A Central Limit Framework for Ring-LWE Decryption

Sean Murphy and Rachel Player

Royal Holloway, University of London, U.K.
s.murphy@rhul.ac.uk, rachel.player@rhul.ac.uk

Abstract. The main contribution of this paper is to develop a statistical framework, based on a Central Limit argument, for analysing the noise in ciphertexts in homomorphic encryption schemes that are based on Ring-LWE. Such an approach is very general: apart from finite variance, no assumption on the distribution of the noise is required (in particular, the noise need not be subgaussian). We demonstrate that such a Central Limit approach can be used to obtain a high-quality approximation of the distribution of the noise in an appropriate decoding basis, even in dimension as small as $n = 100$. We apply our framework and results to a homomorphic Ring-LWE cryptosystem of Lyubashevsky, Peikert and Regev (Eurocrypt 2013, full version) in order to illustrate the benefit of this approach. We show that a Central Limit approach leads to tighter bounds for the probability of decryption failure than have been obtained in prior work.

1 Introduction

The Learning with Errors or *LWE* problem [19, 20] has become a standard hard problem in cryptology that is at the heart of lattice-based cryptography [15, 18]. The Ring Learning with Errors or *Ring-LWE* problem [21, 11] is a generalisation of the LWE problem from the ring of integers to certain other number field rings that potentially give far better efficiency.

A key application area of lattice-based cryptography is (fully, somewhat or levelled) homomorphic encryption [7]. Homomorphic encryption enables an untrusted party to operate meaningfully on encrypted data belonging to a different party, without requiring access to the secret key. A large number of homomorphic encryption schemes have been proposed in the literature, for example [2, 6, 8, 12, 4, 3], many of which [2, 6, 12, 3] are based on Ring-LWE. To illustrate the ideas of this paper, we consider the symmetric key homomorphic cryptosystem given by Lyubashevsky, Peikert and Regev in Section 8.3 of [12] (the full version of [13]), which we term the `SymHom` cryptosystem.

A common feature among all homomorphic encryption schemes is that all ciphertexts have an inherent *noise*. This is typically small in a fresh ciphertext, but the noise grows as homomorphic evaluation operations are performed. If the noise grows too large, then decryption fails. Thus a good understanding of the randomness properties of the noise in a ciphertext is essential to be able to choose appropriate parameters to ensure correctness and efficiency.

1.1 Contributions

The main contribution of this paper is to develop a statistical framework, based on a Central Limit argument, for analysing the noise in ciphertexts in homomorphic encryption schemes that are based on Ring-LWE. This Central Limit framework is essentially based on approximating the mean vector and the covariance matrix of the noise of a ciphertext when embedded into an appropriate complex space and transformed with respect to an appropriate decryption basis. We show that the approximate Normality of this embedded noise when expressed in a decryption basis is fundamentally a Central Limit phenomenon arising from the weighted sum of many random variables, where the weights arise from a change of basis matrix to the decryption basis.

To illustrate the utility of this approach, our second contribution is to apply this framework to the `SymHom` cryptosystem. In Theorems 1 and 2 we present new, tighter bounds for the probabilities of incorrect decryption in degree-1 and degree-2 `SymHom` ciphertexts.

1.2 Motivation for the Central Limit Approach in Ring-LWE

The decryption of ciphertexts in a Ring-LWE-based cryptosystem such as `SymHom` requires us to consider the noise in a ciphertext as a real-valued vector with respect to an appropriate "decoding" basis for the complex space H (see Section 2.3) where the noise in the ciphertext is first obtained as a real-valued vector with respect to a different basis for this complex space H [12]. For example, if $C^{(p\Gamma)}$ is a vector of dimension n expressing the noise in a ciphertext with respect to the decoding $p\Gamma$-basis for H (Section 2.4) and $C^{(T)}$ is a vector of dimension n expressing the noise in a ciphertext with respect to the original T-basis for H (Section 2.3), then $C^{(p\Gamma)} = p\Delta C^{(T)}$ for an appropriate real-valued $n \times n$ change of basis matrix Δ and "scaling prime" p (which is the plaintext modulus in `SymHom`). In particular, this means that we can express a component $C_j^{(p\Gamma)}$ of $C^{(p\Gamma)}$ as

$$C_j^{(p\Gamma)} = p \sum_{k=1}^{n} \Delta_{jk} C_k^{(T)} .$$

The components $C_1^{(T)}, \ldots, C_n^{(T)}$ of $C^{(T)}$ are identically distributed random variables that are uncorrelated and, in general, independent, with mean $\mathbf{E}\left(C_j^{(T)}\right) = 0$ and some finite variance $\mathrm{Var}\left(C_j^{(T)}\right) = \rho^2$. Thus a component $C_j^{(p\Gamma)}$ of a noise vector in the $p\Gamma$-basis is a weighted sum of uncorrelated and in general independent identically distributed random variables. We will show that the weightings $\Delta_{j1}, \ldots, \Delta_{jn}$ are of comparable size, which suggests that a Central Limit argument can be invoked to give a Normal approximation for a component $C_j^{(p\Gamma)}$. For successful decryption, we require each component of $C^{(p\Gamma)}$ to be bounded by an appropriate threshold. A Central Limit approach enables us to bound the probability of incorrect decryption using bounds on the tails of Normal distributions.

Theorems 1 and 2 demonstrate the improvement that can be obtained by using a Central Limit approach in comparison with prior bounds, such as those of [12], obtained using δ-subgaussian random variables [14, 16]. For example, if $\eta_1(n, q, \rho) = \frac{1}{2}(n^{\frac{1}{2}}\rho)^{-1}q$ is moderate or large, Theorem 1 gives a decryption failure probability bound of

$$\frac{2n\exp(-\frac{1}{2}\eta_1^2)}{(2\pi)^{\frac{1}{2}}\eta_1}.$$

This is tighter than the equivalent δ-subgaussian decryption failure probability bound of

$$2n\exp(-\frac{1}{2}\eta_1^2)$$

which is obtained by using the tail bound of [16, Lemma 18] in the manner of [12, Lemma 6.5].

Using such a Central Limit approach has a number of additional advantages over other approaches, such as the subgaussian approaches of [12, 5]. These advantages are listed below and expressed in terms of the above discussion. These advantages are also illustrated by Figure 3 of Example 1 (Section 2.5), which considers a weighted sum of heavy-tailed Laplace distributions that model the distribution of the noise in degree-2 `SymHom` ciphertexts arising as the output of the homomorphic multiplication of two fresh ciphertexts.

1. A Central Limit approach makes no substantive distributional assumption for the components $C_k^{(T)}$ beyond finite variance, so is potentially applicable to $C_k^{(T)}$ that are chosen from heavy-tailed distributions. Thus a Central Limit approach is more generally applicable than other approaches that for example have a subgaussian requirement for such random variables.

2. A Central Limit approach gives an explicit approximating distribution for the cryptographic random variable of interest which can be directly used for general calculation or simulation purposes of use in cryptography. By contrast, a subgaussian approach can never give a explicit approximating distribution and can only give (generally weaker) tail bounds.

3. A Central Limit approach gives not only asymptotically an approximation to a Normal distribution, but also a close approximation concretely, for practically relevant Ring-LWE dimensions n.

1.3 Structure of the Paper

We recall relevant background and introduce new tools in Section 2. We then outline our Central Limit approach in Section 3, and illustrate one aspect of its applicability by considering the `SymHom` cryptosystem in Section 4.

2 Background

In this section, we recall some relevant background and introduce new notation and tools. We give notation and some useful definitions in Section 2.1 and recall algebraic background in Section 2.2. We recall the definition and properties of the complex space H in Section 2.3. We introduce a useful basis for H in Section 2.4 and a useful product for H in Section 2.5. Finally, we recall the statistical background for our Central Limit approach in Section 2.6.

2.1 Notation

We consider the ring $R = \mathbb{Z}[X]/(\Phi_m(X))$, where $\Phi_m(X)$ is the m^{th} cyclotomic polynomial of degree n, and we let R_a denote R/aR for an integer a. For simplicity, we only consider the case where m is a large prime, though our arguments apply more generally. In this case we have $n = \phi(m) = m - 1$, and we also let $n' = \frac{1}{2}n = \frac{1}{2}(m-1)$. We let ζ_m denote a (primitive) m^{th} root of unity, which has minimal polynomial $\Phi_m(X) = 1 + X + \ldots + X^n$. The m^{th} *cyclotomic number field* $K = \mathbb{Q}(\zeta_m)$ is the field extension of the rational numbers $\mathbb{Q}$ obtained by adjoining this m^{th} root of unity ζ_m, so K has degree n. The tensor product $K \otimes_{\mathbb{Q}} \mathbb{R}$ is denoted by $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$.

The value or more formally the coset representative of $(r \bmod q)$ nearest to 0 is denoted by $[\![r]\!]_q = r - q\lfloor q^{-1}r \rceil$, and we use the same notation for a coset of $\mathbb{Z}_q$. We can also extend this idea componentwise to vectors, and we write $[\![\cdot]\!]_q^B$ to indicate such an extension with respect to a basis B. We use † to denote the complex conjugate transpose of a matrix, so $T^\dagger = \overline{T}^T$.

Encryption and decryption in Ring-LWE-based cryptography are inherently statistical processes, and we are giving Central Limit approximations to the distributions of cryptographic random variables of interest. Thus we use the notation $\dot{\sim}$ to denote "is approximately distributed as" in the sense that we may use the approximating distribution for practical purposes without significant error, as is typically done by taking a Central Limit Normal distribution approximation in statistical analysis. Furthermore, whilst Central Limit results are formally asymptotic results concerning sums or means of random variables, such Central Limit approximations usually apply in practice with relatively few summands (except perhaps for pathological distributions) as illustrated for example in Figure 3 of Example 1. We therefore typically use the phrasing "for moderate or large ..." in such a Central Limit context to emphasise that the usual applicability of Central Limit approximations with relatively few summands. Furthermore, we denote by Q the "Q-function" giving the upper tail probability for a standard Normal $\mathrm{N}(0,1)$ distribution, so

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty \exp(-\tfrac{1}{2}z^2)\, dz\,.$$

This tail probability $Q(x)$ is bounded by its asymptotic expansion, so

$$Q(x) \leq (2\pi x^2)^{-\frac{1}{2}} \exp(-\tfrac{1}{2}x^2)\,,$$

and we note that this bound is very tight even for moderate values of $x > 0$.

We now give three definitions relevant to our analysis. Definitions 1 and 2 are used to specify the `SymHom` cryptosystem, whilst Definition 3 specifies the Ring-LWE problem.

Definition 1 ([16]). *The univariate* Balanced Reduction *function $\mathcal{R}$ on $\mathbb{R}$ is the random function* $\mathcal{R}(a) = \begin{cases} 1 - (\lceil a \rceil - a) & \text{with probability} \quad \lceil a \rceil - a \\ -(\lceil a \rceil - a) & \text{with probability} \quad 1 - (\lceil a \rceil - a). \end{cases}$

The multivariate Balanced Reduction *function $\mathcal{R}$ on $\mathbb{R}^l$ with support on $[-1, 1]^l$ is the random function $\mathcal{R} = (\mathcal{R}_1, \ldots, \mathcal{R}_l)$ with component functions $\mathcal{R}_1, \ldots, \mathcal{R}_l$ that are independent univariate Balanced Reduction functions.*

Definition 2 ([16]). *Let B be a (column) basis matrix for the n-dimensional lattice Λ in H. If $\mathcal{R}$ is the Balanced Reduction function, then the* coordinate-wise randomised rounding discretisation *or* CRR discretisation $\lfloor X \rceil^B_{\Lambda+c}$ *of the random variable X on H to the lattice coset $\Lambda + c$ with respect to the basis matrix B is the random variable*

$$\lfloor X \rceil^B_{\Lambda+c} = X + B\, \mathcal{R}\left(B^{-1}(c - X)\right).$$

Definition 3 ([21, 11]). *Let R be the ring of integers of a number field K. Let $q \geq 2$ be an integer modulus. Let $R^\vee$ be the dual fractional ideal of R. Let $R_q = R/qR$ and $R_q^\vee = R^\vee/qR^\vee$. Let $K_\mathbb{R} = K \otimes_\mathbb{Q} \mathbb{R}$.*

Let χ be a distribution over $K_\mathbb{R}$. Let $s \in R_q^\vee$ be a secret. A sample from the Ring-LWE distribution *$A_{s,\chi}$ over $R_q \times K_\mathbb{R}/qR^\vee$ is generated by choosing $a \leftarrow R_q$ uniformly at random, choosing $e \leftarrow \chi$ and outputting*

$$(a, b = (a \cdot s)/q + e \mod qR^\vee).$$

Let Ψ be a family of distributions over $K_\mathbb{R}$. The Search Ring-LWE *problem is defined as follows: given access to arbitrarily many independent samples from $A_{s,\chi}$ for some arbitrary $s \in R_q^\vee$ and $\chi \in \Psi$, find s.*

Let Υ be a distribution over a family of error distributions, each over $K_\mathbb{R}$. The average-case Decision Ring-LWE *problem is to distinguish with non-negligible advantage between arbitrarily many independent samples from $A_{s,\chi}$ for a random choice of $(s, \chi) \leftarrow \mathcal{U}\left(R_q^\vee\right) \times \Upsilon$, and the same number of uniformly random samples from $R_q \times K_\mathbb{R}/qR^\vee$.*

2.2 Cyclotomic Number Fields

There are n *ring embeddings* $\sigma_1, \ldots, \sigma_n \colon K \to \mathbb{C}$ that fix every element of $\mathbb{Q}$. Such a ring embedding σ_k (for $1 \leq k \leq n$) is defined by $\zeta_m \mapsto \zeta_m^k$, so $\sum_{j=1}^n a_j \zeta_m^j \mapsto \sum_{j=1}^n a_j \zeta_m^{kj}$, and such ring embeddings occur in conjugate pairs. The *canonical embedding* $\sigma : K \to \mathbb{C}^n$ is $a \mapsto (\sigma_1(a), \ldots, \sigma_n(a))^T$.

The *ring of integers* $\mathcal{O}_K$ of a number field is the ring of all elements of the number field which are roots of some monic polynomial with coefficients in $\mathbb{Z}$. The ring of integers of the m^{th} cyclotomic number field K is

$$R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[x]/(\Phi_m).$$

The canonical embedding σ embeds R as a lattice $\sigma(R)$. The conjugate dual of this lattice corresponds to the embedding of the dual fractional ideal

$$R^\vee = \{a \in K \mid \mathrm{Tr}(aR) \subset \mathbb{Z}\}.$$

If we define t such that $t^{-1} = m^{-1}(1 - \zeta_m)$, then [12, Lemma 2.16] shows that $R^\vee = \langle t^{-1} \rangle$. We let $(R^\vee)^k$ denote the space of products of k elements of $R^\vee$, that is to say

$$(R^\vee)^k = \{s_1 \dots s_k \mid s_1, \dots, s_k \in R^\vee\} = \left\{t^{-k} r_1 \dots r_k \mid r_1, \dots, r_k \in R\right\}.$$

2.3 The Complex Space H

The ring embeddings $\sigma_1, \dots, \sigma_n$ from K into $\mathbb{C}$ occur in complex conjugate pairs with $\overline{\sigma_k} = \sigma_{m-k}$. Accordingly, much of the analysis of Ring-LWE takes place in a space H of conjugate pairs of complex numbers. It is sometimes convenient to consider such a single conjugate pair in isolation, so giving rise to the space H_2, the 2-dimensional version of H. The *conjugate pair* mappings $\widetilde{\sigma}_i \colon K \to H_2$ for $1 \le i \le n'$ are given for $a \in K$ by

$$\widetilde{\sigma}_i(a) = (\sigma_i(a), \sigma_{m-i}(a))^T,$$

In particular, the canonical embedding actually embeds into $H_2 \times \ldots \times H_2 \cong H$.

Definition 4. The *conjugate pairs matrix* is the complex unitary $n \times n$ matrix T, so $T^{-1} = T^\dagger$, given by

$$T = 2^{-\frac{1}{2}} \begin{pmatrix} 1 & 0 \dots 0 & 0 \dots & 0 & i \\ 0 & 1 \dots 0 & 0 \dots & i & 0 \\ \vdots & \vdots \ddots \vdots & \vdots & \vdots & \vdots \\ 0 & 0 \dots 1 & i \dots & 0 & 0 \\ 0 & 0 \dots 1 & -i \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \ddots & \vdots & \vdots \\ 0 & 1 \dots 0 & 0 \dots & -i & 0 \\ 1 & 0 \dots 0 & 0 \dots & 0 & -i \end{pmatrix}. \qquad \square$$

Definition 5. The complex conjugate pair space $H = T(\mathbb{R}^n)$, where T is the conjugate pairs matrix. In particular, $H_2 = T(\mathbb{R}^2)$. $\square$

Definition 6. The *I-basis* for H is given by the columns of the $n \times n$ identity matrix I, that is to say the I-basis is the standard basis. $\square$

Definition 7. The *T-basis* for H is given by the columns of the conjugate pairs matrix T. $\square$

An element of H is expressed via the I-basis as a vector of n' conjugate pairs. Such an element of H can also be expressed (by construction) in the T-basis as a *real-valued* vector. We also note that the vector representing an element in the T-basis for H has the same norm as an element representing the same element in the I-basis for H, as $|Tv|^2 = |v|^2$ because T is a unitary matrix. Expressing elements of H as vectors in the T-basis therefore gives the isomorphism between H and $\mathbb{R}^n$ as an inner product space.

The canonical embedding under σ of a sum in K gives a componentwise addition in H for any basis for H. Similarly, the canonical embedding under σ of a product in K gives rise to a componentwise $\odot$-product in H when the vectors expressing the embedded elements are in the I-basis for H, when we have $\sigma(aa') = \sigma(a) \odot \sigma(a')$.

2.4 The $p\Gamma$-basis for H

In Definition 8 we specify the *$p\Gamma$-basis* for H in which elements of H are also expressed as real-valued vectors. The $p\Gamma$-basis arises as the embedding of a basis of conjugate pairs for $R^\vee$. The $p\Gamma$-basis is a more convenient basis for H in the case when m is prime, and is a suitable basis for decryption.

Definition 8. The *$p\Gamma$-basis* for H is given by the columns of the matrix $p\Gamma$ (for p prime), where

$$\Gamma = \frac{1}{m} \begin{pmatrix} 1-\zeta_m^1 & 1-\zeta_m^2 & 1-\zeta_m^3 & \dots & 1-\zeta_m^n \\ 1-\zeta_m^2 & 1-\zeta_m^4 & 1-\zeta_m^6 & \dots & 1-\zeta_m^{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1-\zeta_m^n & 1-\zeta_m^{2n} & 1-\zeta_m^{3n} & \dots & 1-\zeta_m^{n^2} \end{pmatrix},$$

and is the embedding of the basis $\left\{\frac{p}{m}(1-\zeta_m^1), \frac{p}{m}(1-\zeta_m^2), \dots, \frac{p}{m}(1-\zeta_m^n)\right\}$ of conjugate pairs for $R^\vee$ in H. □

In Figure 1, we summarise our notation for elements of H expressed with respect to the various bases. If Z is a vector expressing an element of H as a vector of conjugate pairs in the I-basis (or standard basis) for H, then we have real-valued vectors $Z^\ddagger = T^\dagger Z$ and $Z^* = (p\Gamma)^{-1} Z$ expressing this element as a vector in the T-basis and the $p\Gamma$-basis for H respectively.

The change of basis transformations between the T-basis and the $p\Gamma$-basis are summarised in Figure 2, and the relevant properties of the (scaled) change-of-basis matrix $\Delta = \Gamma T^{-1}$ are given in Lemma 1.

Lemma 1. The change of basis matrix from the T-basis to the $p\Gamma$-basis of H is the real invertible matrix $p^{-1}\Delta$, where $\Delta = \Gamma^{-1}T$ satisfies $\Delta\Delta^T = mI - J$. □

Proof. It is clear that $\Delta = \Gamma^{-1}T$ is invertible as both Γ^{-1} and T are invertible. The matrix $\Delta^{-1} = T^{-1}\Gamma = T^\dagger\Gamma$ has matrix entries Δ_{kl}^{-1} satisfying

$$m\Delta_{kl}^{-1} = \begin{cases} 2^{-\frac{1}{2}}\left(\left(1-\zeta_m^{kl}\right) + \left(1-\zeta_m^{-kl}\right)\right) = 2^{\frac{1}{2}}\left(1 - \mathrm{Re}\left(\zeta^{kl}\right)\right) & [1 \le k \le n'] \\ 2^{-\frac{1}{2}}\left(-i\left(1-\zeta_m^{-kl}\right) + i\left(1-\zeta_m^{kl}\right)\right) = 2^{\frac{1}{2}}\mathrm{Im}\left(\zeta^{kl}\right) & [n' < k \le n], \end{cases}$$

Basis for H	I-Basis	T-Basis	$p\Gamma$-Basis
Vector or Random Variable	Z	$Z^{\ddagger}$	Z^{*}
Transformation from the I-Basis	I	$T^{\dagger}$	$p^{-1}\Gamma^{-1}$

Fig. 1. Notation for the expression of an element of H as a vector in the various different vector space bases for H. Note that p is a scaling factor.

$$H \text{ with } T\text{-basis} \quad \underset{p\Delta^{-1} = T^{-1}(p\Gamma)}{\overset{p^{-1}\Delta = (p\Gamma)^{-1}T}{\longleftrightarrow}} \quad H \text{ with } p\Gamma\text{-basis}$$

Fig. 2. Change of Basis Matrices for the T-basis and $p\Gamma$-basis for H in which elements of H are expressed as real-valued vectors.

so Δ^{-1} and hence Δ are real matrices. Thus we have

$$\Delta\Delta^{T} = \Delta\Delta^{\dagger} = (\Gamma^{-1}T)(\Gamma^{-1}T)^{\dagger} = \Gamma^{-1}TT^{\dagger}(\Gamma^{-1})^{\dagger} = \left(\Gamma^{\dagger}\Gamma\right)^{-1}.$$

We note that $\Gamma^{\dagger}_{jk} = m^{-1}(1-\zeta_m^{-jk})$ and that $\sum_{l=1}^{n}\zeta^{l} = -1$ and so on. Thus $\sum_{l=1}^{n}\zeta^{l(j-k)} = n$ if $k = j$ and -1 if $k \neq j$ (for $1 \leq k, j \leq n$), which yields

$$\begin{aligned}
\left(\Gamma^{\dagger}\Gamma\right)_{jk} &= \sum_{l=1}^{n}\Gamma^{\dagger}_{jl}\Gamma_{lk} = \frac{1}{m^2}\sum_{l=1}^{n}(1-\zeta^{-jl})(1-\zeta^{lk}) \\
&= \frac{1}{m^2}\sum_{l=1}^{n}1 - \frac{1}{m^2}\sum_{l=1}^{n}\zeta^{lk} - \frac{1}{m^2}\sum_{l=1}^{n}\zeta^{-jl} + \frac{1}{m^2}\sum_{l=1}^{n}\zeta^{l(k-j)} \\
&= \begin{cases} 2m^{-2}(n+1) = 2m^{-1} & [k=j] \\ m^{-2}(n+1) = m^{-1} & [k \neq j], \end{cases}
\end{aligned}$$

so $\Gamma^{\dagger}\Gamma = m^{-1}(I+J)$. Thus $\Delta\Delta^{T} = (\Gamma^{\dagger}\Gamma)^{-1} = mI - J$. □

2.5 The $\otimes$-product of elements of H

In certain Ring-LWE-based homomorphic encryption schemes, the embedded noise in a fresh ciphertext can be approximated as a Normal random variable. For a ciphertext obtained as the output of a homomorphic multiplication of two fresh ciphertexts, its noise is defined to be the product of the noises in the input ciphertexts. We will therefore be interested in the $\otimes$-product (Definition 9) of two elements of H expressed in the T-basis.

Definition 9. The $\otimes$-*product* of two real vectors $u = (u_{11}, u_{12}, \ldots, u_{n'1}, u_{n'2})^T$ and $v = (v_{11}, v_{12}, \ldots, v_{n'1}, v_{n'2})^T$ of length $n = 2n'$ is

$$u \otimes v = \begin{pmatrix} u_{11} \\ u_{12} \\ \vdots \\ u_{n'1} \\ u_{n'2} \end{pmatrix} \otimes \begin{pmatrix} v_{11} \\ v_{12} \\ \vdots \\ v_{n'1} \\ v_{n'2} \end{pmatrix} = T^{\dagger}\left(Tu \odot Tv\right) = 2^{-\frac{1}{2}} \begin{pmatrix} u_{11}v_{11} - u_{12}v_{12} \\ u_{11}v_{12} + u_{12}v_{11} \\ \vdots \\ u_{n'1}v_{n'1} - u_{n'2}v_{n'2} \\ u_{n'1}v_{n'2} + u_{n'2}v_{n'1} \end{pmatrix}.$$

The $\otimes$-product of two vectors in H expressed in the T-basis is the expression in the T-basis of the componentwise $\odot$-product of those two vectors when expressed in the I-basis. □

In Lemma 2 we consider the bivariate (conjugate pair) case, and show that the resulting $\otimes$-product distribution is a Laplace distribution [10, 17]. The image of this distribution under T then gives the corresponding distribution of the $\odot$-product. The generalisation to the general case with n' conjugate pairs is clear and straightforward.

Lemma 2. Suppose that $W = \begin{pmatrix} W_1 \\ W_2 \end{pmatrix} = U \otimes V = \begin{pmatrix} U_1 \\ U_2 \end{pmatrix} \otimes \begin{pmatrix} V_1 \\ V_2 \end{pmatrix}$ is the $\otimes$-product of the independent random variables $U, V \sim \mathrm{N}(0, I_2)$ with a standard bivariate Normal distribution.
(i) The random variable W has a bivariate Laplace distribution with density function $f_W(w) = \pi^{-1} K_0(2^{\frac{1}{2}}|w|)$ for $w \in \mathbb{R}^2$, where K_0 is the modified Bessel function of the second kind given by $K_0(x) = \displaystyle\int_0^{\infty} \exp(-x \cosh t)\, dt$.
(ii) A component W_j of W has a univariate Laplace distribution with density function $f_{W_j}(w_j) = 2^{-\frac{1}{2}} \exp(-2^{\frac{1}{2}}|w_j|)$, and so has mean $\mathbf{E}(W_j) = 0$, variance $\mathrm{Var}(W_j) = 1$, and tail probability $\mathbf{P}(|W_j| > \theta) = \exp(-2^{\frac{1}{2}}\theta)$. Furthermore, these orthogonal components W_1 and W_2 of W are not independent but are uncorrelated with covariance $\mathrm{Cov}(W_1, W_2) = 0$. □

Proof. Parts (i) and (ii) follow from the Preamble to Part II and from Section 5.1.1 of [10]. Furthermore, it is discussed at the end of Section 1 of [10] how a univariate Laplace distribution arises directly from the components of such an $\otimes$-product $\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} \otimes \begin{pmatrix} V_1 \\ V_2 \end{pmatrix} = 2^{-\frac{1}{2}} \begin{pmatrix} U_1V_1 - U_2V_2 \\ U_1V_2 + U_2V_1 \end{pmatrix}$. □

2.6 Lindeberg's condition for the Central Limit Theorem

To obtain a Normal approximation for a weighted sum $\sum_{j=1}^{n} a_j X_j$ of the form encountered in Ring-LWE, we need a general form of the Central Limit Theorem formally given by the Lindeberg condition [1]. We state such a Central Limit result in Lemma 3. However, Lemma 3 can be informally expressed as that the weighted sum $\sum_{j=1}^{n} a_j X_j$ of the form encountered in Ring-LWE has an approximate Normal distribution for moderate or large n provided that the absolute weights a_j are not dominated by just a few values.

Lemma 3. Suppose $X_1, X_2, \ldots$ are independent and identically distributed continuous random variables that are symmetric about 0 with mean $\mathbf{E}(X_j) = 0$ and variance $\mathrm{Var}(X_j) = 1$, and that have common density function f_{X_j}, and suppose that for constants $a_1, a_2, \ldots$ the sum $\sum_{j=1}^{l} a_j X_j$ has variance function $a(l)^2 = \sum_{j=1}^{l} a_j^2$, and that the functions $\widetilde{a}_j$ are defined by $\widetilde{a}_j(l) = \dfrac{|a_j|}{a(l)}$. In this case, *Lindeberg's condition* is that for any given $\epsilon > 0$, the sum

$$\sum_{j=1}^{l} \widetilde{a}_j(l)^2 \, \Psi_{X_j}\left(\frac{\epsilon}{\widetilde{a}_j(l)}\right) \to 0 \quad \text{as } l \to \infty, \text{ where } \Psi_{X_j}(\theta) = \int_{\theta}^{\infty} x^2 f_{X_j}(x)\, dx.$$

If *Lindeberg's condition* is satisfied, then $a(l)^{-1} \sum_{j=1}^{l} a_j X_j$ tends in distribution to a standard Normal $\mathrm{N}(0,1)$ distribution as $l \to \infty$. □

Proof. We can define a random variable $X_j^{(\alpha)} = \begin{cases} X_j & [|X_j| > \alpha] \\ 0 & [|X_j| \leq \alpha] \end{cases}$ for $\alpha > 0$ obtained by "censoring" X_j at the minimum absolute value α and so on. With this notation, Lindeberg's condition [1] in our case is that for any given $\epsilon > 0$, the sum $\frac{1}{a(l)^2} \sum_{j=1}^{n} \mathbf{E}\left(\left((a_j X_j)^{(\epsilon \alpha(l))}\right)^2\right) \to 0$ as $n \to \infty$. We therefore note that

$$\begin{aligned} \mathbf{E}\left((a_j X_j)^{(\epsilon a(l))^2}\right) &= 2\int_{\epsilon a(l)}^{\infty} x^2 \, f_{a_j X_j}(x)\, dx = 2\int_{\epsilon a(l)}^{\infty} \frac{x^2}{|a_j|} \, f_{X_j}\left(\frac{x}{|a_j|}\right) dx \\ &= 2|a_j|^2 \int_{\epsilon \widetilde{a}_j(l)^{-1}}^{\infty} x'^2 \, f_{X_j}(x')\, dx' = 2|a_j|^2 \, \Psi_{X_j}\left(\frac{\epsilon}{\widetilde{a}_j(l)}\right), \end{aligned}$$

so giving the form of Lindeberg's condition of the Lemma. If Lindeberg's condition is satisfied, then the convergence in distribution to Normality follows from the Lindeberg form of the Central Limit Theorem [1]. □

3 Central Limit Framework

The decryption of ciphertexts in certain Ring-LWE homomorphic encryption schemes requires us to consider the noise in a ciphertext as a real-valued vector in an appropriate "decoding" basis, such as the $p\Gamma$-basis. Let $C^{(p\Gamma)}$ be a vector expressing the noise in such a ciphertext in the $p\Gamma$-basis. In this section, we use a Central Limit approach to approximate the distribution of $C^{(p\Gamma)}$. In Section 3.1 we present our main result, Proposition 3, which gives a good approximation for the distribution of $C^{(p\Gamma)}$ as a Normal random variable. In Section 3.2 we illustrate the closeness of our Central Limit Normal approximation.

3.1 Details of the Central Limit approach

Proposition 1 gives a Central Limit approximation to a weighted sum of the form $\sum_{j=1}^{n} a_j X_j$ for independent and identically distributed random variables

$X_1, \ldots, X_n$. This proposition is a summary of the Lindeberg condition for a Central Limit Theorem (see Section 2.6) and essentially states that a good Normal approximation exists for the weighted sum if enough of the largest weights $|a_j|$ are of comparable size.

Proposition 1. Suppose that $X = (X_1, \ldots, X_n)$ has components $X_1, \ldots, X_n$ that are independent and identically distributed random variables with mean $\mathbf{E}(X_j) = 0$ and finite variance $\mathrm{Var}(X_j) = \rho^2$. For weights $a = (a_1, \ldots, a_n)$, the weighted sum $a^T X = \sum_{j=1}^{n} a_j X_j \dot{\sim} \mathrm{N}(0, |a|^2\rho^2)$ has an approximate Normal distribution for moderate or large n, provided that the weights $a_1, \ldots, a_n$ are not dominated by just a few of these weights. □

Concretely, in a typical parameter situation of Ring-LWE where we have $n > 10^2$, (or $n > 10^3$ in the case of homomorphic encryption), we can expect Proposition 1 to give a good approximation when as few as (for example) about 20 of the largest weights are comparable. We can extend Proposition 1 in the obvious way to give the multivariate case of Proposition 2.

Proposition 2. Suppose that $X = (X_1, \ldots, X_n)$ has components $X_1, \ldots, X_n$ that are independent and identically distributed random variables with mean $\mathbf{E}(X_j) = 0$ and finite variance $\mathrm{Var}(X_j) = \rho^2$, so X has covariance matrix $\rho^2 I_n$. If A is a $n \times n$ matrix whose entries A_{jk} satisfy the Proposition 1 weights criterion, then the transformed random variable $AX \dot{\sim} \mathrm{N}(0, \rho^2 AA^T)$ can be approximated as a multivariate Normal distribution for moderate or large n. □

In Proposition 3, we apply Proposition 2 to approximate the distribution of the noise in a Ring-LWE ciphertext expressed in an appropriate decryption basis. We note the proof of Proposition 3 is complicated by the fact that a pair of random variables in the T-basis arising as the image of a conjugate pair in the I-basis are uncorrelated but not independent (see for example Lemma 1).

Proposition 3. Suppose that $C^{(T)}$ is a vector expressing the noise in a Ring-LWE ciphertext in the T-basis for H, so a component $C_j^{(T)}$ of $C^{(T)}$ has mean $\mathbf{E}\left(C_j^{(T)}\right) = 0$ and finite variance $\mathrm{Var}\left(C_j^{(T)}\right) = \rho^2$. Suppose further that the S-basis given by the columns of the $n \times n$ matrix S is an appropriate basis of H for decryption, and that $\Psi = ST^{-1}$ is the change of basis matrix from the T-basis to the S-basis for H. If the entries Ψ_{jk} of Ψ satisfy the Proposition 1 weights criterion, then the distribution of the noise $C^{(S)}$ in this ciphertext in the (decryption) S-basis for H can be approximated as

$$C^{(S)} \dot{\sim} \mathrm{N}(0; \rho^2 \Psi\Psi^T) \qquad \text{for moderate or large } n.$$

In particular, the $p\Gamma$-basis for H yields $C^{(p\Gamma)} \dot{\sim} \mathrm{N}(0; p^2\rho^2(mI - J))$. □

Proof. We can split $\Psi = (\Psi'|\Psi'')$ into two $n \times n'$ submatrices and we similarly split $C^{(T)} = \left(C^{(T)'} \middle| C^{(T)''}\right)^T$ into the first n' components $C^{(T)'}$ and the final n' components $C^{(T)''}$. Furthermore, their conjugate pairs origin means that $C^{(T)'}$

and $C^{(T)''}$ are uncorrelated (see for example Lemma 2(ii)). The components $C_1^{(T)'}, \ldots, C_{n'}^{(T)'}$ of $C^{(T)'}$ are independent and identically distributed with mean 0 and variance ρ^2, so Proposition 2 gives $\Psi' C^{(T)'} \dot{\sim} \mathrm{N}(0; \rho^2 \Psi' \Psi'^T)$, and we similarly have $\Psi'' C^{(T)''} \dot{\sim} \mathrm{N}(0; \rho^2 \Psi'' \Psi''^T)$. Thus

$$C^{(S)} = \Psi C^{(T)} = \Psi' C^{(T)'} + \Psi'' C^{(T)''} \dot{\sim} \mathrm{N}(0; \rho^2 \Psi \Psi^T)$$

as $C^{(S)}$ is the sum of two uncorrelated approximate multivariate Normal random variables, so has an approximate Normal distribution with covariance matrix $\rho^2 \Psi' \Psi'^T + \rho^2 \Psi'' \Psi''^T = \rho^2 \Psi \Psi^T$. □

3.2 Quality of the CLT approximation

The Central Limit Theorem is formally a statement about the convergence (in distribution) of an appropriate weighted sum of random variables to a Normal distribution in the limit as the number of summands n tends to infinity. When such a result is applied in a concrete setting with a fixed finite n, it is reasonable to question the speed of this convergence, and in particular how accurate the approximation is. This issue is made more precise in a companion work [17], and can be verified empirically.

We illustrate the practical accuracy of our approach in Example 1 below, which models the following situation of interest. Let $C^{(p\Gamma)}$ be the noise vector a Ring-LWE ciphertext expressed in the $p\Gamma$-basis. Then we can express $C^{(p\Gamma)}$ in the T-basis as $C^{(p\Gamma)} = p\Delta C^{(T)}$, where Δ is the change-of-basis matrix of Lemma 1, and the components $C_i^{(T)}$ are identically distributed random variables that are uncorrelated and in general independent having mean 0 and variance ρ^2. We can express a component $C_j^{(p\Gamma)}$ of $C^{(p\Gamma)}$ as

$$C_j^{(p\Gamma)} = p \sum_{k=1}^{n} \Delta_{jk} C_k^{(T)} \tag{1}$$

where the Δ_{jk} are proportional to various sums and differences of m^{th} roots of unity with absolute size about 1, as any row of Δ has squared length $\sum_{k=1}^{n} \Delta_{jk}^2 = n$. Proposition 3 then yields the Normal approximation for a component as $C_j^{(p\Gamma)} \dot{\sim} \mathrm{N}\left(0, np^2\rho^2\right)$. Example 1 illustrates the closeness of such a Central Limit Normal approximation for the situation of Equation 1.

Example 1. Let $m = 101$ and $n = 100$ and let $Y = (Y_1, \ldots, Y_n)^T$ be a vector of independent and identically distributed Laplace random variables $Y_1, \ldots, Y_n$ with mean $\mathbf{E}(Y_j) = 0$ and variance $\rho^2 = \mathrm{Var}(Y_j) = 1$. We take $p = 1$ without loss of generality and consider the distribution of $W = \Delta Y$ where $\Delta = \Gamma^{-1} T$ is the change of basis matrix from the T-basis to the Γ-basis of H. We consider the first component $W_1 = \sum_{j=1}^{n} \Delta_{1k} Y_k$ of $W = \Delta Y$, where the first row $\Delta_1 =$

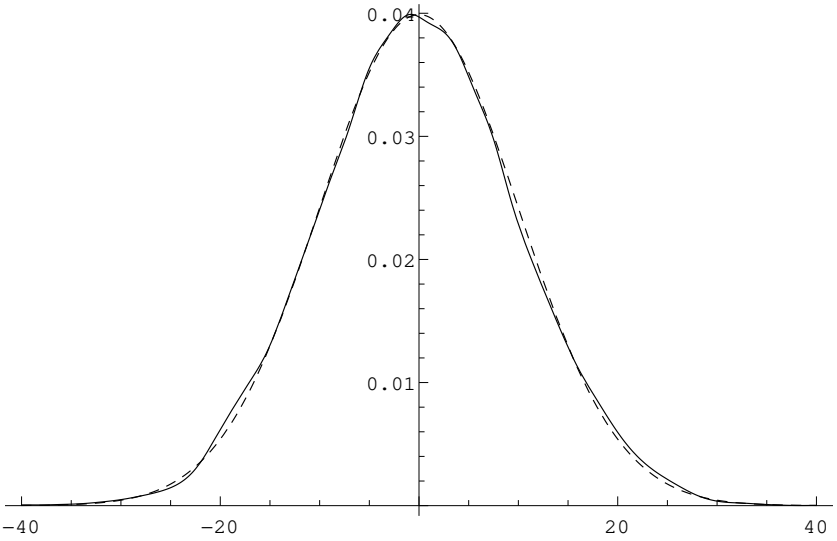

Fig. 3. An empirical density function based on 10^4 realisations of $W_1 = \sum_{j=1}^{100} \Delta_{1k} Y_k$ where $Y_1, \ldots, Y_{100}$ are independent and identically distributed Laplace random variables with variance 1 (solid line) and the density function of the corresponding approximating Normal $\mathrm{N}(0, 10^2)$ distribution (dashed line).

$(\Delta_{11}, \ldots, \Delta_{1n})$ of Δ is given by

$$\Delta_1 = \begin{pmatrix} -1.41, -1.40, -1.39, -1.37, -1.35, -1.32, -1.28, -1.24, -1.20, -1.15, \\ -1.10, -1.04, -0.98, -0.91, -0.84, -0.77, -0.69, -0.62, -0.54, -0.45, \\ -0.37, -0.28, -0.20, -0.11, -0.02, \ 0.07, \ 0.15, \ 0.24, \ 0.33, \ 0.41, \\ \ 0.50, \ 0.58, \ 0.66, \ 0.73, \ 0.81, \ 0.88, \ 0.94, \ 1.01, \ 1.07, \ 1.12, \\ \ 1.17, \ 1.22, \ 1.26, \ 1.30, \ 1.33, \ 1.36, \ 1.38, \ 1.40, \ 1.41, \ 1.41, \\ -0.04, -0.13, -0.22, -0.31, -0.39, -0.47, -0.56, -0.64, -0.71, -0.79, \\ -0.86, -0.93, -0.99, -1.05, -1.11, -1.16, -1.21, -1.25, -1.29, -1.32, \\ -1.35, -1.38, -1.39, -1.41, -1.41, -1.41, -1.41, -1.40, -1.39, -1.37, \\ -1.34, -1.31, -1.27, -1.23, -1.19, -1.14, -1.08, -1.02, -0.96, -0.89, \\ -0.82, -0.75, -0.68, -0.60, -0.52, -0.43, -0.35, -0.26, -0.18, -0.09 \end{pmatrix}.$$

The closeness of the Central Limit approximation for $W_1 = \sum_{j=1}^{n} \Delta_{1k} Y_k$ with mean $\mathbf{E}(W_1) = 0$ and variance $\mathrm{Var}(W_1) = 10^2$ to a Normal $\mathrm{N}(0, 10^2)$ random variable with mean 0 and variance 10^2 is illustrated by the comparison between the empirical density function of W_1 and the $\mathrm{N}(0, 10^2)$ density function shown in Figure 3. □

4 Application to SymHom

In this section, we apply the Central Limit framework developed in Section 3 to the symmetric-key homomorphic cryptosystem presented in [12, Section 8.3],

The SymHom cryptosystem. Let ψ be a continuous LWE error distribution over $K_{\mathbb{R}}$, and let $\lfloor \cdot \rceil$ denote any valid discretisation to cosets of some scaling of $R^\vee$ (e.g. using the decoding basis of $R^\vee$). The cryptosystem is defined formally as follows.

- Gen: choose $s' \leftarrow \lfloor \psi \rceil_{R^\vee}$, and output $s = t \cdot s' \in R$ as the secret key.
- $\mathrm{Enc}_s(\mu \in R_p)$: choose $e \leftarrow \lfloor p\psi \rceil_{t^{-1}\mu + pR^\vee}$. Let $c_0 = -c_1 \cdot s + e \in R_q^\vee$ for uniformly random $c_1 \leftarrow R_q^\vee$, and output the ciphertext $c(S) = c_0 + c_1 S$. The noise in $c(S)$ is defined to be e.
- $\mathrm{Dec}_s(c(S))$ for c of degree k: compute $c(s) \in (R^\vee)_q^k$, and decode it to $e = [\![c(s)]\!] \in (R^\vee)^k$. Output $\mu = t^k \cdot e \bmod pR$.

For ciphertexts c, c' of arbitrary degrees k, k', their homomorphic product is the degree-$(k+k')$ ciphertext $c(S) \boxdot c'(S) = c(S) \cdot c'(S)$, that is to say standard polynomial multiplication. The noise in the result is defined to be the product of the noise terms of c, c'. Similarly, for ciphertexts c, c' of *equal* degree k, their homomorphic sum is $c(S) \boxplus c'(S) = c(S) + c'(S)$, and the noise in the resulting ciphertext is the sum of those of c, c'.

Fig. 4. The SymHom cryptosystem as defined in [12, Section 8.3].

which we refer to as the SymHom cryptosystem. Our analysis enables us to present, in Theorems 1 and 2, new, tighter bounds for the respective probabilities of the incorrect decryption in degree-1 and degree-2 SymHom ciphertexts. Our analysis can similarly be applied for higher-degree ciphertexts [17].

4.1 Noise in SymHom ciphertexts

A description of SymHom cryptosystem, in the notation of [12], is given in Figure 4. We now describe in our notation the relevant parts of the SymHom cryptosystem in order to define the noise in a SymHom ciphertext. We first recall that the SymHom secret key is an element $s \in R$, the plaintext space is R_p, and a plaintext $\mu \in R_p$ is encrypted to give a linear polynomial over $R_q^\vee$.

The first step of the encryption process is to generate a random input for a discretisation process to a coset depending on the plaintext μ. Accordingly, we let Y be a random variable on H such that $TY \sim \mathrm{N}(0; p^2\rho^2 I_n)$ is a spherically symmetric n-dimensional Normal random variable with component variance $p^2\rho^2$ for an appropriately chosen ρ^2. We term Y the *Underlying Noise*, and Y is a complex-valued random vector expressed in the I-basis for H.

Specifically, we discretise Y to the coset $\sigma(pR^\vee) + \sigma(t^{-1}\mu)$ of the lattice $\sigma(pR^\vee)$ obtained by the canonical embedding of the scaled dual fractional ideal $pR^\vee$. We consider the coordinate-wise randomised rounding discretisation with respect to the $p\Gamma$-basis for H, and following Definition 2 we denote this discretisation of Y by $Y'(\mu) = \lfloor\, Y \,\rceil^{p\Gamma}_{\sigma(pR^\vee)+\sigma(t^{-1}\mu)}$.

The *Noise* random variable $Y''(\mu)$ in the encryption of the plaintext μ is then defined to be $Y''(\mu) = \sigma^{-1}(Y'(\mu))$, and is an element of a coset of $pR^\vee + t^{-1}\mu$ containing information about μ. For obvious reasons, we refer to $Y'(\mu) = \sigma(Y''(\mu))$

Description	Random Variable	Range of Random Variable
Underlying Noise	Y	Complex Space H
Embedded Noise	$Y'(\mu)$	Lattice Coset $\sigma(pR^\vee) + \sigma(t^{-1}\mu)$
Noise	$Y''(\mu)$	Number Field Coset $pR^\vee + t^{-1}\mu$

Fig. 5. Notation for the Noise-related quantities used in encryption of the plaintext μ.

as the *Embedded Noise*, and we note that $Y'(\mu)$ expresses the Embedded Noise in the I-basis of H. We summarise this discussion in Figure 5.

In the next step of encryption, we form the ciphertext from the Noise $Y''(\mu)$ and the secret key s in the following way. We choose A uniformly in $R_q^\vee$, and we let $A'(\mu) = -As + Y''(\mu) \in R_q^\vee$. The ciphertext $C(\theta;\mu)$ is the polynomial over $R_q^\vee$ defined as $C(\theta;\mu) = A'(\mu) + A\theta$. We note that this polynomial can be expressed directly in terms of the Noise $Y''(\mu)$ and the secret key s as $C(\theta;\mu) = A(\theta - s) + Y''(\mu)$. A fresh ciphertext is defined to be a degree-1 ciphertext, since the polynomial $C(\theta;\mu)$ is linear.

The output ciphertext of a homomorphic multiplication of two degree-1 ciphertext polynomials is obtained simply by multiplying these polynomials together. Thus we can obtain the degree-2 ciphertext polynomial over $R_q^\vee$ corresponding to the product $\mu_1\mu_2$ of plaintexts μ_1 and μ_2 as $C(\theta;\mu_1,\mu_2) = C(\theta;\mu_1) \boxdot C(\theta;\mu_2)$, where $C(\theta;\mu_1) = A_1'(\mu_1) + A_1\theta$ and $C(\theta;\mu_2) = A_2'(\mu_2) + A_2\theta$. This degree-2 ciphertext polynomial is $C(\theta;\mu_1,\mu_2) = A_1'(\mu_1)A_2'(\mu_2) + (A_2A_1'(\mu_1) + A_1A_2'(\mu_2))\ \theta + A_1A_2\theta^2$, which is given in terms of the secret key s and its constituent Noises $Y_1''(\mu)$ and $Y_2''(\mu)$ by

$$C(\theta;\mu_1,\mu_2) = A_1A_2(\theta - s)^2 + (A_2Y_1''(\mu_1) + A_1Y_2''(\mu_2))\ (\theta - s) + Y_1''(\mu_1)Y_2''(\mu_2).$$

The *Noise* in this degree-2 output ciphertext $C(\theta;\mu_1,\mu_2)$ is defined to be the product $Y_1''(\mu_1)Y_2''(\mu_2)$ of the Noises $Y_1''(\mu_1)$ and $Y_2''(\mu_2)$ of the degree-1 input ciphertexts. This process extends in the obvious way to give ciphertexts of higher degree.

4.2 Decryption using the $p\Gamma$-basis

In this section, we specify in our notation a decryption process for the `SymHom` cryptosystem using the $p\Gamma$-basis of H (though any appropriate basis can be used). We recall (see Figure 1) that we write $Z^\ddagger$ and Z^* to express an element of H as a vector in the T-basis and the $p\Gamma$-basis respectively.

Decryption of a degree-1 ciphertext polynomial $C(\theta;\mu)$ begins by evaluating this polynomial at the secret s. We obtain information about the Noise since $C(s;\mu) = Y''(\mu) \bmod R_q^\vee$. If we embed $C(s;\mu)$ in H under σ and perform a reduction modulo q with respect to to the $p\Gamma$-basis, then we obtain an integer vector $[\![\sigma(C(s;\mu))]\!]_q^{p\Gamma}$ with entries in $[-\frac{1}{2}q, \frac{1}{2}q)$.

The Embedded Noise $Y'(\mu)$ is expressed in the I-basis for H, so $Y'(\mu)$ is expressed with respect to the T-basis of H as the real vector $Y'(\mu)^\ddagger = T^\dagger Y(\mu)$.

However, the change of basis from this T-basis to the $p\Gamma$-basis of H is given by $p^{-1}\Delta = p^{-1}\Gamma^{-1}T$, so there is a real transformation $Y'(\mu)^* = p^{-1}\Delta Y(\mu)^\ddagger$ that gives a real vector $Y'(\mu)^*$ specifying the Embedded Noise expressed in the $p\Gamma$-basis for H. This allows us to write $Y'(\mu)^* = [\![\sigma(C(s,\mu))]\!]_q^{p\Gamma}$ if the Embedded Noise is small enough. In this case, we can recover the real vector $Y'(\mu)^*$ and hence the real Embedded Noise vector $Y'(\mu)^\ddagger$ with respect to the T-Basis. This allows us to determine the coset representative $\sigma(t^{-1}\mu)$ for the coset of the lattice $\sigma(pR^\vee)$ corresponding to the plaintext $\mu \in R_p$. Thus if the Embedded Noise is small enough with high probability, then we can recover the plaintext μ with high probability.

This decryption process generalises to degree-2 and higher degree ciphertexts in a natural way. For example, if $C(\theta;\mu_1)$ and $C(\theta;\mu_2)$ are two degree-1 ciphertexts with respective Embedded Noises $Y_1'(\mu_1)$ and $Y_2'(\mu_2)$, then the degree-2 ciphertext $C(s;\mu_1,\mu_2) = Y''(\mu_1)Y''(\mu_2) = C(s;\mu_1)C(s;\mu_2) \mod (R^\vee)_q^2$, and so we obtain $(Y_1'(\mu_1) \odot Y_2'(\mu_2))^* = [\![\sigma(C(s;\mu_1,\mu_2))]\!]_q^{m^{-1}p\Gamma}$ for small Embedded Noise. Thus if this Embedded Noise is small enough with high probability, we can recover the plaintext product $\mu_1\mu_2 \in R_p$ with high probability.

4.3 Decryption Failure Probabilities in the `SymHom` cryptosystem

We now present in Theorems 1 and 2 our main results, which give (respectively) bounds for the probability of the incorrect decryption of degree-1 and degree-2 `SymHom` ciphertexts. Both results follow from the fact that `SymHom` decryption using (for example) the $p\Gamma$-basis for H fundamentally involves a change of basis transformation between bases for H ultimately to the $p\Gamma$-basis.

Theorem 1. If $\eta_1(n,q,\rho) = \frac{1}{2}(n^{\frac{1}{2}}\rho)^{-1}q$ is moderate or large, then the probability of the incorrect decryption of a `SymHom` degree-1 ciphertext in the $p\Gamma$-basis for H is bounded by

$$\mathbf{P}\left(\begin{array}{l}\text{Incorrect decryption of } \texttt{SymHom} \\ \text{degree-1 ciphertext in } p\Gamma\text{-basis}\end{array}\right) \leq \frac{2n\exp(-\frac{1}{2}\eta_1^2)}{(2\pi)^{\frac{1}{2}}\eta_1}. \qquad \square$$

Proof. The vector expressing the Embedded Noise in the $p\Gamma$-basis for H is of the form $\left(\lfloor Z \rceil_{\Lambda_c}^{p\Gamma}\right)^*$, where $Z = TZ^\ddagger$ and $p^{-1}Z^\ddagger = (p^{-1}T^\dagger)Z \sim \mathrm{N}(0,\rho^2 I_n)$. However, $\left(\lfloor Z \rceil_{\Lambda_c}^{p\Gamma}\right)^* = (p\Gamma)^{-1}\lfloor Z \rceil_{\Lambda+c}^{p\Gamma} \approx \Delta(p^{-1}T^\dagger)Z$, so Lemma 1 shows that

$$\left(\lfloor Z \rceil_{\Lambda_c}^{p\Gamma}\right)^* \dot{\sim} \mathrm{N}(0;\rho^2\Delta\Delta^T) = \mathrm{N}(0;\rho^2(mI - J)).$$

Thus $\left(\lfloor Z \rceil_{\Lambda_c}^{p\Gamma}\right)^*$ is well-approximated by a multivariate Normal random variable $U \sim \mathrm{N}(0;\rho^2(mI-J))$, with components $U_1,\ldots,U_n \sim \mathrm{N}(0,n\rho^2)$. These components therefore have an upper tail probability function given for $\alpha > 0$ by

$$\mathbf{P}(U_j > \alpha) = \mathbf{P}\left((n^{\frac{1}{2}}\rho)^{-1}U_j > (n^{\frac{1}{2}}\rho)^{-1}\alpha\right) = Q\left((n^{\frac{1}{2}}\rho)^{-1}\alpha\right),$$

where the Q-function is as defined in Section 2.1. We can now obtain a bound for the tail probability for the maximum of $|U_1|, \ldots, |U_n|$ for moderate $(n^{\frac{1}{2}}\rho)^{-1}\alpha$ by using the union bound [9] to obtain

$$\begin{aligned}\mathbf{P}\left(\max\{|U_1|,\ldots,|U_n|\} > \alpha\right) = 2\,\mathbf{P}\left(\max\{U_1,\ldots,U_n\} > \alpha\right) &\leq 2n\mathbf{P}(U_j > \alpha) \\ &\leq 2nQ\left((n^{\frac{1}{2}}\rho)^{-1}\alpha\right) \leq \frac{2n^{\frac{3}{2}}\rho}{(2\pi)^{\frac{1}{2}}\alpha}\exp\left(-\frac{\alpha^2}{2n\rho^2}\right).\end{aligned}$$

We can now give a bound for the probability of decryption failure for a degree-1 ciphertext using the Γ-basis. In this case, decryption fails if the absolute size of any component of exceeds $\frac{1}{2}q$, so taking $\alpha = \frac{1}{2}q$ for moderate and large $\eta_1(n,q,\rho) = \frac{1}{2}(n^{\frac{1}{2}}\rho)^{-1}q$ gives

$$\mathbf{P}\left(\begin{array}{l}\text{Incorrect decryption of } \texttt{SymHom} \\ \text{degree-1 ciphertext in } p\Gamma\text{-basis}\end{array}\right) \leq \frac{2n\exp(-\frac{1}{2}\eta_1^2)}{(2\pi)^{\frac{1}{2}}\eta_1}. \qquad \square$$

Theorem 2. If $\eta_2 = \frac{1}{2}(n^{\frac{1}{2}}mp\rho_1\rho_2)^{-1}q$ is moderate or large, then the probability of the incorrect decryption of a `SymHom` degree-2 ciphertext in the $p\Gamma$-basis for H is bounded by

$$\mathbf{P}\left(\begin{array}{l}\text{Incorrect decryption of } \texttt{SymHom} \\ \text{degree-2 ciphertext in } p\Gamma\text{-basis}\end{array}\right) \leq \frac{2n\exp(-\frac{1}{2}\eta_2^2)}{(2\pi)^{\frac{1}{2}}\eta_2}. \qquad \square$$

Proof. The decryption of a `SymHom` degree-2 ciphertext $C(\theta;\mu_1,\mu_2)$ involves processing this ciphertext as $[\![\sigma(C(s;\mu_1,\mu_2))]\!]_q^{m^{-1}p\Gamma}$, that is to say by regarding this Embedded Noise expressed as a vector with respect to the rescaled decoding conjugate pair $m^{-1}p\Gamma$-basis. The processing of a degree-2 ciphertext fundamentally therefore simply involves change of basis transformations for bases for H ultimately to the $m^{-1}p\Gamma$-basis. Thus we can adapt the argument of the proof of Theorem 1 simply by using the appropriate moments, and so we can replace ρ in η_1 with $mp\rho_1\rho_2$ in to give $\eta_2 = \eta_1(n,q,mp\rho_1\rho_2) = \frac{1}{2}(n^{\frac{1}{2}}mp\rho_1\rho_2)^{-1}q$. $\square$

Acknowledgements. Rachel Player was partially supported by an ACE-CSR Ph.D. grant, by the French Programme d'Investissement d'Avenir under national project RISQ P141580, and by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701).

References

1. Billingsley, P.: Probability and Measure, third edn. Wiley (1995)
2. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) Fully Homomorphic Encryption without Bootstrapping. In: Innovations in Theoretical Computer Science 2012, pp. 309–325. ACM (2012)
3. Cheon, J.H., Kim, A., Kim, M., Song, Y.S.: Homomorphic Encryption for Arithmetic of Approximate Numbers. In: T. Takagi, T. Peyrin (eds.) Advances in Cryptology - ASIACRYPT 2017, *LNCS*, vol. 10624, pp. 409–437. Springer (2017)

4. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Faster Fully Homomorphic Encryption: Bootstrapping in Less Than 0.1 Seconds. In: J.H. Cheon, T. Takagi (eds.) Advances in Cryptology - ASIACRYPT 2016, *LNCS*, vol. 10031, pp. 3–33. Springer (2016)
5. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: fast fully homomorphic encryption over the torus. J. Cryptology **33**(1), 34–91 (2020)
6. Fan, J., Vercauteren, F.: Somewhat Practical Fully Homomorphic Encryption. IACR Cryptology ePrint Archive **2012**, 144 (2012)
7. Gentry, C.: Fully Homomorphic Encryption using Ideal Lattices. In: M. Mitzenmacher (ed.) Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, ACM, pp. 169–178 (2009)
8. Gentry, C., Sahai, A., Waters, B.: Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based. In: R. Canetti, J. Garay (eds.) Advances in Cryptology - CRYPTO 2013, *LNCS*, vol. 8042, pp. 75–92. Springer (2013)
9. Grimmett, G., Stirzaker, D.: Probability And Random Processes, 3rd edn. Oxford University Press (2001)
10. Kotz, S., Kozubowski, T., Podórski, K.: The Laplace Distribution and Generalizations. Birkhüser (2001)
11. Lyubashevsky, V., Peikert, C., Regev, O.: On Ideal Lattices and Learning with Errors Over Rings. IACR Cryptology ePrint Archive **2012**, 230 (2012)
12. Lyubashevsky, V., Peikert, C., Regev, O.: A Toolkit for Ring-LWE Cryptography. IACR Cryptology ePrint Archive **2013**, 293 (2013)
13. Lyubashevsky, V., Peikert, C., Regev, O.: A Toolkit for Ring-LWE Cryptography. In: T. Johansson, P. Nguyen (eds.) Advances in Cryptology - EUROCRYPT 2013, *LNCS*, vol. 7881, pp. 35–54. Springer (2013)
14. Micciancio, D., Peikert, C.: Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller. In: D. Pointcheval, T. Johansson (eds.) Eurocrypt 2012, *LNCS*, vol. 7237, pp. 700–718. Springer (2012)
15. Micciancio, D., Regev, O.: Lattice-based Cryptography. In: D.J. Bernstein and J. Buchmann and E. Dahmen (ed.) Post-Quantum Cryptography, pp. 147–191. Springer (2009)
16. Murphy, S., Player, R.: δ-subgaussian Random Variables in Cryptography. In: J. Jang-Jaccard, F. Guo (eds.) ACISP 2019: The 24th Australasian Conference on Information Security and Privacy, *LNCS*, vol. 11547, pp. 251–268. Springer (2019)
17. Murphy, S., Player, R.: Discretisation and Product Distributions in Ring-LWE. MathCrypt 2019, to appear. Available as IACR eprint 2019/596. (2019)
18. Peikert, C.: A decade of lattice cryptography. Foundations and Trends in Theoretical Computer Science **10**(4), 283–424 (2016)
19. Regev, O.: On Lattices, Learning with Errors, Random Linear Codes and Cryptography. In: H. Gabow, R. Fagin (eds.) 37th Annual ACM Symposium of Theory of Computing (2005)
20. Regev, O.: The Learning with Errors Problem (Invited Survey). In: IEEE Conference on Computational Complexity, pp. 191–204 (2010)
21. Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient Public Key Encryption Based on Ideal Lattices. In: M. Matsui (ed.) Advances in Cryptology - ASIACRYPT 2009, *LNCS*, vol. 5912, pp. 617–635 (2009)