A Central Limit Approach for Ring-LWE Noise Analysis

Sean Murphy and Rachel Player

Royal Holloway, University of London, U.K.
s.murphy@rhul.ac.uk, rachel.player@rhul.ac.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 an average-case noise analysis for the BGV scheme. Our approach builds upon the recent work of Costache *et al.* 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 experimentally verify its improvement over prior, worst-case, approaches. 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 have been obtained in prior work.

1 Introduction

The Learning with Errors or *LWE* problem [30, 31] has become a standard hard problem in cryptology that is at the heart of lattice-based cryptography [25, 28]. The Ring Learning with Errors or *Ring-LWE* problem [33, 21] 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 [14]. 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 [4, 12, 17, 22, 6, 5], many of which [4, 12, 22, 5] are based on Ring-LWE. To illustrate the ideas of this paper, we first consider the widely-used BGV scheme [4], that has been implemented in (e.g.) [19, 32]. We also consider the symmetric key homomorphic cryptosystem given by Lyubashevsky, Peikert and Regev in Section 8.3 of [22] (the full version of [23]), which we term the `SymHom` cryptosystem.

Ciphertexts in all homomorphic encryption schemes contain an inherent noise that 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, we would model the noise growth as tightly as possible, so that the most performant parameters that meet the security and correctness requirements can be selected.

Prior approaches for noise analysis in BGV [15, 16, 10, 9] have been 'worst-case': that is, they have modelled the noise growth after every BGV evaluation operation using heuristic worst-case bounds. By tracing through the bounds after each operation, the noise growth incurred by the overall evaluation can also be bounded. However, there can be an unsatisfying gap between the final noise bound and the typical size of the noise as observed in experiments [9], with the gap growing as more computations are performed. In this work, we present for the first time an 'average-case' noise analysis for BGV, where average case is meant in the sense of the noise analysis for the TFHE scheme [6] as presented in [7]. That is, we show how for each homomorphic evaluation operation, the input and output noises can be modelled as a Gaussian 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. We therefore only need to resort to a bound after the evaluation, where the Gaussian distribution of the given variance implies a certain tail bound on the noise (holding with a certain probability). This enables us to set parameters that are still large enough to ensure correctness, but, due to the tighter analysis, 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 `SymHom` cryptosystem is that the noise growth in multiplication is non-linear. 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 [27]. For this reason, an average-case noise analysis for BGV, and related schemes, such as CKKS [5] and BFV [13], was believed until recently to be a challenging open question [9].

In this work, we demonstrate that a Central Limit approach can, under certain assumptions, be used to approximate the output noise of all BGV or `SymHom` operations as a Gaussian. We now expand in more detail on our approach for each scheme.

1.1 A Central Limit Approach for BGV

The first main contribution of this paper is to present an average-case noise analysis for BGV, based on a Central Limit argument.

Average-case analyses for noise growth in FHE schemes have been presented previously, for example for the TFHE scheme [6]. The approach, as presented in [7], is as follows. 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 [7, 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, thus permitting an average-case analysis.

Our approach is built upon the recent work of [8] that develops an average-case noise analysis for the CKKS scheme [5]. To this end, the analysis relies on Theorem 1, developed in [8], that gives the approximation of a polynomial product as a multivariate Normal distribution. Our analyses for the noise polynomials resulting from each BGV homomorphic operation follow from repeated applications of this result and is summarised in Figure 4. We expect that a similar approach could yield an average-case analysis for the BFV homomorphic encryption scheme [13]. Indeed, an analysis of the distributional properties of the multiplication of two polynomial ring elements could also be applicable in wider contexts, such as in analysis of lattice-based key encapsulation mechanisms [2, 11].

We additionally present an experimental verification of the analysis by comparing with practical noise growth in HElib [19] and SEAL [32]. The results are presented in Tables 1, 2, 3 and 4 and show that the average-case approach more tightly models the noise growth. Moreover, we demonstrate the applicability of our analysis by exhibiting specific computations for which the average-case approach predicts lower parameters to support the computation than the worst-case approach, and confirm this by successfully implementing these computations with the smaller parameter set.

1.2 A Central Limit Approach for `SymHom`

The second main contribution of this paper is to develop a statistical framework, based on a Central Limit argument, for analysing the noise in `SymHom` 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 `SymHom` ciphertexts. Our analysis can similarly be applied for higher-degree ciphertexts [27].

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 [22]. 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 SymHom). In particular, this means that we can express a component $c_j^{(p\Gamma)}$ of $C^{(p\Gamma)}$ as

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

The components $c_1^{(T)}, \ldots, c_n^{(T)}$ of $C^{(T)}$ are identically distributed random variables that are uncorrelated and, in general, independent, 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.

Theorem 2 and Corollary 2 demonstrate the improvement that can be obtained by using a Central Limit approach in comparison with prior bounds, such as those of [22], obtained using δ-subgaussian random variables [24, 26]. For example, if $\eta_1(n, q, \rho) = \frac{1}{2}(n^{\frac{1}{2}}\rho)^{-1}q$ is moderate or large, 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 [26, Lemma 18] in the manner of [22, Lemma 6.5].

No concrete parameter recommendations for SymHom are specified in [22], so in contrast to the situation with BGV, it is difficult to quantify the concrete improvement. Asymptotically, ignoring constants, we tighten the bound by a factor of $\omega(\sqrt{\log n})$, for power-of-two n and q following [22, Lemma 8.5]. However, we emphasise that using such a Central Limit approach in analysing SymHom has a number of advantages over other possible approaches, such as the subgaussian approach used in [22]. These advantages are listed below and expressed in terms of the above discussion.

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

2. A Central Limit approach gives an explicit approximating distribution for the cryptographic random variable of interest which can be directly used for general calculation or simulation purposes of use in cryptography. By contrast, a subgaussian approach can never give a explicit approximating distribution and can only give tail bounds. These tail bounds are generally weaker, as is evidenced by comparing our Theorem 2 with the bound that would be obtained following [22].
3. A Central Limit approach gives not only asymptotically an approximation to a Normal distribution, but also a close approximation concretely, for practically relevant Ring-LWE dimensions n.

1.3 Structure of the Paper

We recall relevant background and introduce new tools in Section 2. We outline our Central Limit approach for BGV in Section 3. We then outline our Central Limit approach for the `SymHom` 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

Encryption and decryption in Ring-LWE-based cryptography are inherently statistical processes, and we are giving Central Limit approximations to the distributions of cryptographic random variables of interest. Thus we use the notation $\sim$ 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 by taking a Central Limit Normal distribution approximation in statistical analysis. Furthermore, whilst Central Limit results are formally asymptotic results concerning sums or means of random variables, such Central Limit approximations usually apply in practice with relatively few summands (except perhaps for pathological distributions) as illustrated by the Berry-Esseen conditions [34] and related multidimensional versions [35]. For example, the simplest form of these Berry-Esseen conditions occurs for independent and identically distributed random variables $X_1, X_2, \ldots$ with mean $\mathbf{E}(X_i) = 0$. In this case, if $F_n(x) = \mathbf{P}(Y_n \leq x)$ is the distribution

function of $Y_n = \dfrac{n^{\frac{1}{2}}(X_1 + \ldots + X_n)}{\mathrm{Var}(X_i)^{\frac{1}{2}}}$ and Φ is the distribution function of a standard Normal random variable $Z \sim \mathtt{N}(0,1)$, then

$$|F_n(x) - \Phi(x)| < \frac{\mathbf{E}(|X_i|^3)}{\mathrm{Var}(X_i)^{\frac{3}{2}}}\, n^{-\frac{1}{2}}.$$

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.

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. The tensor product $K \otimes_{\mathbb{Q}} \mathbb{R}$ is denoted by $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$.

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

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

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

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

If we define t such that $t^{-1} = m^{-1}(1 - \zeta_m)$, then [22, Lemma 2.16] shows that $R^\vee = \langle t^{-1} \rangle$. We let $(R^\vee)^k$ denote the space of products of k elements of $R^\vee$, that is to say

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

2.4 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

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

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.

2.5 The BGV scheme

In this section we introduce the BGV scheme [4]. We generally follow the description of BGV given in [9], reproduced in Figure 1, that restricts to a power-of-two cyclotomic ring, $\mathcal{R} = \dfrac{\mathbb{Z}[X]}{(X^n+1)}$ for n a power of two. The plaintext space is given by $\mathcal{R}_t = \dfrac{\mathbb{Z}_t[X]}{(X^n+1)}$ and the ciphertext space is given by $\mathcal{R}_q = \dfrac{\mathbb{Z}_q[X]}{(X^n+1)}$. We generally regard a polynomial element of $\mathcal{R}_q$ as having coefficients in $\{-\frac{1}{2}(q-1), \dots, \frac{1}{2}(q-1)\}$. 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 $h = (h_0, \dots, h_{n-1})$ of coefficients in an appropriate context.

We now describe in our notation the relevant parts of the BGV scheme in order to define the noise in a BGV ciphertext.

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 gaussian with standard deviation $\sigma = 3.2$ [1]. 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 [19, 32]), 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 $\texttt{ct} = ([m + p_0 u + te_1]_q, [p_1 u + te_2]_q)$.
- $\texttt{Decrypt}(\texttt{sk}, \texttt{ct})$: 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 $\texttt{ct} = (c_0', c_1')$.

Fig. 1. The BGV scheme as presented in [9].

SecretKeyGen. For emphasis, we write the secret key as $s \in \{-1, 0, 1\}^n$, a ternary vector of length n, which can more generally be regarded as a polynomial of degree $n - 1$. We regard s as a constant vector known to the genuine receiver. 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. 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. This BGV Criticial Value is (we will show) an n-dimensional multivariate Normal random variable that arises during BGV decryption with secret key s. 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 1 that will be more convenient for our analysis. Lemma 1 can be seen as giving an explicit implementation of the `Scale` operation described in earlier analyses of BGV [10, 15].

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*

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

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 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 `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.

2.6 The `SymHom` scheme

In this section we introduce the `SymHom` cryptosystem. In order to do so, we first need two definitions. A description of `SymHom` cryptosystem, in the notation of [22], is then given in Figure 2.

Definition 5 ([26]). *The univariate* Balanced Reduction *function $\mathcal{R}$ on $\mathbb{R}$ is the random function* $\mathcal{R}(a) = \begin{cases} 1 - (\lceil a \rceil - a) & \textit{with probability} \quad \lceil a \rceil - a \\ -(\lceil a \rceil - a) & \textit{with probability} \quad 1 - (\lceil a \rceil - a). \end{cases}$

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 ([26]). *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 `SymHom` cryptosystem in order to define the noise in a `SymHom` ciphertext. We first recall that the `SymHom` secret key is an element $s \in R$, the plaintext space is R_p, and a plaintext $\mu \in R_p$ is encrypted to give a linear polynomial over $R^\vee_q$.

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

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

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

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

Fig. 2. The SymHom cryptosystem as defined in [22, Section 8.3].

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 $\sigma(pR^\vee) + \sigma(t^{-1}\mu)$ of the lattice $\sigma(pR^\vee)$ obtained by the canonical embedding of the scaled dual fractional ideal $pR^\vee$. We consider the coordinate-wise randomised rounding discretisation with respect to the $p\Gamma$-basis for H, and following Definition 6 we denote this discretisation of Y by $Y'(\mu) = \lfloor\, Y \,\rceil^{p\Gamma}_{\sigma(pR^\vee)+\sigma(t^{-1}\mu)}$.

The *Noise* random variable $Y''(\mu)$ in the encryption of the plaintext μ is then defined to be $Y''(\mu) = \sigma^{-1}(Y'(\mu))$, and is an element of a coset of $pR^\vee + t^{-1}\mu$ containing information about μ. For obvious reasons, we refer to $Y'(\mu) = \sigma(Y''(\mu))$ 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 3.

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 $C(\theta; \mu_1, \mu_2) = A_1'(\mu_1)A_2'(\mu_2) + (A_2A_1'(\mu_1) +$

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

Fig. 3. Notation for the Noise-related quantities used in encryption of the plaintext μ.

$A_1 A_2'(\mu_2))\ \theta + A_1 A_2 \theta^2$, which is given in terms of the secret key s and its constituent Noises $Y_1''(\mu)$ and $Y_2''(\mu)$ by

$$C(\theta; \mu_1, \mu_2) = A_1 A_2 (\theta - s)^2 + (A_2 Y_1''(\mu_1) + A_1 Y_2''(\mu_2))\ (\theta - s) + Y_1''(\mu_1) Y_2''(\mu_2).$$

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 to BGV noise analysis

3.1 BGV Polynomial Multiplication

Many BGV operations involve polynomial multiplication in $\mathcal{R}$ or $\mathcal{R}_q$, that is to say modulo $X^n + 1$, and we express such a polynomial multiplication 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 $\{0, \ldots, n-1\}$.

BGV requires to construct the polynomial product in $\mathcal{R}$ or $\mathcal{R}_q$ of a constant or scalar and a (discretised) multivariate Normal random variable or of two multivariate Normal random variables. We use the following result, developed for the CKKS context in [8].

Theorem 1. *Suppose that $Z \sim \mathtt{N}(\mu; \rho^2 I_n)$ and $Z' \sim \mathtt{N}(\mu'; \rho'^2 I_n)$, then the polynomial product ZZ' (modulo $X^n + 1$) is well-approximated as a multivariate Normal distribution for large n given by*

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

where $\rho_*^2 = n\rho^2\rho'^2 + \rho'^2|\mu|^2 + \rho^2|\mu'|^2$ and S is an off-diagonal matrix with $S_{i,i'} = \rho'^2 \sum_{j=0}^{n-1} \xi(i-j)\xi(i'-j)\mu_{i-j}\mu_{i'-j} + \rho^2 \sum_{j=0}^{n-1} \xi(i-j)\xi(i'-j)\mu'_{i-j}\mu'_{i'-j}$.

Following the approach of [8], we make the *Small-S assumption*: that this off-diagonal matrix S is negligible compared to $\rho_*^2 I_N$ and we disregard it. This assumption is reasonable in many circumstances of interest in BGV as the message vector length is generally bounded.

Corollary 1. *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 for Theorem 1 given in [8]. 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 BGV Noise Analysis

We now give a series of results showing how the noise in a ciphertext output from each BGV operation follows a Gaussian distribution with zero mean and a specified component variance. We begin with Lemma 3 about the noise of a fresh BGV ciphertext, and we note that a similar result can be inferred from Lemma 1 of [9].

Lemma 3. [Fresh] *The noise random variable* V_{fresh} *for a fresh BGV ciphertext has a Normal distribution given by* $V_{\mathsf{fresh}} \sim \mathtt{N}(0; \rho^2_{\mathsf{fresh}} I_N)$*, where the component variance* ρ^2_{fresh} *can be accurately approximated with high probability as*

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

Proof. The first part of the public key $p_0 = [-(as + te)]_q$ (in the notation of Figure 1) 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_{fresh} 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

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

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_{fresh} is

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

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_{fresh} is

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

The random vectors u and s have independent Uniform distributions on $\{-1, 0, 1\}^n$, so squared components u_i^2 and s_i^2 take the value 1 with probability $\frac{2}{3}$ and 0 with probability $\frac{1}{3}$. Thus both $|s|^2, |u|^2 \sim \mathrm{Bin}(n, \frac{2}{3})$ have Binomial distributions, so can be approximated by independent Normal $\mathtt{N}(\frac{2}{3}n, \frac{2}{9}n)$ distributions for large n. The distribution of $\rho^2_{\mathsf{fresh}} = (1 + |u|^2 + |s|^2)t^2\sigma^2$ can therefore be approximated as a Normal $\mathtt{N}((\frac{4}{3}n + 1)t^2\sigma^2, \frac{4}{9}n(t^2\sigma^2)^2)$ distribution. The standard deviation of ρ^2_{fresh} is $\frac{2}{3}(t^2\sigma^2)n^{\frac{1}{2}}$, which is small compared to the mean $(\frac{4}{3}n + 1)t^2\sigma^2$ of ρ^2_{fresh}. Thus ρ^2_{fresh} can be accurately approximated by $(\frac{4}{3}n+1)t^2\sigma^2$ with high probability. If the corresponding component standard deviation $(\frac{4}{3}n+1)^{\frac{1}{2}}t\sigma$ is small compared to q, so it does not generally affect any modular reduction, then the fresh ciphertext noise is $V_{\mathsf{fresh}} \sim \mathtt{N}(0; \rho^2_{\mathsf{fresh}} I_n)$, with noise variance $\rho^2_{\mathsf{fresh}} \approx (\frac{4}{3}n + 1)t^2\sigma^2$.

We now give a series of results about the noise distribution resulting from the application of BGV operations to BGV ciphertexts. We start with Lemma 4 giving the distribution of the noise random variable following the application of the BGV `Add` operation to two BGV ciphertexts.

Lemma 4. [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 $V_{\mathsf{add}} \sim \mathtt{N}(0; \rho^2_{\mathsf{add}} I_n)$, where the component variance ρ^2_{add} is given by*

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

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 BGV `Add` operation can also be used to add a ciphertext to itself. For competeness, we also give (without proof) the distribution of the noise random variable for such an integer multiple of a ciphertext in Lemma 5.

Lemma 5. [Integer Multiple] *Suppose that the noise random variable V of a BGV ciphertext (c_0, c_1) has 0-mean multivariate Normal distribution given by $V \sim \mathtt{N}(0; \rho^2 I_n)$. The noise random variable of the integer multiple $k(c_0, c_1)$ of the BGV ciphertext (c_0, c_1) for an integer k is $kV \sim \mathtt{N}(0; k^2\rho^2 I_n)$.*

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

$$(c^*_0, c^*_1, c^*_2) = (c_0 c'_0, c_0 c'_1 + c_1 c'_0, c_1 c'_1)\,.$$

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

$$\begin{aligned} W_{\mathsf{mult}} &= c^*_0 + s c^*_1 + s^2 c^*_2 = c_0 + s(c'_0, c_0 c'_1 + c_1 c'_0) + s^2 c_1 c'_1 \\ &= (c_0 + s c_1)(c'_0 + s c'_1) = W W', \end{aligned}$$

where $W = c_0 + s c_1$ and $W' = c'_0 + s c'_1$ 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 6.

Lemma 6. [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 distributions $m + V$ and $m' + V'$, 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 $V_{\mathsf{mult}} \sim \mathtt{N}(0; \rho^2_{\mathsf{mult}} I_n)$, where the component variance ρ^2_{mult} is given by*

$$\rho^2_{\mathsf{mult}} = n\rho^2\rho'^2 + \rho'^2|m|^2 + \rho^2|m'|^2.$$

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

$$V = (c_0 + s c_1) - m \sim \mathtt{N}(0; \rho^2 I_n) \;\text{ and }\; V' = (c'_0 + s c'_1) - m' \sim \mathtt{N}(0; \rho'^2 I_n).$$

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

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

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 and Corollary 1.

Remark 1. In practice, to use Lemma 6, we need to approximate $|m|^2$ and $|m'|^2$. If the components of m and m' can be regarded as being independently and uniformly distributed on $T = \{-\frac{1}{2}(t-1), \ldots, \frac{1}{2}(t-1)\}$, 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 the component variance ρ^2_{mult} can be accurately approximated with high probability as

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

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 1 is given in Lemma 7. The result is analogous to prior results [9, 20, 29] 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 7. [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

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

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 $\alpha_i + s\beta_i = s^2 w^i - t d_i$ 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

$$\begin{aligned} W_{\mathsf{relin}} &= \overline{c}_0 + s\overline{c}_1 = c_0^* + \textstyle\sum_{i=0}^{\ell} \alpha_i g_i + s c_1^* + s \sum_{i=0}^{\ell} \beta_i g_i \\ &= c_0^* + s c_1^* + \textstyle\sum_{i=0}^{\ell} (\alpha_i + s\beta_i) g_i = c_0^* + s c_1^* + s^2 \sum_{i=0}^{\ell} w^i g_i - t \sum_{i=0}^{\ell} d_i g_i \\ &= c_0^* + s c_1^* + s^2 c_2^* - t \textstyle\sum_{i=0}^{\ell} d_i g_i = W - t \sum_{i=0}^{\ell} d_i g_i, \end{aligned}$$

where $W = c_0^* + s c_1^* + s^2 c_2^*$ 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

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

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 $d_i g_i \sim \mathtt{N}(0; \frac{1}{12} n w^2 \sigma^2)$, and hence that

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

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

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

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

The BGV `ModSwitch` operation is, as we noted earlier, the key technical tool for noise management in BGV and is used to move from a modulus q to a smaller modulus p. The result of Lemma 8 gives an expression for the component variance $\rho^2_{\mathsf{mod\text{-}sw}}$ of the noise random variable for the BGV `ModSwitch` operation containing two terms. However, the first term $\gamma^2 \rho^2 = \left(\frac{p}{q}\right)^2 \rho^2$ is much smaller than the second term $\frac{1}{12}(\frac{2}{3}n + 1)(t^2 - 1)$ in many situations of interest, in which case the noise component variance $\rho^2_{\mathsf{mod\text{-}sw}} \approx \frac{1}{12}(\frac{2}{3}n + 1)(t^2 - 1)$ is constant and does not depend on the input noise variance ρ^2.

Lemma 8. [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 $p < q$ has noise random variable $V_{\mathsf{mod\text{-}sw}} \sim \mathtt{N}(0; \rho^2_{\mathsf{mod\text{-}sw}} I_n)$, where the component variance $\rho^2_{\mathsf{mod\text{-}sw}}$ can be accurately approximated with high probability in terms of the contraction factor $\gamma = \dfrac{p}{q}$ as*

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

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 $T = \{\frac{1}{2}(t-1), \ldots, \frac{1}{2}(t-1)\}$ of modulo t values (where t is odd) having an almost exactly Uniform distribution on T for BGV moduli $p < q$ (a full justification for this statement is omitted for space reasons).

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

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

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 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

$$W_{\mathsf{mod\text{-}sw}} = \gamma W + (U_0 + s U_1).$$

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

$$V_{\mathsf{mod\text{-}sw}} = \gamma V + (U_0 + s U_1).$$

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 $(U_0 + s U_1)_i$ of the second term $U_0 + s U_1$ 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 $(U_0 + s U_1)_i$ can be regarded as having a Normal distribution with

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	$n\rho^2\rho'^2+\rho'^2\lvert m\rvert^2+\rho^2\lvert m'\rvert^2$
`Relinearize`	ρ^2	$\rho^2+\frac{1}{12}n(\ell+1)w^2t^2\sigma^2$
`ModSwitch`	ρ^2	$\gamma^2\rho^2+\frac{1}{12}\left(\frac{2}{3}n+1\right)(t^2-1)$

Fig. 4. Component variances in the zero-mean Normal random variable giving the noise in the output ciphertext after homomorphic evaluation operations on input ciphertexts with input noises given by the zero-mean Normal random variables of the given component variances. (Other notation as given in Lemmas 3-8.)

$\mathtt{N}(0, \frac{1}{12}(1+|s|^2)(t^2-1)I_n)$ for large n with component variance $\frac{1}{12}(1+|s|^2)(t^2-1)$. Thus the BGV `ModSwitch` operation has a distribution given by

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

with component variance $\rho^2_{\mathsf{mod\text{-}sw}} = \gamma^2\rho^2 + \frac{1}{12}(1+|s|^2)(t^2-1)$. However, $|s|^2$ is the sum of n Uniform Ternary random variable, so is very close to $\frac{2}{3}n$ with high probability for large n, when the component variance can be accuately approximated as

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

3.3 Experimental verification

The analysis of the previous subsections shows that the noise in a BGV ciphertext can, under certain assumptions, be expressed as a multivariate Normal random variable after every homomorphic operation. The analysis is summarised in Figure 4. This enables us to model the noise growth throughout a BGV homomorphic evaluation in an 'average-case' manner, as was done for TFHE in [7]; that is, tracing through the variances at each operation, in order to find the variance of the noise in the output ciphertext. The variance can then be used to determine a bound on the noise in the output ciphertext to set parameters for correctness. Such an average-case approach is in contrast to a 'worst-case' analysis, as employed in prior studies on BGV [16, 10, 9]; that is, tracing through a bound on the noise at each operation in order to determine a bound on the noise in the output ciphertext. Since worst-case bounds are used in these prior studies at each step, we expect that the final bound could be very loose. This was confirmed by experiments in [9], which compared the observed noise growth in the HElib [19] implementation of BGV to the predicted noise growth from worst-case analyses.

In this section, we illustrate the efficacy of our average-case approach by comparing the noise growth predicted by these with observed noise growth in both HElib [19] and SEAL [32] and with the noise growth predicted by worst-case

bounds as developed in [9] following Iliashenko [20]. Our experiments use HElib version 2.2.1 and SEAL version 4.0. We show that our average-case analysis can tightly estimate the practical noise growth, thus closing the gap between worst-case predicted noise and practically observed noise highlighted in [9]. 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.

The first circuit considered is the same circuit as was used in [9]. 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 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 after each operation and output an average over 10000 trials. 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, using the same parameter settings as the previous experiment. The evaluation is as follows in the i-th trial. First, fresh ciphertexts $\mathtt{ct}_1, \ldots, \mathtt{ct}_8$ encrypting $i+1, \ldots, i+8$ respectively are generated. Next, ciphertexts $\mathtt{ct}_9, \ldots, \mathtt{ct}_{12}$ are generated as the multiplication of $\mathtt{ct}_1$ and $\mathtt{ct}_2$; $\ldots$; $\mathtt{ct}_7$ 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 after each multiplication and output an average over 10000 trials. The results for HElib and SEAL are presented in Table 2 and Table 4 respectively.

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 [9]. 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.

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 4. We model the variance after multiplication as in Remark 1. We then translate the variance after each operation into a bound on the noise after each operation following the approach described in [8]. 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.

Lemma 9 ([8]). *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.

Definition 7 ([9]). *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*

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

The HElib results in Tables 1 and 2 show that the average-case approach much more closely models the observed noise growth for fresh ciphertexts, addition, and multiplication. While the average-case modelling does not completely close the heuristic-to-practical gap identified in [9], the improvement is still significant. For example, the gap is reduced by as much as 25 bits in the case of the deeper circuit.

The SEAL results of Tables 3 and 4 are even more promising and show that the average-case heuristics tightly model the observed noise growth for fresh ciphertexts, addition, and multiplication, including deeper multiplication. In most cases, the heuristic-to-practical gap is reduced to only 3-5 bits.

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, the 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 the average-case approach. This may also be due to specificities in the libraries. For example, in our HElib implementation we modulus switch to the 'natural' prime set following the expected usage of the library, whereas the heuristic analyses are for a general situation of modulus switching to any p. Modifying HElib to explore this further is beyond the scope of this work.

Both our worst-case and average-case heuristic estimates assume that the secret distribution is uniform ternary, as is done in our analysis of Section 3.2, and as is the distribution used in SEAL. The secret distribution implemented

n	Enc			Add			Mult			ModSwitch		
	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$
2048	35.0	41.0	48.7	34.0	41.0	48.2	17.0	26.0	39.1	-	-	-
4096	89.0	96.0	104	88.0	95.0	103	70.0	80.0	93.5	39.0	46.0	40.6
8192	199	206	213	198	205	213	179	189	203	148	155	149
16384	417	425	433	416	424	432	396	407	422	366	374	368

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 [9] for parameter sets with dimension $n \in \{2048, 4096, 8192, 16384\}$. The column W gives an estimate of the noise budget using worst-case heuristic bounds following [9]. The column A gives an estimate of the noise budget using an average case approach.

n	Enc			Mult1			Mult2			Mult3		
	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$
4096	89.0	96.0	104	71.0	80.0	94.3	35.0	49.0	75.8	0	0	38.8
8192	199	206	213	180	189	203	142	156	184	66.0	90.0	145
16384	417	425	433	397	407	422	357	372	402	277	302	361

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 W gives an estimate of the noise budget using worst-case heuristic bounds following [9]. The column A gives an estimate of the noise budget using an average case approach.

in HElib is also ternary, but with a slightly different variance[1]. 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.

The results for $n = 4096$ in Table 3 give an interesting example where the worst-case approach 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, the average-case analysis predicts there are 6 bits remaining, and indeed there is an observed average remaining noise budget of 8 bits. To further illustrate the utility of the average-case approach, we now exhibit additional specific computations for which the average-case approach predicts lower parameters to support the computation than the worst-case approach. The examples here are 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 Homomor-

[1] `https://github.com/homenc/HElib/blob/f0e3e010009c592cd411ba96baa8376eb485247a/src/keys.cpp#L1145`

n	Enc			Add			Mult			ModSwitch		
	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$
4096	34.0	40.0	44.0	33.0	40.0	43.0	0	6.00	8.00	0	8.00	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 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 [9] for parameter sets with dimension $n \in \{2048, 4096, 8192, 16384\}$. The column W gives an estimate of the noise budget using worst-case heuristic bounds following [9]. The column A gives an estimate of the noise budget using an average case approach.

n	Enc			Mult1			Mult2			Mult3		
	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$	W	A	$\overline{x}$
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

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 W gives an estimate of the noise budget using worst-case heuristic bounds following [9]. The column A gives an estimate of the noise budget using an average case approach.

phic Encryption Security Standard [1] for error distribution standard deviation $\sigma = 3.2$, uniform ternary secret, and $n \in \{4096, 8192, 16384\}$; 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.

We exhibit in Table 5 an example for $L = 3$, $\zeta = 8$, and $t = 256$, following parameter choices in [9]. It can be seen that the worst-case approach indicates that the $n = 16384$ parameter set is required, while the average-case approach indicates that $n = 8192$ suffices. In Table 6, we see another example, for $L = 2$, $\zeta = 3$, and $t = 257$. In this situation, the average-case approach predicts that the $n = 4096$ parameter set suffices to support the computation, while the worst-case approach suggests $n = 8192$ is required. We implemented this latter circuit in HElib, and found indeed that the computation could be supported with $n = 4096$.

4 A CLT approach to `SymHom` noise analysis

In this section, we present a Central Limit approach to `SymHom` noise analysis. For simplicity, we restrict our discussion to the situation where m is prime, though our arguments apply more generally.

n	W	A
8192	0	17
16384	183	229

Table 5. The column W gives an estimate of the noise budget for the circuit parameterised by $L = 3$, $\zeta = 8$, $t = 256$, for the parameter set determined by the ring dimension n, using worst-case heuristic bounds following [9]. The column A gives an estimate of the noise budget for the same circuit using an average case approach.

n	W	A
4096	0	19
8192	105	124

Table 6. The column W gives an estimate 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, using worst-case heuristic bounds following [9]. The column A gives an estimate of the noise budget for the same circuit using an average case approach.

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

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

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

In Figure 5, we summarise our notation for elements of H expressed with respect to the various bases. If Z is a vector expressing an element of H as a vector of conjugate pairs in the I-basis (or standard basis) for H, then we have real-valued vectors $Z^\ddagger = T^\dagger Z$ and $Z^* = (p\Gamma)^{-1} Z$ expressing this element as a vector in the T-basis and the $p\Gamma$-basis for H respectively.

The change of basis transformations between the T-basis and the $p\Gamma$-basis are summarised in Figure 6, and the relevant properties of the (scaled) change-of-basis matrix $\Delta = \Gamma T^{-1}$ are given in Lemma 10.

Lemma 10. 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 $\Delta\Delta^T = mI - J$.

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

Fig. 5. Notation for the expression of an element of H as a vector in the various different vector space bases for H. Note that p is a scaling factor.

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

Fig. 6. Change of Basis Matrices for the T-basis and $p\Gamma$-basis for H in which elements of H are expressed as real-valued vectors.

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 $k \neq j$ (for $1 \le k, j \le n$), which yields

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

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

The noise in a `SymHom` 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 approach to approximate 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 `SymHom`, we need a general form of the Central Limit Theorem formally given by the Lindeberg condition [3, 34]. We state such a Central Limit result in Lemma 11. However, Lemma 11 can be informally expressed as that the weighted sum $\sum_{j=1}^{n} a_j X_j$ of the form encountered in Ring-LWE has an approximate Normal distribution for moderate or large n provided that the absolute weights a_j are not dominated by just a few values.

Lemma 11. 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 $l \to \infty$.

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 `SymHom` ciphertext expressed in an appropriate decryption basis. We note the proof of Proposition 2 is complicated by the fact that a pair of random variables in the T-basis arising as the image of a conjugate pair in the I-basis are uncorrelated but not independent (see for example Lemma 10).

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

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

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

Proof. We can split $\Psi = (\Psi' | \Psi'')$ into two $n \times n'$ submatrices and we similarly split $C^{(T)} = \left(C^{(T)'} \middle| C^{(T)''}\right)^T$ into the first n' components $C^{(T)'}$ and the final n' components $C^{(T)''}$. Furthermore, their conjugate pairs origin means that $C^{(T)'}$ and $C^{(T)''}$ are uncorrelated. 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 a companion work [27], and can be verified empirically.

4.3 SymHom decryption using the $p\Gamma$-basis

We now specify a decryption process for the SymHom cryptosystem using the $p\Gamma$-basis of H (though any appropriate basis can be used). We recall (see Figure 5) that we write $Z^\ddagger$ and Z^* to express an element of H as a vector in the T-basis and the $p\Gamma$-basis respectively.

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

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

This decryption process generalises to degree-2 and higher degree ciphertexts in a natural way. For example, if $C(\theta;\mu_1)$ and $C(\theta;\mu_2)$ are two degree-1 ciphertexts with respective Embedded Noises $Y_1'(\mu_1)$ and $Y_2'(\mu_2)$, then the degree-2 ciphertext $C(s;\mu_1,\mu_2) = Y''(\mu_1)Y''(\mu_2) = C(s;\mu_1)C(s;\mu_2) \bmod (R^\vee)_q^2$, 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 SymHom 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 SymHom ciphertexts. Both results follow from the fact that SymHom decryption using (for example) the $p\Gamma$-basis for H fundamentally involves a change of basis transformation between bases for H ultimately to the $p\Gamma$-basis.

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

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

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 `SymHom` degree-1 ciphertext in the $p\Gamma$-basis for H is bounded by

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

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 Lemma 10 shows 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 [18] to obtain

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

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

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

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

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

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

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

References

1. Albrecht, M., Chase, M., Chen, H., Ding, J., Goldwasser, S., Gorbunov, S., Halevi, S., Hoffstein, J., Laine, K., Lauter, K., Lokam, S., Micciancio, D., Moody, D., Morrison, T., Sahai, A., Vaikuntanathan, V.: Homomorphic encryption security standard. Tech. rep., HomomorphicEncryption.org (2018)
2. Avanzi, R., Bos, J.W., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Shanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS-Kyber. Tech. rep., National Institute of Standards and Technology (2017). Available at `https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions`
3. Billingsley, P.: Probability and Measure, third edn. Wiley (1995)
4. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) Fully Homomorphic Encryption without Bootstrapping. In: Innovations in Theoretical Computer Science 2012, pp. 309–325. ACM (2012)
5. Cheon, J.H., Kim, A., Kim, M., Song, Y.S.: Homomorphic Encryption for Arithmetic of Approximate Numbers. In: T. Takagi, T. Peyrin (eds.) Advances in Cryptology - ASIACRYPT 2017, *LNCS*, vol. 10624, pp. 409–437. Springer (2017)
6. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Faster Fully Homomorphic Encryption: Bootstrapping in Less Than 0.1 Seconds. In: J.H. Cheon, T. Takagi (eds.) Advances in Cryptology - ASIACRYPT 2016, *LNCS*, vol. 10031, pp. 3–33. Springer (2016)
7. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: fast fully homomorphic encryption over the torus. J. Cryptology **33**(1), 34–91 (2020)
8. Costache, A., Curtis, B., Hales, E., Murphy, S., Ogilvie, T., Player, R.: On the precision loss in approximate homomorphic encryption. In submission (2021)
9. Costache, A., Laine, K., Player, R.: Evaluating the effectiveness of heuristic worst-case noise analysis in FHE. In: L. Chen, N. Li, K. Liang, S.A. Schneider (eds.) Computer Security - ESORICS 2020 - 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14-18, 2020, Proceedings, Part II, *Lecture Notes in Computer Science*, vol. 12309, pp. 546–565. Springer (2020). DOI 10.1007/978-3-030-59013-0_27. URL `https://doi.org/10.1007/978-3-030-59013-0\_27`

10. Costache, A., Smart, N.P.: Which ring based somewhat homomorphic encryption scheme is best? In: K. Sako (ed.) Topics in Cryptology - CT-RSA 2016 - The Cryptographers' Track at the RSA Conference 2016, San Francisco, CA, USA, February 29 - March 4, 2016, Proceedings, *Lecture Notes in Computer Science*, vol. 9610, pp. 325–340. Springer (2016). DOI 10.1007/978-3-319-29485-8_19. URL https://doi.org/10.1007/978-3-319-29485-8_19
11. D'Anvers, J.P., Karmakar, A., Roy, S.S., Vercauteren, F.: SABER. Tech. rep., National Institute of Standards and Technology (2017). Available at https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions
12. Fan, J., Vercauteren, F.: Somewhat Practical Fully Homomorphic Encryption. IACR Cryptology ePrint Archive **2012**, 144 (2012)
13. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144 (2012). https://ia.cr/2012/144
14. Gentry, C.: Fully Homomorphic Encryption using Ideal Lattices. In: M. Mitzenmacher (ed.) Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, ACM, pp. 169–178 (2009)
15. Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. In: D. Pointcheval, T. Johansson (eds.) Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings, *Lecture Notes in Computer Science*, vol. 7237, pp. 465–482. Springer (2012). DOI 10.1007/978-3-642-29011-4_28. URL https://doi.org/10.1007/978-3-642-29011-4_28
16. Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the AES circuit. In: R. Safavi-Naini, R. Canetti (eds.) Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings, *Lecture Notes in Computer Science*, vol. 7417, pp. 850–867. Springer (2012). DOI 10.1007/978-3-642-32009-5_49. URL https://doi.org/10.1007/978-3-642-32009-5_49
17. Gentry, C., Sahai, A., Waters, B.: Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based. In: R. Canetti, J. Garay (eds.) Advances in Cryptology - CRYPTO 2013, *LNCS*, vol. 8042, pp. 75–92. Springer (2013)
18. Grimmett, G., Stirzaker, D.: Probability And Random Processes, 3rd edn. Oxford University Press (2001)
19. HElib. https://github.com/homenc/HElib (2019)
20. Iliashenko, I.: Optimisations of fully homomorphic encryption. Ph.D. thesis, KU Leuven (2019)
21. Lyubashevsky, V., Peikert, C., Regev, O.: On Ideal Lattices and Learning with Errors Over Rings. IACR Cryptology ePrint Archive **2012**, 230 (2012)
22. Lyubashevsky, V., Peikert, C., Regev, O.: A Toolkit for Ring-LWE Cryptography. IACR Cryptology ePrint Archive **2013**, 293 (2013)
23. Lyubashevsky, V., Peikert, C., Regev, O.: A Toolkit for Ring-LWE Cryptography. In: T. Johansson, P. Nguyen (eds.) Advances in Cryptology - EUROCRYPT 2013, *LNCS*, vol. 7881, pp. 35–54. Springer (2013)
24. Micciancio, D., Peikert, C.: Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller. In: D. Pointcheval, T. Johansson (eds.) Eurocrypt 2012, *LNCS*, vol. 7237, pp. 700–718. Springer (2012)
25. Micciancio, D., Regev, O.: Lattice-based Cryptography. In: D.J. Bernstein and J. Buchmann and E. Dahmen (ed.) Post-Quantum Cryptography, pp. 147–191. Springer (2009)

26. Murphy, S., Player, R.: δ-subgaussian Random Variables in Cryptography. In: J. Jang-Jaccard, F. Guo (eds.) ACISP 2019: The 24th Australasian Conference on Information Security and Privacy, *LNCS*, vol. 11547, pp. 251–268. Springer (2019)
27. Murphy, S., Player, R.: Discretisation and Product Distributions in Ring-LWE. Journal of Mathematical Cryptology **15**, 45–59 (2020)
28. Peikert, C.: A decade of lattice cryptography. Foundations and Trends in Theoretical Computer Science **10**(4), 283–424 (2016)
29. Player, R.: Parameter selection in lattice-based cryptography. Ph.D. thesis, Royal Holloway, University of London (2018)
30. Regev, O.: On Lattices, Learning with Errors, Random Linear Codes and Cryptography. In: H. Gabow, R. Fagin (eds.) 37th Annual ACM Symposium of Theory of Computing (2005)
31. Regev, O.: The Learning with Errors Problem (Invited Survey). In: IEEE Conference on Computational Complexity, pp. 191–204 (2010)
32. Microsoft SEAL (release 4.0). `https://github.com/Microsoft/SEAL` (2022). Microsoft Research, Redmond, WA.
33. Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient Public Key Encryption Based on Ideal Lattices. In: M. Matsui (ed.) Advances in Cryptology - ASIACRYPT 2009, *LNCS*, vol. 5912, pp. 617–635 (2009)
34. Stroock, D.: Probability Theory: An Analytic View. Cambridge University Press (2011)
35. Tao, T., Vu, V.: Random matrices: Universality of local eigenvalue statistics. Acta Mathematica **206**, 127–204 (2011)