IACR Communications in Cryptology
ISSN 3006-5496, Vol. 1, No. 2, 41 pages.

https://doi.org/10.62056/ay76c0kr

A Central Limit Approach for Ring-LWE Noise Analysis

Sean Murphy[a] and Rachel Player

Royal Holloway, University of London, Egham, UK

Abstract. This paper develops Central Limit arguments for analysing the noise in ciphertexts in two homomorphic encryption schemes that are based on Ring-LWE. The first main contribution of this paper is to present and evaluate an average-case noise analysis for the BGV scheme. Our approach relies on the recent work of Costache *et al.* (SAC 2023) that gives the approximation of a polynomial product as a multivariate Normal distribution. We show how this result can be applied in the BGV context and evaluate its efficacy. We find this average-case approach can much more closely model the noise growth in BGV implementations than prior approaches, but in some cases it can also underestimate the practical noise growth. Our second main contribution is to develop a Central Limit framework to analyse the noise growth in the homomorphic Ring-LWE cryptosystem of Lyubashevsky, Peikert and Regev (Eurocrypt 2013, full version). Our approach is very general: apart from finite variance, no assumption on the distribution of the noise is required (in particular, the noise need not be subgaussian). We show that our approach leads to tighter bounds for the probability of decryption failure than those of prior work.

Keywords: Ring-LWE · Central Limit Theorem · decryption failure probability · BGV scheme · homomorphic encryption

1 Introduction

The Learning with Errors or *LWE* problem [Reg05, Reg10] has become a standard hard problem in cryptology that is at the heart of lattice-based cryptography [MR09, Pei16]. The Ring Learning with Errors or *Ring-LWE* problem [SSTX09, LPR12] is a generalisation of the LWE problem from the ring of integers to certain other number field rings that potentially give far better efficiency.

A key application area of lattice-based cryptography is (fully, somewhat or levelled) homomorphic encryption [Gen09]. Homomorphic encryption enables an untrusted party to operate meaningfully on encrypted data belonging to a different party, without requiring access to the secret key. A large number of homomorphic encryption schemes have been proposed in the literature, for example [BGV12, FV12, GSW13, LPR13a, CGGI16, CKKS17], many of which [BGV12, FV12, LPR13a, CKKS17] are based on Ring-LWE. In this paper, we consider the widely-used BGV scheme [BGV12], which has been implemented in many major libraries, including HElib [HEl19] and SEAL [SEA22]. We also consider the homomorphic cryptosystem given by Lyubashevsky, Peikert and Regev in Section 8.3 of [LPR13a] (the full version of [LPR13b]), which we term the `LPRHom` cryptosystem.

During parts of the development of this work, 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). The authors thank all reviewers of this work for their valuable feedback and suggestions.

E-mail: s.murphy@rhul.ac.uk (Sean Murphy), rachel.player@rhul.ac.uk (Rachel Player)
[a]Corresponding author

Received: 2024-03-20 Accepted: 2024-06-03

Ciphertexts in all homomorphic encryption schemes contain noise, which is needed for security. As more homomorphic evaluation operations are performed, the noise grows, and if it exceeds a certain threshold, then decryption will fail. It is thus essential to understand the noise growth behaviour in order to choose secure and correct parameters. Ideally, the noise growth behaviour would be modelled tightly, so that the most performant parameters that meet the security and correctness requirements can be selected.

Noise estimation approaches. Existing models for noise growth in the homomorphic encryption literature can be classified according to several strands. A first approach, starting with [GHS12a, GHS12b], seeks to bound the noise in fresh encryption and after each homomorphic operation. Tracing these bounds through leads to a bound on the noise in the output ciphertext. A second approach, used e.g. in [CGGI20, CCH+23], and continued in this work, seeks to obtain distributional results on the noise. For example, the variance of the noise is traced through each homomorphic operation, and the variance in the output noise is then converted to a bound. A third approach, considered e.g. in [LMSS22, ABBB+22], uses empirical measurements of the noise.

The first approach is often described as *worst-case* (see e.g. [CLP20, GNSJ24]) while the second is often described as *average-case* (see e.g. [CCH+23, CNP23]). However, care should be taken here as even the 'worst-case' bounds are typically developed using heuristics (see e.g. [GHS12a, GHS12b, CS16, Ili19]); and moreover, they may fail, although typically with very low probability. For example, prior analyses have assumed that a particular Gaussian random variable is within six [CS16] or ten [HEl19] standard deviations of its mean. Average-case analyses may also rely on heuristics and assumptions, as is the case for [CCH+23] and this work.

Bounds on the noise are often given in the *canonical norm*, i.e. the infinity norm in the canonical embedding (as in e.g. [GHS12a, GHS12b, CS16, CKKS17, HS20, CLP20, KPP22]) but they may also be given in the infinity norm in the ciphertext ring (as in e.g. [KPZ21]). Some works (e.g. [CCH+23, CNP23]) consider both norms. The experiments of [CNP23] for BGV show that using the infinity norm more closely models practical noise growth in HElib than using the canonical norm. In this work, for BGV, we conduct our noise analysis directly in the ciphertext ring and give our eventual noise bounds in the infinity norm. For `LPRHom`, we give our analysis under the canonical embedding, where we consider decryption with respect to an embedding of a particular "decoding basis" (using the terminology of [LPR13a]).

Most approaches for noise analysis are *static*, in the sense that the noise estimates in the given model can be publicly computed based on the scheme parameters. Our work focuses on this context. A *dynamic* approach to noise analysis, in which noise is determined at runtime from a given ciphertext using the secret key, has also been considered in [LMSS22, ABBB+22].

Motivation for our work. Prior to the development of this work, the noise analyses presented for BGV [GHS12a, GHS12b, CS16, CLP20, HS20] have provided bounds for the noise growth after every BGV evaluation operation. By tracing through the bounds after each operation, the noise growth incurred by the overall evaluation can also be bounded. In this sense, these prior analyses can all be classified as worst-case. It was shown in [CLP20] that when using such an approach, there can be an unsatisfying gap between the final noise bound and the typical size of the noise as observed in experiments, with the gap growing as more computations are performed.

We are motivated by the noise analysis for the TFHE scheme [CGGI16] that was presented in [CGGI20]. There, it is assumed that the coefficients of a fresh TFHE ciphertext are independent subgaussians, and that the coefficients of a ciphertext output of the gate bootstrapping operation are also independent subgaussians. The latter assumption

is experimentally verified [CGGI20, Figure 10]. It is shown that every TFHE operation can be implemented via gate bootstrapping on a linear combination of ciphertexts. Thus, by linearity and by the assumption on gate bootstrapping, every TFHE ciphertext noise coefficient can be modelled as a subgaussian. This permits the variance of the noise to be traced through the overall evaluation.

In this work, we show how for each BGV homomorphic evaluation operation, the input and output noises can each be modelled as a Normal random variable. This enables us to trace through the variances of the noise at each operation, and eventually arrive at the variance of the noise after the evaluation. Therefore, we only need to resort to a bound after the evaluation: the (modelled) Normal distribution of the given variance implies a certain tail bound on the noise holds with a certain probability. This can be classified as an average-case approach. We expect that this approach should enable us to set parameters that are still large enough to ensure correctness, but may be smaller (and thus more performant) than those that would be chosen under a worst-case analysis.

The fundamental issue with modelling the noise growth in schemes like BGV or the `LPRHom` cryptosystem is that the noise growth in multiplication is nonlinear. In more detail, if two BGV ciphertexts having noise polynomials v_1 and v_2 are multiplied, then the resulting ciphertext has noise polynomial $v_1 \cdot v_2$. In particular, if X_1 and X_2 are subgaussian random variables arising from such noise polynomials, then the product $X_1 \cdot X_2$ is not necessarily subgaussian and indeed can have a much heavier tail [MP20]. For this reason, developing an analysis for BGV that models the noise as (sub)gaussian (as done in [CGGI20] for TFHE) was believed until recently to be a challenging open question [CLP20]. It was also believed to be challenging for related schemes, such as CKKS [CKKS17] and BFV [FV12], that have a similar multiplication structure to BGV. An important step in resolving this question was made in [CCH+23], which showed how a Central Limit approach could, under certain assumptions, be used to approximate the noise in CKKS ciphertexts.

Our first main contribution is to apply the approach of [CCH+23] to BGV and evaluate its efficacy in the BGV context. Our second main contribution is to show how a Central Limit approach can be used to approximate the output noise of all `LPRHom` operations as a Normal distribution. We now overview these contributions in more detail.

1.1 A Central Limit approach for BGV

The first main contribution of this paper is to present and evaluate an average-case noise analysis for BGV, based on a Central Limit argument. Our approach is built upon the recent work of [CCH+23] that develops an average-case noise analysis for the CKKS scheme [CKKS17]. The CKKS scheme closely follows the BGV scheme, differing mainly in the native plaintext space and in encoding of messages. In particular, the multiplication in CKKS is essentially the same as for BGV, and thus it is expected that analyses for CKKS could be applicable in the BGV setting. Indeed, our results crucially rely on Theorem 1 and Corollary 1, developed in [CCH+23]. Our average-case noise analysis for BGV follows from repeated applications of Corollary 1 and is summarised in Figure 2.

Theorem 1 gives the mean and covariance of a polynomial product $Y := ZZ'$ of two multivariate Normals Z and Z'; and shows that the components Y_i of Y can be well-approximated by a multivariate Normal distribution. As in [CCH+23], we also rely on Heuristic 1, which expresses that Y itself can be approximated as a multivariate Normal distribution of the mean and covariance established in Theorem 1. No detailed justification for Heuristic 1 is given in [CCH+23]. As an additional contribution of this work, we give a partial justification in Lemma 3, which establishes the bivariate Normality of $(Y_i, Y_{i'})$ for any pair of components Y_i and $Y_{i'}$ of Y in the particular case of applying Theorem 1 to determine the noise growth after a multiplication of two fresh ciphertexts. We also discuss how this argument be generalised to finite collections of several components $\{Y_{i_1}, \ldots, Y_{i_k}\}$

and potentially in situtations beyond the multiplication of fresh ciphertexts.

Impact. We conduct an experimental evaluation of the efficacy of our average-case noise approach for BGV by comparing with the practical noise growth in HElib [HEl19] and SEAL [SEA22]. Our findings, presented in Tables 1, 2, 3, and 4, are mixed.

On the one hand, the experiments show that in some cases, our approach much more closely models the noise compared to the prior approach of [CLP20]. Moreover, we show that our average-case approach can lead to practical improvements for parameter selection in some cases. In more detail, we are able to exhibit explicit circuits (Tables 3 and 5) for which the prior approach of [CLP20] implies a larger set of parameters is needed than is suggested by our average-case approach; and in these cases, we find that the practically observed noise is such that the smaller parameter set suggested by our average-case approach is indeed sufficient to support the computation.

On the other hand, for our HElib experiments, we also compare with the inbuilt noise estimations in the library [HS20], and we find that the [HS20] estimates are generally closer than ours to the practically observed noise. In addition, for both HElib and SEAL, our average-case approach sometimes overestimates the remaining noise budget, i.e., underestimates the practical BGV noise growth. This is similar to the findings of [CCH+23] for CKKS and suggests that caution should be employed when relying on this approach for parameter selection.

Discussion. In this work, we present and evaluate an average-case noise analysis for BGV. As the findings of our evaluation are mixed, it is natural to wonder how the negative aspects of our approach can be explained.

Our approach is obtained by directly applying the results of [CCH+23] that were developed for CKKS. As such, we make the same assumptions as were made in [CCH+23], and so our work inherits the limitations of these assumptions. One of these assumptions is that the noise coefficients are independent. This assumption is not experimentally verified in [CCH+23] or the present work. Subsequent work [BMCM23] has found that this independence assumption does not hold for the BFV scheme [FV12], and is the cause for underestimates of noise growth in BFV implementations. Moreover, very recent work suggests this independence assumption does not hold for the BGV scheme in general [BM23]. We have not seen a full version of [BM23], so we do not know if it leads to similar underestimates. However, BGV and CKKS are similar to BFV and so based on the results of [BMCM23] we suspect this may partly explain the discrepancy between our predictions and the observed experimental results.

Notably, our approach leads to overestimates of the remaining noise budget after modulus switching. We did not investigate precisely what causes the analysis to fail for modulus switching. The reason for not doing so is that the subsequent work [CNP23] already provides an improved average-case noise analysis for BGV that is specific to its implementation in HElib, which experiments show closely models HElib noise growth. The main reason why the [CNP23] noise analysis is so effective is the observation (in [CNP23, Lemma 8]) that in HElib, the modulus switch noise is dominated by the rounding term only. In fact, the results of [CNP23] crucially rely on the noise distribution in HElib ciphertexts effectively being reset by every modulus switch. In contrast, our heuristic analysis (Lemma 5) is given for a general situation of modulus switching and including both the rounding term and the term that arises from the scaling of the input noise.

In summary, while the present work gives a first step towards an average-case BGV noise analysis, the findings of this and subsequent works [BMCM23, CNP23, BM23] illustrate that further work is needed to refine BGV noise analyses to tightly model noise growth across different implementations.

1.2 A Central Limit approach for LPRHom

The second main contribution of this paper is to develop a statistical framework, based on a Central Limit argument, for analysing the noise in LPRHom ciphertexts. To illustrate the utility of this approach, we present in Theorem 2 and Corollary 2 new, tighter bounds for the probabilities of incorrect decryption in degree-1 and degree-2 LPRHom ciphertexts. Our analysis can similarly be applied for higher-degree ciphertexts [MP20].

In more detail, the Central Limit framework is essentially based on approximating the mean vector and the covariance matrix of the noise of a ciphertext when embedded into the complex space H and transformed with respect to an appropriate "decoding" basis, that is required during decryption [LPR13a]. We show that the approximate Normality of this embedded noise when expressed in a decoding 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 decoding basis.

For example, if $C^{(p\Gamma)}$ is a vector of dimension n expressing the noise in a ciphertext with respect to the decoding $p\Gamma$-basis for H (Definition 8) and $C^{(T)}$ is a vector of dimension n expressing the noise in a ciphertext with respect to the original T-basis for H (Section 2.4), then $C^{(p\Gamma)} = p\Delta C^{(T)}$ for an appropriate real-valued $n \times n$ change of basis matrix Δ and "scaling prime" p (which is the plaintext modulus in LPRHom). In particular, this means that we can express a component $c_j^{(p\Gamma)}$ of $C^{(p\Gamma)}$ as

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

The components $c_1^{(T)}, \ldots, c_n^{(T)}$ of $C^{(T)}$ are identically distributed random variables that are uncorrelated and, in general, independent, having zero mean $\mathbf{E}\left(c_j^{(T)}\right) = 0$ and some finite variance $\mathrm{Var}\left(c_j^{(T)}\right) = \rho^2$. Thus a component $c_j^{(p\Gamma)}$ of a noise vector in the $p\Gamma$-basis is a weighted sum of uncorrelated and in general independent identically distributed random variables. We will show that the weightings $\Delta_{j1}, \ldots, \Delta_{jn}$ are of comparable size, which suggests that a Central Limit argument can be invoked to give a Normal approximation for a component $c_j^{(p\Gamma)}$. For successful decryption, we require each component of $C^{(p\Gamma)}$ to be bounded by an appropriate threshold. A Central Limit approach enables us to bound the probability of incorrect decryption using bounds on the tails of Normal distributions.

Impact. Our decryption failure probability bounds (Theorem 2 and Corollary 2) are tighter than the prior bounds in the literature [LPR13a]. This demonstrates the improvement that can be obtained by using a Central Limit approach, in comparison with the prior approach of [LPR13a] that uses δ-subgaussian random variables [MP12, MP19]. For example, if $\eta_1(n, q, \rho) = \frac{1}{2}(n^{\frac{1}{2}}\rho)^{-1}q$ is moderate or large (as defined in Section 2.2), Theorem 2 gives a decryption failure probability bound of

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

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

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

which is obtained by using the tail bound of [MP19, Lemma 18] in the manner of [LPR13a, Lemma 6.5]. Asymptotically, ignoring constants, this tightens the bound by a factor of $\omega(\sqrt{\log n})$, for power-of-two n and q following [LPR13a, Lemma 8.5].

No concrete parameter recommendations for LPRHom are specified in [LPR13a], so it is difficult to quantify the concrete improvement. Additionally, to the best of our knowledge,

no implementation of `LPRHom` exists, so we did not consider an experimental verification of our `LPRHom` results.

More generally, we emphasise that a Central Limit approach has the following advantages over an approach which seeks to derive bounds using δ-subgaussian arguments.

- A Central Limit approach makes no substantive distributional assumption for the components $c_k^{(T)}$ beyond finite variance, so is potentially applicable to $c_k^{(T)}$ that are chosen from heavy-tailed distributions. Thus a Central Limit approach is more generally applicable than other approaches that for example have a subgaussian requirement for such random variables.
- A Central Limit approach gives an explicit approximating distribution for the cryptographic random variable of interest which can be directly used for general calculation or simulation purposes of use in cryptography. By contrast, a subgaussian approach can never give an explicit approximating distribution and can only give tail bounds. These tail bounds are generally weaker, as is evidenced by comparing our Theorem 2 with the bound that would be obtained following [LPR13a].
- A Central Limit approach gives not only asymptotically an approximation to a Normal distribution, but also a close approximation concretely, for practically relevant Ring-LWE dimensions.

This indicates that the techniques developed in our work may also lead to improved analyses in other application contexts.

1.3 Structure of the paper

We recall relevant background and introduce new tools in Section 2. We present and evaluate a Central Limit approach for BGV in Section 3. We then present our Central Limit approach for the `LPRHom` cryptosystem in Section 4.

2 Background

2.1 Notation

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

2.2 Central Limit approximations

When giving Central Limit approximations, we use the notation $\sim$ to denote either "is exactly distributed as" or "is approximately distributed as" in the sense that we may use the approximating distribution for practical purposes without significant error, as is typically done in statistical analysis. Furthermore, whilst Central Limit results are formally asymptotic results concerning sums or means of random variables, such Central Limit approximations usually apply in practice with relatively few summands (except perhaps for pathological distributions) as illustrated by the Berry-Esseen conditions [Str11] and related multidimensional versions [TV11]. We therefore typically use the phrasing "for moderate or large ..." in such a Central Limit context to emphasise the usual applicability of Central Limit approximations with relatively few summands. For example, such a Central Limit approximation has been used in a homomorphic encryption context as being empirically justified with as few as 30 summands [LMSS22], in line with routine statistical practice.

2.3 Cyclotomic number fields

We consider the ring $R = \mathbb{Z}[X]/(\Phi_m(X))$, where $\Phi_m(X)$ is the m^{th} cyclotomic polynomial of degree $n = \phi(m)$, and we let R_a denote R/aR for an integer a. We let ζ_m denote a (primitive) m^{th} root of unity. 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. Note that here we are using ζ_m as an abstract root of unity to define K, but in Definition 8 we will abuse notation and use ζ_m as an explicit complex number.

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

The *ring of integers* $\mathcal{O}_K$ of a number field is the ring of all elements of the number field which are roots of some monic polynomial with coefficients in $\mathbb{Z}$. The ring of integers of the m^{th} cyclotomic number field K is $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[x]/(\Phi_m)$. The canonical embedding σ embeds R as a lattice $\sigma(R)$. The conjugate dual of this lattice corresponds to the embedding of the dual fractional ideal $R^\vee = \{a \in K \mid \mathrm{Tr}(aR) \subset \mathbb{Z}\}$.

For m an odd prime, if we define b such that $b^{-1} = m^{-1}(1 - \zeta_m)$, then [LPR13a, Corollary 2.18] shows that $R^\vee = \langle b^{-1} \rangle$. We let $(R^\vee)^k$ denote the space of products of k elements of $R^\vee$, that is to say

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

2.4 The complex space H

The ring embeddings $\sigma_1, \ldots, \sigma_n$ from K into $\mathbb{C}$ occur in complex conjugate pairs with $\overline{\sigma_k} = \sigma_{m-k}$. Accordingly, much of the analysis of Ring-LWE takes place in a space H of conjugate pairs of complex numbers.

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

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

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

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

Definition 4. *The* T-basis *for H is given by the columns of the conjugate pairs matrix T.*

An element of H is expressed via the I-basis as a vector of $n' = \frac{1}{2}n$ conjugate pairs. Such an element of H can also be expressed (by construction) in the T-basis as a *real-valued* vector, giving the isomorphism between H and $\mathbb{R}^n$ as an inner product space. Note that our bases are all for H, whereas the "power basis" etc of [LPR13a] are bases for $R^\vee$.

2.5 The BGV scheme

In this section we introduce the BGV scheme [BGV12]. We generally follow the description of BGV given in [CLP20], reproduced in Figure 3 in Appendix A, that restricts to a power-of-two cyclotomic ring, $\mathcal{R} = \mathbb{Z}[X]/(X^N + 1)$ for N a power of two. The plaintext space is given by $\mathcal{R}_t = \mathbb{Z}_t[X]/(X^N + 1)$, where the integer t denotes the plaintext modulus. The ciphertext space is given by $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^N + 1)$, where the integer q denotes the ciphertext modulus. We generally regard a polynomial element of $\mathcal{R}_q$ as having coefficients in $\{-\frac{1}{2}(q-1), \ldots, \frac{1}{2}(q-1)\}$. A polynomial $h \in \mathcal{R}$ (or $\mathcal{R}_q$ or $\mathcal{R}_t$) is given by

$$h = h(X) = \sum_{i=0}^{n-1} h_j X^j = h_0 + h_1 X + \ldots + h_{n-1} X^{N-1},$$

where this polynomial may also be interpreted as vector $h = (h_0, \ldots, h_{N-1})$ of coefficients in an appropriate context. When multiplying such polynomials in $\mathcal{R}_q$, i.e. modulo $X^N + 1$, we express the result using a modified Sign function ξ on the integers given by $\xi(z) = \mathrm{Sign}(z)$ for $z \neq 0$ with $\xi(0) = 1$. A term of (hh') can then be specified as

$$(hh')_i = \sum_{j=0}^{N-1} \xi(i-j)\, h_{i-j} h'_j \qquad [i = 0, \ldots, N-1].$$

and the subscripts are interpreted modulo N to lie in $\{0, \ldots, N-1\}$.

We now describe in our notation the relevant parts of the BGV scheme in order to define the noise in a BGV ciphertext. In our analysis, to derive expression for the BGV noise growth, we make the following assumptions (which were also made in [CCH+23]) for simplicity. Firstly, we assume that certain quantities are fixed, and we will describe such quantities as 'constant' in these cases. Secondly, we will calculate expressions for the critical quantity in $\mathcal{R}$, rather than $\mathcal{R}_q$.

SecretKeyGen. For emphasis, we write the secret key as $s \in \{-1, 0, 1\}^N$, a ternary vector of length N. We regard s as a constant vector known to the genuine receiver. That is, we will assume that the secret key is fixed, rather than being a random variable. More generally, s can be regarded as a polynomial of degree $N - 1$.

PublicKeyGen. The public key (p_0, p_1) consists of two parts, with the first part p_0 a multivariate random variable and the second part p_1 a constant vector. For the second part p_1, a constant vector $a \in \{-\frac{1}{2}(q-1), \ldots, \frac{1}{2}(q-1)\}^n$ is chosen and p_1 is set to a, so $p_1 = a$. For the first part p_0 with secret key $s \in \{-1, 0, 1\}^N$, we have

$$p_0 = -as - t\epsilon_0, \quad \text{where } \epsilon_0 \sim \mathtt{N}(0; \sigma^2 I_N)$$

is a spherically symmetric multivariate Normal random variable with component variance σ^2, where as denotes the appropriate polynomial product of a and s. The distribution of the public key (p_0, p_1) is therefore given by

$$p_0 \sim \mathtt{N}(-as; t^2\sigma^2 I_N) \quad \text{and} \quad p_1 = a.$$

Noise in BGV. In our analysis, we will give distributions for the multivariate random variables arising in BGV before any reduction modulo q. That is, we will calculate expressions for the critical quantity in R, rather than R_q. For convenience, we approximate discrete random variables in BGV by the obvious appropriate continuous random variable.

For a BGV ciphertext (c_0, c_1) encrypting a message m, our analysis considers the *BGV Critical Value*, W given by

$$W = c_0 + sc_1,$$

where sc_1 denotes the appropriate polynomial product of s and c_1. The Noise V is then given from the Critical Value W by subtracting m.

Modulus switching. The key technical tool for noise management in BGV is modulus switching. In Lemma 1 we give an alternative expression for the BGV `ModSwitch` operation to that given in Figure 3 that will be more convenient for our analysis. Lemma 1 (proved in Appendix B) can be seen as giving an explicit implementation of the `Scale` operation described in earlier analyses of BGV [CS16, GHS12a].

Lemma 1. *Suppose that (c_0, c_1) is a BGV ciphertext with respect to a modulus q and consider a* `ModSwitch` *operation with respect to a new modulus $p < q$. The BGV* `ModSwitch` *operation maps an input ciphertext part c_i to the nearest integer polynomial to $\frac{p}{q}c_i$ having the same value modulo t as c_i. More formally, this output ciphertext (c_0', c_1') after the* `ModSwitch` *operation can be expressed as*

$$c_i' = \left\lfloor \frac{p}{q} c_i \right\rceil + \left(\left(c_i - \left\lfloor \frac{p}{q} c_i \right\rceil \right) \bmod t \right) \qquad [i = 0, 1].$$

2.6 The `LPRHom` scheme

In this section we introduce the `LPRHom` cryptosystem. In order to do so, we first need two definitions. A description of `LPRHom` cryptosystem, in the notation of [LPR13a], is then given in Figure 4 in Appendix C.

Definition 5 ([MP19])**.** The univariate *Balanced Reduction* function $\mathcal{R}$ on $\mathbb{R}$ is the random function $\mathcal{R}(a) = \begin{cases} 1 - (\lceil a \rceil - a) & \text{with probability} \quad \lceil a \rceil - a \\ -(\lceil a \rceil - a) & \text{with probability} \quad 1 - (\lceil a \rceil - a). \end{cases}$

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

Definition 6 ([MP19])**.** Let B be a (column) basis matrix for the n-dimensional lattice Λ in H. If $\mathcal{R}$ is the Balanced Reduction function, then the *coordinate-wise randomised rounding discretisation* or *CRR discretisation* $\lfloor X \rceil^B_{\Lambda+c}$ of the random variable X on H to the lattice coset $\Lambda + c$ with respect to the basis matrix B is the random variable

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

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

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

Specifically, we discretise Y to the coset $\sigma(pR^\vee) + \sigma(b^{-1}\mu)$ of the lattice $\sigma(pR^\vee)$ obtained by the canonical embedding of the scaled dual fractional ideal $pR^\vee$. We consider the coordinate-wise randomised rounding discretisation with respect to the $p\Gamma$-basis for H, and following Definition 6 we denote this discretisation of Y by $Y'(\mu) = \lfloor\, Y \,\rceil^{p\Gamma}_{\sigma(pR^\vee)+\sigma(b^{-1}\mu)}$.

The *Noise* random variable $Y''(\mu)$ in the encryption of the plaintext μ is then defined to be $Y''(\mu) = \sigma^{-1}(Y'(\mu))$, and is an element of a coset of $pR^\vee + b^{-1}\mu$ containing information

Description	Random Variable	Range of Random Variable
Underlying Noise	Y	Complex Space H
Embedded Noise	$Y'(\mu)$	Lattice Coset $\sigma(pR^\vee) + \sigma(b^{-1}\mu)$
Noise	$Y''(\mu)$	Number Field Coset $pR^\vee + b^{-1}\mu$

Figure 1: Notation for the Noise-related quantities used in encryption of the plaintext μ.

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 summarise this discussion in Figure 1.

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

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

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

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

3 A CLT approach for BGV noise analysis

3.1 Distribution of polynomial products in $\mathcal{R}$

BGV noise analysis requires us to construct the polynomial product in $\mathcal{R}$ or $\mathcal{R}_q$, that is to say modulo $X^N + 1$, of a constant or scalar and a (discretised) multivariate Normal random variable or of two multivariate Normal random variables. In this section we present relevant results from [CCH+23], developed for the CKKS context, that we will apply to give an average-case noise analysis for BGV.

Theorem 1 ([CCH+23])**.** *Suppose that $Z \sim \mathtt{N}(\boldsymbol{\mu};\rho^2 I_N)$ and $Z' \sim \mathtt{N}(\boldsymbol{\mu}';\rho'^2 I_N)$, then the polynomial product ZZ' (modulo X^N+1) has mean vector $\mathbf{E}(ZZ')$ and covariance matrix $\mathrm{Cov}(ZZ')$ given by*

$$\mathbf{E}(ZZ') = \boldsymbol{\mu}^* \quad \text{and} \quad \mathrm{Cov}(ZZ') = \rho_*^2 I_N + S,$$

where $\boldsymbol{\mu}^*$ is the polynomial product of $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$, $\rho_*^2 = N\rho^2\rho'^2 + \rho'^2\left\|\boldsymbol{\mu}\right\|_2^2 + \rho^2\left\|\boldsymbol{\mu}'\right\|_2^2$ and S is an off-diagonal matrix with entries

$$S_{i,i'} = \rho'^2 \sum_{j=0}^{N-1} \xi(i-j)\xi(i'-j)\mu_{i-j}\mu_{i'-j} + \rho^2 \sum_{j=0}^{N-1} \xi(i-j)\xi(i'-j)\mu'_{i-j}\mu'_{i'-j},$$

for a modified sign function ξ given by $\xi(z) = Sign(z)$ for $z \neq 0$ and $\xi(0) = 1$. Furthermore, the components $(ZZ')_i$ of this polynomial product can be approximated as a Normal $\mathtt{N}(\mu_i^*, \rho_*^2)$ distribution.

Theorem 1 gives the mean and covariance of the product $Y = ZZ'$, and shows the components Y_i of Y can be well-approximated as Normal. As in [CCH+23], our average-case analysis for BGV will model ZZ' as a multivariate Normal distribution of the established mean and covariance. This is expressed in Heuristic 1.

Heuristic 1 ([CCH+23]). *Suppose that $Z \sim \mathtt{N}(\boldsymbol{\mu}; \rho^2 I_N)$ and $Z' \sim \mathtt{N}(\boldsymbol{\mu}'; \rho'^2 I_N)$. Then, for $\boldsymbol{\mu}^*$, ρ_*^2 and S as specified in Theorem 1, the polynomial product ZZ' (modulo $X^N + 1$) can be approximated as a multivariate Normal distribution as*

$$ZZ' \sim \mathtt{N}\left(\boldsymbol{\mu}^*;\ \rho_*^2 I_N + S\right).$$

Following the approach of [CCH+23], we make the *Small-S assumption*: that the off-diagonal matrix S encountered in Theorem 1 and Heuristic 1 is negligible compared to $\rho_*^2 I_N$ and we disregard it. As noted in [CCH+23], examination of the form of S indicates this assumption may not always hold, for example if the mean vectors have large constant components. However, the Small-S assumption is reasonable in many circumstances of interest in BGV when the message components can be modelled as being uniform modulo t.

Corollary 1 ([CCH+23]). *Suppose that $Z \sim \mathtt{N}(\mu; \rho^2 I_N)$ and $Z' \sim \mathtt{N}(\mu'; \rho'^2 I_N)$ are independent, λ is a constant vector and the Small-S assumption is valid. Approximations to the distribution of λZ, ZZ', Z^2 are then given by:*

$$\begin{aligned} \lambda Z &\sim \mathtt{N}\left(\lambda\mu\ ;\ \rho^2|\lambda|^2 I_N\right), \\ ZZ' &\sim \mathtt{N}\left(\mu\mu'\ ;\ N\rho^2\rho'^2 + \rho'^2|\mu|^2 + \rho^2|\mu'|^2) I_N\right) \\ \text{and } Z^2 &\sim \mathtt{N}\left(\mu^2\ ;\ 2\rho^2(n\rho^2 + 2|\mu|^2) I_N\right). \end{aligned}$$

We also add a further variant of these results, as adapted in a special case for general (i.e., not necessarily Normal) distributions Z and Z', which we use when considering the BGV `ModSwitch` operation.

Lemma 2. *Suppose that $Z = (Z_0, \ldots, Z_{N-1})^T$ and $Z' = (Z'_0, \ldots, Z'_{N-1})^T$ are independent vectors of independent and identically distributed components with mean $\mathbf{E}(Z_i) = \mathbf{E}(Z'_i) = 0$ and respective variances $Var(Z_i) = \rho^2$ and $Var(Z'_i) = \rho'^2$. The polynomial product ZZ' is well-approximated as a multivariate Normal distribution for large N given by*

$$ZZ' \sim \mathtt{N}(0; n\rho^2\rho'^2 I_N).$$

Proof. The proof is similar to that given in [CCH+23] for Theorem 1. A component $(ZZ')_i$ of ZZ' is the sum of N summands of the form $\pm Z_j Z'_{j'}$ with mean $\mathbf{E}(\pm Z_i Z'_{i'}) = 0$ and variance $\mathrm{Var}(\pm Z_i Z'_{i'}) = \rho^2\rho'^2$. Thus the Central Limit Theorem shows that the distribution of this component $(ZZ')_i$ and be approximated for large N as $(ZZ')_i \sim \mathtt{N}(0, N\rho^2\rho'^2)$. Furthermore, distinct components $(ZZ')_i$ and $(ZZ')_{i'}$ $(i \neq i')$ have covariance $\mathrm{Cov}((ZZ')_i, (ZZ')_{i'}) = 0$ (as they have 0 means), which gives the result. □

3.2 Partial justification of Heuristic 1

In [CCH+23] no detailed justification was given for Heuristic 1. Lemma 3 provides a partial justification by establishing the bivariate Normality of $(Y_i, Y_{i'})$ for two components Y_i and $Y_{i'}$ of Y, where $Y_i = (ZZ')_i$ is the i^{th} component of the product of two independent spherically symmetric multivariate Normal random variables Z and Z' for $1 \leq i \leq N$. This lemma is proved in the case that Z and Z' arise in the multiplication of fresh ciphertexts so that certain quantities that arise are bounded and do not depend on N.

Lemma 3. *Suppose that $Z \sim \mathtt{N}(\boldsymbol{\mu}; \rho^2 I_N)$ and $Z' \sim \mathtt{N}(\boldsymbol{\mu}'; \rho'^2 I_N)$ are independent and that $Y_i = (ZZ')_i$ for $1 \leq i \leq N$. Any linear combination $U = \gamma Y_i + \gamma' Y_{i'}$ of Y_i and $Y_{i'}$ can be approximated by a univariate Normal distribution for large N. Thus $(Y_i, Y_{i'})$ can be approximated by a bivariate Normal distribution for large N.*

Proof. (Sketch.) For readability, the full proof of Lemma 3 is provided in Appendix D. The proof relies on Lemmas 8, 9, 10, 11, and 12, which are stated and proved in Appendix D, and we give a brief sketch here. In Lemma 8, we show that Y_i and $Y_{i'}$ can both be expressed with respect to a particular quadratic form. In Lemma 9, we show how the result of Lemma 8 can be generalised to an arbitrary linear combination of Y_i and $Y_{i'}$, expressed as $U = \gamma Y_i + \gamma' Y_{i'}$. In particular Lemma 9 shows how U can be expressed as a sum of two different random variables: the first is a quadratic form in independent and identically distributed Normal $N(0,1)$ random variables, and the second is a Normal random variable. This quadratic form is defined in terms of the eigenvalues of a specified matrix, which arises from the polynomial multiplication of Z and Z'. Thus if we show that this quadratic form has an approximate Normal distribution, then U overall is also approximately Normal. In Lemma 10, we establish results about the powers of these eigenvalues that are needed in the proof of Lemma 11. In Lemma 11, we invoke the Lyapunov Central Limit Theorem [Bil95, Theorem 27.3] to show that a scaled version of the quadratic form random variable of U asymptotically has a Normal distribution. This requires that the Lyapunov quotient tends to zero, which is established using Lemma 12, under an assumption (assured by restricting to Z and Z' arising in the multiplication of fresh ciphertexts) that certain quantities are bounded. Lemma 3 then follows as a corollary of Lemma 11. □

The arguments used in the proof of Lemma 3 can be extended to show that any linear combination $\gamma_{i_1} Y_{i_1} + \ldots + \gamma_{i_k} Y_{i_k}$ of a finite collection $(Y_{i_1}, \ldots, Y_{i_k})$ of k components of Y can be approximated as a univariate Normal distribution for large N, and so this collection $(Y_{i_1}, \ldots, Y_{i_k})$ has an approximate multivariate Normal distribution for large N.

It would also seem possible to adapt the proof of Lemma 3 to situations beyond the multiplication of fresh ciphertexts by modelling the dependence on N of components of μ and μ'. Lemma 12 could then potentially be generalised by using $N^{-\frac{1}{2}(1+d)}$ for an appropriate choice of d, rather than the simpler $N^{-\frac{1}{2}}$ used when modelling fresh messages.

3.3 BGV noise analysis

In this section, we apply results from Section 3.1 to give an average-case noise analysis for BGV. Under the assumptions specified in Section 3.1, these show how the noise in a ciphertext output from each BGV operation follows a Normal distribution with zero mean and a specified component variance. As most of these results are obtained entirely analogously to those for CKKS developed in [CCH+23], we summarise the results in Table 2 and defer their justification to Appendix E.

An exception is the `ModSwitch` operation, which is used to move from a ciphertext modulus q to a smaller modulus p. This operation does not exist for CKKS, although it is similar to the CKKS operation `Rescale`[1]. We first give in Lemma 4 a technical result (proved in Appendix F) that we will rely on in the following lemma. Lemma 5 then shows that a Normal distribution approximates the noise random variable for a ciphertext obtained from a BGV `ModSwitch` operation.

Lemma 4. *Suppose that q and t are odd positive integers with $q \gg t$, and let $\mathcal{Q} = \{-\frac{1}{2}(q-1), \ldots, \frac{1}{2}(q-1)\}$ and $\mathcal{T} = \{-\frac{1}{2}(t-1), \ldots, \frac{1}{2}(t-1)\}$. If $Z \sim Uni(\mathcal{Q})$ and $|\gamma| \ll 1$*

[1]The main difference is that in CKKS the message is in the high-order bits, so `Rescale` also serves the purpose of ensuring the message is scaled up by an appropriate amount. In BGV the message is in the low-order bits, so modulus switching does not affect the encoding of the message.

BGV Operation	Component variance of input noise(s)	Component variance of output noise(s)
`Encrypt`	-	$\left(\frac{4}{3}N+1\right)t^2\sigma^2$
`Add`	ρ^2, ρ'^2	$\rho^2+\rho'^2$
`Multiply`	ρ^2, ρ'^2	$N\rho^2\rho'^2+\rho'^2\lvert\mathtt{m}\rvert^2+\rho^2\lvert\mathtt{m}'\rvert^2$
`Relinearize`	ρ^2	$\rho^2+\frac{1}{12}N(\ell+1)w^2t^2\sigma^2$
`ModSwitch`	ρ^2	$\gamma^2\rho^2+\frac{1}{12}\left(\frac{2}{3}N+1\right)(t^2-1)$

Figure 2: Component variances in the zero-mean Normal random variable giving the noise in the output ciphertext after BGV homomorphic evaluation operations on input ciphertexts with input noises given by the zero-mean Normal random variables of the given component variances. The input ciphertexts to `Multiply` encrypt messages $\mathtt{m}$ and $\mathtt{m}'$. The parameter N denotes the power-of-two cyclotomic ring dimension. The parameter t is the plaintext modulus. The parameter σ is the standard deviation of the Ring-LWE error. In `Relinearize`, a ciphertext component with respect to ciphertext modulus q is decomposed into $\ell+1$ terms with in base w where $\ell := \lfloor \log_q(q) \rfloor$. The parameter $\gamma := p/q$ where q is the original ciphertext modulus and p is the ciphertext modulus after `ModSwitch`.

then the random variable

$$(Z - \lfloor \gamma Z \rceil) \mod t$$

has a distribution very close to Uniform on $\mathcal{T}$.

Lemma 5. **[ModSwitch]** *Suppose that a BGV ciphertext* (c_0, c_1) *with respect to a modulus* q *has a* 0*-mean multivariate Normal noise random variable given by* $V \sim \mathtt{N}(0; \rho^2 I_N)$*. Then the output ciphertext* (c_0', c_1') *after a* `ModSwitch` *operation of this ciphertext to a modulus* $p \ll q$ *has noise random variable* $V_{\mathsf{mod\text{-}sw}(s)} \sim \mathtt{N}(0; \rho^2_{\mathsf{mod\text{-}sw}(s)} I_N)$*, where the component variance* $\rho^2_{\mathsf{mod\text{-}sw}(s)}$ *is given in terms of the contraction factor* $\gamma = \dfrac{p}{q}$ *as*

$$\rho^2_{\mathsf{mod\text{-}sw}} = \gamma^2\rho^2 + \tfrac{1}{12}(\tfrac{2}{3}N+1)(t^2-1).$$

Proof. Lemma 1 shows that the output ciphertext (c_0', c_1') (with modulus p) following the application of the BGV `ModSwitch` to the input ciphertext (c_0, c_1) (with modulus q) is given by

$$c_i' = \lfloor \gamma c_i \rceil + ((c_i - \lfloor \gamma c_i \rceil) \bmod t) \qquad [i = 0, 1].$$

In order to analyse the BGV `ModSwitch` operation, we define

$$U_i = ((c_i - \lfloor \gamma c_i \rceil) \bmod t) = c_i' - \lfloor \gamma c_i \rceil \qquad [i = 0, 1],$$

which we can regard as integer random variables with independent components U_{ij} taking values in the set $\mathcal{T} = \{\frac{1}{2}(t-1), \ldots, \frac{1}{2}(t-1)\}$ of modulo t values (where t is odd). Lemma 4 shows that these integer random variables U_{ij} have a distribution very close to Uniform on T for BGV moduli $p \ll q$ (as $|\gamma| \ll 1$).

The BGV Critical Value $W_{\mathsf{mod\text{-}sw}}$ for the decryption of this ciphertext (c_0', c_1') obtained from the BGV `ModSwitch` operation is

$$\begin{aligned} W_{\mathsf{mod\text{-}sw}(s)} &= c_0' + s c_1' = \lfloor \gamma c_0 \rceil + s \lfloor \gamma c_1 \rceil + (U_0 + s U_1) \\ &= \gamma(c_0 + s c_1) + (U_0 + s U_1) + (\lfloor \gamma c_0 \rceil + s \lfloor \gamma c_1 \rceil - \gamma(c_0 + s c_1)) \\ &= \gamma W + (U_0 + s U_1) + ((\lfloor \gamma c_0 \rceil - \gamma c_0) + s(\lfloor \gamma c_1 \rceil - \gamma c_1)), \end{aligned}$$

We note that the final term $(\lfloor \gamma c_0 \rceil - \gamma c_0) + s(\lfloor \gamma c_1 \rceil - \gamma c_1)$ arises from rounding components to the nearest integers. Thus this term is negligible as each component consists of the

Table 1: The column $\overline{x}$ gives the observed mean of the noise budget in HElib ciphertexts over 10000 trials of the homomorphic evaluation described in the first circuit and in [CLP20] for parameter sets with dimension $N \in \{2048, 4096, 8192, 16384\}$. The column [CLP20] gives an estimate of the noise budget using worst-case heuristic bounds as given in that work. The column 'Ours' gives an estimate of the noise budget using our average case approach. The column [HS20] gives the remaining noise budget estimated using the `getNoiseBound()` function of HElib. The entry '-' denotes that the parameter set was too small to support this operation.

N	Enc				Add				Mult				ModSwitch			
	[CLP20]	Ours	[HS20]	$\overline{x}$	[CLP20]	Ours	[HS20]	$\overline{x}$	[CLP20]	Ours	[HS20]	$\overline{x}$	[CLP20]	Ours	[HS20]	$\overline{x}$
2048	35.0	41.0	43.2	48.7	34.0	41.0	42.2	48.2	17.0	26.0	31.5	39.1	-	-	-	-
4096	89.0	96.0	97.6	104	88.0	95.0	96.6	103	70.0	80.0	85.4	93.5	39.0	46.0	32.4	40.6
8192	199	206	207	213	198	205	206	213	179	189	194	203	148	155	141	149
16384	417	425	426	433	416	424	425	432	396	407	412	422	366	374	358	368

sum of $(1 + |s|) \approx (\frac{2}{3}n + 1)$ $\mathrm{Uni}((-\frac{1}{2}, \frac{1}{2}))$ rounding random variables, and so for practical purposes the BGV `ModSwitch` Critical Value is given by

$$W_{\mathsf{mod\text{-}sw}(s)} = \gamma W + (U_0 + sU_1).$$

The BGV `ModSwitch` noise random variable $V_{\mathsf{mod\text{-}sw}}$ corresponding to this BGV `ModSwitch` Critical Value is given by

$$V_{\mathsf{mod\text{-}sw}(s)} = \gamma V + (U_0 + sU_1).$$

The first term $\gamma V \sim \mathtt{N}(0; \gamma^2\rho^2 I_N)$ in this expression has a symmetric multivariate Normal distribution with mean 0 and component variance $\gamma^2\rho^2$. A component $(U_0 + sU_1)_i$ of the second term $U_0 + sU_1$ is a sum of $(1 + |s|^2)$ independent $\mathrm{Uni}(-\frac{1}{2}t, \frac{1}{2}t)$ random variables, so the Central Limit Theorem shows that the component $(U_0 + sU_1)_i$ can be regarded as having a Normal distribution with $\mathtt{N}(0, \frac{1}{12}(1 + |s|^2)(t^2 - 1)I_N)$ for large N with component variance $\frac{1}{12}(1 + |s|^2)(t^2 - 1)$. Thus the noise random variable $V_{\mathsf{mod\text{-}sw}}$ of BGV `ModSwitch` operation has a distribution given by

$$V_{\mathsf{mod\text{-}sw}(s)} \sim \mathtt{N}\left(0; \rho^2_{\mathsf{mod\text{-}sw}(s)} I_N\right), \text{ where } \rho^2_{\mathsf{mod\text{-}sw}(s)} = \gamma^2\rho^2 + \tfrac{1}{12}(1 + |s|^2)(t^2 - 1).$$

We conclude using the fact that s is chosen from a uniform ternary distribution, similarly to the Justification for Heuristic 2 in Appendix E. □

3.4 Experimental evaluation of efficacy

In this section, we illustrate the efficacy of the average-case approach for BGV noise analysis presented in Section 3.3 by comparing the noise growth predicted by this approach with observed noise growth in both HElib [HEl19] and SEAL [SEA22] and with the noise growth predicted by worst-case bounds as developed in [CLP20] following Iliashenko [Ili19]. Our experiments use HElib version 2.2.1 and SEAL version 4.0. We show that our average-case analysis can more closely estimate the practical noise growth than the prior worst-case approach of [CLP20]. To do so, we consider the homomorphic evaluation of two circuits. The results for HElib are displayed in Tables 1 and 2 respectively. The results for SEAL are displayed in Tables 3 and 4 respectively. For HElib, we also compared our approach to the noise predictions that are inbuilt in HElib using the `getNoiseBound()` function.

Experimental setup. The first circuit considered is the same circuit as was used in [CLP20]. The evaluation is as follows in the i-th trial. First, fresh ciphertexts $\mathtt{ct}_1$ and $\mathtt{ct}_2$ encrypting $i + 1$ and i are generated. Next, $\mathtt{ct}_3$ is generated as the homomorphic

Table 2: The column $\overline{x}$ gives the observed mean of the noise budget in HElib ciphertexts over 10000 trials of the homomorphic evaluation described above in the second circuit for parameter sets with dimension $N \in \{4096, 8192, 16384\}$. The column [CLP20] gives an estimate of the noise budget using worst-case heuristic bounds as given in that work. The column 'Ours' gives an estimate of the noise budget using our average case approach. The column [HS20] gives the remaining noise budget estimated using the `getNoiseBound()` function of HElib.

N	Enc				Mult1				Mult2				Mult3			
	[CLP20]	Ours	[HS20]	$\overline{x}$	[CLP20]	Ours	[HS20]	$\overline{x}$	[CLP20]	Ours	[HS20]	$\overline{x}$	[CLP20]	Ours	[HS20]	$\overline{x}$
4096	89.0	96.0	97.6	104	71.0	80.0	86.4	94.3	35.0	49.0	64.0	75.8	0	0	19.2	38.8
8192	199	206	207	213	180	189	195	203	142	156	171	184	66.0	90.0	125	145
16384	417	425	426	433	397	407	413	422	357	372	389	402	277	302	340	361

Table 3: The column $\overline{x}$ gives the observed mean of the noise budget in SEAL ciphertexts over 10000 trials of the homomorphic evaluation described in the first circuit and in [CLP20] for parameter sets with dimension $N \in \{2048, 4096, 8192, 16384\}$. The column [CLP20] gives an estimate of the noise budget using worst-case heuristic bounds as given in that work. The column 'Ours' gives an estimate of the noise budget using our average case approach.

N	Enc			Add			Mult			ModSwitch		
	[CLP20]	Ours	$\overline{x}$	[CLP20]	Ours	$\overline{x}$	[CLP20]	Ours	$\overline{x}$	[CLP20]	Ours	$\overline{x}$
4096	34.0	40.0	44.0	33.0	40.0	43.0	0	6.00	8.00	0	6.00	2.00
8192	135	142	146	134	141	145	97.0	106	111	95.0	102	95.0
16384	349	357	360	348	356	360	310	321	323	304	312	304
32768	784	792	796	783	792	795	744	755	759	733	741	734

Table 4: The column $\overline{x}$ gives the observed mean of the noise budget in SEAL ciphertexts over 10000 trials of the homomorphic evaluation described above in the second circuit for parameter sets with dimension $N \in \{4096, 8192, 16384\}$. The column [CLP20] gives an estimate of the noise budget using worst-case heuristic bounds as given in that work. The column 'Ours' gives an estimate of the noise budget using our average case approach.

N	Enc			Mult1			Mult2			Mult3		
	[CLP20]	Ours	$\overline{x}$	[CLP20]	Ours	$\overline{x}$	[CLP20]	Ours	$\overline{x}$	[CLP20]	Ours	$\overline{x}$
16384	349	357	361	311	321	325	235	250	252	83.0	108	104
32768	784	792	796	745	756	757	667	683	676	511	537	515

addition of $\mathtt{ct}_1$ and $\mathtt{ct}_2$. Next, $\mathtt{ct}_4$ is generated as the homomorphic multiplication of $\mathtt{ct}_3$ and $\mathtt{ct}_2$. For $N > 2048$, $\mathtt{ct}_5$ is generated by modulus switching $\mathtt{ct}_4$ down to the next prime in the chain (for $N = 2048$ the parameters are too small to support this operation). We measure the noise budget in each of 10000 trials. We then compute the average of these noise budget measurements, and report that value in our results tables. The results for HElib and SEAL are presented in Table 1 and Table 3 respectively.

We also explore the noise growth in a second, deeper, circuit, which is a multiplication tree of eight independent ciphertexts. We used the same parameter settings as the previous experiment. The evaluation is as follows in the i-th trial. First, fresh ciphertexts $\mathtt{ct}_1$, ..., $\mathtt{ct}_8$ encrypting $i+1$, ..., $i+8$ respectively are generated. Next, ciphertexts $\mathtt{ct}_9$, ..., $\mathtt{ct}_{12}$ are generated as the multiplication of $\mathtt{ct}_1$ and $\mathtt{ct}_2$; ...; $\mathtt{ct}_7$ and $\mathtt{ct}_8$ respectively. Next ciphertexts $\mathtt{ct}_{13}$ and $\mathtt{ct}_{14}$ are generated as the multiplication of $\mathtt{ct}_9$ and $\mathtt{ct}_{10}$; and $\mathtt{ct}_{11}$ and $\mathtt{ct}_{12}$ respectively. Finally, ciphertext $\mathtt{ct}_{15}$ is generated as the multiplication of $\mathtt{ct}_{13}$ and $\mathtt{ct}_{14}$. We measure the noise budget in each of 10000 trials. We then compute the average of these noise budget measurements, and report that value in our results tables.The results for HElib and SEAL are presented in Table 2 and Table 4 respectively. Note that in Tables 1, 2, 3 and 4, the columns labelled as "Mult" refer to the noise after

the tensor product (i.e. without ciphertext maintenance operations), so as to compare with Lemma 15.

For both circuits, the HElib parameters were chosen as follows. The standard deviation of the error distribution was set to $\sigma = 3.2$, the ring dimension was set to $N \in \{2048, 4096, 8192, 16384\}$ and the corresponding maximal ciphertext modulus q was set so that $\log q \in \{54, 109, 218, 438\}$. The plaintext modulus was set as $t = 3$. Other parameters are set according to HElib default parameter settings, detailed in [CLP20]. The parameter set $N = 2048$ is omitted in Table 2 as it is too small to support the homomorphic evaluation of the circuit.

For both circuits, the SEAL parameters were chosen as follows. The standard deviation of the error distribution was set to $\sigma = 3.2$, the ring dimension was set to $N \in \{4096, 8192, 16384, 32768\}$ and the corresponding maximal ciphertext modulus q was set so that $\log q \in \{109, 218, 438, 881\}$. The plaintext modulus was set to be a suitable integer of 20 bits, a default choice in the SEAL examples. In SEAL, the parameter sets with $N \in \{4096, 8192\}$ were too small to support the deeper circuit.

Developing bounds from output variance. We present average case bounds for each operation as follows: we trace through the component variance of the noise polynomial after each operation, using the formulae in Figure 2. We model the variance after multiplication as in Heuristic 3. We then translate the variance after each operation into a bound on the noise after each operation following the approach described in [CCH+23]. That is, we allow an error tolerance α (we set $\alpha = 0.001$ in the experiments), such that our noise bound is exceeded with probability α. Recent work [CSBB24, CCP+24, ABMP24] has shown that BGV can be vulnerable to attacks if there are decryption errors, and so in practice, it may be necessary to choose an exponentially small α.

Lemma 6 ([CCH+23])**.** *Suppose a noise polynomial is distributed as* $\mathtt{N}(0, \rho^2 I_N)$*. For a threshold* $T > 0$*, the error tolerance* $\alpha = \mathbf{P}(\|Z\|_\infty > T)$ *satisfies*

$$T = \sqrt{2} \cdot \rho \cdot \mathrm{erf}^{-1}((1-\alpha)^{\frac{1}{N}}) .$$

We express our results in terms of the *noise budget* (Definition 7). Loosely speaking, the noise budget is the number of bits left for homomorphic computation before a wraparound modulo q that would lead to decryption failure. We find the noise budget intuitive to have an idea of "how much space" is left within $\log q$ for the noise to grow. It is also the API for noise provided by SEAL.

Definition 7 ([CLP20])**.** Let `ct` be a BGV ciphertext with respect to modulus q having Critical Value W modulo q. The *noise budget* for this ciphertext is defined as

$$\log_2(q) - \log_2(\|W\|) - 1 .$$

Results. Our findings show mixed results. The HElib results in Tables 1 and 2 show that our average-case approach can much more closely model the observed noise growth for fresh ciphertexts, addition, and multiplication than the prior approach of [CLP20]. The improvement in closing the heuristic-to-practical gap identified in [CLP20] can be significant. For example, the gap is reduced by as much as 25 bits in the case of the deeper circuit. On the other hand, the [HS20] estimates are generally closer than ours to the practically observed noise, although a gap between these estimates and the observed noise remains. The difference between our estimates and the [HS20] estimates can be very minor (e.g. less than two bits for fresh ciphertexts and after addition), but grows with each multiplication. For modulus switching, Table 1 shows that [CLP20] estimates better predict the remaining noise budget than the [HS20] estimates.

The SEAL results of Tables 3 and 4 also show that the average-case heuristics can more closely model the observed noise growth for fresh ciphertexts, addition, and multiplication, including deeper multiplication. In several cases, the heuristic-to-practical gap is reduced to only 3-5 bits. However, in other cases, the average-case heuristic overestimates the remaining noise budget. This is similar to the findings of [CCH+23] for the CKKS scheme, and suggests caution should be employed when relying on this average-case analysis for setting parameters.

Discussion. There are some discrepancies between the SEAL implementation and the heuristic estimates that may account for differences between the observed and predicted behaviour. For example, in Table 4, for $N = 16384$, after the third multiplication, our average-case heuristic overestimates the remaining noise budget by one bit. We do not relinearize (in doing so, diverging from the SEAL recommendations), so by the third multiplication in the second circuit, the ciphertexts are much larger. This introduces additional noise not accounted for in the heuristics. We would expect such an additional noise to increase as N increases, and this expectation is confirmed by the results for $N = 32768$. Moreover, modifying our experiments to relinearize inputs before the next multiplication significantly reduces (but does not totally account for) the overestimation.

For modulus switching, in both libraries, the remaining noise budget is overestimated by our average-case approach. We did not investigate precisely what causes the analysis to fail for modulus switching. The reason for not doing so is that the subsequent work [CNP23] already provides an improved average-case noise analysis for BGV that is specific to its implementation in HElib, which experiments show closely models HElib noise growth. The main reason why the [CNP23] noise analysis is so effective is the observation that in HElib — using the natural prime set in accordance with the expected user behaviour — the mod switch noise is dominated by the rounding term only. In contrast, our heuristic analysis is given for a general situation of modulus switching to any p.

Both the worst-case and average-case heuristic estimates assume that the secret distribution is uniform ternary, as is done in our analysis of Section 3.3, and as is the distribution used in SEAL. The secret distribution implemented in HElib is also ternary, but with a slightly different variance[2]. We found that this discrepancy impacts the heuristic-to-practical gap only minimally. Indeed, adapting the heuristics for the HElib secret distribution made no difference in the predicted remaining average-case noise budget in low-depth computation, while for larger N, and after two or more multiplications, the predicted remaining noise budget was 2 bits closer to the observed remaining noise budget.

Potential practical utility of this average-case approach. The results for $N = 4096$ in Table 3 give an interesting example where the approach of [CLP20] predicts that there is no remaining noise budget after the multiplication, suggesting that the parameter set is too small to support the evaluation of this circuit. In contrast, our average-case analysis predicts there are 6 bits remaining, and indeed there is an observed average remaining noise budget of 8 bits. Following the approach of [CLP20] in this context would suggest to perform the computation in dimension at least $N = 8192$, leading to a consequent performance slowdown, as (e.g.) more coefficients must be processed during each homomorphic operation. In fact, as suggested by the average-case analysis, the computation can be performed for $N = 4096$. This illustrates that the use of an average-case approach can lead to a practical improvement when it comes to parameter selection, e.g., if N were to be chosen in an automated way based on the predicted noise growth by an FHE compiler.

[2] https://github.com/homenc/HElib/blob/f0e3e010009c592cd411ba96baa8376eb485247a/src/keys.cpp

Table 5: Estimates of the noise budget for the circuit parameterised by $L = 2$, $\zeta = 3$, $t = 257$, for the parameter set determined by the ring dimension N, obtained using our average-case approach and the prior approach of [CLP20].

N	[CLP20]	Ours
4096	0	19
8192	105	124

We can exhibit an additional specific computation for which the average-case approach predicts lower parameters to support the computation than the worst-case approach. This example is illustrative and we expect that many other such circuits could be found. To characterise a broad range of circuits, we focus on an L-level circuit with ζ additions and one multiplication at each level. We fix ciphertext moduli q that achieve 128-bit security according to the Homomorphic Encryption Security Standard [ACC+18] for error distribution standard deviation $\sigma = 3.2$, uniform ternary secret, and $N \in \{4096, 8192\}$; and allow to vary the plaintext modulus t. Given a circuit parameterised by L, ζ and t, we investigate the predicted noise growth for different parameter sets according to the average-case and worst-case approaches. Table 5 gives an example for $L = 2$, $\zeta = 3$, and $t = 257$. In this situation, our average-case approach predicts that the $N = 4096$ parameter set suffices to support the computation, while the approach of [CLP20] suggests $N = 8192$ is required. We implemented this circuit in HElib, and found indeed that the computation could be supported with $N = 4096$.

4 A CLT approach for `LPRHom` noise analysis

In this section, we present a Central Limit approach to `LPRHom` noise analysis. An overall summary of our approach is as follows. As done in [LPR13a], we analyse noise with respect to a decoding basis. However, our analysis is in H, whereas the analysis of [LPR13a] is in $R^\vee$; and our analysis uses a Central Limit approach, whereas [LPR13a] uses δ-subgaussians. The hope is, and what we eventually show is, that the use of this Central Limit approach leads to tighter bounds for decryption failure probability.

For simplicity, we restrict our discussion to the situation where m is prime, though our arguments apply more generally.

4.1 Additional background

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

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

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

The $p\Gamma$-basis *is the embedding under* σ *of the "decoding basis" (using the terminology of [LPR13a])* $\left\{\frac{p}{m}(1-\zeta_m^1), \frac{p}{m}(1-\zeta_m^2), \dots, \frac{p}{m}(1-\zeta_m^n)\right\}$ *of conjugate pairs for* $R^\vee$ *in* H.

If Z is a vector expressing an element of H as a vector of conjugate pairs in the I-basis (or standard basis) for H, then we have real-valued vectors $Z^{\ddagger} = T^{\dagger}Z$ and $Z^{*} = (p\Gamma)^{-1}Z$ expressing this element as a vector in the T-basis and the $p\Gamma$-basis for H respectively. The relevant properties of the (scaled) change-of-basis matrix $\Delta = \Gamma T^{-1}$ are given in Lemma 7, which is proved in Appendix G.

Lemma 7. 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 $\Delta\Delta^{T} = mI - J$.

The noise in a `LPRHom` ciphertext obtained as the output of a homomorphic multiplication of two fresh ciphertexts is the product of the noises in the input ciphertexts. We will therefore be interested in the $\otimes$-product (Definition 9) of two elements of H expressed in the T-basis.

Definition 9. *The* $\otimes$-product *of two real vectors* $u = (u_{11}, u_{12}, \ldots, u_{n'1}, u_{n'2})$ *and* $v = (v_{11}, v_{12}, \ldots, v_{n'1}, v_{n'2})$ *of length* $n = 2n'$ *is*

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

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

4.2 A Central Limit approximation of the distribution of $C^{(p\Gamma)}$

To obtain a Normal approximation for a weighted sum $\sum_{j=1}^{n} a_j X_j$ of the form encountered in `LPRHom`, we need a general form of the Central Limit Theorem formally given by the Lindeberg condition [Bil95, Str11]. We state such a Central Limit result in Lemma 17 in Appendix H. However, Lemma 17 can be informally expressed as that the weighted sum $\sum_{j=1}^{n} a_j X_j$ of the form encountered in Ring-LWE has an approximate Normal distribution for moderate or large n provided that the absolute weights a_j are not dominated by just a few values.

Proposition 1 gives a Central Limit approximation to a weighted multivariate sum of the form for independent and identically distributed random variables $X_1, \ldots, X_n$. This proposition is a summary of the Lindeberg condition for a Central Limit Theorem and essentially states that a good Normal approximation exists for the weighted sum if enough of the largest (in absolute value) weights are of comparable size. Concretely, in a typical parameter situation of Ring-LWE where we have $n > 10^2$, (or $n > 10^3$ in the case of homomorphic encryption), we can expect Proposition 1 to give a good approximation when as few as (for example) about 20 of the largest weights are comparable.

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 $\mathrm{Var}(X_j) = \rho^2$, so X has covariance matrix $\rho^2 I_n$. If A is a $n \times n$ matrix whose entries A_{jk} are not dominated by just a few of these entries, then the transformed random variable $AX \sim \mathtt{N}(0, \rho^2 AA^T)$ can be approximated as a multivariate Normal distribution for moderate or large n.

In Proposition 2, we apply Proposition 1 to approximate the distribution of the noise in a `LPRHom` ciphertext expressed in an appropriate decryption basis. We start with $C^{(T)}$, a vector expressing the noise in a `LPRHom` ciphertext in the T-basis for H, and observe that by the structure of H, the components of $C^{(T)}$ can be split into conjugate pairs.

In Proposition 1 we split $C^{(T)}$ into two sets of components where one of each conjugate pair is in each set. We then apply Proposition 1 twice, once to each set. Note that to invoke Proposition 1 requires an independence assumption. Fresh ciphertexts should have independent noise coefficients by construction, but we have not investigated the independence of noise coefficients in more general `LPRHom` ciphertexts.

Proposition 2. Suppose that $C^{(T)}$ is a vector expressing the noise in a `LPRHom` ciphertext in the T-basis for H, so a component $c_j^{(T)}$ of $C^{(T)}$ has mean $\mathbf{E}\left(c_j^{(T)}\right) = 0$ and finite variance $\mathrm{Var}\left(c_j^{(T)}\right) = \rho^2$, and suppose that besides its complex conjugate the component $c_j^{(T)}$ is independent of the other components. Suppose further that the S-basis given by the columns of an $n \times n$ 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 are not dominated by just a few values, then the distribution of the noise $C^{(S)}$ in this ciphertext in the (decryption) S-basis for H can be approximated as

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

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. 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 1 gives $\Psi' C^{(T)'} \sim \mathtt{N}(0; \rho^2 \Psi' \Psi'^T)$, and we similarly have $\Psi'' C^{(T)''} \sim \mathtt{N}(0; \rho^2 \Psi'' \Psi''^T)$. Thus

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

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

The Central Limit Theorem is formally a statement about the convergence (in distribution) of an appropriate weighted sum of random variables to a Normal distribution in the limit as the number of summands n tends to infinity. When such a result is applied in a concrete setting with a fixed finite n, it is reasonable to question the speed of this convergence, and in particular how accurate the approximation is. This issue is made more precise in [MP20], and can be verified empirically.

4.3 `LPRHom` decryption using the $p\Gamma$-basis

We now specify a decryption process for the `LPRHom` cryptosystem using the $p\Gamma$-basis of H (though any appropriate basis can be used). We recall 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 $\sigma(b^{-1}\mu)$ for the coset of the lattice $\sigma(pR^\vee)$ corresponding to the plaintext $\mu \in R_p$. Thus if the Embedded Noise is small enough with high probability, then we can recover the plaintext μ with high probability.

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

4.4 Decryption failure probabilities in the LPRHom cryptosystem

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

In the following, we denote by Q the "Q-function" giving the upper tail probability for a standard Normal $\mathtt{N}(0,1)$ distribution, so

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

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

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

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

Theorem 2. 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 LPRHom degree-1 ciphertext in the $p\Gamma$-basis for H is bounded by

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

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 $p^{-1}Z^\ddagger = (p^{-1}T^\dagger)Z \sim \mathtt{N}(0,\rho^2 I_n)$. 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 Proposition 2 and Lemma 7 show that

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

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

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

where the Q-function is as defined above. 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 [GS01] 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

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

□

Corollary 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 `LPRHom` degree-2 ciphertext in the $p\Gamma$-basis for H is bounded by

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

Proof. The decryption of a `LPRHom` 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 2 simply by using the appropriate moments, and so we can replace ρ in η_1 with $mp\rho_1\rho_2$ in to give $\eta_2 = \eta_1(n, q, mp\rho_1\rho_2) = \frac{1}{2}(n^{\frac{1}{2}}mp\rho_1\rho_2)^{-1}q$. □

References

[ABBB+22] Ahmad Al Badawi, Jack Bates, Flavio Bergamaschi, David Bruce Cousins, Saroja Erabelli, Nicholas Genise, Shai Halevi, Hamish Hunt, Andrey Kim, Yongwoo Lee, Zeyu Liu, Daniele Micciancio, Ian Quah, Yuriy Polyakov, Saraswathy R.V., Kurt Rohloff, Jonathan Saylor, Dmitriy Suponitsky, Matthew Triplett, Vinod Vaikuntanathan, and Vincent Zucca. Openfhe: Open-source fully homomorphic encryption library. In *Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, WAHC'22, pages 53–63, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3560827.3563379.

[ABMP24] Andreea Alexandru, Ahmad Al Badawi, Daniele Micciancio, and Yuriy Polyakov. Application-aware approximate homomorphic encryption: Configuring FHE for practical use. Cryptology ePrint Archive, Paper 2024/203, 2024. https://eprint.iacr.org/2024/203. URL: https://eprint.iacr.org/2024/203.

[ACC+18] M. Albrecht, M. Chase, H. Chen, J. Ding, S. Goldwasser, S. Gorbunov, S. Halevi, J. Hoffstein, K. Laine, K. Lauter, S. Lokam, D. Micciancio, D. Moody, T. Morrison, A. Sahai, and V. Vaikuntanathan. Homomorphic encryption security standard. Technical report, HomomorphicEncryption.org, 2018. URL: http://homomorphicencryption.org/wp-content/uploads/2018/11/HomomorphicEncryptionStandardv1.1.pdf.

[BGV12] 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. doi:10.1145/2090236.2090262.

[Bil95] P. Billingsley. *Probability and Measure*. Wiley, third edition, 1995.

[BM23] Beatrice Biasioli and Chiara Marcolla, 2023. Personal communication.

[BMCM23] Beatrice Biasioli, Chiara Marcolla, Marco Calderini, and Johannes Mono. Improving and automating BFV parameters selection: An average-case approach. *IACR Cryptol. ePrint Arch.*, page 600, 2023. URL: https://eprint.iacr.org/2023/600.

[CCH+23] Anamaria Costache, Benjamin R. Curtis, Erin Hales, Sean Murphy, Tabitha Ogilvie, and Rachel Player. On the precision loss in approximate homomorphic encryption. In Claude Carlet, Kalikinkar Mandal, and Vincent Rijmen, editors, *Selected Areas in Cryptography - SAC 2023 - 30th International Conference, Fredericton, Canada, August 14-18, 2023, Revised Selected Papers*, volume 14201 of *Lecture Notes in Computer Science*, pages 325–345. Springer, 2023. doi:10.1007/978-3-031-53368-6_16.

[CCP+24] Jung Hee Cheon, Hyeongmin Choe, Alain Passelègue, Damien Stehlé, and Elias Suvanto. Attacks against the INDCPA-d security of exact FHE schemes. Cryptology ePrint Archive, Paper 2024/127, 2024. https://eprint.iacr.org/2024/127. URL: https://eprint.iacr.org/2024/127.

[CGGI16] 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. doi:10.1007/978-3-662-53887-6_1.

[CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: fast fully homomorphic encryption over the torus. *J. Cryptology*, 33(1):34–91, 2020. doi:10.1007/s00145-019-09319-x.

[CKKS17] 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. doi:10.1007/978-3-319-70694-8_15.

[CLP20] Anamaria Costache, Kim Laine, and Rachel Player. Evaluating the effectiveness of heuristic worst-case noise analysis in FHE. In Liqun Chen, Ninghui Li, Kaitai Liang, and Steve A. Schneider, editors, *Computer Security - ESORICS 2020 - 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14-18, 2020, Proceedings, Part II*, volume 12309 of *Lecture Notes in Computer Science*, pages 546–565. Springer, 2020. doi:10.1007/978-3-030-59013-0_27.

[CNP23] Anamaria Costache, Lea Nürnberger, and Rachel Player. Optimisations and tradeoffs for HElib. In Mike Rosulek, editor, *Topics in Cryptology - CT-RSA 2023 - Cryptographers' Track at the RSA Conference 2023, San Francisco, CA, USA, April 24-27, 2023, Proceedings*, volume 13871 of *Lecture Notes in Computer Science*, pages 29–53. Springer, 2023. doi:10.1007/978-3-031-30872-7_2.

[CS16] Ana Costache and Nigel P. Smart. Which ring based somewhat homomorphic encryption scheme is best? In Kazue Sako, editor, *Topics in Cryptology - CT-RSA 2016 - The Cryptographers' Track at the RSA Conference 2016, San Francisco, CA, USA, February 29 - March 4, 2016, Proceedings*, volume 9610 of *Lecture Notes in Computer Science*, pages 325–340. Springer, 2016. doi:10.1007/978-3-319-29485-8_19.

[CSBB24] Marina Checri, Renaud Sirdey, Aymen Boudguiga, and Jean-Paul Bultel. On the practical CPAD security of "exact" and threshold FHE schemes and libraries. Cryptology ePrint Archive, Paper 2024/116, 2024. https://eprint.iacr.org/2024/116. URL: https://eprint.iacr.org/2024/116.

[FV12] J. Fan and F. Vercauteren. Somewhat Practical Fully Homomorphic Encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012. URL: https://eprint.iacr.org/2012/144.

[Gen09] 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. doi:10.1145/1536414.1536440.

[GHS12a] Craig Gentry, Shai Halevi, and Nigel P. Smart. Fully homomorphic encryption with polylog overhead. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2012. doi:10.1007/978-3-642-29011-4_28.

[GHS12b] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the AES circuit. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 850–867. Springer, 2012. doi:10.1007/978-3-642-32009-5_49.

[GNSJ24] Qian Guo, Denis Nabokov, Elias Suvanto, and Thomas Johansson. Key recovery attacks on approximate homomorphic encryption with non-worst-case noise flooding countermeasures. In *33rd USENIX Security Symposium (USENIX Security 24). Philadelphia, PA: USENIX Association*, 2024. Preproceedings version available at https://www.usenix.org/system/files/sec24summer-prepub-822-guo.pdf.

[GS01] G. Grimmett and D. Stirzaker. *Probability And Random Processes*. Oxford University Press, 3rd edition, 2001.

[GSW13] 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. doi:10.1007/978-3-642-40041-4_5.

[HEl19] HElib. https://github.com/homenc/HElib, January 2019.

[HS20] Shai Halevi and Victor Shoup. Design and implementation of HElib: a homomorphic encryption library. Cryptology ePrint Archive, Paper 2020/1481,

2020. https://eprint.iacr.org/2020/1481. URL: https://eprint.iacr.org/2020/1481.

[Ili19] I. Iliashenko. *Optimisations of fully homomorphic encryption*. PhD thesis, KU Leuven, 2019.

[KPP22] Andrey Kim, Antonis Papadimitriou, and Yuriy Polyakov. Approximate homomorphic encryption with reduced approximation error. In Steven D. Galbraith, editor, *Topics in Cryptology - CT-RSA 2022 - Cryptographers' Track at the RSA Conference 2022, Virtual Event, March 1-2, 2022, Proceedings*, volume 13161 of *Lecture Notes in Computer Science*, pages 120–144. Springer, 2022. doi:10.1007/978-3-030-95312-6_6.

[KPZ21] Andrey Kim, Yuriy Polyakov, and Vincent Zucca. Revisiting homomorphic encryption schemes for finite fields. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2021 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part III*, volume 13092 of *Lecture Notes in Computer Science*, pages 608–639. Springer, 2021. doi:10.1007/978-3-030-92078-4_21.

[LMSS22] B. Li, D. Micciancio, M. Schutz, and J. Sorrel. Securing Approximate Homomorphic Encryption using Differential Privacy. In Y. Dodis and T. Shrimpton, editors, *Advances in Cryptology - CRYPTO 2022*, volume LNCS 13507, pages 560–589, 2022. doi:10.1007/978-3-031-15802-5_20.

[LPR12] V. Lyubashevsky, C. Peikert, and O. Regev. On Ideal Lattices and Learning with Errors Over Rings. *IACR Cryptology ePrint Archive*, 2012:230, 2012. URL: https://eprint.iacr.org/2012/230.pdf.

[LPR13a] V. Lyubashevsky, C. Peikert, and O. Regev. A Toolkit for Ring-LWE Cryptography. *IACR Cryptology ePrint Archive*, 2013:293, 2013. URL: https://eprint.iacr.org/2013/293.

[LPR13b] 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. URL: https://doi.org/10.1007/978-3-642-38348-9_3.

[MP12] 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. doi:10.1007/978-3-642-29011-4_41.

[MP19] 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*, volume 11547 of *LNCS*, pages 251–268. Springer, 2019. doi:10.1007/978-3-030-21548-4_14.

[MP20] Sean Murphy and Rachel Player. Discretisation and Product Distributions in Ring-LWE. *Journal of Mathematical Cryptology*, 15:45–59, 2020. doi:10.1515/jmc-2020-0073.

[MR09] 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. doi:10.1007/978-3-540-88702-7_5.

[Pei16] Chris Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4):283–424, 2016. doi:10.1561/0400000074.

[Pla18] Rachel Player. *Parameter selection in lattice-based cryptography*. PhD thesis, Royal Holloway, University of London, 2018.

[Reg05] 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. doi:10.1145/1568318.1568324.

[Reg10] O. Regev. The Learning with Errors Problem (Invited Survey). In *IEEE Conference on Computational Complexity*, pages 191–204, 2010. doi:10.1109/CCC.2010.26.

[SEA22] Microsoft SEAL (release 4.0). https://github.com/Microsoft/SEAL, March 2022. Microsoft Research, Redmond, WA.

[SSTX09] 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. doi:10.1007/978-3-642-10366-7_36.

[Str11] D. Stroock. *Probability Theory: An Analytic View*. Cambridge University Press, 2011.

[TV11] T. Tao and Van Vu. Random matrices: Universality of local eigenvalue statistics. *Acta Mathematica*, 206:127–204, 2011. doi:10.1007/s11511-011-0061-3.

A The BGV scheme

Figure 3 outlines the BGV scheme as presented in [CLP20].

B Proof of Lemma 1

Proof. We let $r = \dfrac{q}{p}$, so the integer $r = 1 \bmod t$. The `ModSwitch` operation uses $\delta_i = -c_i \bmod r$ and $\delta_i = 0 \bmod t$ for $i = 0, 1$. The Chinese Remainder Theorem shows that δ_0 and δ_1 are uniquely defined modulo rt, so have coefficients lying between $-\frac{1}{2}rt$ and $\frac{1}{2}rt$. This specification for δ_i also gives

$$c_i + \delta_i = 0 \mod r \quad \text{and} \quad c_i + \delta_i = c_i \mod t \qquad [i = 0, 1].$$

In addition, the Chinese Remainder Theorem shows that $c_0 + \delta_0$ and that $c_1 + \delta_1$ have unique solutions modulo rt given by

$$c_0 + \delta_0 = rc_0 \mod rt \quad \text{and} \quad c_1 + \delta_1 = rc_1 \mod rt.$$

The parts of output ciphertext (c_0', c_1') after the `ModSwitch` operation therefore satisfy

$$c_0' = \frac{c_0 + \delta_0}{r} = c_0 \mod t \quad \text{and} \quad c_1' = \frac{c_1 + \delta_1}{r} = c_1 \mod t,$$

so the output ciphertext parts have the same values modulo t as the input ciphertext parts.

The BGV scheme. BGV is a (levelled) FHE scheme parameterised by N, q, t, χ, S, w, ℓ and λ. Let w be a base, then $\ell + 1 = \lfloor \log_w q \rfloor + 1$ is the number of terms in the decomposition into base w of an integer in base q. The Ring-LWE error distribution is denoted χ and is typically a discrete Normal with standard deviation $\sigma = 3.2$ [ACC+18]. The underlying Ring-LWE problem is parameterised by N, q, σ and S, where the parameter S denotes the secret key distribution. In implementations (e.g [HEl19, SEA22]), S is often chosen as a polynomial that has coefficients in $\{-1, 0, 1\}$. The security parameter is λ.

- $\texttt{SecretKeyGen}(\lambda)$: Sample $s \leftarrow S$ and output $\texttt{sk} = s$.
- $\texttt{PublicKeyGen}(\texttt{sk})$: Set $s = \texttt{sk}$ and sample $a \leftarrow \mathcal{R}_q$ uniformly at random and $e \leftarrow \chi$. Output $\texttt{pk} = ([-(as + te)]_q, a)$.
- $\texttt{EvaluationKeyGen}(\texttt{sk}, w)$: Set $s = \texttt{sk}$. For $i \in \{0, \ldots, \ell\}$, sample $b_i \leftarrow \mathcal{R}_q$ uniformly at random and $d_i \leftarrow \chi$. Output $\texttt{evk} = \big([-(b_i s + t d_i) + w^i s^2]_q, b_i\big)$.
- $\texttt{Encrypt}(\texttt{pk}, m)$: For the message $m \in \mathcal{R}_t$. Let $\texttt{pk} = (p_0, p_1)$, sample $u \leftarrow S$ and $e_1, e_2 \leftarrow \chi$. Output $\texttt{ct} = ([m + p_0 u + t e_1]_q, [p_1 u + t e_2]_q)$.
- $\texttt{Decrypt}(\texttt{sk}, \texttt{ct})$: Let $s = \texttt{sk}$ and $\texttt{ct} = (c_0, c_1)$. Output $m' = [[c_0 + c_1 s]_q]_t$.
- $\texttt{Add}(\texttt{ct}_0, \texttt{ct}_1)$: Output $\texttt{ct} = ([\texttt{ct}_0[0] + \texttt{ct}_1[0]]_q, [\texttt{ct}_0[1] + \texttt{ct}_1[1]]_q)$.
- $\texttt{Multiply}(\texttt{ct}_0, \texttt{ct}_1)$: Set $c_0 = [\texttt{ct}_0[0]\texttt{ct}_1[0]]_q$, $c_1 = [\texttt{ct}_0[0]\texttt{ct}_1[1] + \texttt{ct}_0[1]\texttt{ct}_1[0]]_q$, and $c_2 = [\texttt{ct}_0[1]\texttt{ct}_1[1]]_q$. Output $\texttt{ct} = (c_0, c_1, c_2)$.
- $\texttt{Relinearize}(\texttt{ct}, \texttt{evk})$: Let $\texttt{ct}[0] = c_0$, $\texttt{ct}[1] = c_1$ and $\texttt{ct}[2] = c_2$. Let $\texttt{evk}[i][0] = [-(b_i s + t d_i) + w^i s^2]_q$ and $\texttt{evk}[i][1] = b_i$. Express c_2 in base w as $c_2 = \sum_{i=0}^{\ell} c_2^{(i)} w^i$. Set $c_0' = c_0 + \sum_{i=0}^{\ell} \texttt{evk}[i][0] c_2^{(i)}$, and $c_1' = c_1 + \sum_{i=0}^{\ell} \texttt{evk}[i][1] c_2^{(i)}$. Output $\texttt{ct}' = (c_0', c_1')$.
- $\texttt{ModSwitch}(\texttt{ct}, p)$: For $p = q = 1 \mod t$ with p dividing q. Let $\texttt{ct} = (c_0, c_1)$. Fix δ_i such that $\delta_i = -c_i \pmod{\frac{q}{p}}$ and $\delta_i = 0 \pmod t$. Set $c_0' = \frac{p}{q}(c_0 + \delta_0)$ and $c_1' = \frac{p}{q}(c_1 + \delta_1)$. Output $\texttt{ct} = (c_0', c_1')$.

Figure 3: The BGV scheme as presented in [CLP20].

The output ciphertext parts c_0' and c_1' are "modulo p" polynomials with coefficients lying in $\{-\frac{1}{2}(p-1), \ldots, \frac{1}{2}(p-1)\}$ obtained as the direct contractions of "modulo q" polynomials as

$$c_0' = \frac{c_0 + \delta_0}{r} = \frac{p}{q}(c_0 + \delta_0) \quad \text{and} \quad c_1' = \frac{c_1 + \delta_1}{r} = \frac{p}{q}(c_1 + \delta_1)$$

We note that these new ciphertext parts can also be expressed as

$$c_0' = \frac{c_0}{r} + \frac{\delta_0}{r} \quad \text{and} \quad c_1' = \frac{c_0}{r} + \frac{\delta_1}{r},$$

where $\frac{\delta_0}{r}$ and $\frac{\delta_1}{r}$ are polynomials with coefficients between $-\frac{1}{2}t$ and $\frac{1}{2}t$. Thus the BGV $\texttt{ModSwitch}$ operation maps an input ciphertext part c_i to an output ciphertext part c_i', where c_i' is the nearest integer polynomial to $\frac{c_i}{r} = \frac{p}{q} c_i$ having the same value modulo t as c_i, which gives the expression in the statement of the Lemma. □

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

- Gen: choose $s' \leftarrow \lfloor \psi \rceil_{R^\vee}$, and output $s = b \cdot s' \in R$ as the secret key.
- $\mathrm{Enc}_s(\mu \in R_p)$: choose $e \leftarrow \lfloor p\psi \rceil_{b^{-1}\mu + pR^\vee}$. 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'.

Figure 4: The `LPRHom` cryptosystem as defined in [LPR13a, Section 8.3].

C The `LPRHom` cryptosystem

Figure 4 gives the `LPRHom` cryptosystem as defined in [LPR13a, Section 8.3].

D Proof of Lemma 3

We consider two distinct components $Y_i = (ZZ')_i$ and $Y_{i'} = (ZZ')_{i'}$ for $i \neq i'$ of the polynomial product ZZ' modulo $X^N + 1$. We establish that the bivariate random variable $(Y_i, Y_{i'})$ has an approximate bivariate Normal distribution by showing that any arbitrary linear combination $U = \gamma Y_i + \gamma' Y_{i'}$ has an approximate univariate Normal distribution.

Our approach begins by showing in Lemma 8 that both Y_i and $Y_{i'}$ can be expressed as a quadratic form in the same $2N$ independent standard Normal $\mathtt{N}(0,1)$ univariate random variables.

Lemma 8. *Suppose $Z \sim \mathtt{N}(\mu; \rho^2 I_N)$ and $Z' \sim \mathtt{N}(\mu'; \rho'^2 I_N)$ are independent symmetric N-dimensional random variables (representing polynomials), and that:*

- $\eta = \frac{1}{2}\rho\rho'$,
- $V = \begin{pmatrix} \rho^{-1}(Z-\mu) \\ \rho'^{-1}(Z'-\mu') \end{pmatrix} \sim \mathtt{N}(0; I_{2N})$,
- $\alpha_i = \begin{pmatrix} (-\xi(j)\ \rho\mu'_j) \\ (\xi(i-j)\ \rho'\mu_j) \end{pmatrix}$ *is a vector of dimension* $2N$,
- A_i *is an* $N \times N$ *matrix with* $(A_i)_{j',j} = \begin{cases} \xi(i-j) & [j' = i-j] \\ 0 & [j' \neq i-j], \end{cases}$
- $Q_i = \left(\begin{array}{c|c} 0 & A_i \\ \hline A_i & 0 \end{array}\right)$ *is a symmetric* $2N \times 2N$ *matrix,*
- $\mu^* = \mu\mu'$ *is the polynomial product of* μ *and* μ^*.

A component $Y_i = (ZZ')_i$ (for $i = 0, \ldots N-1$) of the polynomial product ZZ' (modulo $X^N + 1$) can be expressed as a quadratic form in a standard Normal $2N$-dimensional random variable $V \sim \mathtt{N}(0; I_{2N})$ as

$$Y_i = \sum_{j=0}^{N-1} \xi(i-j) Z_{i-j} Z'_j = \eta\, V^T Q_i V + \alpha_i^T V + \mu_i^*.$$

Furthermore, another component $Y_{i'} = (ZZ')_{i'}$ (for $i' \neq i$) of the polynomial product ZZ' (modulo $X^N + 1$) can be expressed as a quadratic form in the same standard Normal $2N$-dimensional random variable $V \sim \mathtt{N}(0; I_{2N})$ with the obvious changes in notation for $\alpha_{i'}$, $A_{i'}$ and $Q_{i'}$ as

$$Y_{i'} = \sum_{j=0}^{N-1} \xi(i'-j) Z_{i'-j} Z'_j = \eta\, V^T Q_{i'} V + \alpha_{i'}^T V + \mu_{i'}^*.$$

Proof. The matrix A_i satisfies $(A_i)_{j',j} = 1$ if $j + j' = i$, $(A_i)_{j',j} = -1$ if $j + j' = i + N$ and $(A_i)_{j',j} = 0$ otherwise. Thus A_i has constant "anti-diagonals" and so Q_i is a real symmetric matrix. We note (interpreting $i - j$ modulo N) that

$$\begin{aligned} Z_{i-j} Z'_j &= \rho\rho' \frac{(Z_{i-j} - \mu_{i-j})}{\rho} \frac{(Z'_j - \mu'_j)}{\rho'} \\ &\qquad + \rho\mu'_j \frac{(Z_{i-j} - \mu_{i-j})}{\rho} + \rho'\mu_{i-j} \frac{(Z'_j - \mu'_j)}{\rho} + \mu_{i-j}\mu'_j \\ &= 2\eta V_{i-j} V_{N+j} + \rho\mu'_j V_{i-j} + \rho'\mu_{i-j} V_{N+j} + \mu_{i-j}\mu'_j. \end{aligned}$$

Thus the quadratic form $Y_i = \sum_{j=0}^{N-1} \xi(i-j) Z_{i-j} Z'_j$ is given by

$$\begin{aligned} Y_i &= \sum_{j=0}^{N-1} \xi(i-j) \left(2\eta V_{i-j} V_{N+j} + \rho \mu'_j V_{i-j} + \rho' \mu_j V_{N+j} + \mu_{i-j} \mu'_j\right) \\ &= \eta\, V^T Q_i V + \sum_{j=0}^{N-1} \left(-\rho \xi(j) \mu'_{i-j} V_j + \rho' \xi(i-j) \mu_j V_{N+j}\right) + \mu^*_i \\ &= \eta\, V^T Q_i V + \alpha_i^T V + \mu^*_i. \end{aligned}$$

By construction, the $2N$-dimensional multivariate random variable $V \sim \mathtt{N}(0; I_{2N})$ has a standard multivariate Normal distribution. The stated result for $Y_{i'}$ then follows immediately. □

We now consider the arbitrary linear combination $U = \gamma Y_i + \gamma' Y_{i'}$, where without loss of generality we assume $\gamma^2 + \gamma'^2 = 1$, and we show that U has an approximate Normal distribution. We use the quadratic forms for Y_i and $Y_{i'}$ given in Lemma 8 to express U as a sum of $2N$ independent univariate random variables, as detailed in Lemma 9.

Lemma 9. *We use the notation of Lemma 8. Suppose that $Z \sim \mathtt{N}(\mu; \rho^2 I_N)$ and $Z' \sim \mathtt{N}(\mu'; \rho'^2 I_N)$ are independent symmetric N-dimensional random variables (representing polynomials), and suppose that $\lambda_0, \ldots, \lambda_{2N-1}$ are the $2N$ real eigenvalues of the $(2N \times 2N)$ matrix $Q = \gamma Q_i + \gamma' Q_{i'}$, where $\gamma^2 + \gamma'^2 = 1$. The linear combination $U = \gamma Y_i + \gamma' Y_{i'}$ of Y_i and $Y_{i'}$ can be expressed in terms of a sum of components of an orthogonal transformation β_i of $\gamma\alpha_i + \gamma'\alpha_{i'}$ and a quadratic form in $2N$ independent standard Normal random variables $W_0, \ldots, W_{2N-1} \sim \mathtt{N}(0,1)$ as*

$$U = \sum_{j=0}^{2N-1} (\eta \lambda_j W_j^2 + \beta_j W_j) + (\gamma \mu^*_i + \gamma' \mu^*_{i'}).$$

Proof. Lemma 8 shows that

$$Y_i = \eta\, V^T Q_i V + \alpha_i^T V + \mu^*_i \quad \text{and} \quad Y_{i'} = \eta\, V^T Q_{i'} V + \alpha_{i'}^T V + \mu^*_{i'}.$$

The linear combination $U = \gamma Y_i + \gamma' Y_{i'}$ is therefore a quadratic form given by

$$\begin{aligned} U &= \eta\, V^T (\gamma Q_i + \gamma' Q_{i'}) V + (\gamma\alpha_i + \gamma'\alpha_{i'})^T V + (\gamma \mu^*_i + \gamma' \mu^*_{i'}) \\ &= \eta\, V^T Q V + (\gamma\alpha_i + \gamma'\alpha_{i'})^T V + (\gamma \mu^*_i + \gamma' \mu^*_{i'}). \end{aligned}$$

The matrix $Q = \gamma Q_i + \gamma' Q_{i'}$ is real and symmetric, so its $2N$ eigenvalues $\lambda_0, \ldots, \lambda_{2N-1}$ are real and Q can be diagonalised. If P is a $2N \times 2N$ matrix of orthonormal column eigenvectors of Q, then P is an orthogonal matrix with $P^T P = P P^T = I_{2N}$ and $P^T Q P = D = \mathrm{Diag}(\lambda_0, \ldots, \lambda_{2N-1})$. If set $W = P^T V$, then W is an orthogonal transformation of the multivariate standard Normal random variable V and so $W \sim \mathtt{N}(0; I_{2N})$. Furthermore, if we set $\beta = P^T(\gamma\alpha_i + \gamma'\alpha_{i'})$, then β is an orthogonal transformation of $(\gamma\alpha_i + \gamma'\alpha_{i'})$, and the quadratic form form Y_i can be expressed as

$$\begin{aligned} U &= \eta\, V^T Q V + (\gamma\alpha_i + \gamma'\alpha_{i'})^T V + (\gamma \mu^*_i + \gamma' \mu^*_{i'}) \\ &= \eta\, V^T (P P^T) Q (P P^T) V + (\gamma\alpha_i + \gamma'\alpha_{i'})^T (P P^T) V + (\gamma \mu^*_i + \gamma' \mu^*_{i'}) \\ &= \eta\, (P^T V)^T (P^T Q P)(P^T V) + ((\gamma\alpha_i + \gamma'\alpha_{i'})^T P)(P^T V) + (\gamma \mu^*_i + \gamma' \mu^*_{i'}) \\ &= \eta\, W^T D W + \beta^T W + (\gamma \mu^*_i + \gamma' \mu^*_{i'}). \end{aligned}$$

The components $W_0, \ldots, W_{2N-1}$ of W are independent and identically distributed standard Normal $\mathtt{N}(0,1)$ random variables, and so the quadratic form U can be expressed as

$$\begin{aligned} U &= \sum_{i=0}^{2N-1} \eta\lambda_j W_j^2 + \sum_{i=0}^{2N-1} \beta_j W_j + (\gamma\mu_i^* + \gamma'\mu_{i'}^*) \\ &= \sum_{j=0}^{2N-1} (\eta\lambda_j W_j^2 + \beta_j W_j) + (\gamma\mu_i^* + \gamma'\mu_{i'}^*). \end{aligned}$$

□

The distribution of the linear combination $U = \gamma Y_i + \gamma' Y_{i'}$ depends on the eigenvalues of the matrix $Q = \gamma Q_i + \gamma' Q_{i'}$, and in particular the sum of squared eigenvalues and the sum of fourth powers of eigenvalues. Lemma 10 gives the relevant results for these sums of powers of eigenvalues of Q.

Lemma 10. *The $2N \times 2N$ matrix $Q = \gamma Q_i + \gamma' Q_{i'}$ (where $i \neq i'$ and $\gamma^2 + \gamma'^2 = 1$) is a real symmetric matrix with real eigenvalues $\lambda_0, \ldots, \lambda_{2N-1}$ satisfying*

$$\sum_{j=0}^{2N-1} \lambda_j^2 = 2N \quad \textit{and} \quad \sum_{j=0}^{2N-1} \lambda_j^4 \leq 3N.$$

Proof. The matrix $Q = \gamma Q_i + \gamma' Q_{i'} = \left(\begin{array}{c|c} 0 & \gamma A_i + \gamma' A_{i'} \\ \hline \gamma A_i + \gamma' A_{i'} & 0 \end{array}\right)$ is a real symmetric matrix, so has real eigenvalues. The matrices A_i and $A_{i'}$ satisfy $A_i^2 = A_{i'}^2 = I_N$, as for example the diagonal entries of A_i^2 are

$$(A_i^2)_{j,j} = \sum_{k=0}^{N-1} (A_i)_{j,k}(A_i)_{k,j} = (A_i)_{j,i-j}\xi(i-j) = \xi(i-j)^2 = 1,$$

and the off-diagonal entries ($j' \neq j$) of A_i^2 are

$$(A_i^2)_{j',j} = \sum_{k=0}^{N-1} (A_i)_{j',k}(A_i)_{k,j} = (A_i)_{j',i-j}\xi(i-j) = 0.$$

Similarly, the entries of the matrix $A_i A_{i'}$ (for $i \neq i'$) are

$$(A_i A_{i'})_{jk} = \sum_{l=0}^{N-1} (A_i)_{j,l}(A_{i'})_{l,k} = \left\{ \begin{array}{cc} \xi(i-j)\,\xi(j-(i-i')) & [(j-k) = (i-i')] \\ 0 & [(j-k) \neq (i-i')], \end{array} \right.$$

so in particular the diagonal entries of the matrix $A_i A_{i'}$ (for $i \neq i'$) are 0. Thus we have $\mathrm{Tr}(A_i^2) = \mathrm{Tr}(A_{i'}^2) = N$ and $\mathrm{Tr}(A_i A_{i'}) = \mathrm{Tr}(A_{i'} A_i) = 0$.

The matrix Q^2 has eigenvalues $\lambda_0^2, \ldots, \lambda_{2N-1}^2$, where

$$Q^2 = \left(\begin{array}{c|c} 0 & \gamma A_i + \gamma' A_{i'} \\ \hline \gamma A_i + \gamma' A_{i'} & 0 \end{array}\right)^2 = \left(\begin{array}{c|c} (\gamma A_i + \gamma' A_{i'})^2 & 0 \\ \hline 0 & (\gamma A_i + \gamma' A_{i'})^2 \end{array}\right).$$

The submatrix $(\gamma A_i + \gamma' A_{i'})^2$ of Q^2 is given by

$$(\gamma A_i + \gamma' A_{i'})^2 = \gamma^2 A_i^2 + \gamma'^2 A_{i'}^2 + \gamma\gamma' A_i A_{i'} + \gamma\gamma' A_i A_{i'} = I_N + \gamma\gamma'(A_i A_{i'} + A_{i'} A_i)$$

and has trace

$$\mathrm{Tr}\left((\gamma A_i + \gamma' A_{i'})^2\right) = \mathrm{Tr}(I_N) + \gamma\gamma'\mathrm{Tr}(A_i A_{i'}) + \gamma\gamma'\mathrm{Tr}(A_{i'} A_i) = N.$$

Thus the sum of squared eigenvalues of Q is given by

$$\sum_{j=0}^{2N-1} \lambda_j^2 = \mathrm{Tr}(Q^2) = 2\mathrm{Tr}\left((\gamma A_i + \gamma' A_{i'})^2\right) = 2N.$$

Results concerning the sum of fourth powers of eigenvalues of Q can be obtained by considering the matrix $Q^4 = \left(\begin{array}{c|c} (\gamma A_i + \gamma' A_{i'})^4 & 0 \\ \hline 0 & (\gamma A_i + \gamma' A_{i'})^4 \end{array}\right)$, which has eigenvalues $\lambda_0^4, \ldots, \lambda_{2N-1}^4$. We note that

$$\begin{aligned} (A_i A_{i'} + A_{i'} A_i)^2 &= (A_i A_{i'})^2 + (A_{i'} A_i)^2 + (A_i A_{i'}^2 A_i) + (A_{i'} A_i^2 A_{i'}) \\ &= 2I_N + (A_i A_{i'})^2 + (A_{i'} A_i)^2 \end{aligned}$$

as $A_i^2 = A_{i'}^2 = I_N$, so the submatrix $(\gamma A_i + \gamma' A_{i'})^4$ of Q^4 is given by

$$\begin{aligned} (\gamma A_i + \gamma' A_{i'})^4 &= (I_N + \gamma\gamma'(A_i A_{i'} + A_{i'} A_i))^2 \\ &= I_N + \gamma^2\gamma'^2(A_i A_{i'} + A_{i'} A_i)^2 + 2\gamma\gamma'(A_i A_{i'} + A_{i'} A_i)) \\ &= (1 + 2\gamma^2\gamma'^2) I_N + \gamma^2\gamma'^2((A_i A_{i'})^2 + (A_{i'} A_i)^2) \\ &\quad + 2\gamma\gamma'(A_i A_{i'} + A_{i'} A_i) \end{aligned}$$

and has trace given by

$$\mathrm{Tr}\left((\gamma A_i + \gamma' A_{i'})^4\right) = \left(1 + 2\gamma^2\gamma'^2\right) N + \gamma^2\gamma'^2 \mathrm{Tr}\left((A_i A_{i'})^2 + (A_{i'} A_i)^2\right).$$

The diagonal entries of the matrix $(A_i A_{i'})^2$ are given by

$$\left((A_i A_{i'})^2\right)_{j,j} = \sum_{k=0}^{N} (A_i A_{i'})_{j,k} (A_i A_{i'})_{k,j}$$

For a summand $(A_i A_{i'})_{j,k}(A_i A_{i'})_{k,j}$ of the above sum to be nonzero, we require both $(j-k) = (i-i')$ and $(k-j) = (i-i')$, so $2(i-i') = N$ giving the necessary condition $(i-i') = \frac{1}{2}N$ (as $i \neq i'$). Thus for the $(i-i') \neq \frac{1}{2}N$ case, we have $\left((A_i A_{i'})^2\right)_{j,j} = 0$ and $\mathrm{Tr}(\left((A_i A_{i'})^2\right) = 0$. In the other case where $(i-i') = \frac{1}{2}N$, we have

$$A_i A_{i+\frac{1}{2}N} = \pm \left(\begin{array}{c|c} 0 & I_{\frac{1}{2}N} \\ \hline -I_{\frac{1}{2}N} & 0 \end{array}\right),$$

where the sign depends on whether $i < \frac{1}{2}N$. Thus $(A_i A_{i+\frac{1}{2}N})^2 = -I_N$ whatever the value of the $\pm$ sign, and so $\mathrm{Tr}\left((A_i A_{i+\frac{1}{2}N})^2\right) = -N$. In summary, we have

$$\mathrm{Tr}\left((A_i A_{i'})^2\right) = \mathrm{Tr}\left((A_{i'} A_i)^2\right) = \left\{\begin{array}{cc} 0 & [(i-i') \neq \frac{1}{2}N] \\ -N & [(i-i') = \frac{1}{2}N]. \end{array}\right.$$

The trace of the submatrix $(\gamma A_i + \gamma' A_{i'})^4$ of Q^4 is therefore given by

$$\mathrm{Tr}\left((\gamma A_i + \gamma' A_{i'})^4\right) = \left\{\begin{array}{cc} (1 + 2\gamma^2\gamma'^2)N & [(i-i') \neq \frac{1}{2}N] \\ N & [(i-i') = \frac{1}{2}N]. \end{array}\right.$$

However, $2\gamma^2\gamma'^2 \leq \frac{1}{2}$ for $\gamma^2 + \gamma'^2 = 1$, so $\mathrm{Tr}\left((\gamma A_i + \gamma' A_{i'})^4\right) \leq \frac{3}{2}N$. Thus the sum of fourth powers of eigenvalues of Q satisfies

$$\sum_{j=0}^{2N-1} \lambda_j^4 = \mathrm{Tr}(Q^4) = 2\mathrm{Tr}\left((\gamma A_i + \gamma' A_{i'})^4\right) \leq 3N.$$

□

We now show that that the arbitrary linear combination $U = \gamma Y_i + \gamma' Y_{i'}$ has an approximate univariate Normal distribution by using the Lyapunov version of the Central Limit Theorem. This version of the Central Limit Theorem is suitable for situations such as this where the summands (of U) are independent but not identically distributed.

Theorem 3 (Lyapunov Central Limit Theorem ([Bil95] Theorem 27.3))**.** *Suppose X_1, X_2 , ..., X_n, ... are a sequence of independent random variables with finite mean κ_j and finite variance τ_j^2, and let $s_n = \sum_{j=1}^n \tau_j^2$. If for some $\delta > 0$, Lyapunov's condition*

$$\lim_{n\to\infty} \frac{1}{s_n^{2+\delta}} \sum_{j=1}^{n} \mathbf{E}\left(|X_j - \kappa_j|^{2+\delta}\right) = 0 \quad \textit{is satisfied,}$$

$$\textit{then} \quad \frac{1}{s_n} \sum_{j=1}^{n} (X_j - \kappa_j) \longrightarrow \mathtt{N}(0,1) \quad \textit{in distribution as } n \to \infty.$$

For simplicity, we assume that Z and Z' arise in the encryption of fresh messages. We make this assumption so that all random variables, and associated quantities such as the components of μ μ' or related vectors α, α' and β can be considered as bounded and all moments are finite, and in particular do not depend on N. We later discuss how to relax this assumption. In this fresh message situation, we consider the independent random variables

$$T_j = (2N)^{-\frac{1}{2}}(\eta\ \lambda_j W_j^2 + \beta_j W_j) \qquad [j = 0, \ldots, 2N-1],$$

and establish an asymptotic Normality in Lemma 11 (also using the technical Lemma 12) for the sum of these random variables by using the Lyapunov Central Limit Theorem with $\delta = 2$.

Lemma 11. *The sum of random variables*

$$\sum_{j=0}^{2N-1} T_j = (2N)^{-\frac{1}{2}} \sum_{j=0}^{2N-1} (\eta\ \lambda_j W_j^2 + \beta_j W_j)$$

tends in distribution to a standard Normal $\mathtt{N}(0,1)$ distribution as $N \to \infty$.

Proof. The mean of the random variable T_j is

$$\kappa_j = \mathbf{E}(T_j) = (2N)^{-\frac{1}{2}} \left(\eta\lambda_j \mathbf{E}(W_j^2) + \beta_j \mathbf{E}(W_j)\right) = (2N)^{-\frac{1}{2}} \eta\lambda_j,$$

so the centred version of this random variable is

$$\begin{aligned} T_j - \kappa_j &= (2N)^{-\frac{1}{2}} \left(\eta\lambda_j W_j^2 + \beta_j W_j - \eta\lambda_j\right) \\ &= (2N)^{-\frac{1}{2}} \left(\eta\lambda_j (W_j^2 - 1) + \beta_j W_j\right). \end{aligned}$$

Squaring this random variable gives

$$(T_j - \kappa_j)^2 = (2N)^{-1} \left(\eta^2\lambda_j^2 (W_j^2 - 1)^2 + \beta_j^2 W_j^2\right) + \text{Odd Degree in } W_j.$$

We note that $\mathbf{E}((W_j^2 - 1)^2) = 2$, $\mathbf{E}(W_j^2) = 1$, and $\mathbf{E}(\text{Odd Degree in } W_j) = 0$, so T_j has variance

$$\begin{aligned} \tau_j^2 &= \mathrm{Var}(T_j) = \mathbf{E}\left((T_j - \kappa_j)^2\right) \\ &= (2N)^{-1} \left(\eta^2\lambda_j^2 \mathbf{E}\left((W_j^2 - 1)^2\right) + \beta_j^2 \mathbf{E}\left(W_j^2\right)\right) \\ &= (2N)^{-1} \left(2\eta^2\lambda_j^2 + \beta_j^2\right) \end{aligned}$$

Lemma 10 shows that the sum of these variances is

$$s_{2N}^2 = \sum_{j=0}^{2N-1} \tau_j^2 = (2N)^{-1}\eta^2 \sum_{j=0}^{2N-1} \lambda_j^2 + (2N)^{-1} \sum_{j=0}^{2N-1} \beta_j^2 = \eta^2 + (2N)^{-1}|\beta|^2,$$

We now consider the Lyapunov Central Limit Theorem with $\delta = 2$ to establish the convergence of

$$\frac{1}{s_{2N}} \sum_{j=0}^{2N-1} (T_j - \kappa_j) \longrightarrow \mathtt{N}(0,1) \quad \text{in distribution as } N \to \infty.$$

This requires us to establish the limit as $N \to \infty$ of the Lyapunov quotient

$$\frac{1}{(s_{2N}^2)^2} \sum_{j=0}^{2N-1} \mathbf{E}\left((X_j - \kappa_j)^4\right) = \frac{\sum_{j=0}^{2N-1} \mathbf{E}\left((X_j - \kappa_j)^4\right)}{\left(\sum_{j=0}^{2N-1} \mathbf{E}\left((X_j - \kappa_j)^2\right)\right)^2}.$$

We first consider the denominator $s_{2N}^2 = \eta^2 + (2N)^{-1}|\beta|^2$ of the Lyapunov quotient. The vector β is an orthogonal transformation of $\gamma\alpha + \gamma'\alpha'$, so $|\beta|^2 = |\gamma\alpha + \gamma'\alpha'|^2$. However, the components of $\gamma\alpha + \gamma'\alpha'$ are bounded (with finite second moment under modelling assumption), and so as $N \to \infty$ we have

$$(s_{2N}^2)^2 = \sum_{j=0}^{2N-1} \left(\mathbf{E}((T_j - \kappa_j)^2)^2\right) = \eta^2 + (2N)^{-1}|\gamma\alpha + \gamma'\alpha'|^2 \to \text{Constant} > 0.$$

Lemma 12 shows that the numerator $\mathbf{E}\left((T_j - \kappa_j)^4\right)$ of the Lyapunov quotient satisfies $\mathbf{E}\left((T_j - \kappa_j)^4\right) \to 0$ as $N \to \infty$, so the Lyapunov quotient satisifies

$$\frac{\sum_{j=0}^{2N-1} \mathbf{E}\left((X_j - \kappa_j)^4\right)}{\left(\sum_{j=0}^{2N-1} \mathbf{E}\left((X_j - \kappa_j)^2\right)\right)^2} \longrightarrow 0 \quad \text{as } N \to \infty.$$

Thus the Lyapunov Central Limit Theorem with $\delta = 2$ shows that

$$\frac{1}{s_{2N}} \sum_{j=0}^{2N-1} T_j \longrightarrow \mathtt{N}(0,1) \text{ in distribution as } N \to \infty.$$

□

Lemma 12. *The numerator of the Lyapunov quotient satifies*

$$\mathbf{E}\left((T_j - \kappa_j)^4\right) \to 0 \quad \textit{as } N \to \infty.$$

Proof. The expression $T_j - \kappa_j = (2N)^{-\frac{1}{2}}\left(\eta\lambda_j(W_j^2 - 1) + \beta_j W_j\right)$ is given in the proof of Lemma 11, so for an appropriate polynomial g of terms of odd degree

$$(T_j - \kappa_j)^4 = (2N)^{-2}\left(\eta^4\lambda_j^4(W_j^2 - 1)^4 + \beta_j^4 W_j^4 + 6\eta^2\lambda_j^2\beta_j^2(W_j^2 - 1)^2 W_j^2\right) + g(W_j).$$

We note that $\mathbf{E}((W_j^2 - 1)^4) = 60$, $\mathbf{E}(W_j^4) = 3$, $\mathbf{E}((W_j^2 - 1)^2 W_j^2) = 10$ and $\mathbf{E}(g(W_j)) = 0$, so

$$\begin{aligned} \mathbf{E}\left((T_j - \kappa_j)^4\right) &= (2N)^{-2}\left(\begin{array}{c} \eta^4\lambda_j^4\mathbf{E}((W_j^2 - 1)^4) + \beta_j^4\mathbf{E}(W_j^4) \\ + 6\eta^2\lambda_j^2\beta_j^2\mathbf{E}((W_j^2 - 1)^2 W_j^2) \end{array}\right) \\ &= (2N)^{-2}\left(60\eta^4\lambda_j^4 + 60\eta^2\lambda_j^2\beta_j^2 + 3\beta_j^4\right) \end{aligned}$$

However, the geometric mean $(\eta^4\lambda_j^4\ \beta_j^4)^{\frac{1}{2}} = \eta^2\lambda_j^2\ \beta_j^2$ of $\eta^4\lambda_j^4$ and β_j^4 is less than or equal to the arithmetic mean $\frac{1}{2}(\eta^4\lambda_j^4 + \beta_j^4)$, so

$$60\ (\eta^2\lambda_j^2)\ \beta_j^2 \leq 30\eta^4\lambda_j^4 + 30\beta_j^4,$$

which shows that $\mathbf{E}\left((T_j - \kappa_j)^4\right)$ can be bounded as

$$\mathbf{E}\left((T_j - \kappa_j)^4\right) \leq (2N)^{-2}\left(90\eta^4\lambda_j^4 + 33\beta_j^4\right).$$

Lemma 10 therefore shows that the sum of the fourth power expectations satisfies

$$\begin{array}{rcl} \sum_{j=0}^{2N-1} \mathbf{E}\left((T_j - \kappa_j)^4\right) & \leq & (2N)^{-2}\sum_{j=0}^{2N-1}\left(90\eta^4\lambda_j^4 + 33\beta_j^4\right) \\ & \leq & (2N)^{-2}90\eta^4\sum_{j=0}^{2N-1}\lambda_j^4 + 33(2N)^{-2}\sum_{j=0}^{2N-1}\beta_j^4 \\ & \leq & (2N)^{-2}90\eta^4(3N) + 33(2N)^{-2}\sum_{j=0}^{2N-1}\beta_j^4 \\ & \leq & (2N)^{-1}135\eta^4 + 33(2N)^{-2}\sum_{j=0}^{2N-1}\beta_j^4. \end{array}$$

Fourth powers of components of β are bounded under the modelling assumption, so $(2N)^{-2}\sum_{j=0}^{2N-1}\beta_j^4 \to 0$. Thus we have shown that the Lyapunov quotient numerator satisfies $\sum_{j=0}^{2N-1}\mathbf{E}\left((T_j - \kappa_j)^4\right) \to 0$ as $N \to \infty$. □

We can finally establish the proof of Lemma 3.

Proof. The arbitrary linear combination $U = (\gamma Y_i + \gamma' Y_{i'})$ can be expressed as

$$U = \sum_{j=0}^{2N-1}(\eta\lambda_j W_j^2 + \beta_j W_j) + (\gamma\mu_i^* + \gamma'\mu_{i'}^*) = (2N)^{\frac{1}{2}}\sum_{j=0}^{2N-1} T_j + (\gamma\mu_i^* + \gamma'\mu_{i'}^*).$$

However, Lemma 11 shows that $\sum_{j=0}^{2N-1} T_j$ has an approximate univariate Normal distribution for large N. Thus any arbitrary linear combination $U = \gamma Y_i + \gamma' Y_{i'}$ of Y_i and $Y_{i'}$ has an approximate Normal distribution for large N, so $(Y_i, Y_{i'})$ has an approximate bivariate Normal distribution for large N. □

E Details of BGV noise analysis

We give a series of results showing how the noise in a ciphertext output from each BGV operation follows a Normal distribution with zero mean and a specified component variance, as summarised in Table 2.

We begin with Lemma 13 about the noise of a fresh BGV ciphertext obtained under a secret key s and an ephemeral encryption key u. Heuristic 2 then gives a Normal distribution that accurately approximates the noise random variable for a fresh BGV ciphertext. We note that a similar result can be inferred from Lemma 1 of [CLP20].

Lemma 13. [Fresh] *The noise random variable $V_{\mathsf{fresh}(s,u)}$ for a fresh BGV ciphertext obtained under a secret key s and an ephemeral encryption key u has a Normal distribution given by $V_{\mathsf{fresh}(s,u)} \sim \mathtt{N}(0; \rho^2_{\mathsf{fresh}(s,u)} I_N)$, where the component variance $\rho^2_{\mathsf{fresh}(s,u)}$ is given by*

$$\rho^2_{\mathsf{fresh}(s,u)} = (|s|^2 + |u|^2 + 1)t^2\sigma^2.$$

Proof. The first part of the public key $p_0 = [-(as + te)]_q$ (in the notation of Figure 3) can be expressed as $p_0 = -as - te + q\alpha$ for an appropriate integer vector α. For the second part of the public key $p_1 = a$, we therefore have $p_0 + sp_1 = -te + q\alpha$. The BGV Critical Value $W_{\mathsf{fresh}(s,u)}$ used for decryption of the fresh ciphertext (c_0, c_1) given by $c_0 = m + p_0 u + te_1$ and $c_1 = p_1 u + te_2$ corresponding to message m is given by

$$\begin{array}{rcl} W_{\mathsf{fresh}(s,u)} & = & c_0 + sc_1 = m + p_0 u + te_1 + s(p_1 u + te_2) \\ & = & m + u(-as - te + q\alpha) + te_1 + s(au + te_2) \\ & = & m + qu\alpha + t(-ue + e_1 + se_2). \end{array}$$

If the standard deviation of $t(-ue + e_1 + se_2)$ is not too large, reducing the BGV Critical Value W modulo q and then modulo t gives the message m. Thus the noise random variable corresponding to the BGV Critical Value $W_{\mathsf{fresh}(s,u)}$ is

$$V_{\mathsf{fresh}(s,u)} = t(-ue + e_1 + se_2).$$

Corollary 1 shows that $-ue \sim \mathtt{N}(0; |u|^2\sigma^2 I_N)$ and that $se_2 \sim \mathtt{N}(0; |s|^2\sigma^2 I_N)$, so the distribution of the fresh noise random variable $V_{\mathsf{fresh}(s,u)}$ is

$$V_{\mathsf{fresh}(s,u)} \sim \mathtt{N}(0; \rho^2_{\mathsf{fresh}(s,u)} I_N), \quad \text{where } \rho^2_{\mathsf{fresh}(s,u)} = (1 + |u|^2 + |s|^2)t^2\sigma^2.$$

□

Heuristic 2. [Fresh] *The noise random variable V_{fresh} for a fresh BGV ciphertext is accurately approximated as $V_{\mathsf{fresh}} \sim \mathtt{N}(0; \rho^2_{\mathsf{fresh}} I_N)$ by a Normal distribution with component variance ρ^2_{fresh} given by*

$$\rho^2_{\mathsf{fresh}} = (\tfrac{4}{3}N + 1)t^2\sigma^2.$$

Justification. Lemma 13 shows that $\rho^2_{\mathsf{fresh}(s,u)} = (1 + |u|^2 + |s|^2)t^2\sigma^2$, where s is the secret key and u is an ephemeral key. However s and u are random vectors of length N with components uniformly distributed in $\{-1, 0, 1\}$, so $|s|^2, |u|^2 \sim \mathrm{Bin}(N, \frac{2}{3})$ are independent Binomial random variables each with mean $\frac{2}{3}N$ and variance $\frac{2}{9}N$, giving $|s|^2 + |u|^2 \sim \mathrm{Bin}(2N, \frac{2}{3})$ with mean $\frac{4}{3}N$ and variance $\frac{4}{9}N$. The fresh noise random variable V_{fresh} is therefore the mixture of Normal distributions

$$\sum_{k=0}^{2N} \mathbf{P}(\mathrm{Bin}(2N, \tfrac{2}{3}) = k) \;\; \mathtt{N}\left(0; (k+1)t^2\sigma^2 I_N\right),$$

weighted by the Binomial probabilities $\mathbf{P}(\mathrm{Bin}(2N, \frac{2}{3}) = k) = \binom{2N}{k}\left(\frac{2}{3}\right)^k \left(\frac{1}{3}\right)^{2N-k}$. Thus V_{fresh} is a mixture of Normal distributions in which the high weight summand distributions are very similar, all having mean 0 and component variance close to $(\frac{4}{3}N + 1)t^2\sigma^2$, as the standard deviation of $|s|^2 + |u|^2$ is $\frac{2}{9}N^{\frac{1}{2}}$. The Normal mixture distribution for V_{fresh} can therefore be approximated as a Normal distribution as

$$V_{\mathsf{fresh}} \sim \mathtt{N}(0; \rho^2_{\mathsf{fresh}} I_N), \quad \text{where } \rho^2_{\mathsf{fresh}} = (\tfrac{4}{3}N + 1)t^2\sigma^2.$$

Next, Lemma 14 gives the distribution of the noise random variable following the application of the BGV `Add` operation to two BGV ciphertexts.

Lemma 14. [Add] *Suppose that the noise random variables V and V' for two independent BGV ciphertexts have 0-mean multivariate Normal distributions given by $V \sim \mathtt{N}(0; \rho^2 I_N)$ and $V' \sim \mathtt{N}(0; \rho'^2 I_N)$. Let V_{add} be the noise random variable for the ciphertext output from the BGV* `Add` *operation applied to these two ciphertexts, then $V_{add} \sim \mathtt{N}(0; \rho^2_{add} I_N)$, where the component variance ρ^2_{add} is given by*

$$\rho^2_{add} = \rho^2 + \rho'^2 .$$

Proof. Suppose that (c_0, c_1) and (c'_0, c'_1) are the independent BGV ciphertexts having respective underlying messages m and m' respectively and having the given noise random variables

$$V = (c_0 + sc_1) - m \sim \mathtt{N}(0; \rho^2 I_N) \;\text{ and }\; V' = (c'_0 + sc'_1) - m' \sim \mathtt{N}(0; \rho'^2 I_N).$$

The BGV `Add` operation gives the new ciphertext $(c_0 + c_1, c'_0 + c'_1)$ with message $m + m'$ and noise random variable

$$V_{\mathsf{add}} = (c_0 + c'_0) + s(c_1 + c'_1) - (m + m') = V + V' \sim \mathtt{N}(0; (\rho^2 + \rho'^2) I_N).$$

□

The application of the BGV `Multiply` operation to the BGV ciphertexts (c_0, c_1) and (c'_0, c'_1) gives a 3-part ciphertext

$$(c^*_0, c^*_1, c^*_2) = (c_0c'_0, c_0c'_1 + c_1c'_0, c_1c'_1)\,.$$

This 3-part ciphertext can potentially be decrypted by considering the 3-part `Multiply` Critical Value

$$\begin{aligned} W_{\mathsf{mult}} &= c^*_0 + sc^*_1 + s^2c^*_2 = c_0 + s(c'_0, c_0c'_1 + c_1c'_0) + s^2c_1c'_1 \\ &= (c_0 + sc_1)(c'_0 + sc'_1) = WW', \end{aligned}$$

where $W = c_0 + sc_1$ and $W' = c'_0 + sc'_1$ are the BGV Critical Values of the original ciphertexts (c_0, c_1) and (c'_0, c'_1). If m and m' are the messages corresponding to the ciphertexts (c_0, c_1) and (c'_0, c'_1), then the message $m \cdot m'$ corresponding to this 3-part ciphertext can be found by reducing this Critical Value W_{mult} modulo q and then modulo t. The distribution of the noise random variable following the application of the BGV `Multiply` operation is given in Lemma 15, and Heuristic 3 then gives a method for using this result by giving an expression for the component variance in circumstances of practical interest.

Lemma 15. [Multiply] *Suppose that the noise random variables V and V' for two independent BGV ciphertexts have 0-mean multivariate Normal distributions given by $V \sim \mathtt{N}(0; \rho^2 I_N)$ and $V' \sim \mathtt{N}(0; \rho'^2 I_N)$. Further suppose that the Small-S assumption is valid for the product distribution $(m + V)(m' + V')$, where m and m' are the underlying messages. Let V_{mult} be the noise random variables for the ciphertext output from the BGV* `Multiply` *operation applied to these two ciphertexts, then $V_{\mathsf{mult}(m,m')} \sim \mathtt{N}(0; \rho^2_{\mathsf{mult}(m,m')} I_N)$, where the component variance $\rho^2_{\mathsf{mult}(m,m')}$ is given by*

$$\rho^2_{\mathsf{mult}(m,m')} = N\rho^2\rho'^2 + \rho'^2|m|^2 + \rho^2|m'|^2.$$

Proof. Suppose that (c_0, c_1) and (c'_0, c'_1) are the independent BGV ciphertexts having respective underlying messages m and m' respectively and having the given noise random variables

$$V = (c_0 + sc_1) - m \sim \mathtt{N}(0; \rho^2 I_N) \;\text{ and }\; V' = (c'_0 + sc'_1) - m' \sim \mathtt{N}(0; \rho'^2 I_N).$$

The BGV multiplication operation gives the new 3-part ciphertext

$$(c_0c'_0, c_0c'_1 + c_1c'_0, c_1c'_1) \quad \text{with corresponding message } m \cdot m'.$$

The corresponding BGV Critical Value is

$$W_{\mathsf{mult}(m,m')} = (c_0 + sc_1)(c'_0 + sc'_1) = (m + V)(m' + V').$$

The corresponding noise random variable V_{mult} therefore has the same covariance matrix as the product of $m + V \sim \mathtt{N}(m; \rho^2 I_N)$ and $m' + V' \sim \mathtt{N}(m'; \rho'^2 I_N)$. The result then follows from Theorem 1, Heuristic 1, and Corollary 1. □

Heuristic 3. [Multiply] *In practice, we need to approximate $|m|^2$ and $|m'|^2$ to use Lemma 15. If the components of m and m' can be regarded as being independently and uniformly distributed on $\mathcal{T} = \{-\frac{1}{2}(t-1), \ldots, \frac{1}{2}(t-1)\}$, then the overall noise random variable V_{mult} for a BGV* `Multiply` *ciphertext is accurately approximated as $V_{\mathsf{mult}} \sim \mathtt{N}(0; \rho^2_{\mathsf{mult}} I_N)$ by a Normal distribution with component variance ρ^2_{mult} given by*

$$\rho^2_{\mathsf{mult}} = N\left(\rho^2\rho'^2 + \tfrac{1}{12}(t^2 - 1)(\rho^2 + \rho'^2)\right).$$

Justification If the components m_i and m'_i are uniformly distributed on $\mathcal{T}$, then $\mathrm{Var}(m_i) = \mathrm{Var}(m'_i) = \frac{1}{12}(t^2-1)$. In this case, we have $|m|^2, |m'|^2 \approx \frac{1}{12}N(t^2-1)$, and so a similar argument to the Justification for Heuristic 2 shows that the Normal mixture distribution for V_{mult} can be accurately approximated as a Normal distribution as

$$V_{\mathsf{mult}} \sim \mathtt{N}(0; \rho^2_{\mathsf{mult}} I_N), \quad \text{where } \rho^2_{\mathsf{mult}} \approx N\left(\rho^2\rho'^2 + \tfrac{1}{12}(t^2-1)(\rho^2+\rho'^2)\right).$$

The BGV `Relinearize` operation is used to convert a 3-part ciphertext arising after a BGV `Multiply` operation to a standard 2-part BGV ciphertext. The distribution of the Noise random variable following the application of a BGV `Relinearize` operation of the form described in Figure 3 is given in Lemma 16. The result is analogous to prior results [CLP20, Ili19, Pla18] about the BGV and BFV `Relinearize` operations.

We note that well-known implementations of BGV use more extensively optimised variants of this basic BGV `Relinearize` operation, so this result may need adapting for such optimised variants.

Lemma 16. [Relinearize] *Suppose that a 3-part BGV ciphertext arising from a BGV* `Multiply` *operation has a 0-mean multivariate Normal noise random variable given by* $V \sim \mathtt{N}(0; \rho^2 I_n)$. *Consider a BGV* `Relinearize` *operation with* $\ell+1$ *terms in the decomposition into base* w *of an integer in base* q *with* $\ell = \lfloor \log_w q \rfloor$ *in which a coefficient in* $\{-\frac{1}{2}(q-1), \ldots, \frac{1}{2}(q-1)\}$ *is represented as vector with* $(\ell+1)$ *components lying between* $-\frac{1}{2}w$ *and* $\frac{1}{2}w$. *Let* V_{relin} *be the noise random variable for the ciphertext output from such a BGV* `Relinearize` *operation, then* $V_{\mathsf{relin}} \sim \mathtt{N}(0; \rho^2_{\mathsf{relin}} I_N)$, *where the component variance* ρ^2_{relin} *is given by*

$$\rho^2_{\mathsf{relin}} = \rho^2 + \tfrac{1}{12}N(\ell+1)w^2t^2\sigma^2\,.$$

Proof. We consider the 3-part ciphertext $(c^*_0, c^*_1, c^*_2) = (c_0c'_0, c_0c'_1 + c_1c'_0, c_1c'_1)$ arising from the application of the BGV `Multiply` operation to the ciphertext (c_0, c_1) and the ciphertext (c'_0, c'_1). For a BGV scheme with parameter ℓ, the ciphertext component c^*_2, a polynomial with coefficients between $\frac{1}{2}(q-1)$ and $\frac{1}{2}(q-1)$, is expressed as

$$c^*_2 = \textstyle\sum_{i=0}^{\ell} g_i w^i, \;\; \text{for decompsition polynomials } g_i(x) = \sum_{j=0}^{N-1} g_{ij}x^j,$$

The integer coefficients g_{ij} of these decomposition polynomials g_i can be regarded as independent random variables lying uniformly between $-\frac{1}{2}w$ and $\frac{1}{2}w$, so we have $\mathbf{E}(g_{ij}) = 0$ and $\mathrm{Var}(g_{ij}) = \frac{1}{12}w^2$.

The BGV `Relinearize` operation transforms this 3-part ciphertext into a standard 2-part BGV ciphertext by using the Evaluation Keys

$$\alpha_i = -(\beta_i s + td_i) + w^i s^2 \quad \text{and} \quad \beta_i \quad [i = 0, \ldots, \ell],$$

where $\beta_0, \ldots, \beta_\ell$ are independent random elements of $\mathcal{R}_q$ and $d_0, \ldots, d_\ell$ are independent random variables with the error distribution χ, and we note that $\alpha_i + s\beta_i = s^2w^i - td_i$ The output of the BGV `Relinearize` operation is the 2-part ciphertext $(\overline{c}_0, \overline{c}_1)$ given by

$$\overline{c}_0 = c^*_0 + \sum_{i=0}^{\ell} \alpha_i g_i \quad \text{and} \quad \overline{c}_1 = c^*_1 + \sum_{i=0}^{\ell} \beta_i g_i.$$

The BGV Critical Value W_{relin} of this 2-part ciphertext $(\overline{c}_0, \overline{c}_1)$ is given by

$$\begin{aligned} W_{\mathsf{relin}} &= \overline{c}_0 + s\overline{c}_1 = c^*_0 + \textstyle\sum_{i=0}^{\ell} \alpha_i g_i + sc^*_1 + s\sum_{i=0}^{\ell} \beta_i g_i \\ &= c^*_0 + sc^*_1 + \textstyle\sum_{i=0}^{\ell} (\alpha_i + s\beta_i) g_i = c^*_0 + sc^*_1 + s^2 \sum_{i=0}^{\ell} w^i g_i - t\sum_{i=0}^{\ell} d_i g_i \\ &= c^*_0 + sc^*_1 + s^2c^*_2 - t\textstyle\sum_{i=0}^{\ell} d_i g_i = W - t\sum_{i=0}^{\ell} d_i g_i, \end{aligned}$$

where $W = c_0^* + sc_1^* + s^2c_2^*$ is the BGV Critical Value for the 3-part ciphertext (c_0^*, c_1^*, c_2^*). Thus the BGV `Relinearize` operation has noise random variable V_{relin} given by

$$V_{\mathsf{relin}} = V - t\sum_{i=0}^{\ell} d_i g_i$$

A component d_{ij} of d has mean $\mathbf{E}(d_{ij}) = 0$ and variance $\mathrm{Var}(d_{ij}) = \sigma^2$, and a component g_{ij} has mean $\mathbf{E}(g_{ij}) = 0$ and variance $\mathrm{Var}(g_{ij}) = \frac{1}{12}w^2$ as g_{ij} is uniformly distributed between $-\frac{1}{2}w$ and $\frac{1}{2}w$. Thus Lemma 2 shows that $d_i g_i \sim \mathtt{N}(0; \frac{1}{12}Nw^2\sigma^2)$, and hence that

$$t\sum_{i=0}^{\ell} d_i g_i \sim \mathtt{N}(0; \tfrac{1}{12}N(\ell+1)w^2t^2\sigma^2\ I_N).$$

Thus the BGV `Relinearize` operation has a noise random variable V_{relin} with a distribution

$$V_{\mathsf{relin}} \sim \mathtt{N}(0; (\rho^2 + \tfrac{1}{12}N(\ell+1)w^2t^2\sigma^2)I_N),$$

with component variance $\rho^2_{\mathsf{relin}} = \rho^2 + \frac{1}{12}N(\ell+1)w^2t^2\sigma^2$. □

F Proof of Lemma 4

Proof. We address the distribution of $((Z - \lfloor \gamma Z \rceil) \mod t)$ by considering the mapping

$$g\colon \mathcal{Q} \to \mathcal{T} \quad \text{given by } g(x) = (x - \lfloor \gamma x \rceil) \bmod t.$$

We first consider the restriction of the mapping g to the set

$$\mathcal{Q}_k = \{z \in \mathcal{Q} \mid z = k \bmod t\}$$

of "modulo q" values that have the value k modulo t, that is to say the mapping

$$g_k\colon \mathcal{Q}_k \to \mathcal{T} \quad \text{given by } g_k(x) = (x - \lfloor \gamma x \rceil) \bmod t.$$

For $|\gamma| \ll 1$, the mapping g_k can be regarded as a random mapping of the finite set $\mathcal{Q}_k$ to the finite set $\mathcal{T}$, so for $x \in \mathcal{Q}_k$ and $y \in \mathcal{T}$ the random variable

$$Z_{k,x,y} = \begin{cases} 1 & \text{if } g_k(x) = y \\ 0 & \text{if } g_k(x) \neq y \end{cases}$$

has a Bernoulli distribution with parameter $\dfrac{1}{\#T} = \dfrac{1}{t}$. The number the pre-images under g_k in $\mathcal{Q}_k$ of a given $y \in \mathcal{T}$ is the random variable

$$R_{k,y} = \sum_{x \in \mathcal{Q}_k} Z_{k,x,y} = \#\{x \in \mathcal{Q}_k \mid g^{-1}(y) = x\}.$$

Thus the number the pre-images of $y \in \mathcal{T}$ under g_k has a Binomial distribution as it is the sum of independent Bernoulli random variables, so we have

$$R_{k,y} \sim \mathrm{Bin}(\#\mathcal{Q}_k, t^{-1}).$$

The number of pre-images under g in $\mathcal{Q}$ of an element $y \in \mathcal{T}$ is the random variable

$$P_y = \sum_{k \in \mathcal{T}} R_{k,y} = \sum_{k \in \mathcal{T}} \sum_{x \in \mathcal{Q}_k} Z_{k,x,y} = \#\{x \in \mathcal{Q} \mid g^{-1}(y) = x\}.$$

The number of pre-images of $y \in \mathcal{T}$ under g in $\mathcal{Q}$ is therefore the sum of independent Binomial random variables with the same probability parameter, so has a Binomial distribution (noting that $\sum_{k \in \mathcal{T}} \#Q_k = \#\mathcal{Q} = q$) given by

$$P_y \sim \mathrm{Bin}(q, t^{-1})$$

with mean $\mathbf{E}(P_y) = \frac{q}{t}$ and variance $\mathrm{Var}(P_y) = \frac{q}{t}\left(1 - \frac{1}{t}\right)$. This Binomial distribution for P_y can be approximated using the standard technique of approximating such a Binomial distribution by an appropriate Normal distribution, given for example by the Central Limit approximation to the sum of $Z_{k,x,y}$. Thus we can write

$$P_y \sim \mathtt{N}\left(\frac{q}{t}, \frac{q}{t}\left(1 - \frac{1}{t}\right)\right).$$

The distribution of $g(Z) = ((Z - \lfloor \gamma Z \rceil) \bmod t)$ for $Z \sim \mathrm{Uni}(\mathcal{Q})$ is therefore given by $\mathbf{P}(g(Z) = k)$ for $k \in \mathcal{T}$, where

$$\mathbf{P}(g(Z) = k) \text{ is given by } \frac{P_y}{q} \sim \mathtt{N}\left(\frac{1}{t}, \frac{1}{tq}\left(1 - \frac{1}{t}\right)\right),$$

This random variable has mean $\mathbf{E}\left(\frac{P_y}{q}\right) = \frac{1}{t}$ and standard deviation satisfying

$$\text{St Dev}\left(\frac{P_y}{q}\right) = \frac{1}{(tq)^{\frac{1}{2}}}\left(1 - \frac{1}{t}\right)^{\frac{1}{2}} < \left(\frac{t}{q}\right)^{\frac{1}{2}}\frac{1}{t}.$$

However, $t \ll q$, so $\left(\frac{t}{q}\right)^{\frac{1}{2}}$ is very small, showing that $\mathbf{P}(g(Z) = k)$ is very close to $\frac{1}{t}$. Thus $g(Z)$ is approximated by a Uniform distribution on $\mathcal{T}$. □

G Proof of Lemma 7

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 $\Delta\Delta^T = (\Gamma^{\dagger}\Gamma)^{-1} = mI - J$. □

H Lindeberg Central Limit Theorem

Lemma 17 ([Bil95, Str11])**.** Suppose $X_1, X_2, \ldots$ are independent and identically distributed continuous random variables that are symmetric about 0 with mean $\mathbf{E}(X_j) = 0$

and variance $\mathrm{Var}(X_j) = 1$, and that have common density function f_{X_j}, and suppose that for constants $a_1, a_2, \ldots$ the sum $\sum_{j=1}^{l} a_j X_j$ has variance function $a(l)^2 = \sum_{j=1}^{l} a_j^2$, and that the functions $\widetilde{a}_j$ are defined by $\widetilde{a}_j(l) = \dfrac{|a_j|}{a(l)}$. In this case, *Lindeberg's condition* is that for any given $\epsilon > 0$, the sum

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

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