computing A'_ℓ used in Extract, so that H and H_E can be programmed independently. Also, since our programming strategy for $A'_\ell = H_E(|ID_1| \dots |ID_\ell|)$ only works for a *single* decryption key \mathbf{t} , we must make Extract deterministic so that it returns the same secret key \mathbf{t} again if queried again at the same $|ID_1| \dots |ID_\ell|$; this is the purpose of the hash function G used to derive the randomness seed for Extract deterministically from $|ID_1| \dots |ID_\ell|$.

4.5 Concrete Parameter Sets Based on Best Known Attacks

The security of each component of LATTE depends on an associated lattice problem and so the computational security of each of these problems must be considered to derive parameters, with the most vulnerable component determining the overall security for a given parameter set. The global parameters for the scheme are dimension N and modulus q , but we will also need to consider level-specific parameters, namely the standard deviation used for sampling at each level, σ_ℓ . The six security constraints to be considered are: (1) Gaussian sampler security (2) Decryption failure (3) Master key recovery (breaking the NTRU problem/finding short vectors in the NTRU lattice) (4) Delegated key recovery (finding short vectors in the lattice) (5) User key recovery (solving closest vector problem) (6) Message recovery (breaking the R-LWE encryption scheme). These are discussed in detail in [1], so here we only state the mathematical conditions which must be satisfied, and compute the security levels using our updated parameters and modifications to the scheme. We first summarise the differences between our security analysis and [1]. Any other differences are negligible and due to precision variations in the attack costing script.

Summary of Differences Compared to [1]. There are three main differences:

- We find that the discrete Gaussian statistical parameter $\varepsilon = 2^{-22.5}/(\ell+1)N$ used by σ_ℓ in [1] was miscalculated. The Kullback-Liebler divergence between the sampled distribution and the ideal discrete Gaussian distribution is bounded by approximately $8((\ell+1)N)^2 \varepsilon^2$. Choosing $\varepsilon =$

Table 6: LATTE Estimated Cost of Master Key Recovery.

Set	β	Classical Security	Quantum Security
LATTE-1	974	301	275
LATTE-2	1501	455	414
LATTE-3	973	301	274
LATTE-4	1501	455	414

$2^{-25.5}/(\ell + 1)N$ ensures the divergence is at most 2^{-48} , as specified by the proposed LATTE specification [1]. If the sampled distribution has a KL-divergence of 2^{-48} from the ideal distribution then using the sampler at most 2^{47} times will only reduce the security of the scheme by up to one bit [47]. However, in [1], the $\varepsilon = 2^{-22.5}/(\ell + 1)N$ would only ensure the KL-divergence is at most 2^{-42} .

- To accommodate the use of the FACCT sampler in Key-Gen, as described in Sec. 5, we modify the value of σ_0 , as displayed in Table 5. This also has an effect on the subsequent σ_ℓ , and therefore difficulty of the underlying lattice problems and success of each attack.
- As our redesign of LATTE discards the polynomial \mathbf{b} in the master public key and reduces the module dimension of the user private key, as described in Section 3.1, we update the attack costings accordingly. First, it reduces decryption failure rate, as there is one less error term. The best user key recovery attack reduces to CVP in the master lattice, so the attack is on the same lattice, but it demands a marginally shorter vector to be successful.

Gaussian Sampler Security. The statistical security of the Gaussian sampler used for sampling short vectors from lattice cosets in extraction and delegation to level ℓ is determined by the standard deviation of the sampler σ_ℓ and its relation to the Gram-Schmidt norm of the input basis. As this property of the basis is determined from the master key generation and any previous delegations, i.e. $\|\tilde{\mathbf{B}}\| \leq \sqrt{(\ell + 2)N} \cdot \sigma_\ell$, we can draw the following condition based on the relationship of the standard deviations at each level:

$$\sigma_\ell \geq \eta_\varepsilon(\mathbb{Z})\sqrt{(\ell + 1)N} \cdot \sigma_{\ell-1}, \quad (7)$$

taking ε as $2^{-25.5}/(\ell + 1)N$ in order to make the Kullback-Leibler divergence of the sampler from the discrete Gaussian is at most 2^{-48} . However, we also require the sampled vectors to be short for the purposes of keeping the underlying lattice problem hard. Therefore, we can set σ_ℓ to be equal to right hand side of Eq. (7), where $\sigma_0 \approx 1.17\sqrt{q/(2N)}$. The quantity σ_0 is chosen to be this as it minimises the Gram-Schmidt norm of the master basis (resulting in short user private keys in the single-level IBE), as deduced in [22]. Table 5 gives the σ_ℓ used for each parameter set, which can be computed indirectly from (N, q) .

Decryption Failure. To protect against attacks which exploit random decryption failures, we must bound the error term incurred in the R-LWE encryption scheme. The probability that the error term is too large is derived in [1], based on the method of [4]. Essentially, the decryption failure rate cannot exceed $2^{-\lambda}$, where λ is

Table 7: LATTE Estimated Cost of Delegated Key Recovery.

Set	ℓ	β	Classical Security	Quantum Security
LATTE-1	1	1020	314	287
LATTE-2	1	1051	323	295
LATTE-3	1	1021	315	287
	2	388	130	119
LATTE-4	1	1051	323	295
	2	907	281	257

Table 8: LATTE Estimated Cost of User Key Recovery.

Set	ℓ	β	Classical Security	Quantum Security
LATTE-1	1	829	258	236
LATTE-2	1	1863	560	510
LATTE-3	1	830	259	236
	2	334	114	105
LATTE-4	1	1864	561	510
	2	799	250	228

the security level in bits of the scheme. For each parameter set and level, we can compute the probability of decryption failure, noting that our design consists of one less ephemeral private key than in [1], reducing the standard deviation τ of the Gaussian distribution of the coefficients of the error term d to $\tau = \sqrt{\sigma_e^2 + (\ell + 1)N\sigma_t^2\sigma_e^2}$, marginally reducing the failure rate. This is given in Table 5.

Master Key Recovery. The security of the master key recovery depends upon the difficulty of finding the short vector (\mathbf{g}, \mathbf{f}) in the lattice, given the public NTRU basis. The attack is successful if the projection of the short vector onto the vector space spanned by the final β Gram-Schmidt vectors is shorter than the length of the $(2N - \beta + 1)^{th}$ Gram-Schmidt vector. This corresponds to minimising block size β , for:

$$\sigma_0\sqrt{\beta} \leq GH(\beta)^{(2\beta-2N)/(\beta-1)} \cdot \det(\Lambda_0)^{1/2N}.$$

The minimum solutions to this inequality for different parameter sets is given in Table 6. The work required to find the shortest vector using this block size with the BKZ2.0 algorithm is also given.

Delegated Key Recovery. For delegated key recovery, the attacker must find a short sequence of vectors in $\Lambda_{\ell-1}$. This can reduce to solving SVP in the master lattice Λ_0 to find a vector of length $\sigma_\ell \cdot \sqrt{2N}$. Table 7 gives the minimum block size β required (as per below Eq. (8)) for a successful attack using BKZ2.0 and the classical and quantum cost of these attacks which depend on N and q .

$$\sigma_\ell \cdot \sqrt{2N} \leq GH(\beta)^{(2N)/(\beta-1)} \cdot \det(\Lambda_0)^{1/2N}. \quad (8)$$

User Key Recovery. User key recovery requires finding a short solution to $t_0 + t_1 \cdot \mathbf{h} + t_2 \cdot \mathbf{A}_1 + \dots + t_\ell \cdot \mathbf{A}_{\ell-1} = \mathbf{A}_\ell$, which reduces to solving the CVP in the master lattice Λ_0 , of the form $t_0 + t_1 \cdot A_0 = A_\ell$. It is enough to find a short (t_0, t_1) with length $\leq \sigma_\ell \cdot \sqrt{2(\ell + 1)} \cdot \sqrt{2N}$. To do this, it is required to minimise Eq. (9) over β . Table 8 gives

Table 9: Cost of Primal Message Recovery Attack.

Set	m	β	Classical Security	Quantum Security
LATTE-1	1018	423	140	128
LATTE-2	1962	967	299	273
LATTE-3	998	232	84	78
LATTE-4	2037	561	180	165

Table 10: Cost of Dual Message Recovery Attack.

Set	m	β	Classical Security	Quantum Security
LATTE-1	1039	422	140	128
LATTE-2	1974	964	298	272
LATTE-3	1005	232	84	78
LATTE-4	2101	560	180	165

Table 11: LATTE Parameters.

Set	Security	N	q
LATTE-1	128	1024	$2^{24} - 2^{14} + 1$
LATTE-2	256	2048	$2^{25} - 2^{12} + 1$
LATTE-3	80	1024	$2^{36} - 2^{20} + 1$
LATTE-4	160	2048	$2^{38} - 2^{26} + 1$

the minimum block size β required for a successful attack, and the classical and quantum cost of these attacks.

$$\sigma_\ell \cdot \sqrt{2(\ell+1)} \cdot \sqrt{2N} \leq GH(\beta)^{(2N)/(\beta-1)} \cdot \det(\Lambda_0)^{1/2N}. \quad (9)$$

Message Recovery. There are two attacks to consider for this. Message recovery depends on solving an extended version of RLWE, which reduces to an instance of the primal-CVP or dual-SVP. In the primal-CVP attack, the ephemeral private keys are recovered via a close vector problem. In the dual-SVP attack, an attempt is made to distinguish the ciphertext elements from uniformly random polynomials in \mathcal{R}_q . In fact, it is enough for the attacker to recover one of the ephemeral private keys, \mathbf{e} and so message recovery cost is not affected by hierarchical level, or by our redesign.

The minimal block size β needed for a successful attack, and the cost of these attacks are given in Tables 9 and 10, depending on (N, q) . The code to populate Tables 9 and 10 is that used in [4]. By considering the cost of all attacks covered in this section, the security levels in Table 11 could be derived.

Setting up Parameters. The parameter sets are given in Table 11. These are the parameters recommended in the original specification [1]. We have extended the security estimates from [1] to give them on a per-level basis. The security decreases as we move down the hierarchy. However, it turns out that each parameter set's security is determined by the message recovery capabilities, which remain constant down the levels. Therefore our parameter security conclusions match that of [1], and furthermore are not affected by our optimisations, as the message recovery attack is independent of the modified parameter ℓ .

Parameter sets LATTE-1 and 2 are only applicable to a single level, essentially an IBE rather than HIBE, version of the scheme. LATTE-3 and 4 can be used for up to two levels. The reason we cannot use these parameters beyond these levels is that the decryption failure rate exceeds the target security level. In fact, the failure rate is so high it renders the scheme completely insecure and also not suitable for use. The key and ciphertext sizes for LATTE are given in Table 13. The method for calculating these is also given in Appendix I.

5 IMPLEMENTATION TECHNIQUES AND OPTIMIZATIONS

In this Section, we discuss the implementation techniques used in our optimised LATTE scheme. First, we present our faster novel fFLDL variant for (Mod)NTRU basis in subsection 5.1. Then, we discuss the techniques adapted from FALCON [51] and ModFALCON [17] in subsection 5.2. After that, we discuss the integer discrete Gaussian sampling techniques in subsection 5.3, including the adaptation of FACCT [56] and COSAC [54, 55] samplers in our LATTE implementation.

Algorithm 1: Optimised fFLDL algorithm for (Mod)NTRU basis in LATTE.

Input: Gram matrix $\mathbf{G} \in (\mathbb{C}[x]/\langle x^n + 1 \rangle)^{d \times d}$ in the FFT domain. $d \in \{2, 3\}$. $\mathbf{D}' \in (\mathbb{R}^+)^{2n}$.

Output: Tree T .

```

1 if  $n = 1$  then
2    $T.value \leftarrow \mathbf{G}_{0,0}$ .
3 else
4    $\mathbf{L} \leftarrow \mathbf{I}_d, \mathbf{D} \leftarrow \mathbf{0}_d$ .
5   for  $j = 0$  to  $n - 1$  do
6      $(\mathbf{D}_{0,0})_j \leftarrow (\mathbf{G}_{0,0})_j, (\mathbf{L}_{1,0})_j \leftarrow \frac{(\mathbf{G}_{1,0})_j}{(\mathbf{D}_{0,0})_j}$ .
7     if  $d = 2$  then
8       if  $n = N$  then  $(\mathbf{D}_{1,1})_j \leftarrow \frac{q^2}{(\mathbf{D}_{0,0})_j}$  else
9          $(\mathbf{D}_{1,1})_j \leftarrow \frac{\mathbf{D}'_{2j} \mathbf{D}'_{2j+1}}{(\mathbf{D}_{0,0})_j}$ ;
10      else if  $d = 3$  then
11         $(\mathbf{D}_{1,1})_j \leftarrow (\mathbf{G}_{1,1})_j - \frac{|(\mathbf{G}_{1,0})_j|^2}{(\mathbf{D}_{0,0})_j}$ .
12         $(\mathbf{D}_{2,2})_j \leftarrow \frac{q^2}{(\mathbf{D}_{0,0})_j (\mathbf{D}_{1,1})_j}, (\mathbf{L}_{2,0})_j \leftarrow \frac{(\mathbf{G}_{2,0})_j}{(\mathbf{D}_{0,0})_j}$ .
13         $(\mathbf{L}_{2,1})_j \leftarrow \frac{(\mathbf{G}_{2,1})_j - (\mathbf{G}_{2,0})_j (\mathbf{L}_{1,0})_j^*}{(\mathbf{D}_{1,1})_j}$ .
14      $T.value \leftarrow \mathbf{L}$ .
15     for  $i = 0$  to  $d - 1$  do
16        $\mathbf{d}_0, \mathbf{d}_1 \leftarrow \text{splitfft}(\mathbf{D}_{i,i})$ .
17        $\mathbf{G}' = \begin{pmatrix} \mathbf{d}_0 & \mathbf{d}_1 \\ \mathbf{d}_1^* & \mathbf{d}_0 \end{pmatrix}$ .
18        $T.child_i \leftarrow \text{fFLDL}(\mathbf{G}', \mathbf{D}_{i,i})$ .
19 return  $T$ .
```

5.1 Improved fFLDL Algorithm for NTRU Basis

The original fFLDL algorithm from FALCON [51] for the Fast Fourier LDL* decomposition is shown in Alg. 8 in Appendix B.1. However,

Table 12: Comparison between Our Optimised LATTE Performance Results (op/s) at 4.2GHz with Original LATTE in [1].

Set	KeyGen	$\ell = 1$				$\ell = 2$		
		Ext	Enc	Dec	Del	Ext	Enc	Dec
Orig. LATTE-1 [1]	-	-	2911	2987	-	-	-	-
Our LATTE-1	9.4	1442.1	16525.2	13140.3	-	-	-	-
Orig. LATTE-2 [1]	-	-	1335	1351	-	-	-	-
Our LATTE-2	3.3	613.1	7692.6	6183.3	-	-	-	-
Orig. LATTE-3 [1]	-	-	1892	1774	-	-	1455	1474
Our LATTE-3	5.7	36.3	8000.9	6548.2	2.4	20.0	6351.9	5258.9
Orig. LATTE-4 [1]	-	-	709	668	-	-	568	541
Our LATTE-4	1.7	17.0	3909.6	3201.6	0.8	9.4	3108.0	2584.6

for the (Mod)NTRU basis S_ℓ in LATTE, we observe the following theorem, which can be adapted to accelerate the computation of the fFLDL algorithm (see Appendix F for the proof):

THEOREM 5.1. *Let S_ℓ be a (Mod)NTRU basis. In fFLDL tree of the matrix $G = S_\ell S_\ell^* \in (\mathbb{C}[x]/\langle x^N + 1 \rangle)^{d \times d}$ in FFT domain, we get:*

- (1) $\forall i \in \{0, \dots, d-1\} : D_{i,i} \in \mathbb{R}^n$ for some $n = 2^k \leq N$ in every node of the tree.
- (2) $\forall j \in \{0, \dots, N-1\} : \prod_{i=0}^{d-1} (D_{i,i})_j = q^2$ in the root of the tree.
- (3) $\forall j \in \{0, \dots, n-1\} : (D_{0,0})_j (D_{1,1})_j = D'_{2j} D'_{2j+1}$ for some $n = 2^k \leq N/2$ in every non-root node of the tree, where $D' \in \{D_{i,i}\}_{i=0}^{d-1}$ is from its parent.
- (4) $\forall i \in \{0, \dots, d-1\}, j \in \{0, \dots, n-1\} : (D_{i,i})_j \in \mathbb{R}^+$ for some $n = 2^k \leq N$ in every node of the tree.

We can utilise Theorem 5.1 when computing D in the fFLDL algorithm, see Alg. 1, for the (Mod)NTRU basis S_ℓ in LATTE with $d \in \{2, 3\}$: $D_{d-1,d-1}$ at the root can be computed by $(D_{d-1,d-1})_j = q^2 / \prod_{i=0}^{d-2} (D_{i,i})_j$ for $0 \leq j \leq N-1$. For all the non-root nodes, we can directly compute $D_{0,0}, D_{1,1}$ by using $(D_{0,0})_j = (G_{0,0})_j$ and $(D_{1,1})_j = D'_{2j} D'_{2j+1} / (D_{0,0})_j$, $0 \leq j \leq n-1$, for some $D' \in \mathbb{R}^{2n}$, $G_{0,0} = d'_0 \in \mathbb{R}^n$ from its parent. Since for all $0 \leq i \leq d-1$, we have $D_{i,i} \in \mathbb{R}^n$ in every node of the tree, D can be computed solely by using the real number arithmetic i.e. without complex number arithmetic. Because every complex number arithmetic computation contains multiple underlying floating-point arithmetic operations, by replacing complex number arithmetic with real number arithmetic when computing D , we reduce the total amount of floating-point arithmetic operations. Therefore, this optimisation technique will accelerate the run-time speed of fFLDL algorithm.

5.2 Techniques from FALCON and MODFALCON

The design of LATTE presented in this paper utilises techniques from the signature scheme FALCON [51]. The two schemes are closely related; they are instantiated over the same type of lattice and share key generation and sampling procedures. FALCON makes use of the “tower of rings” structure to find a solution to the NTRU equation $fG - gF = q \pmod{x^N + 1}$, for a given f and g in the NTRUSolve sub-algorithm of KeyGen, and in the lattice Gaussian sampling (ffSampling) component of LATTE Delegate and Extract. The tower of rings approach utilises the fact that computations over polynomials $f, g \in \mathbb{C}[x]/\langle x^{N/2} + 1 \rangle$ are equivalent to computations

over $f(x^2), g(x^2) \in \mathbb{C}[x]/\langle x^N + 1 \rangle$. When $N = 2^k$, for some $k \in \mathbb{Z}$, this can be applied repeatedly so that computations are performed over polynomials of degree 1. This brings advantages in terms of both memory usage and speed [48].

Furthermore, in LATTE Delegate, to complete the delegated basis S_ℓ for lattice dimension higher than $2N$, we adapt the technique from MODFALCON [17]. Let $S_\ell = \begin{pmatrix} \mathbf{v}^\top & \mathbf{M} \\ \mathbf{G}_\ell & \mathbf{F}'_\ell \end{pmatrix}$ be the delegated basis, where $\mathbf{G}_\ell = (s_{\ell+1,0}, \mathbf{F}'_\ell = (s_{\ell+1,1}, \dots, s_{\ell+1,\ell+1}), \mathbf{v} = (s_{0,0}, s_{1,0}, \dots, s_{\ell,0})$, and $\mathbf{M} = (s_{i,j})$ for $0 \leq i \leq \ell$ and $1 \leq j \leq \ell+1$. By Schur complement, if \mathbf{M} is invertible, we have:

$$\begin{aligned} \det(S_\ell) &= \det(\mathbf{M}) \cdot \det(\mathbf{G}_\ell - \mathbf{F}'_\ell \mathbf{M}^{-1} \mathbf{v}^\top) \\ &= \det(\mathbf{M}) (\mathbf{G}_\ell - \mathbf{F}'_\ell \mathbf{M}^{-1} \mathbf{v}^\top) = \det(\mathbf{M}) \mathbf{G}_\ell - \mathbf{F}'_\ell \text{adj}(\mathbf{M}) \mathbf{v}^\top. \end{aligned}$$

Since one can choose any $(S_\ell)_{\ell+1} = (\mathbf{G}_\ell, \mathbf{F}'_\ell)$ such that $\det(S_\ell) = q$ when filling the bottom row $(S_\ell)_{\ell+1}$ of S_ℓ , let \mathbf{F}'_ℓ have the form $(\mathbf{F}'_\ell, \mathbf{0}, \dots, \mathbf{0})$. We have $\det(S_\ell) = \det(\mathbf{M}) \cdot \mathbf{G}_\ell - \mathbf{F}'_\ell \cdot \mathbf{u}_0$ where \mathbf{u}_0 is the first coordinate of $\mathbf{u} = \text{adj}(\mathbf{M}) \cdot \mathbf{v}^\top$. In order to fill the bottom row $(s_{\ell+1,0}, \dots, s_{\ell+1,\ell+1})$ of S_ℓ , if \mathbf{M} is invertible, we can use the same NTRUSolve algorithm as in LATTE KeyGen to find $\mathbf{F}_\ell, \mathbf{G}_\ell$ such that $\det(\mathbf{M}) \cdot \mathbf{G}_\ell - \mathbf{F}_\ell \cdot \mathbf{u}_0 = q$, and we simply resample when $\det(\mathbf{M}) = 0$.

However, since the NTRUSolve algorithm [48] performs the length reduction based on the size of the coefficients in the input, the coefficient size of the output $\mathbf{F}_\ell, \mathbf{G}_\ell$ will be approximately the coefficient size of the input $\det(\mathbf{M}), \mathbf{u}_0$. Since \mathbf{M} is an $(\ell+1) \times (\ell+1)$ sub-matrix of S_ℓ with coordinate sizes being in the order of q among each element, the size of coefficients of $\det(\mathbf{M}), \mathbf{u}_0, \mathbf{F}_\ell$, and \mathbf{G}_ℓ is in the order of $q^{\ell+1}$. To make the infinity norm of S_ℓ less than q , one employs length reduction using Cramer’s rule (see Appendix H).

5.3 Discrete Gaussian Sampling over Integers

In LATTE KeyGen, f, g may need to be resampled multiple times due to the norm check and possible failure to find solutions of the NTRU equation. In order to sample $2N$ coordinates efficiently from \mathcal{D}_{σ_0} , we employ the FACCT sampler [56], which is fast and compact even for larger σ_0 used in LATTE-3 and 4. However, since the FACCT sampler can only sample with $\sigma = k\sqrt{1/(2 \ln 2)}$ where k is a positive integer, we slightly increase $\sigma_0 \approx 1.17\sqrt{q/(2N)}$ in LATTE parameters by setting $k = \lceil 1.17\sqrt{q/(2N)} / \sqrt{1/(2 \ln 2)} \rceil$.

Let $S_\ell = \mathbf{L} \cdot \tilde{S}_\ell$ be the GSO decomposition of the delegated basis $S_\ell \in \mathcal{R}^{(\ell+2) \times (\ell+2)}$, where \mathbf{L} is a unit lower triangular and rows \tilde{s}_i of \tilde{S}_ℓ are pairwise orthogonal. We find that the Euclidean norm of the last GSO vector $\tilde{s}_{\ell+1}$ is very small compared to $\tilde{s}_0, \dots, \tilde{s}_\ell$. This is

Table 13: LATTE and Original LATTE Master Key, Delegated Key and Ciphertext Sizes (Bytes).

Set	Master Public Key	Master Private Key	User Private Key		Ciphertext		Delegated Public Key	Delegated Private Key
			$\ell = 1$	$\ell = 2$	$\ell = 1$	$\ell = 2$		
Orig. LATTE-1 [1]	6144	12288	9216	-	9248	-	-	-
Our LATTE-1	3072	12288	3072	-	6176	-	-	-
Orig. LATTE-2 [1]	12800	25600	19200	-	19232	-	-	-
Our LATTE-2	6400	25600	6400	-	12832	-	-	-
Orig. LATTE-3 [1]	9216	18432	13824	18432	13856	18464	9216	41472
Our LATTE-3	4608	18432	4608	9216	9248	13856	9216	41472
Orig. LATTE-4 [1]	19456	38912	29184	38912	29216	38944	19456	87552
Our LATTE-4	9728	38912	9728	19456	19488	29216	19456	87552

because rows s_0, \dots, s_ℓ of S_ℓ are sampled with a large σ_ℓ but $\det(S_\ell \cdot S_\ell^*) = \prod_{i=0}^{\ell+1} \langle \tilde{s}_i, \tilde{s}_i \rangle$ is constant and always equal to q^2 [17]. The experiment results in Fig.3 of [15] also verified that $\|\tilde{s}_{\ell+1}\|$ decreases significantly by increasing $\|s_0\|$ for $S_\ell \in \mathcal{R}^{3 \times 3}$. In this case, the ratio between the maximal and minimal standard deviation σ' used by the integer discrete Gaussian sampling subroutine in `ffSampling` is very large and the isochronous sampler [32] used by `FALCON` [51] will be inefficient for our scheme, since the rejection rate of [32] is proportional to $\max(\sigma')/\min(\sigma')$. In order to sample with σ' in a broad range, we employ a variant [54] of the `COSAC` sampler [55] instead, which is scalable to large σ' without sacrificing efficiency.

The precision analysis in Sec. 4.3 requires the bound B on Rényi divergence between a single sample from `COSAC` and an ideal Gaussian \mathbb{Z} sample. In [55] it is shown that $B \leq 1 + 4\sigma^2 e_x^2 \lambda$, where e_x denotes the absolute error of the underlying Box-Muller continuous Gaussian sampler used by the `COSAC` sampler and σ denotes the upper bound of the integer Gaussian standard deviation σ . For $\ell = 1$, we can derive $\sigma \leq \eta_\epsilon(\mathbb{Z}) \cdot (\sigma_0 \sqrt{2N})^2 / q$ from [32] by symplecticity of the basis. For $\ell = 2$, we can derive $\sigma \leq \sigma_2 \cdot 9\sigma_1^2 N^2 / q$ from the analysis in Appendix G. We use double precision i.e. 53-bit floating-point arithmetic precision in the `COSAC` sampler for `LATTE-1` and `2`, which provides $e_x \leq 2^{-48}$ [55]. Since the run-time speed of the underlying Box-Muller continuous Gaussian sampler is critical for the speed of the `COSAC` sampler [55], for the `COSAC` implementation in `LATTE-3` and `4`, we use `binary128` i.e. 113-bit floating-point arithmetic precision and reduce the absolute precision of uniform sampling in the underlying Box-Muller continuous Gaussian sampler to 96 bits. This will make e_x less than approximately 2^{-96} .

To accelerate the `LATTE` Encrypt and Decrypt speed, we sample the ephemeral keys $\mathbf{e}, \mathbf{e}_1, \dots, \mathbf{e}_\ell, \mathbf{e}_h$ from a binomial distribution with center 0 and small standard deviation $\sigma_e = 2.0$ instead of \mathcal{D}_{σ_e} used by [1]. Sampling from a binomial distribution is much faster than sampling from \mathcal{D}_{σ_e} and the impact on security is negligible in the encryption [4].

6 PERFORMANCE RESULTS

The first published specification of `LATTE` [1] only provided the Encrypt and Decrypt performance results, as displayed in “Orig. `LATTE`” rows in Table 12, scaled and converted into op/s at 4.2GHz. Here, we give the first full performance results for our optimised variant of `LATTE`, including KeyGen, Extract, and Delegate.

We employ the `gmp` [28] library for multiprecision integer arithmetic. For precisions of floating-point and complex number arithmetic, we use 53 bits i.e. double precision for `LATTE-1` and `2`, and we use 113 bits i.e. `binary128` for `LATTE-3` and `4`, respectively. We use the `__float128` and `__complex128` variable types from `gcc` to implement the 113-bit floating-point and complex number arithmetic for `LATTE-3` and `4`, respectively. Although the error analysis in Sec. 4.3 indicates that the arithmetic precisions for `LATTE-3` and `4` can be further reduced, however, the generic multiprecision floating-point library such as `MPFR` [24] is not optimised for less than 1000-bit precision in terms of the run-time speed [35]. We will leave using hand-optimised floating-point arithmetic routines with lower precision as future works.

We use AES-256 CTR mode with hardware AES-NI instructions as the pseudorandom generator, and use `SHAKE-256` [43] as the KDF in `LATTE` Encrypt and Decrypt. The performance results have been obtained from a desktop machine with an Intel i7-7700K CPU at 4.2GHz, with both hyper-threading and TurboBoost disabled. We use `gcc 11.2.0` compiler with compiling options `-O3 -march=native` enabled. Results are given as “Our `LATTE`” rows in Table 12.

As expected, the KeyGen, Extract, and Delegate processes are the most time consuming components of the scheme, and this increases as security and therefore lattice dimension increase. The trend down the hierarchical levels is that the Extract, Encrypt, and Decrypt all become more time consuming as hierarchical level increases. For Extract in `LATTE-3` and `4`, this corresponds to about 45% decrease in op/s from level 1 to level 2. On the other hand, for the Encrypt and Decrypt, our implementation is 3.6x–5.8x faster compared to the previous performance results from [1]. The speedup might be due to: (1) We change the distribution of the ephemeral keys from discrete Gaussian distribution to binomial distribution. (2) We only perform NTT for the ephemeral keys and \mathbf{m} during the Encrypt and Decrypt, since other inputs are already in the NTT domain. (3) Since we reduce the dimension of extracted user keys by 1, there is also 1 less ephemeral key in Encrypt/Decrypt. In addition, our optimised `LATTE` Delegate only takes about 0.4–1.3 seconds on a desktop machine at 4.2GHz, which is practical and much faster than the estimated run-time (in the order of minutes) for the Delegate in [1]. For comparisons with other (H)IBE and the `FALCON` signature scheme, please refer to Appendix J.

In addition, since we reduce the dimension of extracted user keys by 1 in our improved LATTE scheme, here we compare the key and ciphertext sizes of our improved LATTE scheme with the original LATTE in [1]. The key/ciphertext sizes are summarised in Table 13, with detailed analysis given in Appendix I. The key/ciphertext sizes of the original LATTE [1] are shown in rows with “Orig. LATTE”, and the key/ciphertext sizes of our improved LATTE scheme are shown in rows with “Our LATTE”, respectively. From Table 13, our improved LATTE scheme reduces the key/ciphertext sizes by 25%–67% among all LATTE parameter sets.

Our current implementation is not constant-time, since the gmp multiprecision integer arithmetic library [28] and the gcc runtime library for the binary128 floating-point and complex number arithmetic are unlikely to be constant-time [45]. We will leave the constant-time implementation of our optimised LATTE scheme as future works.

REFERENCES

- [1] 2019. *Quantum-Safe Identity-based Encryption*. Technical Report. The European Telecommunications Standards Institute, Sophia-Antipolis, France.
- [2] 2020. *LTE;Security of the mission critical service (3GPP TS 33.180 version 14.8.0 Release 14)*.
- [3] Masayuki Abe, Rosario Gennaro, and Kaoru Kurosawa. 2005. Tag-KEM/DEM: A New Framework for Hybrid Encryption. *IACR Cryptol. ePrint Arch.* (2005), 27.
- [4] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. 2016. Post-quantum Key Exchange - A New Hope. In *USENIX Security Symposium*. USENIX Association, 327–343.
- [5] S. Bai, T. Lepoint, A. Roux-Langlois, A. Sakzad, D. Stehlé, and R. Steinfeld. 2018. Improved Security Proofs in Lattice-Based Cryptography: Using the Rényi Divergence Rather than the Statistical Distance. *J. Cryptology* 31, 2 (2018), 610–640.
- [6] Dan Boneh and Xavier Boyen. 2004. Efficient Selective-ID Secure Identity-Based Encryption Without Random Oracles. In *EUROCRYPT (Lecture Notes in Computer Science, Vol. 3027)*. Springer, 223–238.
- [7] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. 2005. Hierarchical Identity Based Encryption with Constant Size Ciphertext. In *EUROCRYPT (Lecture Notes in Computer Science, Vol. 3494)*. Springer, 440–456.
- [8] Dan Boneh and Matthew K. Franklin. 2001. Identity-Based Encryption from the Weil Pairing. In *CRYPTO (Lecture Notes in Computer Science, Vol. 2139)*. Springer, 213–229.
- [9] Nicolas Brisebarre, Mioara Joldes, Jean-Michel Muller, Ana-Maria Nanes, and Joris Picot. 2020. Error Analysis of Some Operations Involved in the Cooley-Tukey Fast Fourier Transform. *ACM Trans. Math. Softw.* 46, 2 (2020), 11:1–11:27.
- [10] Peter Campbell and Michael Groves. 2017. Practical post-quantum hierarchical identity-based encryption. available at <https://www.qub.ac.uk/csit/FileStore/Fileupload,785752,en.pdf>.
- [11] Ran Canetti, Shai Halevi, and Jonathan Katz. 2003. A Forward-Secure Public-Key Encryption Scheme. In *EUROCRYPT (Lecture Notes in Computer Science, Vol. 2656)*. Springer, 255–271.
- [12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. 2010. Bonsai Trees, or How to Delegate a Lattice Basis. In *EUROCRYPT (Lecture Notes in Computer Science, Vol. 6110)*. Springer, 523–552.
- [13] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. 2012. Bonsai trees, or how to delegate a lattice basis. *Journal of cryptology* 25, 4 (2012), 601–639.
- [14] Yuanmi Chen and Phong Q. Nguyen. 2011. BKZ 2.0: Better Lattice Security Estimates. In *ASIACRYPT (Lecture Notes in Computer Science, Vol. 7073)*. Springer, 1–20.
- [15] Jung Hee Cheon, Duhyeon Kim, Taechan Kim, and Yongha Son. 2019. A New Trapdoor over Module-NTRU Lattice and its Application to ID-based Encryption. *IACR Cryptol. ePrint Arch.* 2019 (2019), 1468.
- [16] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. August 2016. TFHE: Fast Fully Homomorphic Encryption Library. <https://tfhe.github.io/tfhe/>.
- [17] HITCHANOK CHUENGSAITANSUP, THOMAS PREST, DAMIEN STEHLÉ, ALEXANDRE WALLET, and KEITA XAGAWA. 2020. ModFalcon: Compact Signatures Based On Module-NTRU Lattices. In *AsiaCCS*. ACM, 853–866.
- [18] Urgent Comm. 2022. ESN sees ‘good progress’ but challenges remain in the UK, director says. <https://urgentcomm.com/2022/03/19/esn-sees-good-progress-but-challenges-remain-in-the-uk-director-says/>.
- [19] Cramer and Gabriel. 1750. *Introduction a l’analyse des lignes courbes algébriques par Gabriel Cramer*. chez les freres Cramer & Cl. Philibert.
- [20] Eloïsa Díaz-Francés and Francisco J Rubio. 2013. On the existence of a normal approximation to the distribution of the ratio of two independent normal random variables. *Statistical Papers* 54, 2 (2013), 309–323.
- [21] Yevgeniy Dodis and Nelly Fazio. 2002. Public Key Broadcast Encryption for Stateless Receivers. In *Digital Rights Management Workshop (Lecture Notes in Computer Science, Vol. 2696)*. Springer, 61–80.
- [22] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. 2014. Efficient Identity-Based Encryption over NTRU Lattices. In *ASIACRYPT (2) (Lecture Notes in Computer Science, Vol. 8874)*. Springer, 22–41.
- [23] Léo Ducas and Thomas Prest. 2016. Fast Fourier Orthogonalization. In *ISSAC*. ACM, 191–198.
- [24] Laurent Fousse, Guillaume Hanrot, Vincent Lefèvre, Patrick Pélicissier, and Paul Zimmermann. 2007. MPFR: A multiple-precision binary floating-point library with correct rounding. *ACM Trans. Math. Softw.* 33, 2 (2007), 13.
- [25] Eiichiro Fujisaki and Tatsuaki Okamoto. 1999. Secure Integration of Asymmetric and Symmetric Encryption Schemes. In *CRYPTO (Lecture Notes in Computer Science, Vol. 1666)*. Springer, 537–554.
- [26] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. 2008. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*. ACM, 197–206.
- [27] Craig Gentry and Alice Silverberg. 2002. Hierarchical ID-Based Cryptography. In *ASIACRYPT (Lecture Notes in Computer Science, Vol. 2501)*. Springer, 548–566.
- [28] Torbjørn Granlund and Gmp Development Team. 2015. *GNU MP 6.0 Multiple Precision Arithmetic Library*. Samurai Media Limited, London, GBR.
- [29] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. 1998. NTRU: A Ring-Based Public Key Cryptosystem. In *ANTS (Lecture Notes in Computer Science, Vol. 1423)*. Springer, 267–288.
- [30] Jeremy Horwitz and Ben Lynn. 2002. Toward Hierarchical Identity-Based Encryption. In *EUROCRYPT (Lecture Notes in Computer Science, Vol. 2332)*. Springer, 466–481.
- [31] James Howe, Thomas Prest, Thomas Ricosset, and Mélissa Rossi. 2020. Isochronous Gaussian Sampling: From Inception to Implementation. In *PQCrypto*. 53–71.
- [32] James Howe, Thomas Prest, Thomas Ricosset, and Mélissa Rossi. 2020. Isochronous Gaussian Sampling: From Inception to Implementation. In *PQCrypto (Lecture Notes in Computer Science, Vol. 12100)*. Springer, 53–71.
- [33] Takeshi Koshihara and Katsuyuki Takashima. 2018. New Assumptions on Isogenous Pairing Groups with Applications to Attribute-Based Encryption. In *ICISC (Lecture Notes in Computer Science, Vol. 11396)*. Springer, 3–19.
- [34] Adeline Langlois, Damien Stehlé, and Ron Steinfeld. 2014. GGHLite: More Efficient Multilinear Maps from Ideal Lattices. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings (Lecture Notes in Computer Science, Vol. 8441)*, Phong Q. Nguyen and Elisabeth Oswald (Eds.). Springer, 239–256. https://doi.org/10.1007/978-3-642-55220-5_14
- [35] Christoph Quirin Lauter. 2015. Easing development of precision-sensitive applications with a beyond-quad-precision library. In *ACSSC*. IEEE, 742–746.
- [36] Arjen Lenstra, Hendrik Lenstra, and László Lovász. 1982. Factoring Polynomials With Rational Coefficients. *Math. Ann.* 261 (1982), 515–534.
- [37] Granit Luzhnica. 2011. *Pairing based cryptography and implementation in Java*. Master Thesis.
- [38] Vadim Lyubashevsky. 2012. Lattice Signatures without Trapdoors. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7237)*, David Pointcheval and Thomas Johansson (Eds.). Springer, 738–755. https://doi.org/10.1007/978-3-642-29011-4_43
- [39] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. 2010. On Ideal Lattices and Learning with Errors over Rings. In *EUROCRYPT (Lecture Notes in Computer Science, Vol. 6110)*. Springer, 1–23.
- [40] Sarah McCarthy, Neil Smyth, and Elizabeth O’Sullivan. 2017. A Practical Implementation of Identity-Based Encryption Over NTRU Lattices. In *IMACC (Lecture Notes in Computer Science, Vol. 10655)*. Springer, 227–246.
- [41] Daniele Micciancio and Oded Regev. 2004. Worst-Case to Average-Case Reductions Based on Gaussian Measures. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy. Proceedings*. IEEE Computer Society, 372–381. <https://doi.org/10.1109/FOCS.2004.72>
- [42] Daniele Micciancio and Michael Walter. 2017. Gaussian Sampling over the Integers: Efficient, Generic, Constant-Time. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 10402)*, Jonathan Katz and Hovav Shacham (Eds.). Springer, 455–485. https://doi.org/10.1007/978-3-319-63715-0_16
- [43] NIST. 2015. SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions. <https://doi.org/10.6028/NIST.FIPS.202>.
- [44] NIST. 2016. Post-Quantum Crypto Project. <https://csrc.nist.gov/Projects/post-quantum-cryptography/Post-Quantum-Cryptography-Standardization>.
- [45] Tobias Oder, Julian Speith, Kira Hölftgen, and Tim Güneysu. 2019. Towards Practical Microcontroller Implementation of the Signature Scheme Falcon. In

- PQCrypto (Lecture Notes in Computer Science, Vol. 11505)*. Springer, 65–80.
- [46] UK Home Office. 2022. Emergency Services Network: overview. <https://www.gov.uk/government/publications/the-emergency-services-mobile-communications-programme/emergency-services-network>.
- [47] Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. 2014. Enhanced Lattice-Based Signatures on Reconfigurable Hardware. In *CHES (Lecture Notes in Computer Science, Vol. 8731)*. Springer, 353–370.
- [48] Thomas Pornin and Thomas Prest. 2019. More Efficient Algorithms for the NTRU Key Generation Using the Field Norm. In *Public Key Cryptography (2) (Lecture Notes in Computer Science, Vol. 11443)*. Springer, 504–533.
- [49] Thomas Prest. 2017. Sharper bounds in lattice-based cryptography using the Rényi divergence. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 347–374.
- [50] Thomas Prest. 2021. Rényi Divergence Analysis. Unpublished. Private communication.
- [51] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. 2020. *Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU*. Technical Report. National Institute of Standards and Technology. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.
- [52] Markku-Juhani O Saarinen. 2015. Gaussian sampling precision in lattice cryptography. *IACR ePrint* 953 (2015), 2015.
- [53] Claus-Peter Schnorr and M. Euchner. 1994. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.* 66 (1994), 181–199.
- [54] Shuo Sun, Yongbin Zhou, Yunfeng Ji, Rui Zhang, and Yang Tao. 2021. Generic, Efficient and Isochronous Gaussian Sampling over the Integers. *IACR Cryptol. ePrint Arch.* 2021 (2021), 199.
- [55] Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad. 2020. COSAC: CCompact and Scalable Arbitrary-Centered Discrete Gaussian Sampling over Integers. In *PQCrypto (Lecture Notes in Computer Science, Vol. 12100)*. Springer, 284–303.
- [56] Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad. 2020. FACCT: FAsT, Compact, and Constant-Time Discrete Gaussian Sampler over Integers. *IEEE Trans. Computers* 69, 1 (2020), 126–137.

A HIBE SCHEME DESCRIPTION

This section presents the components of an HIBE scheme and how they interact. They are as follows:

- (1) **KeyGen**: The master key generator establishes the master public and private keys.
- (2) **Delegate**: Through a delegation function, the master key generator creates a public/private key pair for the sub key manager. This gives it the ability to delegate further key pairs, and extract user private keys at that level.
- (3) **Delegate**: The sub key manager delegates a further public/private key to the next level of the hierarchy.
- (4) **Extract**: The extractor uses their public/private key pair to extract and share user public/private keys, as in the single-level IBE scheme.
- (5) **Encrypt/Decrypt**: Encryption/decryption works as a regular encryption scheme, such as R-LWE encryption.

Fig. 1 depicts a diagram of a 2-level HIBE scheme with all its algorithms.

An HIBE scheme is said to be IND-CCA-secure if it is indistinguishable under chosen ciphertext attacks; that is, an adversary with the ability to decrypt any other ciphertext does not possess an advantage in decrypting the challenge ciphertext. ID-IND-CCA further implies the adversary has access to an extraction oracle which allows them to extract keys for other identities before committing to the challenge identity, yet gains no advantage. The challenge consists of the ciphertext *and* the identity under which it is encrypted.

Table 14 indicates the notational practises used to identify each level of the hierarchy.

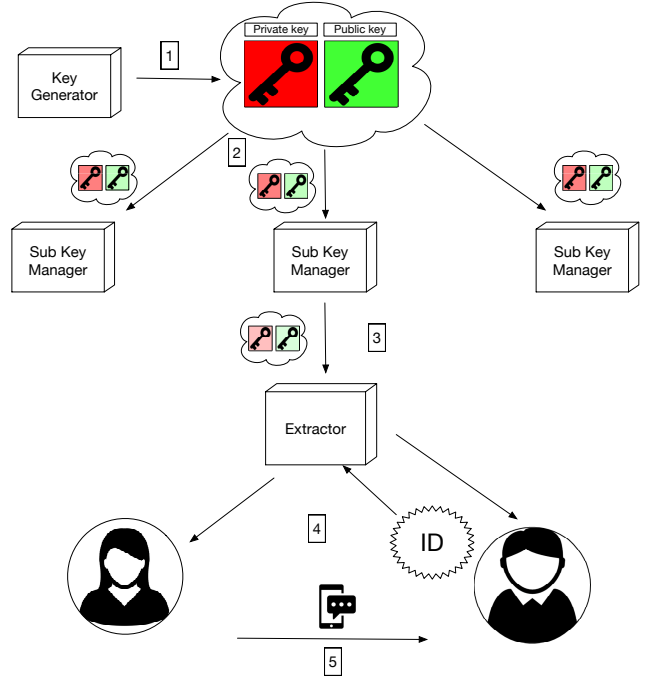


Figure 1: A 2-level HIBE scheme.

Table 14: Explanation of Notational Practice of HIBE Functions.

Level	Function
Level 0	Master KeyGen $2N \times 2N$
Level 1	Extracting with $2N \times 2N \rightarrow \text{Enc/Dec}$
	Delegating to $3N \times 3N$
Level 2	Extracting with $3N \times 3N \rightarrow \text{Enc/Dec}$

B LATTE HIBE SCHEME

We present the full pseudocode of LATTE in Alg. 2–6.

Algorithm 2: The LATTE KeyGen algorithm.

Input: N, q, σ_0 .
Output: $S_0 \in \mathcal{R}_q^{2 \times 2}, \mathbf{h} \in \mathcal{R}_q$.

- 1 $\mathbf{f}, \mathbf{g} \leftarrow \mathcal{D}_{\sigma_0}^N$.
- 2 $Norm \leftarrow \max\left(\|\mathbf{g}, -\mathbf{f}\|, \left\| \left(\frac{q \cdot \mathbf{f}^*}{\mathbf{f} \cdot \mathbf{f}^* + \mathbf{g} \cdot \mathbf{g}^*}, \frac{q \cdot \mathbf{g}^*}{\mathbf{f} \cdot \mathbf{f}^* + \mathbf{g} \cdot \mathbf{g}^*} \right) \right\| \right)$.
- 3 **if** $Norm > \sigma_0 \cdot \sqrt{2N}$ **then** go to Step 1;
- 4 $\mathbf{F}, \mathbf{G} \leftarrow \text{NTRUSolve}_{N,q}(\mathbf{f}, \mathbf{g})$.
- 5 **if** *NTRUSolve is aborted* **then** go to Step 1;
- 6 **if** \mathbf{f} is not invertible on \mathcal{R}_q **then** go to Step 1;
- 7 $\mathbf{h} \leftarrow \mathbf{g} \cdot \mathbf{f}^{-1} \pmod{q}$ in NTT domain.
- 8 **return** $S_0 = \begin{pmatrix} \mathbf{g} & -\mathbf{f} \\ \mathbf{G} & -\mathbf{F} \end{pmatrix}, \mathbf{h}$.

Algorithm 3: The LATTE Delegate algorithm (from level $\ell - 1$ to ℓ).

Input: $N, q, \sigma_\ell, S_{\ell-1}, H : \{0, 1\}^* \rightarrow \mathcal{R}_q, \text{ID}_\ell$.
Output: $S_\ell \in \mathcal{R}_q^{(\ell+2) \times (\ell+2)}$.

- 1 $A_\ell \leftarrow H(\text{ID}_1 | \dots | \text{ID}_\ell)$ in NTT domain.
- 2 $T_{\ell-1} \leftarrow \text{Tree}(S_{\ell-1}, \sigma_\ell)$.
- 3 $seed \leftarrow \{0, 1\}^{256}$.
- 4 **for** $i \in \{0, \dots, \ell\}$ **do**
- 5 $s_{i,\ell+1} \leftarrow \mathcal{D}_{\sigma_\ell}^N$.
- 6 $\mathbf{t} \leftarrow (-s_{i,\ell+1} \cdot A_\ell, \mathbf{0}, \dots, \mathbf{0}) \cdot S_{\ell-1}^{-1}$.
- 7 $\mathbf{z} \leftarrow \text{FFT}^{-1}(\text{ffSampling}(\mathbf{t}, T_{\ell-1}, seed))$.
- 8 $(s_{i,0}, s_{i,1}, \dots, s_{i,\ell}) \leftarrow \lfloor \bar{\mathbf{z}} \rfloor$, where $\bar{\mathbf{z}} \leftarrow (\mathbf{t} - \mathbf{z}) \cdot S_{\ell-1}$.
- 9 **if** $\| (s_{i,0}, s_{i,1}, \dots, s_{i,\ell}, s_{i,\ell+1}) \| > \sqrt{(\ell+2)N} \cdot \sigma_\ell$ **then**
 resample;
- 10 Set $\mathbf{M} = (s_{i,j})$, for $0 \leq i \leq \ell$ and $1 \leq j \leq \ell + 1$.
- 11 **if** \mathbf{M} is not invertible **then** go to Step 3;
- 12 $\mathbf{u} \leftarrow \text{adj}(\mathbf{M}) \cdot (s_{0,0}, s_{1,0}, \dots, s_{\ell,0})^\top$.
- 13 $(F_\ell, G_\ell) \leftarrow \text{NTRUSolve}_{N,q}(\det(\mathbf{M}), \mathbf{u}_0)$ where \mathbf{u}_0 is the first coordinate of \mathbf{u} .
- 14 **if** NTRUSolve is aborted **then** go to Step 3;
- 15 $(s_{\ell+1,0}, \dots, s_{\ell+1,\ell+1}) \leftarrow (G_\ell, F_\ell, \mathbf{0}, \dots, \mathbf{0})$.
- 16 Set $\mathbf{C} = (c_{i,j})$, where $c_{i,j} = s_{j,0} \cdot s_{i,0}^* + \dots + s_{j,\ell+1} \cdot s_{i,\ell+1}^*$,
 $0 \leq i, j \leq \ell$.
- 17 Let $\mathbf{k} = (k_i)_{0 \leq i \leq \ell}$ be the solution to $\mathbf{C} \cdot \mathbf{k} = \mathbf{d}$. By Cramer's rule, $k_i = \frac{\det(C_i(\mathbf{d}))}{\det(\mathbf{C})}$, where $C_i(\mathbf{d})$ is the matrix \mathbf{C} with its i^{th} column replaced by \mathbf{d} .
 $\mathbf{d}_i = s_{\ell+1,0} \cdot s_{i,0}^* + \dots + s_{\ell+1,\ell+1} \cdot s_{i,\ell+1}^*$.
- 18 **for** $i \in \{0, \dots, \ell\}$ **do**
- 19 $(s_{\ell+1,0}, \dots, s_{\ell+1,\ell+1}) =$
 $(s_{\ell+1,0}, \dots, s_{\ell+1,\ell+1}) - \lfloor k_i \rfloor \cdot (s_{i,0}, \dots, s_{i,\ell+1})$.
- 20 **return** $S_\ell = (s_{i,j})$, for $0 \leq i, j \leq \ell + 1$.

Algorithm 4: The LATTE Extract algorithm (from level $\ell - 1$ to user at level ℓ).

Input: $N, q, \sigma_\ell, S_{\ell-1}, H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^N, G : \{0, 1\}^* \rightarrow \{0, 1\}^{256}, \text{ID}_\ell$.
Output: $\mathbf{t}_0, \dots, \mathbf{t}_\ell \in \mathcal{R}_q$.

- 1 $A'_\ell \leftarrow H("E" | \text{ID}_1 | \dots | \text{ID}_\ell)$ in NTT domain.
- 2 $seed \leftarrow G(\text{ID}_1 | \dots | \text{ID}_\ell)$.
- 3 $T_{\ell-1} \leftarrow \text{Tree}(S_{\ell-1}, \sigma_\ell)$.
- 4 $\mathbf{t} \leftarrow (A'_\ell, \mathbf{0}, \dots, \mathbf{0}) \cdot S_{\ell-1}^{-1}$.
- 5 $\mathbf{z} \leftarrow \text{FFT}^{-1}(\text{ffSampling}(\mathbf{t}, T_{\ell-1}, seed))$.
- 6 $(\mathbf{t}_0, \mathbf{t}_1, \dots, \mathbf{t}_\ell) \leftarrow \lfloor \bar{\mathbf{z}} \rfloor$, where $\bar{\mathbf{z}} \leftarrow (\mathbf{t} - \mathbf{z}) \cdot S_{\ell-1}$.
- 7 **return** $\mathbf{t}_1, \dots, \mathbf{t}_\ell \in \mathcal{R}_q$ in NTT domain.

B.1 Sub-Algorithms from FALCON

This subsection presents sub-algorithms (Alg. 7–10) from FALCON [51]. Readers may refer to the FALCON specification [51] for sub-routines (ffLDL, splitfft, mergefft, etc.) used by these algorithms.

Algorithm 5: The LATTE Encrypt algorithm (at level ℓ).

Input: $N, q, \sigma_e, \mathbf{h}, \text{KDF}, \text{ID}_\ell, \mu \in \{0, 1\}^{256}$.
Output: $Z \in \{0, 1\}^{256}, C_1, \dots, C_\ell, C_h \in \mathcal{R}_q$.

- 1 $seed \leftarrow \{0, 1\}^{256}$.
- 2 $Z \leftarrow \mu \oplus \text{KDF}(seed)$.
- 3 Sample $\mathbf{e}, \mathbf{e}_1, \dots, \mathbf{e}_\ell, \mathbf{e}_h$ from a binomial distribution with center 0 and standard deviation σ_e using the seed $\text{KDF}(seed|Z)$.
- 4 **for** $i \in \{1, \dots, \ell - 1\}$ **do**
- 5 $C_i \leftarrow A_i \cdot \mathbf{e} + \mathbf{e}_i$ where $A_i = H(\text{ID}_1 | \dots | \text{ID}_i)$ in NTT domain.
- 6 $\mathbf{m} \leftarrow \text{Encode}(seed)$.
- 7 $C_\ell \leftarrow A'_\ell \cdot \mathbf{e} + \mathbf{e}_\ell + \mathbf{m}$ where $A'_\ell = H("E" | \text{ID}_1 | \dots | \text{ID}_\ell)$ in NTT domain.
- 8 $C_h \leftarrow \mathbf{h} \cdot \mathbf{e} + \mathbf{e}_h$.
- 9 **return** $Z \in \{0, 1\}^{256}, C_1, \dots, C_\ell, C_h \in \mathcal{R}_q$ in NTT domain.

Algorithm 6: The LATTE Decrypt algorithm (at level ℓ).

Input: $N, q, \sigma_e, \mathbf{h}, \text{KDF}, \text{ID}_\ell, Z, (C_1, \dots, C_\ell, C_h)$,
 $(\mathbf{t}_0, \dots, \mathbf{t}_\ell)$.
Output: μ' .

- 1 $\mathbf{V} \leftarrow C_\ell - C_h \cdot \mathbf{t}_1 - C_1 \cdot \mathbf{t}_2 - \dots - C_{\ell-1} \cdot \mathbf{t}_\ell$.
- 2 $seed' \leftarrow \text{Decode}(\mathbf{V})$.
- 3 Sample $\mathbf{e}', \mathbf{e}'_1, \dots, \mathbf{e}'_\ell, \mathbf{e}'_h$ from a binomial distribution with center 0 and standard deviation σ_e using the seed $\text{KDF}(seed'|Z)$.
- 4 **for** $i \in \{1, \dots, \ell - 1\}$ **do**
- 5 $C'_i \leftarrow A_i \cdot \mathbf{e}' + \mathbf{e}'_i$ where $A_i = H(\text{ID}_1 | \dots | \text{ID}_i)$ in NTT domain.
- 6 $\mathbf{m}' \leftarrow \text{Encode}(seed')$.
- 7 $C'_\ell \leftarrow A'_\ell \cdot \mathbf{e}' + \mathbf{e}'_\ell + \mathbf{m}'$ where $A'_\ell = H("E" | \text{ID}_1 | \dots | \text{ID}_\ell)$ in NTT domain.
- 8 $C'_h \leftarrow \mathbf{h} \cdot \mathbf{e}' + \mathbf{e}'_h$.
- 9 Check $(C'_1, \dots, C'_\ell, C'_h)$ agrees with $(C_1, \dots, C_\ell, C_h)$, else return \perp .
- 10 **return** $\mu' = Z \oplus \text{KDF}(seed')$.

C PROOF OF LEMMA 4.2

PROOF. We have that $\frac{D_3(z)}{D_2(z)} = \frac{\rho_{t,\sigma}(I)}{\rho_{i,\sigma}(I)} \cdot \exp(u(z))$, where, following assumption (3) in the statement of Lemma and the notations of the proof of [49, Lemma 7], we note Eq. (10).

$$u(z) = \frac{(z-t)^2}{2\sigma^2} - \frac{(z-\bar{t})^2}{2\bar{\sigma}^2} \quad (10)$$

$$= -\frac{(t-\bar{t})^2 + 2(t-\bar{t})(z-t) - (2\delta_\sigma + \delta_\sigma^2)(z-t)^2}{2(1-\delta_\sigma)^2\sigma^2}. \quad (11)$$

Algorithm 7: NTRUSolve $_{N,q}$ ([51]).

Input: $f, g \in \mathbb{Z}[x] / \langle x^N + 1 \rangle$.

Output: $F, G \in \mathbb{Z}[x] / \langle x^N + 1 \rangle$ such that $fG - gF = q \pmod{x^N + 1}$.

- 1 **if** $N = 1$ **then**
- 2 Compute $u, v \in \mathbb{Z}$ such that $uf - vg = \gcd(f, g)$.
- 3 **if** $\gcd(f, g) \neq 1$ **then abort**;
- 4 $(F, G) \leftarrow (vq, uq)$.
- 5 **return** (F, G) .
- 6 **else**
- 7 $f' \leftarrow N(f)$.
- 8 $g' \leftarrow N(g)$.
- 9 $(F', G') \leftarrow \text{NTRUSolve}_{N/2,q}(f', g')$.
- 10 $F \leftarrow F'(x^2) \cdot f'(x^2) / f(x)$.
- 11 $G \leftarrow G'(x^2) \cdot g'(x^2) / g(x)$.
- 12 $\mathbf{k} \leftarrow \begin{bmatrix} F \cdot f^* + G \cdot g^* \\ f \cdot f^* + g \cdot g^* \end{bmatrix} \in \mathcal{R}$.
- 13 $F \leftarrow F - \mathbf{k} \cdot f$ and $G \leftarrow G - \mathbf{k} \cdot g$.
- 14 **return** (F, G) .

Algorithm 8: The original fLDL algorithm from [23, 51].

Input: Gram matrix $G \in (\mathbb{C}[x] / \langle x^n + 1 \rangle)^{d \times d}$ in the FFT domain.

Output: Tree T .

- 1 **if** $n = 1$ **then**
- 2 $T.\text{value} \leftarrow G_{0,0}$.
- 3 **else**
- 4 $L \leftarrow I_d, D \leftarrow 0_d$.
- 5 **for** $i = 0$ **to** $d - 1$ **do**
- 6 **for** $j = 0$ **to** $i - 1$ **do**
- 7 $L_{i,j} \leftarrow \frac{1}{D_{j,j}} \left(G_{i,j} - \sum_{k < j} L_{i,k} \odot L_{j,k}^* \odot D_{k,k} \right)$.
- 8 $D_{i,i} \leftarrow G_{i,i} - \sum_{j < i} L_{i,j} \odot L_{i,j}^* \odot D_{j,j}$.
- 9 $T.\text{value} \leftarrow L$.
- 10 **for** $i = 0$ **to** $d - 1$ **do**
- 11 $\mathbf{d}_0, \mathbf{d}_1 \leftarrow \text{splitfft}(D_{i,i})$.
- 12 $G' = \begin{pmatrix} \mathbf{d}_0 & \mathbf{d}_1 \\ \mathbf{d}_1^* & \mathbf{d}_0 \end{pmatrix}$.
- 13 $T.\text{child}_i \leftarrow \text{fLDL}(G')$.
- 14 **return** T .

We first bound $\frac{\rho_{t,\sigma}(I)}{\rho_{t,\sigma}(I)}$. Let $\text{erfc}(\alpha) := \frac{1}{\sqrt{2\pi}} \int_{\alpha}^{\infty} \exp(-y^2/2) dy$ be the complementary error function; then we get:

$$\begin{aligned} \rho_{t,\sigma}(I) &= \rho_{t,\sigma}(\mathbb{Z}) - 2 \cdot \text{erfc}(\tau \cdot \sigma) \\ &= \rho_{t,\sigma}(\mathbb{Z}) \cdot \left(1 - \frac{2 \cdot \text{erfc}(\tau \cdot \sigma)}{\rho_{t,\sigma}(\mathbb{Z})} \right) \geq \rho_{t,\sigma}(\mathbb{Z}) \cdot (1 - 1/Q), \end{aligned}$$

where the last inequality uses $\rho_{t,\sigma}(\mathbb{Z}) \sim \sqrt{2\pi}(1-\varepsilon)\sigma$, $\text{erfc}(\tau\sigma) \leq \frac{\exp(-(\tau\sigma)^2/2)}{\sqrt{2\pi}\tau\sigma}$, and $Q \leq \exp(\tau^2/2)/2 \leq \exp(\sigma^2\tau^2/2)/(2\tau)$. Deriving

Algorithm 9: The ffSampling Tree computation algorithm ([51]).

Input: S_{ℓ}, σ_{ℓ} .

Output: Tree T_{ℓ} .

- 1 $G_{\ell} \leftarrow S_{\ell} \cdot S_{\ell}^*$.
- 2 $T \leftarrow \text{fLDL}(\text{FFT}(G_{\ell}))$.
- 3 **For** each leaf of T_{ℓ} , $\text{leaf.value} \leftarrow \sigma_{\ell} / \sqrt{\text{leaf.value}}$.
- 4 **return** T_{ℓ} .

Algorithm 10: The ffSampling algorithm ([51]).

Input: $\mathbf{t} = (t_0, t_1, \dots, t_{\ell})$ in FFT format, tree T , $\text{seed} \in \{0, 1\}^{256}$.

Output: $\mathbf{z} = (z_0, z_1, \dots, z_{\ell})$ in FFT format.

- 1 **if** $n = 1$ **then**
- 2 $\sigma' \leftarrow T.\text{value}$.
- 3 Sample $z_0 \leftarrow \mathcal{D}_{\sigma', t_0}$ using seed .
- 4 Sample $z_1 \leftarrow \mathcal{D}_{\sigma', t_1}$ using seed .
- 5 **return** $\mathbf{z} = (z_0, z_1)$.
- 6 **else**
- 7 $m \leftarrow$ number of children of T .
- 8 **for** $j \leftarrow m, \dots, 0$ **do**
- 9 $T_j \leftarrow j$ -th child of T .
- 10 $\mathbf{t}'_j \leftarrow \mathbf{t}_j + \sum_{i=j+1}^m (t_i - z_i) \cdot T.\text{value}_{i,j}$.
- 11 $\mathbf{t}'_j \leftarrow \text{splitfft}(\mathbf{t}'_j)$.
- 12 $\mathbf{z}'_j \leftarrow \text{ffSampling}(\mathbf{t}'_j, T_j)$.
- 13 $\mathbf{z}_j \leftarrow \text{mergefft}(\mathbf{z}'_j)$.
- 14 **return** $\mathbf{z} = (z_0, z_1, \dots, z_m)$.

a similar inequality for $\rho_{\bar{t}, \bar{\sigma}}(I)$, we have:

$$1 - 2/Q \lesssim \frac{\rho_{t,\sigma}(I)}{\rho_{\bar{t}, \bar{\sigma}}(I)} \left/ \frac{\rho_{t,\sigma}(\mathbb{Z})}{\rho_{\bar{t}, \bar{\sigma}}(\mathbb{Z})} \right. \lesssim 1 + 2/Q. \quad (12)$$

Let $n := ((\Delta'_t)^2 + 2\Delta'_t(z-t)/\sigma - (2\delta_{\sigma} + \delta_{\sigma}^2))((z-t)/\sigma)^2$. By applying [49, Lemma 7], followed by Eq. (11), followed by [41, Lemma 4.4] and finally using $\max(\delta_{\sigma}, \varepsilon) \leq \delta$, it follows that:

$$\ln \left(\frac{\rho_{t,\sigma}(\mathbb{Z})}{\rho_{\bar{t}, \bar{\sigma}}(\mathbb{Z})} \right) \leq \left| \mathbb{E}_{z \leftarrow D_1} [u] \right| \leq \left| \mathbb{E}_{z \leftarrow D_1} \left[\frac{n}{2(1 - \delta_{\sigma})^2} \right] \right| \quad (13)$$

$$\leq \frac{1}{2(1 - \delta_{\sigma})^2} \left| \mathbb{E}_{z \leftarrow D_1} [n] \right| \quad (14)$$

$$\leq \text{num}(\Delta'_t, \varepsilon, \delta_{\sigma}) \quad (15)$$

Similarly, we bound $\exp(u(z))$ over I by using Eq. (11):

$$\max_I |u| \leq \frac{1}{1 - \delta_{\sigma}} \cdot (\tau \Delta'_t + \tau^2 \delta_{\sigma}). \quad (16)$$

Combining Eq. (12), Eq. (15), and Eq. (16), we bound the relative error:

$$\begin{aligned} \left| \ln \left(\frac{D_3}{D_2} \right) \right| &\leq \ln(1 + 2/Q) + \text{num}(\Delta'_t, \varepsilon, \delta_{\sigma}) + \frac{1}{1 - \delta_{\sigma}} \cdot (\tau \Delta'_t + \tau^2 \delta_{\sigma}) \\ &\leq 2/Q + \text{num}(\Delta'_t, \varepsilon, \delta_{\sigma}) + \frac{1}{1 - \delta_{\sigma}} \cdot (\tau \Delta'_t + \tau^2 \delta_{\sigma}) = \text{ub}. \end{aligned} \quad (17)$$

Combining Eq. (17) and [49, Lemma 3], the Rényi divergence between D_3 and D_2 is derived as Eq. (5). Finally, we combine the first weak triangle inequality of [34, Lemma 4.1] with remark 1 to obtain the Rényi divergence between D_3 and D_1 as in Eq. (6). \square

D PROOF OF LEMMA 4.1

PROOF. Consider the ffSampling Alg. 10 in Appendix B.1. We employ a sequence of games G_0, \dots, G_5 , and track the probability of the events E and B_U over those games using a Rényi divergence approach. Let E_i and $B_{U,i}$ denote the events E and B_U in game i for $i = 0, \dots, 5$. The games REAL and IDEAL are defined in the Lemma statement. The sequence of games is as follows:

- G_0 : Game REAL.
- G_1 : G_0 , but we change the 1-dimensional \mathbb{Z} -sampler from the finite precision sampler distribution $\tilde{\mathcal{D}}$ to infinite precision sampler distribution \mathcal{D} .
- G_2 : G_1 , but we abort the game if B_U happens, meaning either there exists i such that the errors $\Delta'_{t^{(i)}}$ (relative to $\sigma^{(i)}$) in centers $t^{(i)}$ exceed $\Delta'_{t^{(i)}}$ or relative errors $\delta_{\sigma^{(i)}}$ in standard deviations $\sigma^{(i)}$ exceed $\delta_{\sigma^{(i)}}$ or there exists j such that the infinity-norm absolute errors $\Delta_{\bar{z}^{(j)}}$ in $\bar{z}^{(j)}$ exceed $\Delta_{\bar{z}^{(j)}}$.
- G_3 : G_2 , but we restrict the 1-dimensional \mathbb{Z} samplers \mathcal{D} to the corresponding τ -bounded distribution \mathcal{D}^τ .
- G_4 : G_3 , but changing arithmetic from finite precision to infinite precision, and removing the τ -tailcut on the 1-dimensional \mathbb{Z} samplers to return to the ideal Gaussian distribution \mathcal{D} . This game is identical to $Game_{IDEAL}$, except for the abort condition introduced in the previous game.
- G_5 : G_4 , but remove the abort introduced in G_2 . This game is identical to $Game_{IDEAL}$.

$G_0 \rightarrow G_1$. Changing the 1-D \mathbb{Z} -sampler. Let $(\bar{\sigma}^{(i)}, \bar{t}^{(i)}) = (\sigma^{(i)}(1 + \delta_{\sigma^{(i)}}, t^{(i)} + \Delta_{t^{(i)}})$ denote the i 'th query to the 1-D sampler in the execution of these games, and denote by $\zeta^{(i)}$ the output integer returned by the sampler for the i 'th query. We apply Proposition 4, with x_0 denoting the remaining source of randomness in the game (i.e. the random coins of \mathcal{A} and the hash function H), and we let $x_i := \zeta^{(i)}$ for $i = 1, \dots, M_{\mathbb{Z}}$. Consider $(x_i | x_{j < i})$, the conditional distribution of x_i , conditioned on all previous x_j , for $j < i$, and the Rényi divergence between this distribution in G_0 and G_1 . Observe that conditioned on the same value of $x_{j < i}$, the values of the following query $\bar{t}^{(i)}$ and std. deviation $\bar{\sigma}^{(i)}$ are identical in both G_0 and G_1 since they both use the same finite precision arithmetic. We therefore have:

$$R_a\left((x_i | x_{j < i})_{G_0}, (x_i | x_{j < i})_{G_1}\right) = R_a\left(\tilde{\mathcal{D}}_{\bar{\sigma}^{(i)}, \bar{t}^{(i)}}, \mathcal{D}_{\bar{\sigma}^{(i)}, \bar{t}^{(i)}}\right) \leq B,$$

then Proposition 2.2 implies:

$$R_a\left((x_0, \dots, x_K)_{G_0}, (x_0, \dots, x_K)_{G_1}\right) \leq B^{M_{\mathbb{Z}}} := B_T.$$

By the data processing and probability preservation properties of Rényi divergence, $\Pr[E_1] \geq \Pr[E_0]^{a/(a-1)}/B_T$.

$G_1 \rightarrow G_2$. Adding a τ tailcut to the \mathbb{Z} Gaussian samplers. By a standard tail bound [38, Lemma 4.4], the statistical distance between this game and the previous one is $\leq M_{\mathbb{Z}} \cdot 2 \exp(-\tau^2/2) \leq 1/Q$. Hence, we have $\Pr[E_2] \geq \Pr[E_1] - 1/Q$.

$G_2 \rightarrow G_3$. Aborting the game if the errors exceed the bounds. Recall that $B_{U,2}$ denotes the event B_U in G_2 that the errors exceed the bounds in the Lemma statement. If the event $B_{U,2}$ does not occur, games G_2 and G_3 proceed identically. Hence, we have $\Pr[E_3] \geq \Pr[E_2] - \Pr[B_{U,2}]$ and $\Pr[B_{U,2}] = \Pr[B_{U,3}]$.

$G_3 \rightarrow G_4$. Changing finite precision arithmetic to infinite precision and removing the τ -tailcut on the Gaussians. We again apply Proposition 2.2, except that this time $x_i := \zeta^{(i)}$ for $i = 1, \dots, M_{\mathbb{Z}}$ except if the event B_U occurs at the i 'th query to the \mathbb{Z} sampler (determined by x_0, \dots, x_{i-1}), in which case $x_i := \perp$, and all subsequent $x_j := \perp$ for $j > i$. As in the previous game, we consider the conditional distribution $(x_i | x_{j < i})$, of x_i conditioned on all previous x_j for $j < i$, and the Rényi divergence between this conditional distribution in G_3 and G_4 . When the event B_U occurs at the at (or before) the i 'th query to the \mathbb{Z} sampler, the conditional distribution $(x_i | x_{j < i})$ is identical in both games (as both conditional distributions return \perp with probability 1) and have Rényi divergence 0. Whereas, if the event B_U does not occur at (or before) the i 'th query to the \mathbb{Z} conditioned on the same fixed value of $x_{j < i}$ in the support of the j 'th 1-D \mathbb{Z} -samplers, we have $\Delta'_{t^{(i)}} \leq \Delta'_{t^{(i)}}$. Also, the query std deviation values $\sigma^{(i)}$ in G_4 and $\bar{\sigma}^{(i)}$ in G_3 have a relative error $\delta_{\sigma^{(i)}} \leq \delta_{\sigma^{(i)}}$ by definition of event B_U . We therefore have:

$$R_a\left((x_i | x_{j < i})_{G_3}, (x_i | x_{j < i})_{G_4}\right) \quad (18)$$

$$\leq R_a\left(\mathcal{D}_{\sigma^{(i)} \cdot (1 \pm \delta_{\sigma^{(i)}}), t^{(i)} + \Delta'_{t^{(i)}} \cdot \sigma^{(i)}}, \mathcal{D}_{\bar{\sigma}^{(i)}, t^{(i)}}\right) \leq C^{(i)}, \quad (19)$$

where in the last inequality we used Lemma 4.2. Then Proposition 2.2 above implies:

$$R_a\left((x_0, \dots, x_{M_{\mathbb{Z}}})_{G_2}, (x_0, \dots, x_{M_{\mathbb{Z}}})_{G_3}\right) \leq \prod_{i < M_{\mathbb{Z}}} C^{(i)} := C_T.$$

Due to the abort condition, we have that conditioned on the same fixed value of x_i 's that do not cause an abort, the values $\bar{z}^{(j)}$ in G_3 and G_4 differ by an absolute error at most $\Delta_{\bar{z}} < 1/2$, and therefore, observing that in G_4 the $\bar{z}^{(j)}$ has integer coordinates (due to the infinite precision), the rounded $\bar{z}^{(j)}$ values in G_4 are identical to those in G_3 conditioned on the same x_i 's. Since the adversary's view in the game depends on x_i 's only via the rounded $\bar{z}^{(j)}$, we conclude by the Rényi probability preservation property that

$$\Pr[E_4] \geq \Pr[E_3]^{a/(a-1)}/C_T,$$

and

$$\Pr[B_{U,4}] \geq \Pr[B_{U,3}]^{a/(a-1)}/C_T.$$

$G_4 \rightarrow G_5$. In this game, we remove the abort introduced in G_2 . Since the games G_4 and G_5 proceed identically until an abort occurs, we have $\Pr[B_{U,5}] = \Pr[B_{U,4}]$ and $\Pr[E_5] \geq \Pr[E_4] - \Pr[B_{U,4}]$. Furthermore, by the Lemma hypothesis, we have $\Pr[B_{U,5}] := p_U$.

Putting together the above bounds, we obtain that the probability of E_5 (i.e. event E in IDEAL) is lower bounded by

$$\Pr[E_5] \geq \Pr[E_3]^{a/(a-1)} / C_T - p_U \\ \geq (\Pr[E_0]^{a/(a-1)} / B_T - (\Pr[B_{U,3}] + 1/Q))^{a/(a-1)} / C_T - p_U$$

and using $p_U = \Pr[B_{U,5}] = \Pr[B_{U,4}] \geq \Pr[B_{U,3}]^{a/(a-1)} / C_T$, we get the claimed bound on $\Pr[E_4] := \Pr[E_{IDEAL}]$. \square

E SUB-ALGORITHMS IN THE STATISTICAL MODEL

We present the sub-algorithms (Alg. 11–16) of our statistical model in Sec. 4.3, including fLDLBounds, ffSamplingBounds, splitfftBounds, mergefftBounds, FFTBounds, and FFTInvBounds.

Algorithm 11: The fLDLBounds algorithm based on statistical model.

Input: G, D' .

Output: Tree T .

```

1 if  $n = 1$  then
2    $(T.value)_0 \leftarrow (\mu_{G_{0,0},R}, \sigma_{G_{0,0},R}^2, 0, 0)$ .
3   Output  $\mu_{leaf,R} = \mu(T.value)_{0,R}$ ,  $\sigma_{leaf,R} = \sigma(T.value)_{0,R}$ .
4   return
5 for  $j \in \{0, \dots, n-1\}$  do
6    $(D_{0,0})_j \leftarrow (\mu_{(G_{0,0})_j,R}, \sigma_{(G_{0,0})_j,R}^2, 0, 0)$ .
7    $(L_{1,0})_j \leftarrow \text{DivBound}((G_{1,0})_j, (D_{0,0})_j)$ .
8   if  $d = 2$  then
9     if  $n = N$  then
10       $(D_{1,1})_j \leftarrow \text{DivBound}(q^2, (D_{0,0})_j)$ .
11     else
12       $x \leftarrow \text{MultBound}(D'_{2j}, D'_{2j+1})$ .
13       $(D_{1,1})_j \leftarrow \text{DivBound}(x, (D_{0,0})_j)$ .
14     else if  $d = 3$  then
15       $x \leftarrow \text{DivBound}(\text{AbsSqrBound}((G_{1,0})_j), (D_{0,0})_j)$ .
16       $(D_{1,1})_j \leftarrow \text{SubBound}((G_{1,1})_j, x)$ .
17       $x \leftarrow \text{MultBound}((D_{0,0})_j, (D_{1,1})_j)$ .
18       $(D_{2,2})_j \leftarrow \text{DivBound}(q^2, x)$ .
19       $(L_{2,0})_j \leftarrow \text{DivBound}((G_{2,0})_j, (D_{0,0})_j)$ .
20       $x \leftarrow \text{MultBound}((G_{2,0})_j, (L_{1,0})_j^*)$ .
21       $y \leftarrow \text{SubBound}((G_{2,1})_j, x)$ .
22       $(L_{2,1})_j \leftarrow \text{DivBound}(y, (D_{1,1})_j)$ .
23  $T.value \leftarrow L$ .
24 for  $i \in \{0, \dots, d-1\}$  do
25    $d_0, d_1 \leftarrow \text{splitfftBounds}(D_{i,i}, n)$ .
26    $G' = \begin{pmatrix} d_0 & d_1 \\ d_1^* & d_0 \end{pmatrix}$ .
27    $T.child_i \leftarrow \text{fLDL}(G', D_{i,i})$ .
28 return  $T$ .
```

Algorithm 12: The ffSamplingBounds algorithm based on statistical model.

Input: $t = (t_0, \dots, t_\ell)$ in FFT format, tree T .

Output: $z = (z_0, \dots, z_\ell)$ in FFT format.

```

1 if  $n = 1$  then
2    $z_0 \leftarrow (\mu_{(t_0)_{0,R}}, 0, 0, 0)$ .
3    $z_1 \leftarrow (\mu_{(t_1)_{0,R}}, 0, 0, 0)$ .
4   Output  $\Delta'_{t_0} = \sigma_{(t_0)_{0,R}} / (\sigma_\ell / \sqrt{\mu(T.value)_{0,R}})$ .
5   Output  $\Delta'_{t_1} = \sigma_{(t_1)_{0,R}} / (\sigma_\ell / \sqrt{\mu(T.value)_{0,R}})$ .
6   return  $z = (z_0, z_1)$ .
7 else
8    $m \leftarrow$  number of children of  $T$ .
9   for  $j \leftarrow m-1, \dots, 0$  do
10     $T_j \leftarrow$   $j$ -th child of  $T$ .
11     $t'_j \leftarrow t_j$ .
12    for  $i \leftarrow j+1, \dots, m-1$  do
13      for  $k \leftarrow 0, \dots, n-1$  do
14         $x \leftarrow \text{SubBound}((t_i)_k, (z_i)_k)$ .
15         $y \leftarrow \text{MultBound}(x, (T.value)_{i,j})_k$ .
16         $(t'_j)_k \leftarrow \text{AddBound}((t'_j)_k, y)$ .
17     $f_0, f_1 \leftarrow \text{splitfftBounds}(t'_j, n)$ .
18     $z'_{j,0}, z'_{j,1} \leftarrow \text{ffSamplingBounds}((f_0, f_1), T_j)$ .
19     $z_j \leftarrow \text{mergefftBounds}(z'_{j,0}, z'_{j,1}, n)$ .
20 return  $z$ .
```

Algorithm 13: The splitfftBounds algorithm based on statistical model.

Input: a, n .

Output: f_0, f_1 .

```

1  $(f_0)_0 \leftarrow (\mu_{a_0,R}, \sigma_{a_0,R}^2, 0, 0)$ .
2  $(f_1)_0 \leftarrow (\mu_{a_0,I}, \sigma_{a_0,I}^2, 0, 0)$ .
3 for  $k \leftarrow 0, \dots, n/2-1$  do
4    $(f_0)_k \leftarrow \text{MultBound}(1/2, \text{AddBound}(a_{2k}, a_{2k+1}))$ .
5    $x \leftarrow \text{SubBound}(a_{2k}, a_{2k+1})$ .
6    $y \leftarrow \text{MultBound}(x, \omega^{-\text{bitrev}(n/2+k)})$ .
7    $(f_1)_k \leftarrow \text{MultBound}(1/2, y)$ .
8 return  $f_0, f_1$ .
```

F PROOF OF THEOREM 5.1

PROOF. (1) From Alg. 8 in Appendix B.1, we have $(D_{0,0})_j = (G_{0,0})_j$ and $(D_{1,1})_j = (G_{1,1})_j - |(L_{1,0})_j|^2 (G_{0,0})_j$, $0 \leq j \leq n-1$, for some input matrix G in the FFT domain in every node of the tree. In addition, we have $(D_{i,i})_j = (G_{i,i})_j - \sum_{k < i} (|(L_{i,k})_j|^2 (D_{k,k})_j)$ at the root when $d > 2$. Therefore, we have $D_{i,i} \in \mathbb{R}^n$ assuming that $G_{i,i} \in \mathbb{R}^n$ for all $i \in \{0, \dots, d-1\}$. To show that latter assumption is true, we observe that at the root we have the input $G = S_\ell S_\ell^*$ in the FFT domain, $G_{i,i} \in \mathbb{R}^N$ for $0 \leq i \leq d-1$. Thus, $D_{i,i} \in \mathbb{R}^N$ for $0 \leq i \leq d-1$ at the root. Assuming $D_{i,i} \in \mathbb{R}^n$ for $0 \leq i \leq d-1$ at a non-leaf node, for its i -th child, we have the

Algorithm 14: The mergeffBounds algorithm based on statistical model.

Input: $\mathbf{f}_0, \mathbf{f}_1, n$.
Output: \mathbf{a} .

- 1 $\mathbf{a}_0 \leftarrow (\mu_{(\mathbf{f}_0)_0, R}, \sigma_{(\mathbf{f}_0)_0, R}^2, \mu_{(\mathbf{f}_1)_0, R}, \sigma_{(\mathbf{f}_1)_0, R}^2)$.
- 2 **for** $k \leftarrow 0, \dots, n/2 - 1$ **do**
- 3 $u \leftarrow \text{MultBound}((\mathbf{f}_1)_k, \omega^{\text{bitrev}(n/2+k)})$.
- 4 $\mathbf{a}_{2k} \leftarrow \text{AddBound}((\mathbf{f}_0)_k, u)$.
- 5 $\mathbf{a}_{2k+1} \leftarrow \text{SubBound}((\mathbf{f}_0)_k, u)$.
- 6 **return** \mathbf{a} .

Algorithm 15: The FFTBounds algorithm based on statistical model.

Input: \mathbf{a} .

- 1 $m = 1$.
- 2 $t = n$.
- 3 **while** $m < n$ **do**
- 4 $t \leftarrow t/2$.
- 5 **for** $i = 0$ **to** $m - 1$ **do**
- 6 $j_1 = 2it$.
- 7 $j_2 = j_1 + t - 1$.
- 8 **for** $j = j_1$ **to** j_2 **do**
- 9 $u \leftarrow \mathbf{a}_j$.
- 10 $v \leftarrow \text{MultBound}(\mathbf{a}_{j+t}, \omega^{\text{bitrev}(m+i)})$.
- 11 $\mathbf{a}_j \leftarrow \text{AddBound}(u, v)$.
- 12 $\mathbf{a}_{j+t} \leftarrow \text{SubBound}(u, v)$.
- 13 $m \leftarrow 2m$.

ffLDL input $\mathbf{G}'_{0,0} = \mathbf{G}'_{1,1} = \mathbf{d}_0$, where $(\mathbf{d}_0)_j = \frac{1}{2}[(\mathbf{D}_{i,i})_{2j} + (\mathbf{D}_{i,i})_{2j+1}] \in \mathbb{R}$ for $j \in \{0, \dots, n/2 - 1\}$. Thus, $\mathbf{D}'_{0,0}, \mathbf{D}'_{1,1} \in \mathbb{R}^{n/2}$ in this child and we can deduce the conclusion by induction.

- (2) Since by definition of the LDL* decomposition [23], \mathbf{L} is a lower triangular matrix with 1 on its diagonal and \mathbf{D} is a diagonal matrix, we have $\det(\mathbf{D}) = \prod_{i=0}^{d-1} \mathbf{D}_{i,i} = \det(\mathbf{G})$. Because $\mathbf{G} = \mathbf{S}_\ell \mathbf{S}_\ell^*$ at the root and the determinant of a (Mod)NTRU basis \mathbf{S}_ℓ is q , we have $\prod_{i=0}^{d-1} (\mathbf{D}_{i,i})_j = q^2$ in the FFT domain at the root for $0 \leq j \leq N - 1$.
- (3) For the i -th child of a non-leaf node, we have the ffLDL input $\mathbf{G}' = \begin{pmatrix} \mathbf{d}_0 & \mathbf{d}_1 \\ \mathbf{d}_1^* & \mathbf{d}_0 \end{pmatrix}$ for $\mathbf{d}_0, \mathbf{d}_1 \leftarrow \text{splitfft}(\mathbf{D}_{i,i}), 0 \leq i \leq d - 1$ (see [51] for the definition of splitfft). By the definition of the LDL* decomposition, for this child, we have $\mathbf{D}'_{0,0} \mathbf{D}'_{1,1} = \det(\mathbf{G}') = \mathbf{d}_0^2 - \mathbf{d}_1 \mathbf{d}_1^*$. Thus, in the FFT domain, we have:

$$(\mathbf{D}'_{0,0})_j (\mathbf{D}'_{1,1})_j = (\mathbf{d}_0)_j^2 - |(\mathbf{d}_1)_j|^2 = \left(\frac{1}{2} [(\mathbf{D}_{i,i})_{2j} + (\mathbf{D}_{i,i})_{2j+1}] \right)^2 - \left| \frac{1}{2} [(\mathbf{D}_{i,i})_{2j} - (\mathbf{D}_{i,i})_{2j+1}] \omega^{-\text{bitrev}(n/2+j)} \right|^2,$$

for $0 \leq j \leq n/2 - 1$. Since $(\mathbf{D}_{i,i})_{2j}, (\mathbf{D}_{i,i})_{2j+1} \in \mathbb{R}$ and $|\omega| = 1$, we get $(\mathbf{D}'_{0,0})_j (\mathbf{D}'_{1,1})_j = (\mathbf{D}_{i,i})_{2j} (\mathbf{D}_{i,i})_{2j+1}$.

Algorithm 16: The FFTInvBounds algorithm based on statistical model.

Input: \mathbf{a} .

- 1 $t = 1$.
- 2 $h = n$.
- 3 $m = n$.
- 4 **while** $m > 1$ **do**
- 5 $j_1 = 0$.
- 6 $h \leftarrow h/2$.
- 7 **for** $i = 0$ **to** $h - 1$ **do**
- 8 $j_2 = j_1 + t - 1$.
- 9 **for** $j = j_1$ **to** j_2 **do**
- 10 $u \leftarrow \text{AddBound}(\mathbf{a}_j, \mathbf{a}_{j+t})$.
- 11 $x \leftarrow \text{SubBound}(\mathbf{a}_j, \mathbf{a}_{j+t})$.
- 12 $\mathbf{a}_j \leftarrow u$.
- 13 $\mathbf{a}_{j+t} \leftarrow \text{MultBound}(x, \omega^{-\text{bitrev}(h+i)})$.
- 14 $j_1 \leftarrow j_1 + 2t$.
- 15 $t \leftarrow 2t$.
- 16 **for** $i = 0$ **to** $n - 1$ **do**
- 17 $\mathbf{a}_i \leftarrow \text{MultBound}(1/n, \mathbf{a}_i)$.

- (4) The ffLDL algorithm computes the LDL* decomposition in the FFT domain. Let $\mathbf{S}_\ell = \mathbf{L} \cdot \tilde{\mathbf{S}}_\ell$ be the GSO decomposition of $\mathbf{S}_\ell \in \mathcal{R}^{d \times d}$ where rows of $\tilde{\mathbf{S}}_\ell$ are pairwise orthogonal. For the input $\mathbf{G} = \mathbf{S}_\ell \mathbf{S}_\ell^*$ at the root, we have $\mathbf{G} = \mathbf{LDL}^*$ where $\mathbf{D} = \tilde{\mathbf{S}}_\ell \tilde{\mathbf{S}}_\ell^*$ [23]. Thus, in the FFT domain, $\mathbf{D}_{i,i} \in (\mathbb{R}^+)^N$ at the root. Assuming $\mathbf{D}_{i,i} \in (\mathbb{R}^+)^n$ for some $i \in \{0, \dots, d - 1\}$ at an non-leaf node, for the i -th child of this node, we have $(\mathbf{D}'_{0,0})_j (\mathbf{D}'_{1,1})_j = (\mathbf{D}_{i,i})_{2j} (\mathbf{D}_{i,i})_{2j+1} \in \mathbb{R}^+$ for $0 \leq j \leq n/2 - 1$. Because $(\mathbf{D}'_{0,0})_j = (\mathbf{d}_0)_j = \frac{1}{2} [(\mathbf{D}_{i,i})_{2j} + (\mathbf{D}_{i,i})_{2j+1}] \in \mathbb{R}^+$ due to the ffLDL input $\mathbf{G}' = \begin{pmatrix} \mathbf{d}_0 & \mathbf{d}_1 \\ \mathbf{d}_1^* & \mathbf{d}_0 \end{pmatrix}$ where $\mathbf{d}_0, \mathbf{d}_1 \leftarrow \text{splitfft}(\mathbf{D}_{i,i})$, we get $\mathbf{D}'_{0,0}, \mathbf{D}'_{1,1} \in (\mathbb{R}^+)^{n/2}$. Thus, we deduce the conclusion by induction. \square

G UPPER BOUNDS FOR THE STANDARD DEVIATION OF INTEGER GAUSSIAN IN FFSAMPLING

When $\ell = 1$, we have $\sigma \leq \sigma_{\min} \cdot \frac{\max_i \|(\tilde{\mathbf{S}}_0)_i\|}{\min_i \|(\tilde{\mathbf{S}}_0)_i\|}$, $0 \leq i \leq 2N - 1$, for σ of the integer Gaussian in ffSampling [32]. From Sec. 4.5, we have $\sigma_{\min} = \eta_\epsilon(\mathbb{Z})$ and $\max_i \|(\tilde{\mathbf{S}}_0)_i\| \leq \sigma_0 \sqrt{2N}$. By symplecticity of \mathbf{S}_0 [32], we have $\min_i \|(\tilde{\mathbf{S}}_0)_i\| \geq q/(\sigma_0 \sqrt{2N})$. Therefore, we get $\sigma \leq \eta_\epsilon(\mathbb{Z}) \cdot (\sigma_0 \sqrt{2N})^2 / q$. In order to analyse the upper bound of σ when $\ell = 2$, first we introduce the following Lemma:

LEMMA G.1. *For every non-root, non-leaf node in an ffLDL tree, we have:*

$$\min_{k=0}^{2n-1} \mathbf{D}'_k \leq (\mathbf{D}_{i,i})_j \leq \max_{k=0}^{2n-1} \mathbf{D}'_k,$$

$i \in \{0, 1\}, 0 \leq j \leq n - 1$, for some $\mathbf{D}' \in (\mathbb{R}^+)^{2n}$ from its parent.

PROOF. From Theorem 5.1, for a non-root, non-leaf node, since $(\mathbf{D}_{0,0})_j = \frac{1}{2}(\mathbf{D}'_{2j} + \mathbf{D}'_{2j+1})$, $0 \leq j \leq n-1$, for some $\mathbf{D}' \in (\mathbb{R}^+)^{2n}$, $(\mathbf{D}_{0,0})_j$ gets the minimal value $\min_{k=0}^{2n-1} \mathbf{D}'_k$ when both \mathbf{D}'_{2j} and \mathbf{D}'_{2j+1} are equal to $\min_{k=0}^{2n-1} \mathbf{D}'_k$. Similarly, $(\mathbf{D}_{0,0})_j$ gets the maximal value $\max_{k=0}^{2n-1} \mathbf{D}'_k$ when both \mathbf{D}'_{2j} and \mathbf{D}'_{2j+1} are equal to $\max_{k=0}^{2n-1} \mathbf{D}'_k$. For $(\mathbf{D}_{1,1})_j = \mathbf{D}'_{2j} \mathbf{D}'_{2j+1} / (\mathbf{D}_{0,0})_j = \frac{\mathbf{D}'_{2j} \mathbf{D}'_{2j+1}}{1/2 \cdot (\mathbf{D}'_{2j} + \mathbf{D}'_{2j+1})}$, it gets the minimal value $\min_{k=0}^{2n-1} \mathbf{D}'_k$ when both \mathbf{D}'_{2j} and \mathbf{D}'_{2j+1} are equal to $\min_{k=0}^{2n-1} \mathbf{D}'_k$ and $(\mathbf{D}_{1,1})_j$ gets the maximal value $\max_{k=0}^{2n-1} \mathbf{D}'_k$ when both \mathbf{D}'_{2j} and \mathbf{D}'_{2j+1} are equal to $\max_{k=0}^{2n-1} \mathbf{D}'_k$ for $\mathbf{D}' \in (\mathbb{R}^+)^{2n}$. \square

From Lemma G.1, if the ancestor of a non-root, non-leaf node is the m -th child of the root, $0 \leq m \leq d-1$, then $(\mathbf{D}_{i,i})_j$ of this node has the minimal value $\min_{k=0}^{N-1} (\mathbf{D}'_{m,m})_k$ and the maximal value $\max_{k=0}^{N-1} (\mathbf{D}'_{m,m})_k$, $i \in \{0, 1\}$, $0 \leq j \leq n-1$, for $\mathbf{D}'_{m,m}$ from the root, respectively. The leaf value of an fFLDL tree is $\sigma = \sigma_\ell / \sqrt{(\mathbf{G}_{0,0})_0}$, where $(\mathbf{G}_{0,0})_0 = \frac{1}{2}(\mathbf{D}'_0 + \mathbf{D}'_1)$ for some \mathbf{D}' from its parent. Following a similar approach in the proof of Lemma G.1, we have:

$$\min\{\mathbf{D}'_0, \mathbf{D}'_1\} \leq (\mathbf{G}_{0,0})_0 \leq \max\{\mathbf{D}'_0, \mathbf{D}'_1\}.$$

Therefore, similar to a non-root, non-leaf node, if the ancestor of a leaf node is the m -th child of the root, then the leaf value σ has the minimal value $\sigma_\ell / \sqrt{\max_{k=0}^{N-1} (\mathbf{D}'_{m,m})_k}$ and the maximal value $\sigma_\ell / \sqrt{\min_{k=0}^{N-1} (\mathbf{D}'_{m,m})_k}$.

In order to analyse the minimal and maximal values of $\mathbf{D}'_{m,m}$ from the root, we introduce the following Lemma:

LEMMA G.2. *For the Gram matrix $\mathbf{G} = \mathbf{S}_{\ell-1} \mathbf{S}_{\ell-1}^* \in (\mathbb{C}[x] / \langle x^N + 1 \rangle)^{(\ell+1) \times (\ell+1)}$ in the FFT domain, we have $|(\mathbf{G}_{i,i})_j| \leq \sigma_{\ell-1}^2 N^2 (\ell+1)^2$ for $0 \leq i \leq \ell-1$, $0 \leq j \leq N-1$.*

PROOF. We have $\mathbf{G}_{i,i} = \sum_{k=0}^{\ell} \text{FFT}(\mathbf{S}_{\ell-1})_{i,k} \odot \text{FFT}(\mathbf{S}_{\ell-1}^*)_{k,i}$, and thus $|(\mathbf{G}_{i,i})_j| = \sum_{k=0}^{\ell} |(\text{FFT}(\mathbf{S}_{\ell-1})_{i,k})_j|^2$. For N -point FFT result \mathbf{z} of scalar \mathbf{a} , we have $|\mathbf{z}_i| \leq \|\mathbf{z}\| = \sqrt{N} \|\mathbf{a}\|$ for $0 \leq i \leq N-1$ [9]. Thus, we have $|(\mathbf{G}_{i,i})_j| \leq (\ell+1) \cdot N \cdot \|(\mathbf{S}_{\ell-1})_{i,k}\|^2 \leq \sigma_{\ell-1}^2 N^2 (\ell+1)^2$, since $\|(\mathbf{S}_{\ell-1})_{i,k}\| \leq \sigma_{\ell-1} \cdot \sqrt{(\ell+1)N}$. \square

For the root of an fFLDL tree when $\ell = 2$, we have $(\mathbf{D}_{0,0})_j = (\mathbf{G}_{0,0})_j \leq 9\sigma_1^2 N^2$ by Lemma G.2. For $(\mathbf{D}_{1,1})_j = (\mathbf{G}_{1,1})_j - \frac{|(\mathbf{G}_{1,0})_j|^2}{(\mathbf{D}_{0,0})_j}$, since $(\mathbf{D}_{0,0})_j \in \mathbb{R}^+$ from Theorem 5.1, we have $(\mathbf{D}_{1,1})_j \leq (\mathbf{G}_{1,1})_j \leq 9\sigma_1^2 N^2$. By Theorem 5.1, we have $(\mathbf{D}_{2,2})_j = \frac{q^2}{(\mathbf{D}_{0,0})_j (\mathbf{D}_{1,1})_j} \geq \frac{q^2}{81\sigma_1^4 N^4}$, by taking the upper bound $9\sigma_1^2 N^2$ of $(\mathbf{D}_{0,0})_j$, $(\mathbf{D}_{1,1})_j$. Thus, for the leaf values σ , we have $\sigma \leq \sigma_2 / \sqrt{q^2 / (81\sigma_1^4 N^4)} = \sigma_2 \cdot 9\sigma_1^2 N^2 / q$.

H CRAMER'S RULE

Cramer's rule [19] is used for solving systems of linear equations. Considering a system of N equations with N unknowns \mathbf{x} , represented as $\mathbf{A}\mathbf{x} = \mathbf{b}$. Cramer's rule states that the solution can be written as $\mathbf{x}_i = \frac{\det(\mathbf{A}_i)}{\det(\mathbf{A})}$, where \mathbf{A}_i is the matrix formed by replacing the i -th column of \mathbf{A} by the column vector \mathbf{b} .

The formulae for the reduction coefficients in the KeyGen and Delegate process come directly from Cramer's Rule applied to the system $\mathbf{A}\mathbf{x} = \mathbf{b}$, where, in the first level, \mathbf{A} is the 2×2 matrix whose (i, j) -entry is the Hermitian product $\langle s_i, s_j \rangle$ of the i^{th} and j^{th} rows of the delegation matrix, and where \mathbf{b} is the two-dimensional column vector whose i^{th} coefficient is $\langle s_2, s_i \rangle$. This result generalises to arbitrary levels; i.e., for any given number of levels $\ell \geq 1$, the reduction of the vector $\mathbf{s}_{\ell+1}$ is effected by replacing it with $\mathbf{s}_{\ell+1} - [\mathbf{k}_0] \mathbf{s}_0 - \dots - [\mathbf{k}_\ell] \mathbf{s}_\ell$, where the \mathbf{k}_i are the coefficients of the solution \mathbf{x} to the system $\mathbf{A}\mathbf{x} = \mathbf{b}$, where \mathbf{A} is the $(\ell+1) \times (\ell+1)$ matrix whose (i, j) -entry is the Hermitian product $\langle s_i, s_j \rangle$ of the i^{th} and j^{th} rows of the delegation matrix, and where \mathbf{b} is the $(\ell+1)$ -dimensional column vector whose i -th coefficient is $\langle s_{\ell+1}, s_i \rangle$.

I KEY AND CIPHERTEXT SIZE CALCULATIONS

The LATTE keys and ciphertexts are mainly collections of polynomials in \mathcal{R} . The degree of each polynomial is N and the number of bits in each coefficient is $\kappa = \lceil \log_2 q \rceil$. The parameters N and q are dependent on the security level required, and values for these are given in Table 13 in Sec. 6. The key/ciphertext bit-size is equal to $N \cdot \kappa \cdot \text{number of polynomials}$, plus any additional bit strings sent, in the case of the ciphertext. Furthermore, we usually consider the key and ciphertext sizes in bytes, and so when the total bit-size is computed, it will be divided by 8 to give the size in bytes.

Master Keys. The master public key consists of a polynomial $\mathbf{h} \in \mathcal{R}_q$. Therefore the bit-size is $N \cdot \kappa$. The master private key \mathbf{S}_0 consists of $(\mathbf{f}, \mathbf{g}, \mathbf{F}, \mathbf{G})$. However, \mathbf{F} and \mathbf{G} can be recomputed on the fly from \mathbf{f} and \mathbf{g} using NTRUSolve. The solution is not unique but as long as it is a short solution, it will suffice. However, this is not efficient and so this research considers the entire $(\mathbf{f}, \mathbf{g}, \mathbf{F}, \mathbf{G})$ to be stored as the private key. Therefore, the master private key is of size $4N \cdot \kappa$.

Delegated Keys. The delegated public key can be straightforwardly generated using the master public key and the chain of user IDs along which the delegation process is happening. Although this can be efficiently generated on the fly, given the user ID chain, we will consider it being stored as the polynomials $\mathbf{h}, \mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_\ell$, which translates as $(\ell+1)$ polynomials in \mathcal{R} , and so the total bit-size of the delegated public key is $(\ell+1) \cdot N \cdot \kappa$. The delegated private key generated from level $\ell-1$ to level ℓ , to be passed onto users at level $\ell+1$, is a $(\ell+2) \times (\ell+2)$ matrix of polynomials in \mathcal{R}_q . Its size is therefore $(\ell+2) \cdot (\ell+2) \cdot N \cdot \kappa$.

User Private Keys. The user public key is entirely dependent on the identity, and so we only examine the size of the user private key. In LATTE for a user at level ℓ , this is a tuple of $(\ell+1)$ polynomials in \mathcal{R}_q . However, we only need to store ℓ of these polynomials (disregarding t_0) and so the user private key is of bit size $\ell \cdot N \cdot \kappa$.

Ciphertexts. Let's consider the ciphertext at level ℓ . This consists of $\ell+1$ polynomials $C_1, \dots, C_\ell, C_h \in \mathcal{R}_q$ along with a 256-bit string Z (which is essentially the encrypted message). Therefore, at level ℓ , the bit-size of the full ciphertext is $(\ell+1) \cdot N \cdot \kappa + 256$.

Table 15: Performance Results (op/s) for the DLP IBE Scheme from [1] (Scaled to 4.2GHz).

Set	Security	n	$\log_2 q$	KeyGen	Ext	Enc	Dec
DLP-0	80	512	22	14.7	873.2	8731.8	6202.9
DLP-3	192	1024	22	4.9	454.1	2639.8	1621.6

Table 16: Performance Results (op/s) for the FALCON [51] (Scaled to 4.2GHz).

Set	Sec.	n	$\log_2 q$	KeyGen	Sign	Verify
FALCON-512	128	512	14	211.4	10861.7	51008.1
FALCON-1024	256	1024	14	66.5	5319.4	24926.1

J COMPARISON TO OTHER SCHEMES

Comparison to DLP IBE. Performance results of the single-level DLP IBE scheme from [1] (converted to op/s at 4.2GHz) are given in Table 15. Since the decryption in the DLP IBE did not include ciphertext validation, for a fair comparison with LATTE, we use the sum of DLP encryption and decryption run-time to compute the op/s of decryption in Table 15. We focus on the comparison between LATTE-1 and DLP-3, since the sizes of parameters N and q are similar. The KeyGen speed of our LATTE-1 implementation is 1.9x faster than DLP-3, and the speed of our LATTE-1 Extract implementation is about 3.2x faster than DLP-3 extraction. This is mainly because we adapt the faster NTRUSolve and lattice Gaussian sampling procedure from FALCON [51]. In addition, the Encrypt/Decrypt speed is 6.3x–8.1x faster in our implementation.

Comparison to FALCON. After adapting the NTRUSolve and lattice Gaussian sampling procedures from FALCON [51], our optimised LATTE KeyGen becomes similar to the FALCON KeyGen, and our optimised LATTE Extract becomes similar to the FALCON Sign, in terms of the operations used by these algorithms. Therefore, here we compare the run-time speed of our optimised LATTE KeyGen/Extract against the FALCON KeyGen/Sign, respectively. The performance results of the FALCON is summarised in Table 16. We focus on the comparison between LATTE-1 and FALCON-1024, since the size of parameter N is the same. The KeyGen speed of our LATTE-1 implementation is about 7.1x slower than FALCON-1024, and the speed of our LATTE-1 Extract implementation is about 3.7x slower than FALCON-1024 Sign. This is mainly because: (1) The size of q in LATTE is much larger than FALCON (24 bits for LATTE-1 compared to 14 bits for FALCON-1024), which will significantly increase the maximal integer size in NTRUSolve as well as the run-time overhead in KeyGen [48]. (2) FALCON computes the fFLDL Tree during the KeyGen, while the fFLDL Tree is computed during the Extract in our LATTE scheme. This difference will add overhead to the run-time speed of our LATTE Extract implementation. (3) From the Falcon specification [51], the AVX2 and FMA instructions were used in the source code during the benchmark. However, these instructions are not used in the source code of our LATTE implementation.

Comparison to Non-quantum Safe HIBE. Pairings-based HIBE scheme performance results from [37] (converted to op/s at 4.2GHz) are given in Table 17. Parameter s is the bit size of the field, which is comparable to the bit-length of RSA modulus by providing the

Table 17: Performance Results (op/s) for the Gentry-Silverberg HIBE Scheme (2-level) Using Java v1.6 [37] (Scaled to 4.2GHz).

s	Security	length of m	Enc	Dec
1024	80	160	11.5	6.7
3072	128	256	0.8	0.5

same security. Although not directly comparable, these results give a good indication of the feasibility of LATTE at levels 1 or 2.

K LATTE VARIANT COMPUTING FFLDL TREE IN KEYGEN/DELEGATE

If the key extraction speed is critical for the application, similar to FALCON [51], we can move the fFLDL Tree computation from the LATTE Extract (Line 3 in Alg. 4 in Appendix B) to the LATTE KeyGen/Delegate when generating a master/delegated private key S_ℓ , at the expense of significantly larger master/delegated private key size. The KeyGen/Delegate/Extract speed of this LATTE variant is shown in Table 18. The LATTE Extract in this variant is about 1.3x–1.7x faster than the run-time speed in Table 12, while the KeyGen/Delegate is at most 6% slower.

Here we also analyse the overhead in the master/delegated private key size of this variant due to the fFLDL Tree T . Assuming a floating-point value has p bytes, the size of T consists of the following 3 parts:

- For a $d \times d$ basis S_ℓ , the root of T stores $d(d-1)/2$ components of L from the LDL^* decomposition, with each component in $\mathbb{C}[x]/\langle x^N + 1 \rangle$. Thus, the root of T has $d(d-1)/2 \cdot 2Np = Npd(d-1)$ bytes.
- The root of T has d sub-trees. The i -th non-leaf level of a sub-tree has 2^i nodes, $0 \leq i \leq \log_2 N - 2$. Each node at i -th level of a sub-tree stores $L_{1,0} \in \mathbb{C}[x]/\langle x^n + 1 \rangle$ from the LDL^* decomposition, where $n = N/2^{i+1}$. Therefore, the total size of i -th level of a sub-tree is $2^i \cdot 2(N/2^{i+1})p = Np$ bytes, and the total size of all non-leaf nodes in a sub-tree is $Np(\log_2 N - 1)$ bytes.
- A sub-tree has $N/2$ leaf nodes. Each leaf node stores a p -byte floating-point value. Therefore, the total size of all leaf nodes in a sub-tree is $Np/2$ bytes.

Table 18: Performance Results (op/s) and fFLDL Tree Size (Bytes) of LATTE Variant at 4.2GHz.

Set	KeyGen	$\ell = 1$			$\ell = 2$	
		Ext	Del	Tree Size	Ext	Tree Size
LATTE-1	9.4	1865.1	-	172032	-	-
LATTE-2	3.3	783.6	-	376832	-	-
LATTE-3	5.4	54.7	2.3	344064	33.7	565248
LATTE-4	1.6	25.9	0.8	753664	15.9	1228800

Thus, the total size of T is $Npd(d-1) + d(Np(\log_2 N - 1) + Np/2) = Npd(\log_2 N + d - 3/2)$ bytes. Columns “Tree Size” in Table 18

summarise the fFLDL tree size for the parameters and floating-point precisions in our LATTE implementation.