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 and rachel.player@rhul.ac.uk

Abstract. The purpose of this paper is to use a Central Limit approach to develop a statistical framework for analysing ciphertexts in Ring-LWE homomorphic encryption schemes. This statistical framework gives rise to Normal approximations for ciphertext random variables, and we show that this allows probabilities to be determined more accurately and hence enables better bounds for decryption failure probabilities than the widely used existing approach based on δ-subgaussian random variables. To demonstrate the benefit of the Central Limit approach, we apply our framework and results to a homomorphic Ring-LWE cryptosystem of Lyubashevsky, Peikert and Regev (Eurocrypt 2013, full version).

Keywords. Ring-LWE, Central Limit Theorem, δ-subgaussian.

1 Introduction

The Learning with Errors or *LWE* problem [21, 22] has become a standard hard problem in cryptology that is at the heart of lattice-based cryptography [16, 20]. The Ring Learning with Errors or *Ring-LWE* problem [23, 12] is a generalisation of the LWE problem from the ring of integers to certain other number field rings. Both the LWE problem and the Ring-LWE problem are related to well-studied lattice problems that are believed to be hard [3, 12, 13, 18, 21].

A key application of lattice-based cryptography is the ability to achieve (fully, somewhat or levelled) homomorphic encryption. Using homomorphic encryption means that one party (the server) can operate meaningfully on encrypted data belonging to a different party (the client), and the server does not need access to the secret key in order to do this. Constructing a *fully homomorphic encryption* scheme was a longstanding open problem until it was resolved in Gentry's seminal work [7]. Gentry's original scheme specifies a *somewhat homomorphic encryption* scheme, which is transformed into a fully homomorphic encryption scheme using a technique known as bootstrapping.

A large number of somewhat homomorphic cryptosystems have been proposed in the literature, for example [2, 6, 8, 13, 5, 4], many of which [2, 6, 13, 4] 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 [13] (the full version of [14]), 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 of the Paper

The first contribution of this paper is to develop a statistical framework, based on a Central Limit argument, for ciphertexts in homomorphic encryption schemes that are based on Ring-LWE. To illustrate its utility, our second contribution is to apply this Central Limit framework to the `SymHom` cryptosystem. We give results on the probability of incorrect decryption for freshly encrypted (degree-1) `SymHom` ciphertexts in Theorem 1 and for degree-2 `SymHom` ciphertexts formed as a result of homomorphic multiplication in Theorem 2.

This Central Limit analysis of the `SymHom` cryptosystem is essentially based on approximating the mean vector and the covariance matrix of the noise of a ciphertext when embedded into the complex space and transformed into 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.

The third contribution of this paper is to contrast our Central Limit approach to evaluating decryption failure probabilities in the `SymHom` cryptosystem with the approach used in the original analysis [13]. The analysis in [13] is based on δ-subgaussian (for $\delta \geq 0$) random variables [15, 17], which are a generalisation of subgaussian random variables [10].

Our main conclusion is that the Central Limit approach allows us to determine probabilities more accurately and so gives rise to better bounds for decryption failure probabilities in Ring-LWE than a δ-subgaussian approach. Such a Central Limit approach could therefore lead to the design of more efficient Ring-LWE cryptosystems, and should be used in place of a δ-subgaussian approach.

1.2 Structure of the Paper

We review the algebraic and statistical background for Ring-LWE, and in particular the `SymHom` cryptosystem, in Section 2. We outline the Central Limit approach to evaluating decryption failure probabilities in Ring-LWE in Section 3, and illustrate this approach by considering the `SymHom` cryptosystem in Section 4. We contrast the Central Limit approach with the δ-subgaussian approach to evaluating decryption failure probabilities in Ring-LWE in Section 5.

2 Mathematical Background for Ring-LWE

In this section, we give the necessary algebraic and statistical background. The algebraic background has its origins in [13] and in part follows [13]. 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, so $n = \phi(m) = m - 1$, and we also let $n' = \frac{1}{2}n = \frac{1}{2}(m-1)$, though our arguments apply more generally.

There are three natural algebraic settings for the discussion of Ring-LWE: the n-dimensional real vector space $\mathbb{R}^n$, the m^{th} cyclotomic number field K (Section 2.1) and the complex space H (Section 2.2). We move between these settings at different places in our discussion of Ring-LWE.

2.1 Cyclotomic Number Fields

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.

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 $a \mapsto (\sigma_1(a), \ldots, \sigma_n(a))^T$.

We can define an induced geometry on K which has an ℓ_2-norm given by $\|a\|_2 = \|\sigma(a)\|_2 = \sum_{j=1}^n |\sigma_j(a)|^2 = 2\sum_{j=1}^{n'} |\sigma_j(a)|^2$ and an ℓ_∞-norm $\|\cdot\|_\infty$ given by $\|a\|_\infty = \|\sigma(a)\|_\infty = \max\{|\sigma_1(a)|, \ldots, |\sigma_n(a)|\}$.

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

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

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 [13, 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

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

2.2 The Complex Space H

As we noted in Section 2.1, the ring embeddings $\sigma_1, \ldots, \sigma_n$ 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 of conjugate pairs of complex numbers. We now specify the appropriate complex space for analysing Ring-LWE, which following [13] we denote by H. In order to do so, we first define the conjugate pairs matrix T. We use † to denote the complex conjugate transpose of a matrix, so $T^\dagger = \overline{T}^T$ and so on.

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

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

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

Definition 3. The *I-basis* for H is given by the columns of the $n \times n$ identity matrix I, that is to say by standard basis vectors. $\square$

Definition 4. The *T-basis* for H is given by the columns of the conjugate pair matrix T. $\square$

The I-basis and T-basis for H give two different ways of expressing an element of H as a vector:

$$\begin{aligned} &H\text{-vectors in the } I\text{-basis} \quad \left\{ (z_1, \ldots, z_{n'}, \overline{z}_{n'}, \ldots, \overline{z}_1)^T \,\middle|\, z_1, \ldots, z_{n'} \in \mathbb{C} \right\}, \\ &H\text{-vectors in the } T\text{-basis} \quad \left\{ (w_1, \ldots, w_{n'}, w_{n'+1}, \ldots, w_n)^T \,\middle|\, w_1, \ldots, w_n \in \mathbb{R} \right\}. \end{aligned}$$

An element of H is expressed as a vector in 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. Thus the T-basis for H is a very natural basis to use for the analysis of Ring-LWE.

We also specify the *$p\Gamma$-basis* for H in Definition 5 in which elements of H are also expressed as real-valued vectors, where the $p\Gamma$-basis arises as the embedding

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

Fig. 1. 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.

of a basis of conjugate pairs for $R^\vee$. This $p\Gamma$-basis for H is a more convenient basis in the case when m is prime, and is a suitable basis for decryption. The change of basis transformations between the T-basis and the $p\Gamma$-basis are summarised in Figure 1, and the relevant properties of the matrix $\Delta = \Gamma T^{-1}$ are given in Lemma 1.

Definition 5. The *$p\Gamma$-basis* for H is given by the columns of the matrix $p\Gamma$ (for p prime), where

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

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. □

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

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

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 $\Gamma_{jk}^\dagger = m^{-1}(1-\zeta_m^{-jk})$ 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 $k \ne j$ (for $1 \le k, j \le n$), which yields

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

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

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. 2. 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.

At various times in our discussion of Ring-LWE, we consider the expression of an element of H as a real-valued vector with respect to the T-basis and the $p\Gamma$-basis for H. We therefore introduce the notation of Figure 2 for the purposes of clarity when dealing with an element of H expressed with respect to the various different bases for H. Thus 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.

2.3 Products of Elements of H

We have seen that the expression of an element of H in the I-basis gives a vector of complex conjugate pairs. It is sometimes convenient to consider such a single conjugate pair in isolation, so giving rise to the space $H_2 = T(\mathbb{R}^2)$. The *conjugate pair* mappings $\widetilde{\sigma}_i$ on K for $1 \leq i \leq n'$ are given by

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

where σ_i are the ring embeddings defined in Section 2.1. The conjugate pair mappings are each (by definition) an embedding $\widetilde{\sigma}_i \colon K \to H_2$. The canonical embedding σ can therefore be regarded as essentially the concatenation of the n' conjugate pair embeddings $\widetilde{\sigma}_1, \ldots, \widetilde{\sigma}_{n'}$, In particular, the canonical embedding actually embeds K into $H_2 \times \ldots \times H_2 \cong H \subset \mathbb{C}^n$, and such an embedded element is expressed as a vector (with appropriate component re-ordering) with respect to the I-basis for H.

The canonical embedding under σ of a sum in the cyclotomic number field gives a componentwise addition in H of the vectors expressing the embedded elements for any basis for H. Similarly, the canonical embedding under σ of a product in the cyclotomic number field 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 $\sigma(aa') = \sigma(a) \odot \sigma(a')$.

The canonical embedding of a product under σ gives other forms of "product" for the corresponding vectors expressing elements of H when other bases are used. The appropriate notion of a product of two elements of the complex space H when these elements are expressed as real vectors in the T-basis for H is given by Definition 6, which specifies the $\otimes$-product of two real vectors.

Definition 6. 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. □

2.4 The $\otimes$-product of Normal Random Variables

In Ring-LWE cryptosystems such as the `SymHom` cryptosystem, random variables relating to the ciphertexts can be closely approximated by gaussian random variables. Thus Definition 6 shows that the distribution of the $\otimes$-product of Normal random variables is of fundamental interest when considering the product of such random variables expressed in an appropriate basis as a real-valued vector. Lemma 2 considers the bivariate (conjugate pair) case, and shows that the resulting $\otimes$-product distribution is a Laplace distribution [11]. The image of this distribution under T then gives the corresponding distribution of the $\odot$-product. The generalisation to the general case with many 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 [11]. Furthermore, it is discussed at the end of Section 1 of [11] how a univariate Laplace distribution arises directly from the components of such an $\otimes$-product $\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} \otimes \begin{pmatrix} V_1 \\ V_2 \end{pmatrix} = 2^{-\frac{1}{2}} \begin{pmatrix} U_1V_1 - U_2V_2 \\ U_1V_2 + U_2V_1 \end{pmatrix}$. □

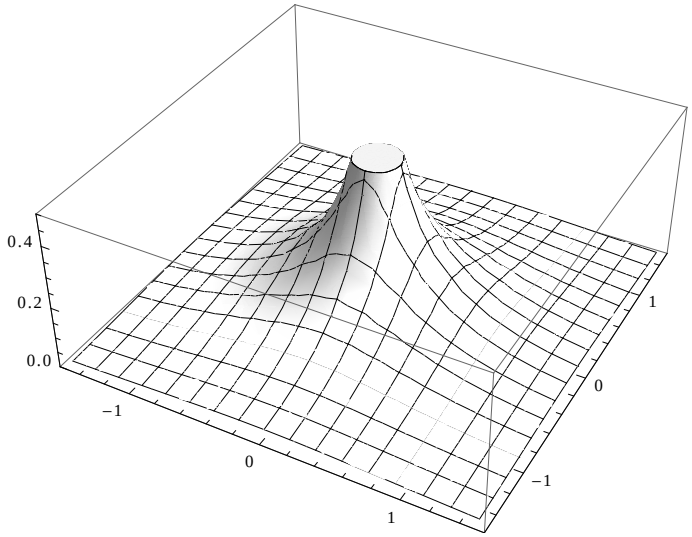

Fig. 3. The density function of a standard bivariate Laplace distribution.

The density function of a standard bivariate Laplace random variable arising as the $\otimes$-product of standard bivariate Normal random variables is illustrated in Figure 3. The density function of the standard univariate Laplace random variable, essentially two "back-to-back" exponential random variables, arising as a component of such an $\otimes$-product is illustrated in Figure 4. For comparison, the density function of a standard Normal $\mathrm{N}(0,1)$ random variable having the same mean 0 and variance 1 is also shown Figure 4. It can be seen that this univariate Laplace distribution arising from an $\otimes$-product has a far heavier tail than the corresponding Normal distribution.

3 A Central Limit Approach to Ring-LWE Decryption

In a Ring-LWE decryption process such as that used in the `SymHom` cryptosystem, we have to consider the ciphertext as a real-valued vector in an appropriate basis for H to allow for decryption, such as the $p\Gamma$-basis. If $C^{(p\Gamma)}$ is a vector expressing such a ciphertext in the $p\Gamma$-basis, the we require all components of $C^{(p\Gamma)}$ to be bounded by an appropriate threshold for a successful decryption. In order to evaluate or bound a decryption failure probability, we therefore have to determine or approximate the distribution of $C^{(p\Gamma)}$. In this Section, we consider a Central Limit approach to such a determination.

3.1 Motivation for the Central Limit Approach in Ring-LWE

In the `SymHom` cryptosystem, the vector $C^{(p\Gamma)}$ expressing the ciphertext in the $p\Gamma$-basis appropriate for decryption is a linear transformation $C^{(p\Gamma)} = p\Delta C^{(T)}$

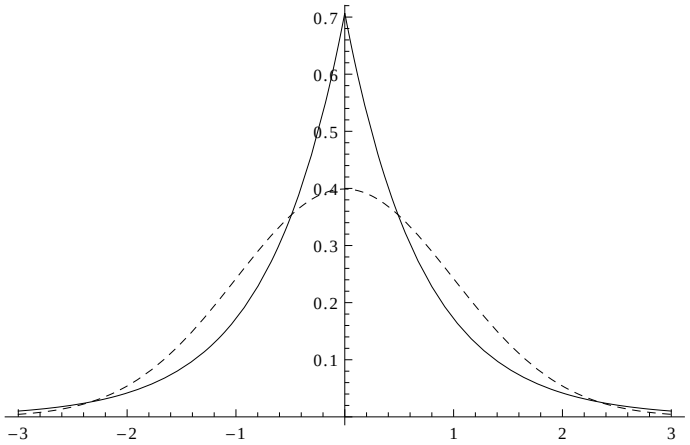

Fig. 4. The density function $2^{-\frac{1}{2}} \exp(-2^{\frac{1}{2}}|z|)$ of a standard univariate Laplace random variable (solid line) and the density function $(2\pi)^{-\frac{1}{2}} \exp(-\frac{1}{2}z^2)$ of a standard Normal $\mathrm{N}(0,1)$ random variable with the same mean 0 and variance 1 (dashed line).

of the vector $C^{(T)}$ expressing the ciphertext in the T-basis, where $\Delta = \Gamma^{-1}T$ is the real change of basis matrix of Lemma 1 and Figure 1. In particular, this means that we can express a component $C_j^{(p\Gamma)}$ of $C^{(p\Gamma)}$ as

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

However, the components $C_1^{(T)}, \ldots, C_n^{(T)}$ of $C^{(T)}$ are identically distributed random variables that are uncorrelated and in general independent having 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 ciphertext vector in the $p\Gamma$-basis for decryption is a weighted sum of uncorrelated and in general independent identically distributed random variables. As Lemma 1 shows that $\Delta\Delta^T = mI - J$, so $\sum_{k=1}^{n} \Delta_{jk}^2 = m - 1 = n$, we obtain

$$\mathbf{E}\left(C_j^{(p\Gamma)}\right) = 0 \quad \text{and} \quad \mathrm{Var}\left(C_j^{(p\Gamma)}\right) = np^2\rho^2.$$

The weightings $\Delta_{j1}, \ldots, \Delta_{jn}$ in the sum $C_j^{(p\Gamma)} = p\sum_{k=1}^{n} \Delta_{jk} C_k^{(T)}$ are of comparable size, as they are proportional to various sums and differences of m^{th} roots of unity with absolute size about 1 as any row of the $n \times n$ matrix Δ has squared length $\sum_{k=1}^{n} \Delta_{jk}^2 = n$. This suggests that a Central Limit argument (detailed below) gives a Normal approximation for a component $C_j^{(p\Gamma)}$ of $C^{(p\Gamma)}$. Such a Central Limit argument yields the Normal approximation

$$C_j^{(p\Gamma)} \dot{\sim} \mathrm{N}\left(0, np^2\rho^2\right) \qquad \text{for moderate or large } n,$$

where $\dot{\sim}$ denotes "is approximately distributed as". In particular, we note such a Central Limit argument makes no distributional assumption for $C_1^{(T)}, \ldots, C_n^{(T)}$ (beyond finite variance), so is potentially applicable to heavy-tailed distributions. The closeness of such a Central Limit Normal approximation for a sum such as $C_j^{(p\Gamma)} = p \sum_{k=1}^{n} \Delta_{jk} C_k^{(T)}$ is illustrated by Example 1.

Example 1. We consider a situation where $m = 101$ and $n = 100$ and we let $Y = (Y_1, \ldots, Y_n)^T$ be a vector of independent and identically distributed (heavy-tailed) Laplace random variables $Y_1, \ldots, Y_n$ with mean $\mathbf{E}(Y_j) = 0$ and variance $\rho^2 = \mathrm{Var}(Y_j) = 1$, as arises when considering the $\otimes$-product of Normal random variables. We consider the distribution of $W = \Delta Y$ (taking $p = 1$ without loss of generality), where $\Delta = \Gamma^{-1} T$ is the change of basis matrix from the T-basis to the Γ-basis of H given in Lemma 1.

We consider the first component $W_1 = \sum_{j=1}^{n} \Delta_{1k} Y_k$ of $W = \Delta Y$, where the first row $(\Delta_{11}, \ldots, \Delta_{1n})$ of Δ is given by

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

The closeness of the Central Limit approximation of a Normal $\mathrm{N}(0, 10^2)$ of variance $n = 10^2$ for $W_1 = \sum_{j=1}^{n} \Delta_{1k} Y_k$ to of variance $n = 10^2$ is illustrated in Figure 5. □

3.2 Details of the Central Limit Approach to Ring-LWE

We now outline further details justifying the potential use of a Central Limit approach in Ring-LWE. We begin with Proposition 1 concerning 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 states that a good Normal approximation exists if enough of the largest weights $|a_j|$ are of comparable size and is a summary of the Lindeberg condition for a Central Limit Theorem in this case (see Appendix A). In this situation, Proposition 1 is essentially a perturbation of the basic Central Limit Theorem. A comparison with this basic Central Limit Theorem would therefore suggest that Proposition 1 is usually engaged to give a good approximation when as few as about 20 of the largest weights are comparable. We can then extend Proposition 1 in the obvious way to give the multivariate case of Proposition 2.

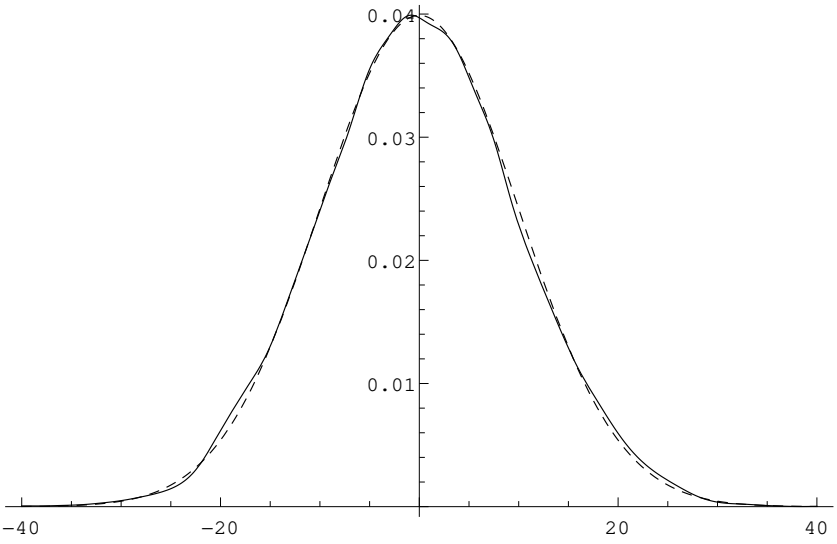

Fig. 5. An empirical density function based on 10^4 realisations of $W_1 = \sum_{j=1}^{100} \Delta_{1k} Y_k$ where $Y_1, \ldots, Y_{100}$ are independent and identically distributed Laplace random variables with variance 1 (solid line) and the density function of the corresponding approximating Normal $\mathrm{N}(0, 10^2)$ distribution (dashed line).

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. □

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 $AX \dot{\sim} \mathrm{N}(0, \rho^2 A A^T)$ can be approximated as a multivariate Normal distribution for moderate or large n. □

Proposition 3 now gives a good Normal approximation for the distribution of a vector expressing the ciphertext in an appropriate basis for decryption in a `SymHom` cryptosystem in most cases of practical interest. 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 a ciphertext in the `SymHom` cryptosystem 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 $C^{(S)}$ this ciphertext in the (decryption) S-basis for H can be approximated as

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

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

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)). It is also the 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 $\Psi' C^{(T)'} \mathrel{\dot\sim} \mathrm{N}(0; \rho^2 \Psi' \Psi'^T)$, and we similarly have $\Psi'' C^{(T)''} \mathrel{\dot\sim} \mathrm{N}(0; \rho^2 \Psi'' \Psi''^T)$. Thus

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

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

3.3 Evaluating or Bounding the Decryption Failure Probability

For a Ring-LWE cryptosystem such as the `SymHom` cryptosystem, the decryption failure probability is directly given by the probability that a vector expressing the ciphertext in an appropriate basis for decryption lies outside a particular box, so motivating Definition 7. For example, the decryption failure probability for a ciphertext vector $C^{(p\Gamma)}$ expressed in the $p\Gamma$-basis for H is the box-out probability function $\overline{\mathbf{B}}_{C^{(p\Gamma)}}(\theta)$ for an appropriate choice of θ.

Definition 7. The *box* $\mathcal{B}(\theta) = \left\{ (v_1, \ldots, v_l)^T \in \mathbb{R}^l \middle| \, |v_1|, \ldots, |v_l| \le \theta \right\}$ is a subset of $\mathbb{R}^l$ defined for $\theta \ge 0$. The *box-in* probability function $\mathbf{B}_V(\theta)$ and the *box-out* probability function $\overline{\mathbf{B}}_V(\theta)$ for a real-valued multivariate random variable $V = (V_1, \ldots, V_l)^T$ on $\mathbb{R}^l$ are given for $\theta \ge 0$ by

$$\begin{aligned} \mathbf{B}_V(\theta) &= \mathbf{P}(V \in \mathcal{B}(\theta)) = \mathbf{P}\left(\max(|V_1|, \ldots, |V_l|) \le \theta\right) \\ \text{and } \overline{\mathbf{B}}_V(\theta) &= \mathbf{P}(V \notin \mathcal{B}(\theta)) = \mathbf{P}\left(\max(|V_1|, \ldots, |V_l|) > \theta\right) = 1 - \mathbf{B}_V(\theta). \end{aligned}$$ □

The box-in probability function and the box-out probability function for an l-dimensional random variable $V = (V_1, \ldots, V_l)^T$ are both simple functions of the distribution function F_V of V, which is given by

$$F_V(v_1, \ldots, v_l) = F_{(V_1, \ldots, V_l)}(v_1, \ldots, v_l) = \mathbf{P}(V_1 \le v_1, \ldots, V_l \le v_l).$$

For example, the box-in probability function is $\mathbf{B}_V(\theta) = F_V(\theta) - F_V(-\theta)$ for a univariate ($l = 1$) random variable V, and for a bivariate ($l = 2$) random variable V we have

$$\mathbf{B}_V(\theta) = \mathbf{B}_{(V_1,V_2)}(\theta) = F_V(\theta,\theta) - F_V(\theta,-\theta) - F_V(-\theta,\theta) + F_V(-\theta,-\theta).$$

More generally, the box-in probability function $\mathbf{B}_Y(\theta)$ is an appropriate sum or difference (on an inclusion/exclusion basis) of distribution function evaluations of the form $F_V(\pm\theta,\ldots,\pm\theta)$, and the box-out probability function $\overline{\mathbf{B}}_V(\theta)$ can also be similarly expressed in terms of distribution function evaluations.

A Central Limit approach to approximate the asymptotic forms of the box-in and box-out probability functions, is a natural and direct approach for the following reasons. In its most basic form the Central Limit Theorem [9] states that if $X_1, X_2, \ldots$ are independent and identically distributed random variables with mean $\mathbf{E}(U_j) = 0$ and variance $\mathrm{Var}(U_j) = 1$, then $Q_l = l^{-1}\sum_{j=1}^{l} U_j$ tends *in distribution* to a standard Normal $\mathrm{N}(0,1)$ distribution. Such "convergence in distribution" is formally a statement about convergence of distribution functions, that is to say that $F_{Q_l} \to \Phi$, where Φ is the distribution function of a standard Normal $\mathrm{N}(0,1)$ random variable. Thus, for example, we can obtain the limiting function for the box-in probability function of Q_l as $\mathbf{B}_{Q_l}(\theta) \to \Phi(\theta) - \Phi(-\theta)$ as $l \to \infty$. More generally, a multivariate Central Limit approach shows that of a box-in or box-out probability function can in principle be expressed to high degree of accuracy in terms of sums and differences of the evaluations of the distribution function of a multivariate Normal random variable. This potentially allows us to evaluate a decryption failure probability for the `SymHom` cryptosystem in practical cases directly (for example by numerical integration) or to obtain good bounds for a a decryption failure probability.

4 The `SymHom` Cryptosystem

We analyse the `SymHom` cryptosystem of Section 8.3 of [13], which is described in Figure 6. In this section, we denote by $[\![r]\!]_q = r - q[q^{-1}r]$ for $r \in \mathbb{Z}$ the coset representative of ($r \bmod q$) nearest to 0, 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 the basis B. Our analysis enables us to obtain good bounds for the probabilities of the incorrect decryption of a degree-1 ciphertext (Theorem 1) and a degree-2 ciphertext (Theorem 2) in the `SymHom` cryptosystem.

4.1 Encryption and Homomorphic Multiplication

We first give a description of the relevant parts of the encryption process of the `SymHom` cryptosystem. The 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 encryption process for a plaintext μ requires us to generate a *Noise*

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 $\overrightarrow{d}$ 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. 6. The `SymHom` Cryptosystem as stated in Section 8.3 of [13]

random variable that is the result of a discretisation process involving μ and some random input. We summarise notation and terminology relating to the Noise in Figure 7.

The first step of the encryption process is to generate a random input for the discretisation process involving 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.

In order to encrypt $\mu \in R_p$, we have to 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* or CRR-discretisation [13, 17] with respect to the $p\Gamma$-basis for H. We denote the discretisation of Y by $Y'(\mu) = \lfloor\!\lfloor Y \rceil\!\rceil^{p\Gamma}_{\sigma(pR^\vee)+\sigma(t^{-1}\mu)}$.

The *Noise* random variable $Y''(\mu)$ in the encryption of the plaintext μ is then defined in [13] 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 $Y'(\mu) = \sigma(Y''(\mu))$ as the *Embedded Noise*, and we note that $Y'(\mu)$ expresses the Embedded Noise in the I-basis of H.

We then form the ciphertext as a linear polynomial over $R_q^\vee$ from the Noise $Y''(\mu)$ that depends on 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 polynomial $C(\theta; \mu)$ is then 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.

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. 7. Notation for the Noise-related quantities used in the `SymHom` encryption of the plaintext μ.

The ciphertext given by the homomorphic product 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 ciphertext polynomial $C(\theta;\mu_1,\mu_2)$ 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 constituent degree-1 ciphertexts, that is to say the constant term in the above formulation of $C(\theta;\mu_1,\mu_2)$. This process extends in the obvious way to give ciphertexts of higher degree.

4.2 Decryption

We specify a decryption process for the `SymHom` cryptosystem using the $p\Gamma$-basis of H (though any appropriate basis can be used), which is essentially that given in Figure 6. We recall (see Figure 2) 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 Noise $Y_1'(\mu_1)$ 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 illustrate the Central Limit approach to evaluating decryption failure probabilities by giving bounds for the incorrect decryption of degree-1 and degree-2 ciphertexts for the SymHom cryptosystem. The results follow from the fact that the SymHom decryption process using the $p\Gamma$-basis for H (for example) 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

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

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 Q is the "Q-function" giving the upper tail probability for a standard Normal $\mathrm{N}(0,1)$ distribution. This tail probability $Q(x)$ is bounded by its asymptotic expansion as $Q(x) \leq (2\pi x^2)^{-\frac{1}{2}}\exp(-\frac{1}{2}x^2)$, and we note that this bound is extremely tight even for moderate values of $x > 0$. 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] (also used in a similar way in Lemma 6.5 of [13]) 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 $\eta_2 = \eta_1(n,q,mp\rho_1\rho_2) = \frac{1}{2}(n^{\frac{1}{2}}mp\rho_1\rho_2)^{-1}q$. $\square$

5 A δ-subgaussian Approach to Ring-LWE Decryption

A δ-subgaussian random variable [15] is a relaxation of a subgaussian random variable [10], which is a random variable with mean 0 that is bounded in a particular technical sense by a Normal random variable. Statistical arguments based on δ-subgaussian random variables have been used in Ring-LWE cryptography for example in [13, 15, 19], and further properties of δ-subgaussian random variables are given in [17]. For completeness, we note that the class of δ-subgaussian random variables includes Normal random variables. By contrast, a Laplace random variable (see Lemma 2), which arises for example when considering a degree-2 `SymHom` ciphertext, is not a δ-subgaussian random variable and has a far heavier tail than that given by a δ-subgaussian tail bound.

The δ-subgaussian approach used to obtain a decryption failure probability bound for `SymHom` cryptosystem given in [13] is described in Figure 8 in terms of

Let V be a real-valued multivariate random variable. The following process is a suggested approach for bounding the box-out probability function $\overline{\mathbf{B}}_V(\theta) = \mathbf{P}(V \notin \mathcal{B}(\theta))$.
• Find a δ-subgaussian random variable V' such that V is "well-approximated" by V'.
• Find a bounding function for the box-out probability function $\overline{\mathbf{B}}_{V'}(\theta) = \mathbf{P}(V' \notin \mathcal{B}(\theta))$ for V' by using the δ-subgaussian properties of V'.
• Use the bounding function for the box-out probability function $\overline{\mathbf{B}}_{V'}(\theta)$ for V' as a bounding function for the box-out probability function $\overline{\mathbf{B}}_V(\theta)$ for V.

Fig. 8. Approach for obtaining a bounding function for a box-out probability function $\overline{\mathbf{B}}_V$ for a random variable V by approximation with a δ-subgaussian random variable.

the equivalent box-out probability function. The idea is to show that a vector expressing a `SymHom` ciphertext random variable in a suitable basis for decryption is close to some δ-subgaussian random variable and then to use a δ-subgaussian tail bound to obtain a bound for the decryption failure probability. We discuss possible issues with this δ-subgaussian approach in Section 5.1 and contrast such a δ-subgaussian approach with a Central Limit approach in Section 5.2.

5.1 Proximity to a δ-subgaussian Random Variable

The δ-subgaussian approach outlined in Figure 8 to evaluating a box-out probability function $\overline{\mathbf{B}}_Y(\theta)$ for a random variable Y, which is used to find a decryption failure probability, is based on finding a δ-subgaussian random variable close to Y. However, Example 2 shows that proximity to a δ-subgaussian random variable is an extremely wide ranging property encompassing many random variables.

Example 2. Suppose $Y_1, \ldots, Y_{2^{2l'}}$ are independent and identically distributed random variables with mean $\mathbf{E}(Y_j) = 0$ and finite variance $\mathrm{Var}(Y_j) = \rho^2$. We consider the orthogonal transformation given by a scaled $2^{2l'} \times 2^{2l'}$ Hadamard matrix S with entries $\pm 2^{-l'}$, so $SS^T = I_{2^{2l'}}$ and S is an orthogonal transformation. The real-valued multivariate random variable $W = SY$ is an orthogonal transformation of Y, with a generic component W_j of $W = (W_1, \ldots, W_{2^{2l'}})^T$ is given by

$$W_j = 2^{-\frac{1}{2}l'} \left((\pm Y_1) + (\pm Y_2) + \ldots + (\pm Y_{2^{l'}}) \right).$$

However, $\pm Y_1, \pm Y_2, \ldots, \pm Y_{2^{2l'}}$ are independent and identically distributed random variables with mean $\mathbf{E}(\pm Y_j) = 0$, so $\mathbf{E}(W) = 0$. Furthermore, the covariance matrix of $W = SY$ is given by

$$\mathrm{Cov}(W) = \mathbf{E}\left(SY(SY)^T\right) = S\mathbf{E}\left(YY^T\right)S^T = \rho^2 S I_{2^{2l'}} S^T = \rho^2 I_{2^{2l'}}.$$

Thus a component W_j of W has variance $\mathrm{Var}(W_j) = \rho^2$ and two distinct components W_j and $W_{j'}$ of W have covariance $\mathrm{Cov}(W_j, W_{j'}) = 0$. The Central Limit Theorem shows that that $W \dot{\sim} \mathrm{N}(0, \rho^2 I_{2^{2l'}})$ is well-approximated by a multivariate Normal $\mathrm{N}(0, \rho^2 I_{2^{2l'}})$ distribution for moderate and large $2^{2l'}$. Thus W is close to a spherically symmetric multivariate Normal distribution, but $Y = S^{-1}W$ is an orthogonal transformation of W, so Y has a distribution that is close to the δ-subgaussian $\mathrm{N}(0, \rho^2 I_{2^{2l'}})$ distribution. □

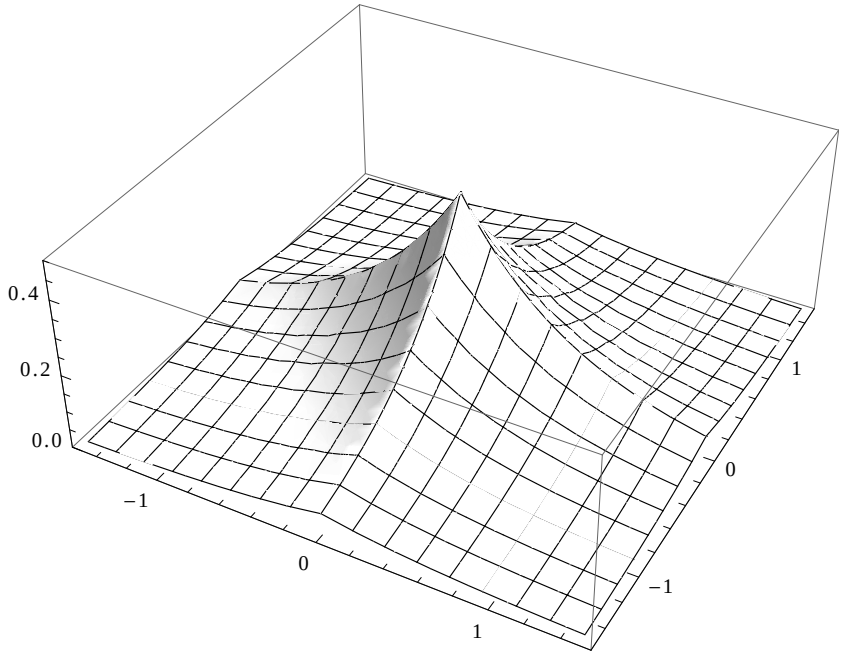

Fig. 9. The joint density function $f_{(Y_1,Y_2)}(y_1, y_2) = \frac{1}{2}\exp(-2^{\frac{1}{2}}(|y_1| + |y_2|))$ of the two independent Laplace random variables Y_1 and Y_2 with mean $\mathbf{E}(Y_j) = 0$, variance $\mathrm{Var}(Y_j) = 1$ and density function $f_{Y_j}(y_j) = 2^{-\frac{1}{2}}\exp(-2^{\frac{1}{2}}|y_j|)$.

Example 2 illustrates that any moderate or high dimension random variable Y with independent and identically distributed components having finite variance can be considered as being close to some multivariate Normal random variable and hence to some δ-subgaussian random variable. However, it would clearly be inaccurate to calculate a box-out probability function $\overline{\mathbf{B}}_Y(\theta)$ using a δ-subgaussian tail bound for a typical component Y_j in a case where Y_j is a very heavy-tailed random variable that is far from satisfying such a δ-subgaussian tail bound. The fundamental issue with the δ-subgaussian approach outlined in Figure 8 is that the $\overline{\mathbf{B}}_Y(\theta)$ can be far from invariant under orthogonal transformations of a spherically symmetric random variable Y. More generally $\overline{\mathbf{B}}_Y(\theta)$ does not transform in an appropriate way under linear transformations of Y to allow the universal use of such a "δ-subgaussian proximity" approach.

5.2 Contrast with the Central Limit Approach

Figure 9 is a two dimensional illustration of the general pyramidal shape arising as the Laplace density function for a real-valued high-dimensional random vector expressing a degree-2 `SymHom` ciphertext. Such a density function consists of ridges along the axes of the T-basis for H, with corries or depressions between the ridges. Univariate marginal random variables in directions "along" or "close

to" to the ridges cannot be approximated by a δ-subgaussian random variable. By contrast, univariate marginal random variables in directions "away" from the ridges that go "through the corries" can be approximated by a Normal random variable. Furthermore, in high dimensions, "most" directions (loosely speaking) go through the corries and stay away from the ridges, so in general univariate marginal random variables can be approximated by a Normal random variable. Such an argument can loosely be thought of as the geometric expression of a Lindeberg Central Limit result, such as Propostion 3. Thus Figure 9 illustrates that the density function of a vector whose independent components have finite variance is close to that of some Normal random variable "away" from the ridges. Furthermore, we also note that these comments also apply to higher degree `SymHom` ciphertext vectors, whose density functions have the same general pyramidal shape but with more pronounced ridges.

In addressing `SymHom` decryption, we have to approximate the distribution of the vector $C^{(S)} = \Psi C^{(T)} = ST^{-1}C^{(T)}$ expressing a `SymHom` ciphertext in an appropriate decryption S-basis in the situation of Proposition 3. The δ-subgaussian approach is always to approximate each component $C_j^{(S)} = \sum_{j=1}^n \Psi_{jk} C_k^{(T)}$ of $C^{(S)}$ as a δ-subgaussian random variable, even though such an approximation may not be valid, for example "along the ridges". More generally, Proposition 3 and the associated Lindeberg Central Limit theory show that the natural approach for *any* distribution of $C_k^{(T)}$ with finite component variance is to approximate the distribution of $C_j^{(S)} = \sum_{j=1}^n \Psi_{jk} C_k^{(T)}$ as a Normal distribution for appropriate weights Ψ_{jk}. Indeed, this theory makes it difficult to conceive of a general situation in which the entries Ψ_{jk} of the matrix Ψ would mean that the distribution of the ciphertext vector $C^{(S)}$ in the decryption S-basis for H could be approximated by some δ-subgaussian distribution, but that $C^{(S)}$ could not be approximated by the Normal $\mathrm{N}(0; \Psi\Psi^T)$ distribution in the manner of Proposition 3.

Another advantage of the Central Limit approach is that is allows the use of better tail bounds than a δ-subgaussian approach. For example, Theorem 1 gives a decryption failure probability bound for a degree-1 `SymHom` cryptosystem ciphertext of $\dfrac{2n\exp(-\frac{1}{2}\eta_1^2)}{(2\pi)^{\frac{1}{2}}\eta_1}$ for moderate or large $\eta_1(n,q,\rho) = \frac{1}{2}(n^{\frac{1}{2}}\rho)^{-1}q$, which is tighter than the equivalent δ-subgaussian decryption failure probability bound of $2n\exp(-\frac{1}{2}\eta_1^2)$ obtained by using the tail bound of [17, Lemma 18] in the manner of [13, Lemma 6.5]. More generally, such a Central Limit framework is particularly suited for making a concrete determination of or finding a good bound for a decryption failure probability in a Ring-LWE cryptosystem.

In summary, the Central Limit approach is to be preferred to a δ-subgaussian approach. The Central Limit approach has a basic theoretical foundation, gives a specified Normal approximating distribution (addressing component correlations), and allows the use of better tail bounds than a δ-subgaussian approach.

References

1. P. Billingsley. *Probability and Measure*. Wiley, third edition, 1995.
2. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) Fully Homomorphic Encryption without Bootstrapping. In *Innovations in Theoretical Computer Science 2012*, pages 309–325. ACM, 2012.
3. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical Hardness of Learning with Errors. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th Annual ACM Symposium on Theory of Computing*, 2013.
4. J. H. Cheon, A. Kim, M. Kim, and Y. S. Song. Homomorphic Encryption for Arithmetic of Approximate Numbers. In T. Takagi and T. Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017*, volume 10624 of *LNCS*, pages 409–437. Springer, 2017.
5. I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène. Faster Fully Homomorphic Encryption: Bootstrapping in Less Than 0.1 Seconds. In J. H. Cheon and T. Takagi, editors, *Advances in Cryptology - ASIACRYPT 2016*, volume 10031 of *LNCS*, pages 3–33. Springer, 2016.
6. J. Fan and F. Vercauteren. Somewhat Practical Fully Homomorphic Encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.
7. C. Gentry. Fully Homomorphic Encryption using Ideal Lattices. In M. Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, ACM, pages 169–178, 2009.
8. C. Gentry, A. Sahai, and B. Waters. Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based. In R. Canetti and J.A. Garay, editors, *Advances in Cryptology - CRYPTO 2013*, volume 8042 of *LNCS*, pages 75–92. Springer, 2013.
9. G. Grimmett and D. Stirzaker. *Probability And Random Processes*. Oxford University Press, 3rd edition, 2001.
10. J. Kahane. Propriétés locales des fonctions à séries de Fourier aléatoires. *Studia Mathematica*, 19:1–25, 1960.
11. S. Kotz, T. Kozubowski, and K. Podórski. *The Laplace Distribution and Generalizations*. Birkhüser, 2001.
12. V. Lyubashevsky, C. Peikert, and O. Regev. On Ideal Lattices and Learning with Errors Over Rings. *IACR Cryptology ePrint Archive*, 2012:230, 2012.
13. V. Lyubashevsky, C. Peikert, and O. Regev. A Toolkit for Ring-LWE Cryptography. *IACR Cryptology ePrint Archive*, 2013:293, 2013.
14. V. Lyubashevsky, C. Peikert, and O. Regev. A Toolkit for Ring-LWE Cryptography. In T. Johansson and P. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, 2013.
15. D. Micciancio and C. Peikert. Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller. In D. Pointcheval and T. Johansson, editors, *Eurocrypt 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, 2012.
16. D. Micciancio and O. Regev. Lattice-based Cryptography. In D.J. Bernstein and J. Buchmann and E. Dahmen, editor, *Post-Quantum Cryptography*, pages 147–191. Springer, 2009.
17. S. Murphy and R. Player. δ-subgaussian Random Variables in Cryptography. In J. Jang-Jaccard and F. Guo, editors, *ACISP 2019: The 24th Australasian Conference on Information Security and Privacy,* to appear. Available as *IACR Cryptology eprint Archive 2017:698*, 2019.

18. C. Peikert. Public-Key Cryptosystems from the worst-case Shortest Vector Problem. In M. Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Comuting*, 2009.
19. C. Peikert. Lattice Cryptography for the Internet. In M. Mosca, editor, *PQCrypto 2014*, volume 8772 of *LNCS*, pages 197–219, Springer, 2014.
20. C. Peikert. A Decade of Lattice Cryptography. *IACR Cryptology ePrint Archive*, 2015:939, 2016.
21. O. Regev. On Lattices, Learning with Errors, Random Linear Codes and Cryptography. In H. Gabow and R. Fagin, editors, *37th Annual ACM Symposium of Theory of Computing*, 2005.
22. O. Regev. The Learning with Errors Problem (Invited Survey). In *IEEE Conference on Computational Complexity*, pages 191–204, 2010.
23. D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient Public Key Encryption Based on Ideal Lattices. In M. Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 617–635, 2009.

Acknowledgements

Rachel Player was 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).

A The Lindeberg Central Limit Theorem

The basic univariate Central Limit Theorem [9] is concerned with the distribution of the sum $\sum_{j=1}^{n} X_j$ of independent and identically distributed random variables $X_1, \ldots, X_n$. When considering the distribution of a component of vector expressing the ciphertext in a Ring-LWE cryptosystem such as `SymHom` cryptosystem, we need to consider the distribution of the weighted sum $\sum_{j=1}^{n} a_j X_j$ for an appropriate choice of weights $a_1, \ldots, a_n$. To obtain a Normal approximation for such a weighted sum $\sum_{j=1}^{n} a_j X_j$ we need a more general form of the Central Limit Theorem formally given by the Lindeberg condition [1]. Such a result is informally summarised in Proposition 1, as stated in Section 3.2.

Proposition 1. Suppose $X_1, X_2, \ldots$ are independent and identically distributed random variables with mean $\mathbf{E}(X_j) = 0$ and finite variance $\mathrm{Var}(X_j) = \sigma^2$. For weights $a = (a_1, \ldots, a_n)$, the weighted sum

$$\sum_{j=1}^{n} a_j X_j \mathrel{\dot{\sim}} \mathrm{N}(0, |a|^2\sigma^2) \qquad \text{for moderate or large } n$$

has an approximate Normal distribution provided that the weights $a_1, \ldots, a_n$ are not dominated by just a few weights. □

The informal summary of the Central Limit Theorem for a weighted sum given by Proposition 1 can be justified by such a Lindeberg Central Limit Theorem as stated for a weighted sum in Lemma 3. We also give such a result for the particular case of the Laplace distribution in Lemma 4.

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}. 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 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 $X_j^{(\alpha)} = \begin{cases} X_j & [|X_j| > \alpha] \\ 0 & [|X_j| \le \alpha] \end{cases}$ 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]. □

Lemma 4. Suppose $X_1, X_2, \ldots$ are independent Laplace random variables with mean $\mathbf{E}(X_j) = 0$ and variance $\mathrm{Var}(X_j) = 1$, so have common density function $f_{X_j}(x_j) = 2^{-\frac{1}{2}} \exp(-2^{\frac{1}{2}}|x_j|)$. For constants $a_1, a_2, \ldots$, define the functions $a(l)^2 = \sum_{j=1}^{l} a_j^2$ and $\widetilde{a}_j(l) = \dfrac{|a_j|}{a(l)}$. If $\sum_{j=1}^{l} \exp\left(-\epsilon \widetilde{a}_j(l)^{-1}\right) \to 0$ as $l \to \infty$ for any given $\epsilon > 0$, then $a(l)^{-1} \sum_{j=1}^{l} a_j X_j$ tends in distribution to a standard Normal $\mathrm{N}(0,1)$ as $l \to \infty$. □

Proof. When X_j has a Laplace distribution, the function Ψ_{X_j} of the Lindeberg condition of Lemma 3 evaluates to

$$\begin{aligned} \Psi_{X_j}(\theta) &= \int_{\theta}^{\infty} x^2 f_{X_j}(x)\, dx = \int_{\theta}^{\infty} x^2 \, 2^{-\frac{1}{2}} \exp(-2^{\frac{1}{2}}|x|)\, dx \\ &= (\theta^2 + 2^{\frac{1}{2}}\theta + 1)\exp(-2^{\frac{1}{2}}\theta). \end{aligned}$$

In this case, the Lindeberg condition of Lemma 3 is therefore given by

$$\begin{aligned}\sum_{j=1}^{l} \widetilde{a}_j(l)^2 \, \Psi_{X_j}\left(\frac{\epsilon}{\widetilde{a}_j(l)}\right) &= \sum_{j=1}^{l} \widetilde{a}(l)^2 \left(\frac{\epsilon^2}{\widetilde{a}_j(l)^2} + \frac{2^{\frac{1}{2}}\epsilon}{\widetilde{a}_j(l)} + 1\right) \exp\left(-\frac{2^{\frac{1}{2}}\epsilon}{\widetilde{a}_j(l)}\right) \\ &= \sum_{j=1}^{l} \left(\epsilon^2 + 2^{\frac{1}{2}}\epsilon\widetilde{a}_j(l) + \widetilde{a}_j(l)^2\right) \exp\left(-\frac{2^{\frac{1}{2}}\epsilon}{\widetilde{a}_j(l)}\right).\end{aligned}$$

We recall that $0 \leq \widetilde{a}_j(l) \leq 1$, so the convergence of this sum depends only on the exponential term. Upon re-scaling ϵ, we thus obtain a Lindeberg's condition of $\sum_{j=1}^{l} \exp\left(-\epsilon\widetilde{a}_j(l)^{-1}\right) \to 0$ as $l \to \infty$ for convergence of $a(l)^{-1}\sum_{j=1}^{l} a_j X_j$ to a standard Normal $\mathrm{N}(0,1)$ distribution. □