Thorns in Polynomial Convolution: Correlation, Large Deviations, and Applications

Dongshu Cai[1,2], Yijian Liu[1,2], Jiabo Wang[1,2](✉), and Xianhui Lu[1,2]

[1] Key Laboratory of Cyberspace Security Defense, Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China;
[2] School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China
{caidongshu, wangjiabo}@iie.ac.cn

Abstract. When estimating the decryption failure rate (DFR) of structured lattice-based cryptography, some schemes implicitly assume that the coefficients of the decryption noise are independent. In practice, however, the decryption noise typically contains terms arising from convolutions of small polynomials, which introduce correlations among coefficients. These correlations can create a non-negligible gap between independence-based estimates and empirical failure rates, leading to underestimated DFRs, overestimated security levels, and exploitable attack surfaces. They also obscure the effect of error-correcting mechanisms in structured lattice-based encryption designs. To date, there has been no practical framework for characterizing such correlations.

In this paper, we give the first systematic characterization of correlations among the coefficients of convolved polynomials with Gaussian coefficients, using the canonical embedding as the central viewpoint. We establish large-deviation results for the coefficients of the resulting polynomial. Our analysis shows that, as the norm grows, convolutional polynomials asymptotically concentrate near a finite set of fixed two-dimensional planes. This gives rise to directional tail structures in the n-dimensional joint probability density, which we call *thorns*.

As a direct application, we prove that existing decryption-failure attacks succeed precisely by forcing the noise to lie on these thorns. This phenomenon endows the noise with extremely strong correlations, ultimately triggering decryption failures. Furthermore, adopting the canonical embedding perspective allows us to comprehensively illustrate how the independence assumption distorts the true noise distribution. We prove that the independence assumption systematically underestimates the noise norm, and we derive an analytic expression for the probability density function of the Euclidean norm of the decryption noise.

Keywords: Decryption failure rate · Convolution · Canonical embedding.

1 Introduction

In the field of post-quantum cryptography, lattice-based primitives stand out prominently and have attracted widespread attention throughout standardiza-

tion efforts, exemplified by NIST-selected CRYSTALS-Kyber for key encapsulation [2]. Despite years of optimization, lattice-based encryption schemes still incur significantly greater bandwidth and storage overhead than classical cryptosystems such as ECC.

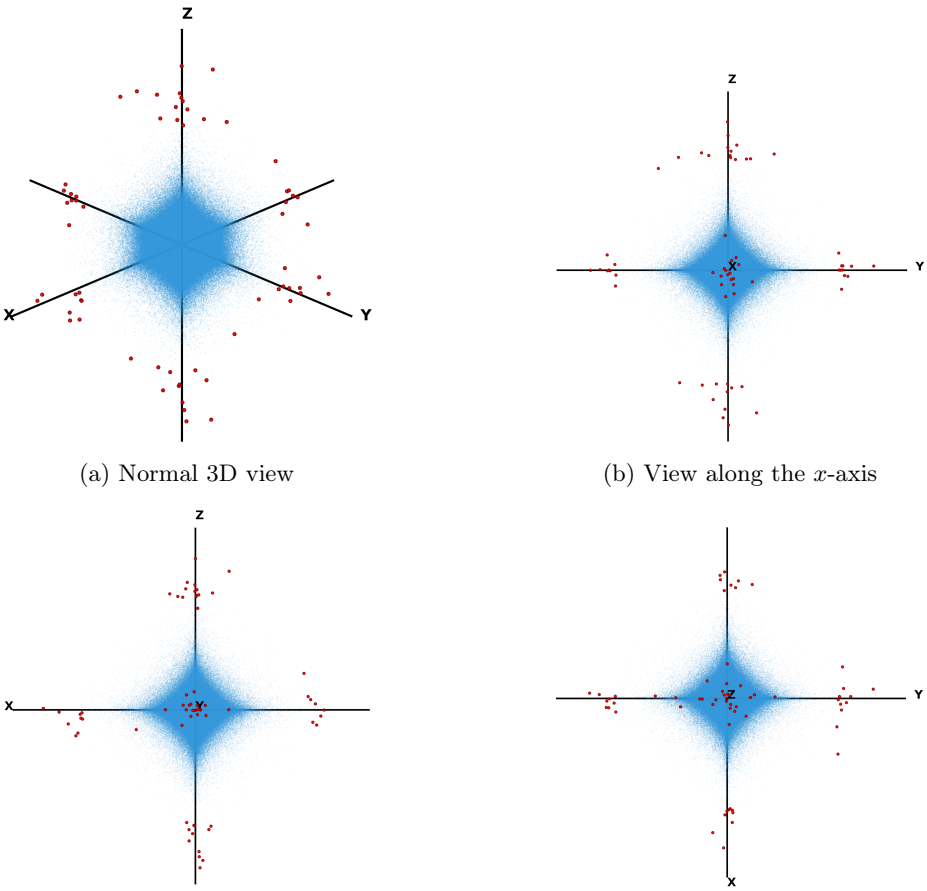
Such overhead mainly stems from polynomial convolution components within decryption noise, which cause noise expansion. To guarantee correct decryption, practitioners are forced to adopt relatively large moduli, further increasing overall ciphertext and key sizes. To mitigate this issue, a practical approach is to equip encryption schemes with stronger error-correcting structures, enabling the use of smaller moduli while maintaining a satisfactory decryption failure rate. Representative instantiations include LAC [19] and DAWN [17]. Also, in FHE, several implementations use the average-case noise analysis with the independence assumption (e.g., BFV [1] and HElib [5]).

Nevertheless, when error-correcting mechanisms are introduced, the inherent correlation among noise components arising from polynomial convolution invalidates the simple, explicit decryption failure-rate estimation used in Kyber. Statistically characterizing such a coefficient correlation remains challenging. Prior work by D'Anvers et al. has made extensive explorations along this line, yet has not revealed the fundamental origin of such a correlation [6,8,7]. We note that the decryption failure not only affects functional correctness but also compromises security under the chosen-ciphertext attack (CCA) model. The ciphertext that fails to decrypt would leak information about the secret key, and there have been a lot of failure-based CCA attacks on lattice-based KEMs [13,14,10,9,11] and FHE schemes [15,3,4,18]. Therefore, current CCA lattice-based KEMs always limit the lower bound of decryption failure rate (DFR) as 2^{-128} (e.g., $< 2^{-139}$ in Kyber [2]).

Conversely, blindly adopting the component-wise independence assumption in practical design deviates from the real noise property and inevitably underestimates the actual decryption failure probability, thereby undermining cryptographic security. A typical instance is the decryption-failure attack against the error-correction-aided LAC proposed by Guo et al. [13]. Some works have studied this effect; for example, [8] proposed a correlation-aware method to estimate the DFR of LAC more accurately, but it still relies on an assumption. Guo and Zheng analyzed the effects of higher moments (kurtosis) arising from correlations and reported a non-negligible bias, which was supported by experiments. A very recent work [16] analyzed the correlation using the ℓ_2 norm decomposition (essentially the autocorrelations) and proposed a generic framework to mitigate its effect via rejection sampling, though its estimate still relies on heuristics. Therefore, accurately characterizing the correlation between decryption noise and enabling reliable estimation of the decryption failure rate have become essential research problems in lattice-based cryptography.

1.1 Contributions

In previous studies, when analyzing the correlation among coefficients of polynomial convolution outputs, researchers commonly expanded all coefficients di-



(a) Normal 3D view

(b) View along the x-axis

(c) View along the y-axis

(d) View along the z-axis

Fig. 1: Hadamard product of two three-dimensional Gaussian vectors. The four panels show the normal 3D view and projections along the three coordinate axes. These samples visually exhibit distinct thorn-shaped structures. We mark the top 64 samples with the largest norms using red dots, all of which approximately distribute along the directions of these thorns.

rectly in terms of the coefficients of the original polynomials. Since different coefficients of the convolution result are computed from identical variables with distinct matching rules, it has been widely acknowledged that convolution inevitably introduces correlations among the resulting coefficients. Nevertheless, analyzing the explicit correlation characteristics via such variable decomposition approaches remains highly intricate.

Instead of adopting this variable splitting perspective, we investigate this problem from the viewpoint of canonical embeddings. For the first time, we clarify the essential origin of such correlation. While components of the product of

two Gaussian polynomials stay independent in the embedding space, they obey non-spherical distributions and form prominent thorn structures, as shown in Figure 1. Such structural features are further converted into coefficient correlation on polynomial rings via inverse canonical embedding.

Based on the above observations, we conduct comprehensive analyses on polynomial convolution from multiple perspectives, and our main contributions are summarized as follows:

First, We establish large deviation results on the angular distribution of vectors obtained by convolving two polynomials with independent Gaussian coefficients. We prove that such convolution vectors asymptotically converge to several fixed thorn directions as the vector scale increases, which endows decryption noise with strong inherent correlation.

Second, We revisit existing decryption-failure-based CCA attacks, and verify that the feasibility of these attacks essentially relies on noise samples falling close to the aforementioned thorns and exhibiting strong correlation. Meanwhile, we propose a general approach to generalize these attack strategies conveniently.

Third, We derive the probability density function of the ℓ_2-norm of decryption noise for the first time. Experimental and theoretical results confirm that the norm values calculated under the independence assumption systematically underestimate the true noise magnitude.

In addition, we pinpoint the sole essential discrepancy between the independence assumption and the actual noise distribution. Specifically, the independence assumption only distorts the marginal distribution of component magnitudes under canonical embedding, while keeping other statistical properties unchanged.

1.2 Techniques

The decryption noise of structured lattice-based cryptography is generally composed of the sum of multiple components, which can be written as

$$\sum_i s_i \cdot e_i + \sum_j e'_j \in \mathcal{R},$$

where s_i, e_i, e'_j are small-coefficient polynomials over the ring $\mathcal{R}$. A typical example is the Ring-LWE encryption scheme presented in the **Applications** section of Section 1.1 in [20], whose decryption error takes the form $r \cdot e - s \cdot e_1 + e_2$.

For the noise decomposition $\sum_i s_i \cdot e_i + \sum_j e'_j$, each entry of every e'_j is sampled in an i.i.d. manner and introduces no correlation among polynomial coefficients. Consequently, the coefficient correlation of decryption noise originates entirely from the term $\sum_i s_i \cdot e_i$. Since each summand in this sum is mutually independent, we focus our analysis on a single representative term denoted by $s \cdot e$.

Cryptographic schemes based on both NTRU and algebraic LWE adopt a decryption framework similar to the Ring-LWE construction in [20]. This yields a clear paradigm for the distributions of s and e. In [20], e is required to follow a spherical Gaussian distribution in the embedded space of the underlying number

field, with mutually independent entries. When expressed in polynomial form, the coefficients of e are continuous; for the number field $\mathbb{Q}[x]/(x^{2^k}+1)$, each polynomial coefficient of e obeys an i.i.d. Gaussian distribution.

Lyubashevsky et al. further provided several useful technical tools in [21]. Specifically, Lemma 2.23 justifies the rounding operation on e, and Lemma 2.24 guarantees that s can be independently sampled from the same distribution as e, while preserving the provable security of the resulting cryptosystem. Although s and e are integer-coefficient polynomials in practical implementations, they still inherit essential characteristics of the spherical Gaussian model in [20], which underpins the security of concrete lattice-based schemes.

To simplify theoretical analysis, we omit the rounding operation described in Lemma 2.23 of [21], while retaining the setup from Lemma 2.24 that polynomials s and e follow identical distributions. Accordingly, when analyzing coefficient correlation of the product $s \cdot e$, we assume both s and e follow i.i.d. continuous spherical Gaussian distributions in the canonical embedding space.

We acknowledge that the discrete distribution of e adopted in practical schemes is not strictly equivalent to the ideal spherical Gaussian model used throughout our analysis. Nevertheless, the reduced canonical embedding of the cyclotomic field $\mathbb{Q}[x]/(x^{2^k}+1)$ we consider forms a unitary transformation, which introduces neither cross-covariance nor pseudo-covariance between distinct frequency coordinates; this property can be observed from Lemma 1. Combined with the multivariate central limit theorem, the embedding components of discrete sub-Gaussian coefficients asymptotically obey i.i.d. circularly symmetric complex Gaussian distributions. To further validate the soundness of our theoretical claims, we supplement numerical experiments demonstrating that when s and e are drawn from discrete sub-Gaussian distributions, the coefficients of $s \cdot e$ exhibit identical correlation behavior as observed under the spherical Gaussian setting.

We now formally characterize the non-independent behavior among the coefficients of $s \cdot e$. Let $\sigma(s), \sigma(e) \in \mathbb{C}^n$ follow the spherical Gaussian condition specified in Lemma 5.16 of [20]. After removing the conjugate part, $\sigma(s)$ and $\sigma(e)$ are i.i.d. complex Gaussian vectors[3]. A key property reads

$$\sigma(s \cdot e) = \sigma(s) \odot \sigma(e) = \big(\sigma_i(s) \cdot \sigma_i(e)\big)_i,$$

where $\odot$ denotes the Hadamard product. Since each entry of $\sigma(s)$ and $\sigma(e)$ is independent, the entries of $\sigma(s \cdot e)$ remain mutually independent in the embedded space.

Each component $\sigma_i(s) \cdot \sigma_i(e)$ of $\sigma(s \cdot e)$ is the product of two independent Gaussian variables, and such a product no longer follows a Gaussian distribution. The probability density function of a standard Gaussian decays as $e^{c|x|^2}$, whereas the tail of the product distribution asymptotically behaves like $e^{c|x|}$ far from the origin. This implies that $\sigma_i(s) \cdot \sigma_i(e)$ possesses heavier tails than the Gaussian

[3] Real and complex embeddings are each i.i.d. with identical standard deviation, and all embedding components are mutually independent.

counterpart under the same standard deviation. As a result, the joint distribution of $\sigma(s \cdot e)$ exhibits heavy tails concentrated around the coordinate axes, forming directional tail patterns we refer to as *thorns*.

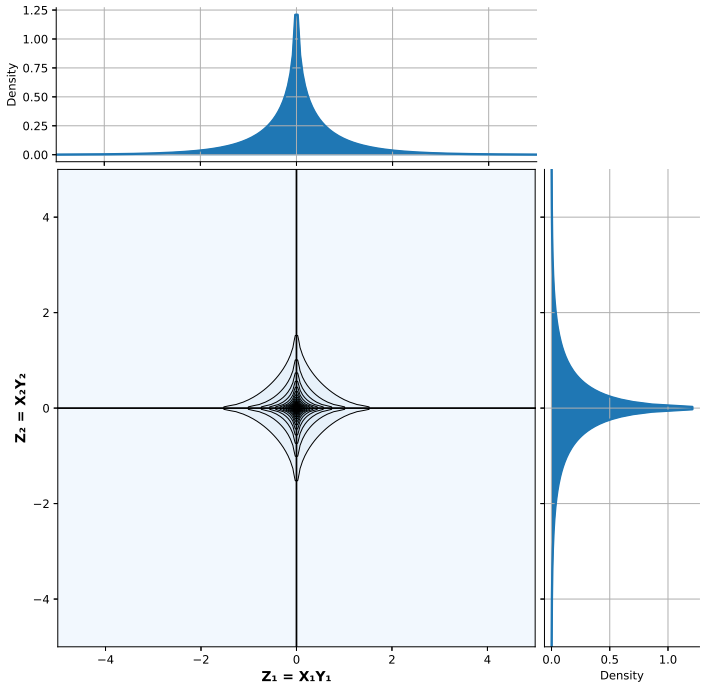


Fig. 2: Joint probability density and marginal densities of the Hadamard product of two 2-dimensional spherical Gaussian vectors. This distribution exhibits heavier tails than a standard Gaussian distribution.

Figure 2 plots the probability density function of (X_1Y_1, X_2Y_2), where X_1, X_2, Y_1, Y_2 are i.i.d. Gaussian variables. Although $Z_1 = X_1Y_1$ and $Z_2 = X_2Y_2$ are independent, their joint distribution clearly displays the aforementioned thorn structure. Similar thorn phenomena persist in higher dimensions. Figure 1 illustrates the corresponding three-dimensional sample distribution with a sufficiently large number of samples; the subsequent three subfigures are three orthogonal views of the original plot. We mark samples with relatively large norm using red dots, which demonstrates that points far from the origin tend to concentrate near the three coordinate axes. Analogous conclusions hold for complex Gaussian cases.

In summary, as the ℓ_2-norm increases, samples of $\sigma(s \cdot e)$ asymptotically concentrate near the coordinate axes in the embedded space with high probability. We formalize this claim for complex Gaussian variables in Corollary 2.

Denote the coefficient vectors of polynomials s and e by $\boldsymbol{s}$ and $\boldsymbol{e}$, respectively, and write the polynomial multiplication $s \cdot e$ as the convolution $\boldsymbol{s} * \boldsymbol{e}$. Since the canonical embedding is a linear transformation acting on coefficient vectors, we write $\boldsymbol{s}V = \sigma(s)$. As shown above, the entries of $\sigma(s \cdot e)$ are independent but non-spherical in the embedded space. The correlation among polynomial coefficients is therefore introduced by the inverse embedding transformation V^{-1}.

For the number field $\mathbb{Q}[x]/(x^n+1)$ with n a power of two, the matrix $V/\sqrt{n}$ is unitary. This means that any vector in the embedded space can be transformed into the polynomial coefficient domain by scaling by $1/\sqrt{n}$ followed by a unitary rotation. Consequently, the thorn directions in the embedded space are directly inherited by the coefficient space. Notably, each coordinate axis in the complex embedded space corresponds to a two-dimensional plane in the coefficient domain. Hence, as its ℓ_2-norm grows, the convolution result $\boldsymbol{s} * \boldsymbol{e}$ asymptotically concentrates near one of these fixed two-dimensional planes with high probability. We rigorously state this result in Theorem 2.

We show that existing decryption-failure CCA attacks precisely exploit the correlation properties demonstrated above. We take the attack on LAC proposed by Guo et al. [13] as an example, and again focus our analysis on the term $s \cdot e$. Triggering a decryption failure requires the error vector to have a relatively large ℓ_2-norm. As established earlier, samples of $\sigma(s \cdot e)$ with large ℓ_2-norm concentrate near the coordinate axes. Without loss of generality, suppose $\sigma_i(s \cdot e)$ is large while all other components remain at a normal magnitude; this condition can be satisfied simply by letting $\sigma_i(s)$ and $\sigma_i(e)$ be large respectively.

In [13], Guo et al. constructed two corresponding pattern pairs for e_1 and the secret key s: the first pair consists of e_1 bearing a long consecutive run of all 1 or all -1 (with optional zeros inserted) and s satisfying a large absolute sum $|\sum_{i=0}^{n-1} s_i|$; the second pair consists of e_1 following an alternating pattern 1 on even indices and -1 on odd indices (or vice versa, with optional zeros inserted) and s achieving a large combined even–odd absolute sum $|\sum_{i=0}^{n/2-1} s_{2i}| + |\sum_{i=0}^{n/2-1} s_{2i+1}|$. One can directly verify that these two paired families exactly render $\sigma_1(e)$ (resp. $\sigma_{n/2}(e)$) and $\sigma_1(s)$ (resp. $\sigma_{n/2}(s)$) large. Conversely, we prove that large values in these canonical embeddings correspond precisely back to the two paired pattern families above.

As observed in practice, finding such patterns requires only moderate computational effort, and the resulting samples typically have large ℓ_2-norm, making decryption failures far more likely to occur.

This phenomenon can be interpreted from two complementary perspectives. First, from the viewpoint of correlation: when $s \cdot e$ has a large ℓ_2-norm, it lies near fixed low-dimensional planes. Thus, enlarging only one component of $s \cdot e$ will cause other polynomial coefficients to increase simultaneously due to non-independence, since these planes are not aligned with the coordinate axes of the coefficient domain. Second, from the viewpoint of ℓ_2-norm concentration:

the distribution of $\sigma(s \cdot e)$ inherently exhibits thorns (as illustrated in Figure 2 and Figure 1). Although large-norm samples are restricted to the vicinity of these thorns, sampling such samples does not require significant computational resources.

Thus, our theoretical framework enables a straightforward generalization of the attack proposed in [13]. Specifically, by finding an index i such that both $|\sigma_i(s)|$ and $|\sigma_i(e)|$ are large, one can achieve a decryption failure probability significantly higher than the value predicted under the independence assumption, thereby making the attack effective.

From the perspective of the canonical embedding, we obtain additional insights that advance the understanding of decryption noise distributions. For the number field $\mathbb{Q}[x]/(x^n + 1)$, the canonical embedding and coefficient vectors are closely related. This correspondence allows us to easily derive the ℓ_2-norm of the coefficient vector from that of the embedded representation of $s \cdot e$. Since the components of $s \cdot e$ in the canonical embedding are independent and identically distributed, we can obtain an analytical expression for the distribution of the ℓ_2-norm of the coefficient vector $\boldsymbol{s} * \boldsymbol{e}$. Such a derivation would be significantly more challenging without the canonical embedding framework.

The canonical embedding also enables a clear comparison between the independence assumption and the true distribution. Under this embedding, the entries of $\sigma(s \cdot e)$ are already independent, and the magnitude and phase of each component are also mutually independent. The independence assumption implicitly treats $\sigma(s \cdot e)$ as an i.i.d. complex Gaussian vector. By contrast, the only deviation from the true distribution lies in the marginal distribution of the magnitudes of these components. Under practical parameter choices, this mischaracterization consistently underestimates the magnitude of each component of $\sigma(s \cdot e)$.

1.3 Not Only Lattice

Our proposed correlation analysis is not limited to lattice-based cryptography, and its core methodology is applicable to arbitrary polynomial quotient rings $\mathbb{Z}[x]/\langle f(x)\rangle$. For instance, the code-based cryptosystem HQC adopts the quotient ring $\mathbb{Z}[x]/\langle x^p - 1\rangle$. In this scenario, the canonical embedding perspective adopted in this paper can be naturally transformed into the discrete Fourier transform (DFT) perspective. The frequency domain of $\mathbb{Z}[x]/\langle x^p - 1\rangle$ under DFT corresponds to the embedding space in the canonical embedding framework. Accordingly, analogous correlation conclusions can be derived for the product of two Gaussian polynomials over $\mathbb{Z}[x]/\langle x^p - 1\rangle$.

In particular, the DFT matrix is a fixed scalar multiple of a unitary matrix, which shares the same property with the canonical embedding matrix of the number field $\mathbb{Z}[x]/\langle x^n + 1\rangle$ investigated in this work. This implies that time-domain independent Gaussian variables over $\mathbb{Z}[x]/\langle x^p - 1\rangle$ remain independent in the frequency domain, which greatly facilitates quantitative correlation analysis.

Guo et al. also constructed CCA attacks against HQC [12], which essentially exploit the thorn structure of decryption noise to trigger decryption failures in

the same manner. The main distinction lies in the coefficient range constraint: in lattice-based schemes, the coefficients of the product of two small polynomials rarely exceed the modulus q, while HQC employs modulus 2. As a result, all components forming the thorn structure are taken modulo 2, which may weaken the inherent correlation to a certain extent.

2 Preliminaries

For $\mathbf{x} \in \mathbb{R}^n/\mathbb{C}^n$ and $p \in [1, \infty]$, the ℓ_p-norm is $\|\mathbf{x}\|_p = \left(\sum_{i\in[n]} |x_i|^p\right)^{1/p}$ for $p < \infty$, and $\|\mathbf{x}\|_\infty = \max_{i\in[n]} |x_i|$, where we typically omit the subscript for $p = 2$ and write $\|\mathbf{x}\|$ instead.

For any vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n/\mathbb{C}^n$, the Hadamard product (element-wise product) $\mathbf{x} \odot \mathbf{y}$ is given by $(\mathbf{x} \odot \mathbf{y})_i = x_i y_i$ for all $i \in [n]$, where we use $\odot$ exclusively to denote this operation throughout this paper.

2.1 Gaussian

A zero-mean circularly symmetric complex Gaussian variable $Z \sim \mathcal{CN}(0, \sigma^2)$ can be written as $Z = X + iY$ where $X, Y \sim \mathcal{N}(0, \sigma^2/2)$ are independent real Gaussians. Throughout this work, we only consider zero-mean Gaussian and complex Gaussian distributions.

If $X_1, X_2 \sim \mathcal{N}(0, \sigma^2)$ are i.i.d. real Gaussians, their product $P = X_1 X_2$ follows a product-normal distribution with density

$$f_P(t) = \frac{1}{\pi\sigma^2} K_0\left(\frac{|t|}{\sigma^2}\right),$$

where K_0 denotes the modified Bessel function of the second kind of order zero. It has no closed-form expression in terms of elementary functions, and is defined by the integral representation

$$K_0(z) = \int_0^\infty e^{-z\cosh t}\, dt, \quad z > 0.$$

It satisfies the asymptotic approximations

$$K_0(z) \sim -\ln\left(\frac{z}{2}\right) - \gamma \quad (z \to 0^+), \qquad K_0(z) \sim \sqrt{\frac{\pi}{2z}} e^{-z} \quad (z \to \infty),$$

where $\gamma \approx 0.5772$ is the Euler–Mascheroni constant.

For i.i.d. zero-mean complex Gaussians $Z_1, Z_2 \sim \mathcal{CN}(0, \sigma^2)$, their product $W = Z_1 Z_2$ has the distribution

$$f_W(z) = \frac{2}{\pi\sigma^4} K_0\left(\frac{2|z|}{\sigma^2}\right), \qquad z \in \mathbb{C}.$$

The distribution of W can be fully decomposed into statistically independent modulus and argument components:

– The modulus $|W| = |Z_1| \cdot |Z_2|$ has probability density function

$$f_{|W|}(r) = \frac{4r}{\sigma^4} K_0\left(\frac{2r}{\sigma^2}\right), \quad r \geq 0.$$

– The argument $\arg(W) = (\arg(Z_1) + \arg(Z_2)) \bmod 2\pi$ is uniformly distributed over $[0, 2\pi)$.

2.2 Algebraic Number Theory Background

Let $K = \mathbb{Q}(\zeta)$ be a number field of degree n, and let $f(x) \in \mathbb{Q}[x]$ be the minimal polynomial of ζ with degree n. Then we have the isomorphism $K \simeq \mathbb{Q}[x]/(f(x))$. For the scalar extension $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$, we similarly have the isomorphism $K_{\mathbb{R}} \simeq \mathbb{R}[x]/(f(x))$. For any element $a = \sum_{j=0}^{n-1} a_j \zeta^j \in K_{\mathbb{R}}$ with $a_j \in \mathbb{R}$, there exists a coefficient embedding

$$\tau \colon K_{\mathbb{R}} \to \mathbb{R}^n, \qquad \tau(a) = (a_0, a_1, \ldots, a_{n-1}).$$

On the other hand, let $a = \sum_{j=0}^{n-1} a_j \zeta^j$ be the polynomial $g_a(\zeta)$ of a with respect to ζ, and let $\zeta_1, \zeta_2, \ldots, \zeta_n$ be the n roots of the minimal polynomial $f(x)$ of ζ. Then we have n embeddings

$$\sigma_k \colon K_{\mathbb{R}} \to \mathbb{C}, \qquad \sigma_k(a) = g_a(\zeta_k).$$

Suppose $f(x)$ has s_1 real roots and $2s_2$ complex roots, then there is a canonical embedding

$$\sigma \colon K_{\mathbb{R}} \to H, \qquad \sigma(a) = \big(g_a(\zeta_1), g_a(\zeta_2), \ldots, g_a(\zeta_n)\big),$$

where the space H is defined by

$$H = \left\{(z_1, z_2, \ldots, z_n) \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} : z_j = \overline{z_{j+s_2}},\ j = s_1 + 1, \ldots, s_1 + s_2\right\}.$$

From the definitions of the two embeddings, for any $a \in K_{\mathbb{R}}$, we have $\tau(a) \cdot V = \sigma(a)$, where $V = (v_{ij})_{n \times n} = (\zeta_j^{i-1})_{n \times n}$ is the Vandermonde matrix generated by $\zeta_1, \zeta_2, \ldots, \zeta_n$. Since $f(x)$ is the minimal polynomial of ζ, it has no repeated roots, which implies that V is invertible.

For any $a, b \in K_{\mathbb{R}}$ with their product $c = ab$, the product corresponds to the convolution $\tau(a) * \tau(b) = \tau(c)$ under the coefficient embedding[4], and corresponds to the Hadamard product $\sigma(a) \odot \sigma(b) = \sigma(c)$ under the canonical embedding.

[4] This operation is essentially the standard convolution combined with reduction modulo $f(x)$. Since the computations in this paper do not involve field switching, we use $*$ directly to denote this operation.

3 Hadamard Product of Gaussian Distributions

In this section, we focus on the setting in the embedding space, corresponding to the Hadamard product of two Gaussians. In fact, the conditional spherical large deviation principle below applies more generally to any distribution whose density satisfies the exponential tail decay 1, as well as to the joint distribution of multiple independent samples from such a distribution.

Theorem 1 (Conditional spherical large deviations in $\mathbb{C}^n$ under exponential radial tails). *Let X be a complex-valued random variable with density f on $\mathbb{C}$ with respect to the two-dimensional Lebesgue measure. Assume that there exists a constant $c > 0$ such that for every $\eta > 0$ there exist constants $M_\eta > 0$ and $0 < C_\eta^- \le C_\eta^+ < \infty$ satisfying*

$$C_\eta^- e^{-(c+\eta)|z|} \le f(z) \le C_\eta^+ e^{-(c-\eta)|z|}, \qquad |z| \ge M_\eta. \tag{1}$$

Let

$$Z = (X_1, \dots, X_n) \in \mathbb{C}^n,$$

where $X_1, \dots, X_n$ are independent copies of X. For $R > 0$, define the conditional directional law μ_R on the complex unit sphere

$$S_{\mathbb{C}}^{n-1} := \{u \in \mathbb{C}^n : \|u\|_2 = 1\}$$

by

$$\mu_R(B) := \Pr\left(\frac{Z}{\|Z\|_2} \in B \,\middle|\, \|Z\|_2 > R\right), \qquad B \subseteq S_{\mathbb{C}}^{n-1} \text{ Borel.}$$

Then $(\mu_R)_{R>0}$ satisfies a large deviation principle on $S_{\mathbb{C}}^{n-1}$ with speed R and good rate function

$$I(u) := c(\|u\|_1 - 1), \qquad u \in S_{\mathbb{C}}^{n-1},$$

where

$$\|u\|_1 := \sum_{j=1}^n |u_j|, \qquad \|u\|_2 := \Big(\sum_{j=1}^n |u_j|^2\Big)^{1/2}.$$

Equivalently, for every closed set $F \subseteq S_{\mathbb{C}}^{n-1}$,

$$\limsup_{R\to\infty} \frac{1}{R} \log \mu_R(F) \le - \inf_{u\in F} I(u),$$

and for every open set $G \subseteq S_{\mathbb{C}}^{n-1}$,

$$\liminf_{R\to\infty} \frac{1}{R} \log \mu_R(G) \ge - \inf_{u\in G} I(u).$$

Proof. We first show an exponential scale for the denominator:

$$\lim_{R\to\infty} \frac{1}{R} \log \Pr(\|Z\|_2 > R) = -c. \tag{2}$$

For the upper bound, fix $\eta \in (0, c)$. By (1), for every $t \in (0, c-\eta)$,

$$\mathbb{E}[e^{t|X|}] = \int_{\mathbb{C}} e^{t|z|} f(z)\, dz < \infty.$$

Indeed, the tail integral is bounded by

$$\int_{|z|\geq M_\eta} e^{t|z|} C_\eta^+ e^{-(c-\eta)|z|}\, dz \ll \int_{M_\eta}^{\infty} r\, e^{-(c-\eta-t)r}\, dr < \infty.$$

Since $\|Z\|_2 \leq \|Z\|_1 = \sum_{j=1}^n |X_j|$, for every $t \in (0, c)$,

$$\Pr(\|Z\|_2 > R) \leq \Pr\left(\sum_{j=1}^n |X_j| > R\right) \leq \left(\mathbb{E}[e^{t|X|}]\right)^n e^{-tR}.$$

Hence

$$\limsup_{R\to\infty} \frac{1}{R} \log \Pr(\|Z\|_2 > R) \leq -t.$$

Letting $t \uparrow c$ yields

$$\limsup_{R\to\infty} \frac{1}{R} \log \Pr(\|Z\|_2 > R) \leq -c.$$

For the lower bound, choose $a > 0$ such that

$$p_a := \Pr(|X| \leq a) > 0.$$

Then

$$\Pr(\|Z\|_2 > R) \geq \Pr(R \leq |X_1| \leq R+1,\ |X_2| \leq a, \ldots, |X_n| \leq a) = \Pr(R \leq |X_1| \leq R+1)\, p_a^{n-1}.$$

Fix $\eta > 0$. By (1), for all sufficiently large R,

$$\Pr(R \leq |X_1| \leq R+1) = \int_{R\leq|z|\leq R+1} f(z)\, dz \geq C_\eta^- \int_{R\leq|z|\leq R+1} e^{-(c+\eta)|z|}\, dz.$$

Using polar coordinates,

$$\int_{R\leq|z|\leq R+1} e^{-(c+\eta)|z|}\, dz = 2\pi \int_R^{R+1} r e^{-(c+\eta)r}\ dr,$$

hence

$$\liminf_{R\to\infty} \frac{1}{R} \log \Pr(R \leq |X_1| \leq R+1) \geq -(c+\eta).$$

Therefore

$$\liminf_{R\to\infty} \frac{1}{R} \log \Pr(\|Z\|_2 > R) \geq -(c+\eta).$$

Letting $\eta \downarrow 0$ proves (2).

We next derive an upper bound for closed sets. Let $F \subseteq S_{\mathbb{C}}^{n-1}$ be closed, and define

$$\alpha_F := \inf_{u\in F} \|u\|_1.$$

On the event

$$\left\{ \frac{Z}{\|Z\|_2} \in F,\ \|Z\|_2 > R \right\},$$

we have

$$\|Z\|_1 = \|Z\|_2 \left\| \frac{Z}{\|Z\|_2} \right\|_1 \geq \alpha_F R.$$

Hence, for every $t \in (0, c)$,

$$\Pr\left(\frac{Z}{\|Z\|_2} \in F,\ \|Z\|_2 > R \right) \leq \Pr\left(\sum_{j=1}^{n} |X_j| \geq \alpha_F R \right) \leq \left(\mathbb{E}[e^{t|X|}] \right)^n e^{-t\alpha_F R}.$$

Thus

$$\limsup_{R\to\infty} \frac{1}{R} \log \Pr\left(\frac{Z}{\|Z\|_2} \in F,\ \|Z\|_2 > R \right) \leq -c\,\alpha_F.$$

Combining this with (2), we obtain

$$\limsup_{R\to\infty} \frac{1}{R} \log \mu_R(F) \leq -c(\alpha_F - 1) = -\inf_{u\in F} I(u).$$

Finally, we give a lower bound for open sets. Let $G \subseteq S_{\mathbb{C}}^{n-1}$ be open, and define

$$\alpha_G := \inf_{u\in G} \|u\|_1.$$

Fix $\eta > 0$. Since points in $S_{\mathbb{C}}^{n-1}$ with all coordinates nonzero are dense, we may choose

$$u^\star = (u_1^\star, \ldots, u_n^\star) \in G$$

such that

$$u_j^\star \neq 0 \quad \forall j, \qquad \|u^\star\|_1 \leq \alpha_G + \eta.$$

Because G is open, there exists $\delta > 0$ such that

$$B_{S_{\mathbb{C}}^{n-1}}(u^\star, \delta) \subseteq G.$$

Fix $\varepsilon > 0$ sufficiently small. For all sufficiently large R, the polydisc

$$Q_R := \prod_{j=1}^{n} \{ z_j \in \mathbb{C} :\ |z_j - R u_j^\star| \leq \varepsilon \}$$

is contained in

$$\left\{ z \in \mathbb{C}^n : \ \|z\|_2 > R, \ \frac{z}{\|z\|_2} \in G \right\}.$$

Indeed, Q_R/R converges to $\{u^\star\}$ in Hausdorff distance, so for large R every $z \in Q_R$ has direction in $B_{S^{n-1}_{\mathbb{C}}}(u^\star, \delta) \subseteq G$, and also has Euclidean norm strictly larger than R.

Fix now $\rho > 0$. Since each $u_j^\star \neq 0$, we have $|z_j| \asymp R$ uniformly on Q_R for all large R. By (1), for all sufficiently large R and all $z = (z_1, \dots, z_n) \in Q_R$,

$$f(z_j) \geq C_\rho^- e^{-(c+\rho)|z_j|}, \qquad j = 1, \dots, n.$$

Writing $f_Z(z) = \prod_{j=1}^n f(z_j)$, we obtain

$$f_Z(z) \geq (C_\rho^-)^n \exp\left(-(c+\rho) \sum_{j=1}^n |z_j| \right) = (C_\rho^-)^n e^{-(c+\rho)\|z\|_1}.$$

Since uniformly on Q_R,

$$\|z\|_1 \leq R\|u^\star\|_1 + O(1),$$

it follows that

$$\inf_{z \in Q_R} f_Z(z) \geq \widetilde{C}_\rho \exp(-(c+\rho)R\|u^\star\|_1)$$

for some constant $\widetilde{C}_\rho > 0$ independent of R. Therefore,

$$\Pr\left(\frac{Z}{\|Z\|_2} \in G, \ \|Z\|_2 > R \right) \geq \Pr(Z \in Q_R) \geq \mathrm{Vol}(Q_R) \inf_{z \in Q_R} f_Z(z).$$

Since

$$\mathrm{Vol}(Q_R) = (\pi\varepsilon^2)^n,$$

we get

$$\liminf_{R\to\infty} \frac{1}{R} \log \Pr\left(\frac{Z}{\|Z\|_2} \in G, \ \|Z\|_2 > R \right) \geq -(c+\rho)\|u^\star\|_1 \geq -(c+\rho)(\alpha_G + \eta).$$

Combining this with (2),

$$\liminf_{R\to\infty} \frac{1}{R} \log \mu_R(G) \geq -(c+\rho)(\alpha_G + \eta) + c.$$

Letting first $\rho \downarrow 0$ and then $\eta \downarrow 0$, we obtain

$$\liminf_{R\to\infty} \frac{1}{R} \log \mu_R(G) \geq -c(\alpha_G - 1) = -\inf_{u \in G} I(u).$$

This completes the proof. □

From the conditional spherical large deviation principle in Theorem 1, we immediately obtain the following characterization of the limiting behavior of typical directions under large deviations.

Since the rate function is minimized precisely for directions concentrated on a single coordinate, the exponentially most probable directions lie exactly on the complex axes. This result is formalized in the Corollary 1.

Corollary 1 (Concentration near the complex coordinate axes). *Under the assumptions of Theorem 1, define*

$$A_{\mathbb{C}} := \bigcup_{j=1}^{n} \{\omega e_j : \ |\omega| = 1\} \subset S_{\mathbb{C}}^{n-1},$$

where $e_1, \dots, e_n$ are the standard basis vectors of $\mathbb{C}^n$. Then

$$I(u) \geq 0 \qquad \forall\, u \in S_{\mathbb{C}}^{n-1},$$

and

$$I(u) = 0 \iff u \in A_{\mathbb{C}}.$$

Consequently, for every $\varepsilon > 0$,

$$\Pr\left(\operatorname{dist}\left(\frac{Z}{\|Z\|_2}, A_{\mathbb{C}}\right) > \varepsilon \,\middle|\, \|Z\|_2 > R\right) \leq \exp(-c_\varepsilon R + o(R))$$

for some constant $c_\varepsilon > 0$, and in particular

$$\Pr\left(\operatorname{dist}\left(\frac{Z}{\|Z\|_2}, A_{\mathbb{C}}\right) > \varepsilon \,\middle|\, \|Z\|_2 > R\right) \xrightarrow[R\to\infty]{} 0.$$

Proof. For every $u \in S_{\mathbb{C}}^{n-1}$,

$$\|u\|_1 \geq \|u\|_2 = 1,$$

hence $I(u) = c(\|u\|_1 - 1) \geq 0$. Moreover,

$$\|u\|_1 = \|u\|_2$$

holds if and only if u has at most one nonzero coordinate. Since $\|u\|_2 = 1$, this is equivalent to $u = \omega e_j$ for some j and some $\omega \in \mathbb{C}$ with $|\omega| = 1$, i.e. $u \in A_{\mathbb{C}}$.

Now apply Theorem 1 to the closed set

$$F_\varepsilon := \left\{u \in S_{\mathbb{C}}^{n-1} : \operatorname{dist}(u, A_{\mathbb{C}}) \geq \varepsilon\right\}.$$

Since F_ε is compact and disjoint from $A_{\mathbb{C}}$, continuity of I implies that

$$c_\varepsilon := \inf_{u \in F_\varepsilon} I(u) > 0.$$

Therefore,

$$\limsup_{R\to\infty} \frac{1}{R} \log \Pr\left(\operatorname{dist}\left(\frac{Z}{\|Z\|_2}, A_{\mathbb{C}}\right) \geq \varepsilon \,\middle|\, \|Z\|_2 > R\right) \leq -c_\varepsilon,$$

which proves the claim. □

Corollary 2 establishes that the distribution corresponding to the Hadamard product of two Gaussian vectors on $\mathbb{C}^n$ satisfies the conditional spherical large deviation principle in Theorem 1.

As a consequence, the directions of samples with sufficiently large Euclidean norm asymptotically concentrate on the complex coordinate axes in probability.

Corollary 2 (Complex Gaussian product case). *Let*

$$U, V \overset{\text{i.i.d.}}{\sim} \mathcal{CN}(0, \tau^2),$$

where $\mathcal{CN}(0, \tau^2)$ denotes the circularly symmetric complex Gaussian distribution with density

$$z \mapsto \frac{1}{\pi\tau^2} e^{-|z|^2/\tau^2}.$$

Set

$$X := UV, \qquad Z = (X_1, \dots, X_n) \in \mathbb{C}^n,$$

where $X_1, \dots, X_n$ are independent copies of X. Then the assumptions of Theorem 1 hold with

$$c = \frac{2}{\tau^2}.$$

Consequently, the conditional directional laws

$$\mu_R(B) = \Pr\left(\frac{Z}{\|Z\|_2} \in B \,\middle|\, \|Z\|_2 > R \right)$$

satisfy an LDP on $S_{\mathbb{C}}^{n-1}$ with speed R and rate function

$$I(u) = \frac{2}{\tau^2}(\|u\|_1 - 1).$$

In particular, conditioned on $\|Z\|_2 > R$, the direction $Z/\|Z\|_2$ concentrates asymptotically near the complex coordinate axes $A_{\mathbb{C}}$.

Proof. The density of $X = UV$ on $\mathbb{C}$ is

$$f_X(z) = \frac{2}{\pi\tau^4} K_0\left(\frac{2|z|}{\tau^2} \right), \qquad z \in \mathbb{C},$$

where K_0 is the modified Bessel function of the second kind. Using the standard asymptotic

$$K_0(t) = \sqrt{\frac{\pi}{2t}}\, e^{-t}\big(1 + o(1)\big), \qquad t \to \infty,$$

we obtain

$$\log f_X(z) = -\frac{2}{\tau^2}|z| + O(\log |z|), \qquad |z| \to \infty.$$

Hence, for every $\eta > 0$, there exist constants $M_\eta > 0$ and $0 < C_\eta^- \le C_\eta^+ < \infty$ such that

$$C_\eta^- e^{-(2/\tau^2+\eta)|z|} \le f_X(z) \le C_\eta^+ e^{-(2/\tau^2-\eta)|z|}, \qquad |z| \ge M_\eta.$$

Thus Theorem 1 applies with $c = 2/\tau^2$, proving the corollary. □

4 Convolution of Gaussian Distributions

In this section, we establish a spherical large deviation principle for the product of two independent Gaussian variables under the coefficient embedding. Throughout this chapter, we focus on cyclotomic fields of power-of-two order $K \simeq \mathbb{Q}[x]/(x^n+1)$. Our results show that if one takes independent and identically distributed Gaussian elements in $K_{\mathbb{R}}$ under the coefficient embedding and multiplies them, the resulting samples asymptotically lie on one of $n/2$ real 2-dimensional planes as the Euclidean 2-norm grows large. These planes represent intrinsic directions imposed by the number field structure, a phenomenon that is completely neglected under the standard independence assumption.

Lemma 1. *Let $K = \mathbb{Q}(\zeta) \simeq \mathbb{Q}[x]/(x^n+1)$, where n is a power of two, and let $\zeta_j = \exp\Big(\frac{(2j-1)\pi \mathrm{i}}{n}\Big)$ for $j = 1, \ldots, n$ be all roots of x^n+1. Define the reduced canonical embedding $\rho : K_{\mathbb{R}} \to \mathbb{C}^{n/2}$ by*

$$\rho(x) := \big(g_x(\zeta_1), \ldots, g_x(\zeta_{n/2})\big).$$

If a random polynomial $a \in K_{\mathbb{R}}$ satisfies $\tau(a) \sim \mathcal{N}(0, \sigma^2 I_n)$ under the coefficient embedding τ, i.e., its coefficients $a_0, \ldots, a_{n-1}$ are i.i.d. real Gaussian variables with zero mean and variance σ^2, then the entries of $\rho(a)$ are independent identically distributed circularly symmetric complex Gaussian variables:

$$\rho(a)_1, \ldots, \rho(a)_{n/2} \overset{\text{i.i.d.}}{\sim} \mathcal{CN}(0, n\sigma^2).$$

Proof. Write $\tau(a) = (a_0, \ldots, a_{n-1})$. For each $1 \le i \le n/2$, we have

$$\rho(a)_i = \sum_{s=0}^{n-1} a_s \zeta_i^s.$$

Since $a_0, \ldots, a_{n-1}$ are i.i.d. real Gaussian variables, $\rho(a)$ is a complex Gaussian vector.

We first compute the covariance. For any $1 \le i, j \le n/2$,

$$\mathbb{E}\Big[\rho(a)_i \overline{\rho(a)_j}\Big] = \sigma^2 \sum_{s=0}^{n-1} \zeta_i^s \overline{\zeta_j^s} = \sigma^2 \sum_{s=0}^{n-1} (\zeta_i \overline{\zeta_j})^s.$$

When $i \neq j$, $\zeta_i \overline{\zeta_j}$ is a nontrivial n-th root of unity, so the geometric sum vanishes; when $i = j$, the sum equals n. Hence

$$\mathbb{E}\Big[\rho(a)_i \overline{\rho(a)_j}\Big] = n\sigma^2 \, \delta_{ij}.$$

We then compute the pseudo-covariance:

$$\mathbb{E}[\rho(a)_i \rho(a)_j] = \sigma^2 \sum_{s=0}^{n-1} (\zeta_i \zeta_j)^s.$$

For all $1 \le i, j \le n/2$, $\zeta_i \zeta_j$ remains a nontrivial n-th root of unity, so the sum is zero.

Combining the above, the entries of $\rho(a)$ are uncorrelated with zero pseudo-covariance. Together with the complex Gaussian property, they are independent circularly symmetric complex Gaussian variables with common variance $n\sigma^2$.

□

Theorem 2 (Spherical concentration on cyclotomic Fourier planes). *Let $K = \mathbb{Q}(\zeta) \simeq \mathbb{Q}[x]/(x^n + 1)$, where n is a power of two and*

$$\zeta = \exp\left(\frac{\pi \mathrm{i}}{n}\right), \qquad \zeta_j = \exp\left(\frac{(2j-1)\pi \mathrm{i}}{n}\right), \quad j = 1, \dots, n,$$

so that $\zeta_1, \dots, \zeta_n$ are the n roots of $x^n + 1$. Let $a, b \in K_{\mathbb{R}}$ be independent random elements such that

$$\tau(a), \tau(b) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2 I_n)$$

under the coefficient embedding τ, and let

$$c = ab.$$

For each $i = 1, \dots, n/2$, define

$$\theta_i := \frac{(2i-1)\pi}{n},$$

and the vectors $u_i, v_i \in \mathbb{R}^n$ by

$$u_i[s] = \cos(s\theta_i), \qquad v_i[s] = \sin(s\theta_i), \qquad 0 \le s \le n-1.$$

Let

$$P_i := \mathrm{span}_{\mathbb{R}}\{u_i, v_i\} \subset \mathbb{R}^n.$$

Then for every $\varepsilon > 0$,

$$\Pr\left(\mathrm{dist}\left(\frac{\tau(c)}{\|\tau(c)\|_2}, \bigcup_{i=1}^{n/2} (P_i \cap S^{n-1}) \right) > \varepsilon \;\middle|\; \|\tau(c)\|_2 > R \right) \xrightarrow[R \to \infty]{} 0. \tag{3}$$

More precisely, there exists a constant $\kappa_\varepsilon > 0$ such that

$$\Pr\left(\mathrm{dist}\left(\frac{\tau(c)}{\|\tau(c)\|_2}, \bigcup_{i=1}^{n/2} (P_i \cap S^{n-1}) \right) > \varepsilon \;\middle|\; \|\tau(c)\|_2 > R \right) \le \exp(-\kappa_\varepsilon R + o(R)). \tag{4}$$

In particular, conditioned on a large coefficient-embedding norm, the product $c = ab$ asymptotically lies on one of the $n/2$ real two-dimensional planes $P_1, \dots, P_{n/2}$.

Proof. We first characterize the product in the frequency domain via the reduced canonical embedding. Since $x^n + 1$ has no real roots and its roots form complex conjugate pairs

$$\zeta_{n+1-i} = \overline{\zeta_i}, \qquad i = 1, \ldots, n/2,$$

the reduced canonical embedding ρ captures exactly the nonredundant part of the full canonical embedding.

By Lemma 1, the entries of $\rho(a)$ and $\rho(b)$ are i.i.d. $\mathcal{CN}(0, n\sigma^2)$, and $\rho(a)$ is independent of $\rho(b)$. Since multiplication in $K_{\mathbb{R}}$ corresponds to coordinatewise multiplication under the canonical embedding, we have

$$\rho(c) = \rho(a) \odot \rho(b).$$

Therefore, the coordinates of $\rho(c)$ are independent and identically distributed as the product of two independent $\mathcal{CN}(0, n\sigma^2)$ variables.

We then apply Corollary 1. Let

$$A_{\mathbb{C}} = \bigcup_{i=1}^{n/2} \{\omega e_i : \ |\omega| = 1\} \subset S_{\mathbb{C}}^{n/2-1}$$

denote the complex coordinate-axis set in $\mathbb{C}^{n/2}$. By Corollary 1, for every $\varepsilon > 0$,

$$\Pr\left(\operatorname{dist}\left(\frac{\rho(c)}{\|\rho(c)\|_2}, A_{\mathbb{C}} \right) > \varepsilon \ \middle| \ \|\rho(c)\|_2 > R \right) \leq \exp(-\lambda_\varepsilon R + o(R)) \tag{5}$$

for some constant $\lambda_\varepsilon > 0$. In particular, the conditional probability above tends to 0 as $R \to \infty$.

We further characterize the coefficient embedding via inverse Fourier representation. We now identify the preimage of each complex coordinate axis under the inverse canonical transform. For $i = 1, \ldots, n/2$, define the normalized vectors

$$\widehat{u}_i := \sqrt{\frac{2}{n}}\, u_i, \qquad \widehat{v}_i := \sqrt{\frac{2}{n}}\, v_i.$$

By the usual discrete Fourier orthogonality relations, the family

$$\{\widehat{u}_1, \widehat{v}_1, \ldots, \widehat{u}_{n/2}, \widehat{v}_{n/2}\}$$

is an orthonormal basis of $\mathbb{R}^n$.

For any $x \in K_{\mathbb{R}}$, write

$$\rho(x)_i = \alpha_i + \mathrm{i}\beta_i, \qquad i = 1, \ldots, n/2.$$

Since $\sigma(x)$ consists of the conjugate pairs

$$(\rho(x)_1, \ldots, \rho(x)_{n/2}, \overline{\rho(x)_{n/2}}, \ldots, \overline{\rho(x)_1}),$$

the inverse Vandermonde transform gives, for each coefficient index $s = 0, \dots, n-1$,

$$\begin{aligned}\tau(x)[s] &= \frac{1}{n}\sum_{j=1}^{n} \sigma_j(x)\zeta_j^{-s} \\ &= \frac{1}{n}\sum_{i=1}^{n/2}\Big(\rho(x)_i\zeta_i^{-s} + \overline{\rho(x)_i}\zeta_i^{s}\Big) \\ &= \frac{2}{n}\sum_{i=1}^{n/2}\big(\alpha_i\cos(s\theta_i) + \beta_i\sin(s\theta_i)\big).\end{aligned}$$

Equivalently,

$$\tau(x) = \sqrt{\frac{2}{n}}\sum_{i=1}^{n/2}\big(\alpha_i\widehat{u}_i + \beta_i\widehat{v}_i\big). \tag{6}$$

Define

$$T : \mathbb{C}^{n/2} \to \mathbb{R}^n, \qquad T(z_1, \dots, z_{n/2}) := \sum_{i=1}^{n/2}\big(\Re(z_i)\widehat{u}_i + \Im(z_i)\widehat{v}_i\big).$$

Since $\{\widehat{u}_i, \widehat{v}_i\}$ is orthonormal, T is an isometry:

$$\|T(z)\|_2 = \|z\|_2, \qquad z \in \mathbb{C}^{n/2}.$$

By (6),

$$\tau(x) = \sqrt{\frac{2}{n}}\,T(\rho(x)), \qquad x \in K_{\mathbb{R}}.$$

Hence

$$\|\tau(x)\|_2 = \sqrt{\frac{2}{n}}\,\|\rho(x)\|_2. \tag{7}$$

Moreover, for each fixed i,

$$T(\omega e_i) = \Re(\omega)\widehat{u}_i + \Im(\omega)\widehat{v}_i, \qquad |\omega| = 1,$$

so

$$T\big(\{\omega e_i : \ |\omega| = 1\}\big) = P_i \cap S^{n-1}.$$

Therefore,

$$T(A_{\mathbb{C}}) = \bigcup_{i=1}^{n/2}(P_i \cap S^{n-1}). \tag{8}$$

Since T is an isometry, for every nonzero $z \in \mathbb{C}^{n/2}$,

$$\operatorname{dist}\left(\frac{T(z)}{\|T(z)\|_2},\ \bigcup_{i=1}^{n/2}(P_i \cap S^{n-1})\right) = \operatorname{dist}\left(\frac{z}{\|z\|_2},\ A_{\mathbb{C}}\right). \tag{9}$$

Finally, we transfer the results back to the coefficient embedding. Applying (9) with $z = \rho(c)$, and using $\tau(c) = \sqrt{2/n}\,T(\rho(c))$, we get

$$\operatorname{dist}\left(\frac{\tau(c)}{\|\tau(c)\|_2},\ \bigcup_{i=1}^{n/2}(P_i \cap S^{n-1})\right) = \operatorname{dist}\left(\frac{\rho(c)}{\|\rho(c)\|_2},\ A_{\mathbb{C}}\right).$$

Likewise, by (7),

$$\|\rho(c)\|_2 = \sqrt{\frac{n}{2}}\,\|\tau(c)\|_2.$$

Therefore, for every $\varepsilon > 0$,

$$\begin{aligned}&\Pr\left(\operatorname{dist}\left(\frac{\tau(c)}{\|\tau(c)\|_2},\ \bigcup_{i=1}^{n/2}(P_i \cap S^{n-1})\right) > \varepsilon \;\middle|\; \|\tau(c)\|_2 > R\right)\\ &\quad = \Pr\left(\operatorname{dist}\left(\frac{\rho(c)}{\|\rho(c)\|_2},\ A_{\mathbb{C}}\right) > \varepsilon \;\middle|\; \|\rho(c)\|_2 > \sqrt{\frac{n}{2}}\,R\right).\end{aligned}$$

Now (3) follows immediately from (5), and (4) follows as well after replacing R there by $\sqrt{n/2}\,R$ and absorbing the factor into κ_ε. □

5 Applications

In this section, we present several useful conclusions by adopting canonical embeddings and leveraging the large deviation results established previously.

5.1 CCA Attacks: Explanation and Extension

We focus on the most classical decryption-failure attack against LAC proposed by Guo et al. [13]. For the convolution term $s \cdot e$, their work reveals that imposing proper constraints on s and e can induce strong correlation within decryption noise. Let $s = \sum_i s_i x^i$ and $e = \sum_i e_i x^i$. They put forward two main patterns that trigger decryption failures with high probability.

The first pattern is categorized as Type 1: it requires $|\sum_i s_i|$ to be sufficiently large, together with a consecutive segment consisting of 1 or -1 in the sequence $\{e_i\}$. If several zeros are allowed to be inserted into such segments, this pattern is defined as Type 2.

The second pattern is named Type 1b, which demands a sufficiently large value of $|\sum_i s_{2i}| + |\sum_i s_{2i+1}|$, and requires $\{e_i\}$ to contain specific alternating segments where odd-indexed entries are 1 and even-indexed entries are -1, or vice versa[5]. Correspondingly, the version permitting inserted zeros in such alternating segments is referred to as Type 2b.

[5] The original Type 1b and Type 2b patterns are originally designed for polynomials with fixed Hamming weights. We argue that these patterns can still effectively induce decryption failures even when the Hamming weight is not fixed.

To trigger decryption failures, it is necessary to maximize the norm of $s \cdot e$. According to our analysis, a large norm of $s \cdot e$ indicates that it lies on the thorns, which equivalently requires the existence of some index i such that both $|\sigma_t(s)|$ and $|\sigma_t(e)|$ take relatively large values.

We demonstrate that the constraints of Type 1 and Type 2 naturally lead to large values of $|\sigma_1(s)|$ and $|\sigma_1(e)|$ simultaneously. Similarly, the constraints corresponding to Type 1b and Type 2b result in large magnitudes of $|\sigma_{n/2}(s)|$ and $|\sigma_{n/2}(e)|$. In addition, any polynomial e with a large $\sigma_t(e)$ tends to possess coefficient characteristics consistent with the aforementioned patterns, and the same property holds for polynomial s.

Accordingly, we can readily extend existing decryption-failure-based attacks. There is no need to explore specific coefficient patterns of s and e for triggering decryption failures. It suffices to guarantee an index t such that both $|\sigma_t(s)|$ and $|\sigma_t(e)|$ achieve large values. The actual decryption failure rate turns out to be considerably higher than the estimation under the independence assumption, owing to the strong correlation among polynomial coefficients.

Canonical embeddings and cyclic signed partial sums. Consider the number field $\mathbb{Q}(\zeta) \cong \mathbb{Q}[x]/\langle x^n + 1\rangle$, where n is a power of two. For each canonical embedding $\sigma_t(e)$, the coefficients $\{e_i\}$ are regularly permuted and assigned alternating signs according to the index t, yielding a new sequence $\{b_j^{(t)}\}$. The embedding value is then defined as $\sigma_t(e) = \sum_j b_j^{(t)} \zeta^j$. By applying Abel's summation formula to the right-hand side, we observe that partial sums derived from such permutation and sign assignment significantly affect the magnitude $|\sigma_t(e)|$. This property is formally stated in Theorem 3.

All constraints imposed on s and e by Guo et al. induce large partial sums at $t = 1$ or $t = n/2$, which simultaneously enlarge $\sigma_t(s)$ and $\sigma_t(e)$. Specifically, the two proposed patterns require that the partial sums of $\{s_i\}$ are sufficiently large either under the original order corresponding to $t = 1$ or after rearrangement for $t = n/2$, which raises $|\sigma_t(s)|$. For polynomial e, Type 1 and Type 2 demand sufficiently long consecutive segments of ± 1 under the original order for $t = 1$, while Type 1b and Type 2b require such segments after rearrangement for $t = n/2$. Both settings produce large partial sums under corresponding permutations and consequently increase $|\sigma_t(e)|$.

Theorem 3. *Let $n = 2^m$, let $K = \mathbb{Q}(\zeta)$ where ζ is a primitive $2n$-th root of unity, and let*

$$R = \mathbb{Z}[\zeta] \simeq \mathbb{Z}[x]/\langle x^n + 1\rangle.$$

For

$$e = \sum_{k=0}^{n-1} e_k \zeta^k \in K_{\mathbb{R}},$$

and each odd integer $t \in \{1, 3, \ldots, 2n-1\}$, define

$$\sigma_t(e) = \sum_{k=0}^{n-1} e_k \zeta^{tk}.$$

Let

$$q_t(k) \equiv tk \pmod{n}, \qquad 0 \le q_t(k) \le n-1,$$

and

$$\varepsilon_t(k) = (-1)^{\lfloor tk/n \rfloor} \in \{\pm 1\}.$$

Then

$$\zeta^{tk} = \varepsilon_t(k)\zeta^{q_t(k)}.$$

Since $\gcd(t,n)=1$*, the map* q_t *is a permutation of* $\{0,\dots,n-1\}$*. Let* $p_t = q_t^{-1}$*, and define*

$$b_j^{(t)} = \varepsilon_t\big(p_t(j)\big)e_{p_t(j)}, \qquad 0 \le j \le n-1.$$

Then

$$\sigma_t(e) = \sum_{j=0}^{n-1} b_j^{(t)}\zeta^j.$$

Now fix an arbitrary cut position $s \in \{0,\dots,n-1\}$*, and define the cyclically cut signed sequence*

$$d_r^{(t,s)} = \begin{cases} b_{s+r}^{(t)}, & 0 \le r \le n-1-s, \\ -b_{s+r-n}^{(t)}, & n-s \le r \le n-1. \end{cases}$$

Equivalently,

$$d_r^{(t,s)} = (-1)^{\mathbf{1}_{\,s+r\ge n}}\, b^{(t)}_{\,s+r \bmod n}.$$

Let

$$T_m^{(t,s)} = \sum_{r=0}^{m} d_r^{(t,s)}, \qquad 0 \le m \le n-1.$$

Then

$$\zeta^{-s}\sigma_t(e) = \sum_{r=0}^{n-1} d_r^{(t,s)}\zeta^r,$$

and hence, by Abel summation,

$$\zeta^{-s}\sigma_t(e) = T_{n-1}^{(t,s)}\zeta^{n-1} + \sum_{m=0}^{n-2} T_m^{(t,s)}\big(\zeta^m - \zeta^{m+1}\big).$$

In particular,

$$|\sigma_t(e)| = \left| T_{n-1}^{(t,s)}\zeta^{n-1} + \sum_{m=0}^{n-2} T_m^{(t,s)}\big(\zeta^m - \zeta^{m+1}\big) \right|.$$

Proof. Since $\zeta^n = -1$, for each k we may write

$$tk = n\left\lfloor \frac{tk}{n} \right\rfloor + q_t(k), \qquad 0 \le q_t(k) \le n-1,$$

and therefore

$$\zeta^{tk} = (\zeta^n)^{\lfloor tk/n \rfloor} \zeta^{q_t(k)} = (-1)^{\lfloor tk/n \rfloor} \zeta^{q_t(k)} = \varepsilon_t(k) \zeta^{q_t(k)}.$$

Thus

$$\sigma_t(e) = \sum_{k=0}^{n-1} e_k \zeta^{tk} = \sum_{k=0}^{n-1} \varepsilon_t(k) e_k \zeta^{q_t(k)}.$$

Since $\gcd(t, n) = 1$, multiplication by t modulo n is bijective, so q_t is a permutation. Re-indexing with $j = q_t(k)$, namely $k = p_t(j)$, gives

$$\sigma_t(e) = \sum_{j=0}^{n-1} b_j^{(t)} \zeta^j.$$

Next, fix $s \in \{0, \ldots, n-1\}$. Then

$$\zeta^{-s} \sigma_t(e) = \sum_{j=0}^{n-1} b_j^{(t)} \zeta^{j-s}.$$

Split the indices according to whether $j \geq s$ or $j < s$. If $j = s + r$ with $0 \leq r \leq n - 1 - s$, then $\zeta^{j-s} = \zeta^r$. If $j = s + r - n$ with $n - s \leq r \leq n - 1$, then

$$\zeta^{j-s} = \zeta^{r-n} = \zeta^r \zeta^{-n} = -\zeta^r,$$

since $\zeta^n = -1$. Therefore

$$\zeta^{-s} \sigma_t(e) = \sum_{r=0}^{n-1} d_r^{(t,s)} \zeta^r.$$

Finally, applying Abel summation to the sequence $d_r^{(t,s)}$ gives

$$\sum_{r=0}^{n-1} d_r^{(t,s)} \zeta^r = T_{n-1}^{(t,s)} \zeta^{n-1} + \sum_{m=0}^{n-2} T_m^{(t,s)} (\zeta^m - \zeta^{m+1}),$$

which proves the desired identity. Taking absolute values yields the final formula for $|\sigma_t(e)|$.

Remark 1. The theorem separates two distinct mechanisms.

First, the choice of t specifies which canonical embedding coordinate $\sigma_t(e)$ is under consideration; algebraically, this produces the signed permutation $b^{(t)}$ of the original coefficient sequence.

Second, the choice of s specifies where one starts to take partial sums in that signed permutation. Since multiplication by ζ^{-s} does not change the modulus, every cut position s yields an exact representation of the same quantity $|\sigma_t(e)|$.

Therefore, different canonical embedding coordinates correspond to different signed summation orders, and for each such order one may further choose an arbitrary starting position for the partial sums.

Canonical embeddings and coherent $\{-1, 0, 1\}$-blocks. In LAC, all polynomial coefficients are restricted to the set $\{-1, 0, 1\}$. The constraints on polynomial e adopted in the attack by Guo et al. require the rearranged sequence $\{e_i\}$ to contain consecutive segments of ones with several zeros inserted inside. Different from the analysis via Abel's summation formula, we provide a dedicated interpretation for this specific scenario. In this case, $|\sigma_t(e)|$ can be decomposed into two components. The dominant component corresponds to the sum of a geometric sequence, while the remaining part serves as minor random perturbation that barely affects the overall magnitude of $|\sigma_t(e)|$.

In this subsection, we still specialize to the 2-power cyclotomic case with

$$K = \mathbb{Q}(\zeta), \qquad f(x) = x^n + 1,$$

where ζ is a primitive $2n$-th root of unity. We write

$$e = \sum_{r=0}^{n-1} e_r \zeta^r \in K_{\mathbb{R}}, \qquad e_r \in \{-1, 0, 1\}.$$

We order the roots of f as

$$\zeta_k = \zeta^{2k-1} \quad (1 \le k \le n/2), \qquad \zeta_{k+n/2} = \overline{\zeta_k} \quad (1 \le k \le n/2),$$

so that, in particular, $\zeta_{n/2+1} = \overline{\zeta_1}$. For $1 \le k \le n$, the k-th coordinate of the canonical embedding is

$$\sigma_k(e) = g_e(\zeta_k).$$

For each $1 \le k \le n/2$, let

$$t_k = 2k - 1, \qquad q_k(r) \equiv t_k r \pmod{n}, \quad 0 \le q_k(r) \le n - 1,$$

and

$$\varepsilon_k(r) = (-1)^{\lfloor t_k r/n \rfloor} \in \{\pm 1\}.$$

Since $\gcd(t_k, n) = 1$, the map q_k is a permutation of $\{0, \ldots, n-1\}$. Let $p_k = q_k^{-1}$, and define

$$b_j^{(k)} = \varepsilon_k\big(p_k(j)\big) e_{p_k(j)}, \qquad 0 \le j \le n - 1.$$

Then $b_j^{(k)} \in \{-1, 0, 1\}$, and

$$\sigma_k(e) = \sum_{j=0}^{n-1} b_j^{(k)} \zeta^j.$$

Moreover,

$$|\sigma_{k+n/2}(e)| = |\sigma_k(e)|, \qquad 1 \le k \le n/2.$$

Thus, for each embedding position k, the corresponding canonical embedding coordinate is obtained from the original coefficient sequence $(e_0, \ldots, e_{n-1})$ by a signed permutation

$$(e_0, \ldots, e_{n-1}) \longmapsto (b_0^{(k)}, \ldots, b_{n-1}^{(k)}).$$

The relevant structured patterns are therefore not necessarily contiguous in the original coefficient order, but contiguous after passing to the embedding-adapted sequence $b^{(k)}$.

Definition 1 (k-coherent interval). *Fix $1 \le k \le n/2$. Let*

$$I = \{s, s+1, \dots, s+L-1\} \subset \mathbb{Z}/n\mathbb{Z}$$

be a cyclic interval of length L, and let

$$A = \{j \in I : \ b_j^{(k)} \neq 0\}.$$

We say that I is a k-coherent interval of sign $\eta \in \{\pm 1\}$ and weight ℓ for e, if

$$b_j^{(k)} \in \{0, \eta\} \qquad \text{for all } j \in I,$$

and

$$|A| = \ell.$$

Equivalently, after passing to the embedding-adapted sequence $b^{(k)}$, all nonzero entries on I have the same sign η, and there are exactly ℓ such nonzero entries.

The point of this definition is that a k-coherent interval forces the k-th canonical embedding coordinate to contain a deterministic coherent contribution.

Proposition 1 (Deterministic contribution of a k-coherent interval). *Let $1 \le k \le n/2$, and suppose that $I = \{s, \dots, s+L-1\}$ is a k-coherent interval of sign η and weight ℓ for e. Let*

$$u = \sum_{j \in I} b_j^{(k)} \zeta^j, \qquad v = \sum_{j \notin I} b_j^{(k)} \zeta^j,$$

so that

$$\sigma_k(e) = u + v.$$

Then

$$u = \eta \zeta^s \sum_{m \in B} \zeta^m,$$

where

$$B = \{j - s \bmod n : \ j \in I, \ b_j^{(k)} = \eta\} \subseteq \{0, \dots, L-1\}, \qquad |B| = \ell.$$

Hence

$$|u| = \left| \sum_{m \in B} \zeta^m \right|.$$

If $\ell = L$, namely if $b_j^{(k)} = \eta$ for every $j \in I$, then

$$|u| = \left| \sum_{m=0}^{L-1} \zeta^m \right| = \frac{\sin(L\pi/(2n))}{\sin(\pi/(2n))}.$$

More generally, for arbitrary $\ell \le L$,

$$|u| \ge \ell \cos\left(\frac{(L-1)\pi}{2n}\right).$$

Consequently,

$$|\sigma_k(e)| \ge \ell \cos\left(\frac{(L-1)\pi}{2n}\right) - |v|.$$

Proof. By definition of a k-coherent interval, every nonzero coefficient of $b^{(k)}$ on I equals η. Therefore

$$u = \sum_{j \in I,\ b_j^{(k)} = \eta} \eta \zeta^j = \eta \zeta^s \sum_{m \in B} \zeta^m,$$

which proves the first identity.

If $\ell = L$, then $B = \{0, 1, \ldots, L-1\}$, and thus

$$|u| = \left|\sum_{m=0}^{L-1} \zeta^m\right| = \left|\frac{1-\zeta^L}{1-\zeta}\right| = \frac{\sin(L\pi/(2n))}{\sin(\pi/(2n))}.$$

For the general case, the points $\{\zeta^m : m \in B\}$ lie on an arc of the unit circle of angular length

$$\frac{(L-1)\pi}{n} < \pi.$$

Projecting onto the bisector of this arc yields

$$\left|\sum_{m \in B} \zeta^m\right| \ge \ell \cos\left(\frac{(L-1)\pi}{2n}\right).$$

Finally, since $\sigma_k(e) = u + v$, the reverse triangle inequality gives

$$|\sigma_k(e)| \ge |u| - |v|,$$

which concludes the proof.

Remark 2 (Interpretation). For a fixed embedding position k, the sequence $b^{(k)}$ is the embedding-adapted coefficient order. A long interval in $b^{(k)}$ whose nonzero entries are all equal to $+1$ or all equal to -1 produces a large deterministic contribution to $|\sigma_k(e)|$. The contribution is exactly a geometric sum when the interval is full, and admits the explicit lower bound

$$\ell \cos\left(\frac{(L-1)\pi}{2n}\right)$$

when zeros are inserted inside the interval.

The contribution of the remaining coefficients is isolated in the error term $|v|$.

Remark 3 (The four motivating patterns). The general formulation above contains the four motivating patterns as special cases.

For $k = 1$, one has $t_1 = 1$, hence

$$q_1(r) = r, \qquad \varepsilon_1(r) \equiv 1,$$

so that

$$b_j^{(1)} = e_j.$$

Therefore a 1-coherent interval is exactly a contiguous interval in the original coefficient order in which all nonzero entries have the same sign. This contains both: (i) a contiguous block of $+1$'s or -1's; and (ii) a contiguous interval containing ℓ entries equal to $+1$ (or all equal to -1) and $L - \ell$ zeros in arbitrary order. Its conjugate embedding coordinate is $\sigma_{n/2+1}$.

For $k = n/2$, one has $t_{n/2} = n-1 \equiv -1 \pmod{n}$. In this case, the embedding-adapted order $b^{(n/2)}$ corresponds, after a cyclic cut, to the alternating sign rule in the original coefficient order. Therefore a $n/2$-coherent interval corresponds to: (i) a contiguous alternating block

$$+1, -1, +1, -1, \ldots$$

or its negation; and (ii) the same pattern with zeros inserted, while preserving the alternating sign rule on the nonzero positions. Its conjugate embedding coordinate is σ_n.

5.2 Experiments

Throughout our theoretical analysis of the correlation among the coefficients of $s \cdot e$, we assume that the coefficients of s and e follow continuous Gaussian distributions. Nevertheless, practical lattice-based schemes adopt discrete sub-Gaussian distributions instead. This section conducts numerical experiments to verify that the correlation characteristics derived in our theoretical analysis still hold even under realistic discrete distributions.

We again take the patterns defined in [13] as illustrative examples. When both s and e are of Type 1, i.e., they contain consecutive runs of 1 or -1, the coefficients of $s \cdot e$ concentrate near the plane P_1 according to Theorem 2. When both s and e belong to Type 1b, characterized by alternating sequences of 1 and -1, the coefficients of $s \cdot e$ asymptotically lie close to the plane $P_{n/2}$.

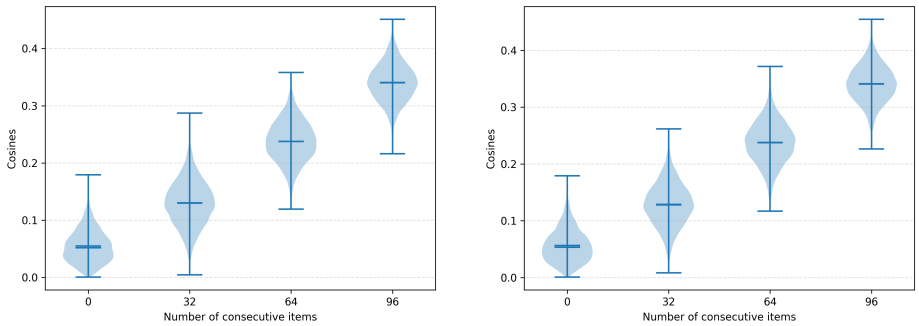
The experimental results are presented in Figure 3. In our sampling setup, each entry of s and e is independently assigned a value of 1 or -1 with probability $1/4$ each, and 0 with probability $1/2$. For each test case, we first fix special segments with specified lengths and random starting positions, and then sample the remaining entries at random. A total of 2^{12} samples are collected for each violin plot.

The violin plot at the horizontal axis value 0 corresponds to the fully random case[6]. During the sampling of Type 1b, we do not impose a fixed Hamming weight

[6] Due to inherent randomness, fully random sampling may occasionally generate special subsequences that match predefined patterns, such as consecutive 1s.

constraint, which explains the nearly identical appearance of the two plots. This observation further validates our theoretical claim that all thorns of $s \cdot e$ are structurally equivalent.

The experimental outcomes demonstrate that as the length of the predefined special segment increases, sampled vectors gradually concentrate closer to the corresponding planes. By contrast, fully random samples are almost orthogonal to the planes associated with these thorns.



(a) Type 1 (b) Type 1b

Fig. 3: Violin plots of the cosine values of the angles between the coefficient vector of $s \cdot e$ and the corresponding planes, where both s and e are sampled under Type 1 (or Type 1b). The horizontal axis denotes the length of consecutive $1/-1$ entries (or alternating 1 and -1 entries).

5.3 Distribution of noise length

Theorem 4. *Consider a number field $K \cong \mathbb{Q}[x]/(x^n+1)$ where n is a power of two. For $a, b \in K_{\mathbb{R}}$, suppose that $\tau(a)$ and $\tau(b)$ are independently distributed as $\mathcal{N}(0, \sigma^2 I_n)$. Then the probability density function of the ℓ_2-norm of the coefficient vector of $a \cdot b$ is given by*

$$f_{\|\tau(a\cdot b)\|}(r) = nr \cdot g^{*(n/2)}\left(\frac{nr^2}{2}\right), \quad r \geq 0,$$

where

$$g(r) := \frac{2}{n^2\sigma^4} K_0\left(\frac{2\sqrt{r}}{n\sigma^2}\right), \quad r \geq 0.$$

Here $f^{*\ell}(x)$ denotes the ℓ-fold convolution of $f(x)$, which is explicitly defined as

$$f^{*\ell}(y) = \int_{\substack{x_i \geq 0 \\ x_1+\cdots+x_{\ell-1} \leq y}} \left(\prod_{i=1}^{\ell-1} f(x_i)\right) f\left(y - \sum_{i=1}^{\ell-1} x_i\right) dx_1 \cdots dx_{\ell-1}, \quad y > 0.$$

Proof. Let $c = a \cdot b \in \mathbb{R}[x]/(x^n+1)$. The first $n/2$ entries of $\sigma(a)$ and $\sigma(b)$ follow independent and identically distributed circularly symmetric complex Gaussian distribution $\mathcal{CN}(0, n\sigma^2)$, while the remaining $n/2$ entries are the complex conjugates of the first half. Following the proof of Theorem 2, $\tau(a)$ and $\tau(b)$ obey $\mathcal{N}(0, \sigma^2 I_n)$ and are mutually independent.

Define the embedding matrix V such that $\tau(c)V = \sigma(c)$. For the number field $K = \mathbb{Q}[x]/(x^n+1)$, one can verify that $V^H V = nI_n$, which yields

$$\|\tau(c)\|_2 = \|\sigma(c)\|_2 / \sqrt{n}.$$

We first derive the probability density function of $T_{\text{cano-half}} = \sum_{i=1}^{n/2} \sigma_i(c)^2$. Recall that the probability density function of $|\sigma_i(c)|$ is

$$f_{|\sigma_i(c)|}(r) = \frac{4r}{n^2\sigma^4} K_0\left(\frac{2r}{n\sigma^2}\right), \quad r \geq 0.$$

The probability density function of $\sigma_i(c)^2$ is then given by

$$g(r) := f_{\sigma_i(c)^2}(r) = \frac{2}{n^2\sigma^4} K_0\left(\frac{2\sqrt{r}}{n\sigma^2}\right), \quad r \geq 0.$$

Accordingly, the probability density function of $T_{\text{cano-half}}$ corresponds to the $(n/2)$-fold convolution of g, denoted as $g^{*(n/2)}(r)$.

Since $\|\sigma(c)\|_2^2 = 2T_{\text{cano-half}}$, we have

$$\|\tau(c)\|_2 = \sqrt{\frac{2T_{\text{cano-half}}}{n}}.$$

As a result, the probability density function of the ℓ_2-norm of the coefficient vector $\tau(a \cdot b)$ is

$$f_{\|\tau(a\cdot b)\|}(r) = nr \cdot g^{*(n/2)}\left(\frac{nr^2}{2}\right), \quad r \geq 0.$$

□

5.4 Real Scenario versus Independence Assumption

Characteristic functions or Laplace transforms can be adopted to derive explicit closed-form expressions for the norm of $\tau(s \cdot e)$, which facilitate practical theoretical analysis. This work presents a more concise argument to illustrate that the independence assumption inevitably underestimates the actual ℓ_2-norm of $\tau(s \cdot e)$. This implies that if the decryption procedure essentially depends on the ℓ_2-norm of the noise vector, the decryption failure rate evaluated under the independence assumption provably underestimates the true decryption failure rate from a theoretical perspective.

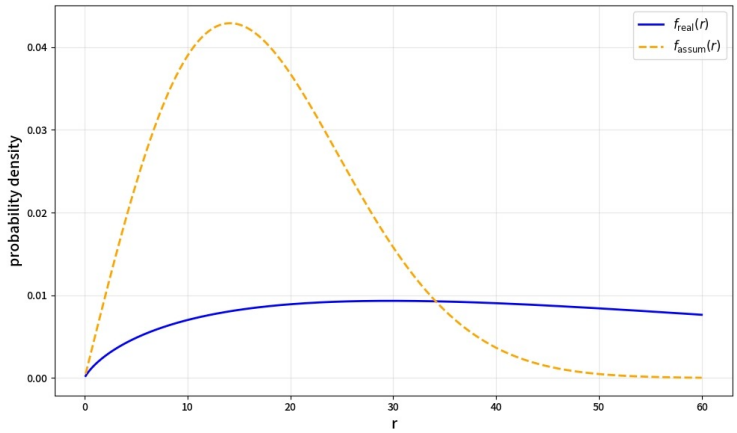


Fig. 4: This figure plots $f_{\text{assum}}(r)$ and $f_{\text{real}}(r)$ under $\sigma = 10$. It can be observed that the independence assumption misjudges $|\sigma(s \cdot e)|$ to take small values with relatively high probability. In contrast, practical results show that larger values occur with notably higher probabilities in the real distribution.

We still focus on the ring $\mathbb{R}[x]/(x^n + 1)$. Let $s = \sum_i s_i x^i$ and $e = \sum_i e_i x^i$. The ℓ-th entry of $\tau(s \cdot e)$ satisfies

$$\tau(s \cdot e)_\ell = \sum_{i+j \equiv \ell \pmod n} (-1)^{\left\lfloor \frac{i+j-\ell}{n} \right\rfloor} s_i e_j.$$

Each individual entry $\tau(s \cdot e)_\ell$ is a linear combination of numerous independent and identically distributed random variables, and thus asymptotically follows a Gaussian distribution by the central limit theorem. If one only adopts the marginal distribution of each $\tau(s \cdot e)_\ell$ and further imposes the independence assumption, it implicitly regards $\tau(s \cdot e)$ as a spherical Gaussian vector.

Following the proof of Theorem 1, this treatment is equivalent to assuming that the first $n/2$ entries of $\sigma(s \cdot e)$ are i.i.d. complex Gaussian variables, while the remaining $n/2$ entries are their complex conjugates. Under the component-wise independence assumption in the canonical embedding domain, the induced correlation structure coincides with the real setting: the first $n/2$ embedding components are mutually independent, and the latter half are conjugate to the first half.

We analyze an arbitrary embedding component $\sigma_i(s \cdot e)$. In the real setting, its amplitude follows the distribution

$$f_{\text{real}}(r) = \frac{4r}{\sigma^4} K_0\left(\frac{2r}{\sigma^2}\right), \quad r \geq 0.$$

Under the independence assumption, the amplitude is instead governed by the marginal distribution of a circularly symmetric complex Gaussian with identical

standard deviation:

$$f_{\text{assum}}(r) = \frac{r}{2\sigma^2} e^{-\frac{r^2}{4\sigma^2}}, \quad r \geq 0.$$

The two densities $f_{\text{assum}}(r)$ and $f_{\text{real}}(r)$ intersect at only two points, and the first intersection point approaches zero as σ grows.

In practical lattice-based schemes, σ is scaled by $\sqrt{n}$ relative to individual coefficients and typically lies on the order of 10. Figure 4 plots $f_{\text{assum}}(r)$ and $f_{\text{real}}(r)$ for $\sigma = 10$. In this case, the first intersection is nearly invisible near the origin, demonstrating that the independence assumption almost always underestimates the noise norm under practical parameter choices.

Finally, we emphasize that the discrepancy between $f_{\text{real}}(r)$ and $f_{\text{assum}}(r)$ constitutes the only essential difference between the real noise distribution and its independence-assumed approximation. Nevertheless, this single mischaracterization of the true distribution fundamentally enables the emergence of practical decryption-failure-based attacks.

6 Discussion

We characterize the origin of correlation among the entries of the coefficient embedding $\tau(c)$ for the product $c = s \cdot e$, where s, e are spherical Gaussian random elements of $K_{\mathbb{R}}$, from the perspective of canonical embeddings. Our analysis focuses on the special case $K \cong \mathbb{Q}[x]/\langle x^n+1\rangle$. The findings established in this paper provide instructive guidance for both cryptanalysis and cryptosystem design.

Taking algebraic-structure-based LWE public-key encryption schemes as an illustrative example, we split the implications into two aspects. From a cryptanalytic standpoint, an efficient strategy for locating extreme samples—either those exhibiting strong correlation across entries of $\tau(c)$ or those with a large ℓ_2-norm of $\tau(c)$—is to search for pairs (s, e) such that the corresponding coordinate magnitudes of $\sigma(s)$ and $\sigma(e)$ are simultaneously large at the same embedding index.

From a cryptographic design perspective, we prove that the independence assumption strictly underestimates the norm of decryption noise. Consequently, as previously observed empirically in numerous prior works, one critical takeaway is to avoid relying solely on the independence assumption, especially when computing the decryption failure rate of a scheme. Furthermore, to mitigate the occurrence of extreme products c, controlling the infinity norms of $\sigma(s)$ and $\sigma(e)$ separately constitutes a reliable countermeasure. This follows from our theoretical result that large-norm products c arise with far higher probability from multiplying pairs (s, e) that feature a single coordinate with amplified magnitude in their canonical embeddings while all remaining coordinates retain typical magnitudes.

References

1. Biasioli, B., Marcolla, C., Calderini, M., Mono, J.: Improving and Automating BFV Parameters Selection: An Average-Case Approach. Cryptology ePrint Archive, Paper 2023/600 (2023), https://eprint.iacr.org/2023/600
2. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehle, D.: CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM. In: 2018 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 353–367 (2018). https://doi.org/10.1109/EuroSP.2018.00032
3. Checri, M., Sirdey, R., Boudguiga, A., Bultel, J.P.: On the Practical CPAD Security of "exact" and Threshold FHE Schemes and Libraries. In: Advances in Cryptology – CRYPTO 2024: 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2024, Proceedings, Part III. p. 3–33. Springer-Verlag, Berlin, Heidelberg (2024). https://doi.org/10.1007/978-3-031-68382-4_1, https://doi.org/10.1007/978-3-031-68382-4_1
4. Cheon, J.H., Choe, H., Passelègue, A., Stehlé, D., Suvanto, E.: Attacks Against the IND-CPAD Security of Exact FHE Schemes. In: Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security. p. 2505–2519. CCS '24, Association for Computing Machinery, New York, NY, USA (2024). https://doi.org/10.1145/3658644.3690341, https://doi.org/10.1145/3658644.3690341
5. Costache, A., Nürnberger, L., Player, R.: Optimisations and Tradeoffs for HElib. In: Rosulek, M. (ed.) Topics in Cryptology – CT-RSA 2023. pp. 29–53. Springer International Publishing, Cham (2023)
6. D'Anvers, J., Guo, Q., Johansson, T., Nilsson, A., Vercauteren, F., Verbauwhede, I.: Decryption failure attacks on IND-CCA secure lattice-based schemes. In: Lin, D., Sako, K. (eds.) Public-Key Cryptography - PKC 2019 - 22nd IACR International Conference on Practice and Theory of Public-Key Cryptography, Beijing, China, April 14-17, 2019, Proceedings, Part II. pp. 565–598. Lecture Notes in Computer Science, Springer (2019). https://doi.org/10.1007/978-3-030-17259-6_19, https://doi.org/10.1007/978-3-030-17259-6_19
7. D'Anvers, J., Rossi, M., Virdia, F.: (one) failure is not an option: Bootstrapping the search for failures in lattice-based encryption schemes. In: Canteaut, A., Ishai, Y. (eds.) Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part III. pp. 3–33. Lecture Notes in Computer Science, Springer (2020). https://doi.org/10.1007/978-3-030-45727-3_1, https://doi.org/10.1007/978-3-030-45727-3_1
8. D'Anvers, J., Vercauteren, F., Verbauwhede, I.: The impact of error dependencies on ring/mod-lwe/lwr based schemes. In: Ding, J., Steinwandt, R. (eds.) Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019, Chongqing, China, May 8-10, 2019 Revised Selected Papers. pp. 103–115. Lecture Notes in Computer Science, Springer (2019). https://doi.org/10.1007/978-3-030-25510-7_6, https://doi.org/10.1007/978-3-030-25510-7_6
9. D'Anvers, J.P., Batsleer, S.: Multitarget Decryption Failure Attacks and Their Application to Saber and Kyber. In: Public-Key Cryptography – PKC 2022: 25th IACR International Conference on Practice and Theory of Public-Key Cryptography, Virtual Event, March 8–11, 2022, Proceedings, Part I. p. 3–33. Springer-Verlag, Berlin, Heidelberg (2022). https://doi.org/10.1007/978-3-030-97121-2_1, https://doi.org/10.1007/978-3-030-97121-2_1

10. D'Anvers, J.P., Guo, Q., Johansson, T., Nilsson, A., Vercauteren, F., Verbauwhede, I.: Decryption Failure Attacks on IND-CCA Secure Lattice-Based Schemes. In: Public-Key Cryptography – PKC 2019: 22nd IACR International Conference on Practice and Theory of Public-Key Cryptography, Beijing, China, April 14-17, 2019, Proceedings, Part II. p. 565–598. Springer-Verlag, Berlin, Heidelberg (2019). https://doi.org/10.1007/978-3-030-17259-6_19, https://doi.org/10.1007/978-3-030-17259-6_19
11. D'Anvers, J.P., Rossi, M., Virdia, F.: (One) Failure Is Not an Option: Bootstrapping the Search for Failures; Lattice-Based Encryption Schemes. In: Advances in Cryptology – EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part III. p. 3–33. Springer-Verlag, Berlin, Heidelberg (2020). https://doi.org/10.1007/978-3-030-45727-3_1, https://doi.org/10.1007/978-3-030-45727-3_1
12. Guo, Q., Johansson, T.: A new decryption failure attack against hqc. In: Moriai, S., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2020. pp. 353–382. Springer International Publishing, Cham (2020)
13. Guo, Q., Johansson, T., Yang, J.: A novel cca attack using decryption errors against lac. In: Galbraith, S.D., Moriai, S. (eds.) Advances in Cryptology – ASIACRYPT 2019. pp. 82–111. Springer International Publishing, Cham (2019)
14. Howgrave-Graham, N., Nguyen, P.Q., Pointcheval, D., Proos, J., Silverman, J.H., Singer, A., Whyte, W.: The Impact of Decryption Failures on the Security of NTRU Encryption. In: Boneh, D. (ed.) Advances in Cryptology - CRYPTO 2003. pp. 226–246. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
15. Li, B., Micciancio, D.: On the Security of Homomorphic Encryption on Approximate Numbers. In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology – EUROCRYPT 2021. pp. 648–677. Springer International Publishing, Cham (2021)
16. Liu, Y., Ge, J., Zhang, Y., Wang, J., Lu, X.: A Primer on Dependency in Polynomial Product: Identify, Exploit, and Trim. Cryptology ePrint Archive, Paper 2026/802 (2026), https://eprint.iacr.org/2026/802
17. Liu, Y., Zhang, Y., Lu, X., Cheng, Y., Yin, Y.: Dawn: Smaller and faster ntru encryption via double encoding. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025. pp. 396–427. Springer Nature Singapore, Singapore (2026)
18. Liu, Z., Wang, Y., Fisch, B.: IND-CPA-D of Relaxed Functional Bootstrapping: A New Attack, A General Fix, and A Stronger Model. In: Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security. p. 2907–2921. CCS '25, Association for Computing Machinery, New York, NY, USA (2025). https://doi.org/10.1145/3719027.3765170, https://doi.org/10.1145/3719027.3765170
19. Lu, X., Liu, Y., Zhang, Z., Jia, D., Xue, H., He, J., Li, B., Wang, K.: LAC: Practical ring-LWE based public-key encryption with byte-level modulus. Cryptology ePrint Archive, Paper 2018/1009 (2018), https://eprint.iacr.org/2018/1009
20. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) Advances in Cryptology – EUROCRYPT 2010. pp. 1–23. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
21. Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-lwe cryptography. In: Johansson, T., Nguyen, P.Q. (eds.) Advances in Cryptology – EUROCRYPT 2013. pp. 35–54. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)