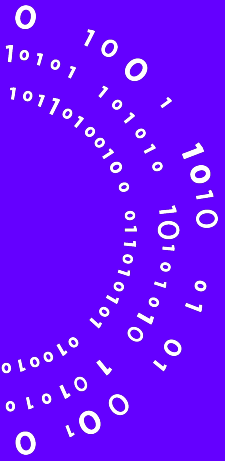

Cryptography Research Centre

Why It Still Makes Sense to Talk About Average-Case Noise Analysis in FHE: The BGV Example

FHE.org 2026

8 March 2026

Beatrice Biasioli, Chiara Marcolla, Nadir Murru, Matilda Urani
beatrice.biasioli@ibm.com
chiara.marcolla@tii.ae
nadir.murru@unitn.it
matilda.urani@polito.it

TII Technology Innovation Institute

Motivation

Does it still make sense to talk about average-case approaches in FHE?

Motivation

Does it still make sense to talk about average-case approaches in FHE?

Yes!

(At least for BGV)

The context

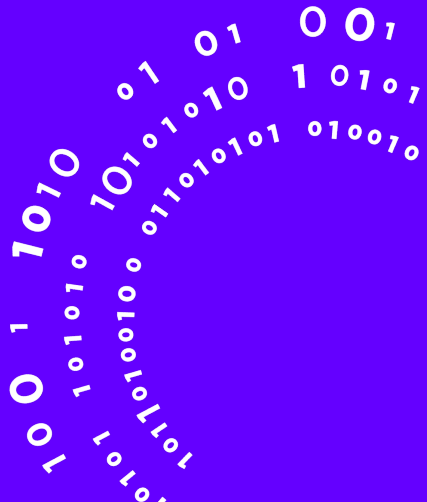

What are the parameters of an FHE scheme?

Almost all FHE schemes used today are based on the Learning with Errors (LWE) problem (or its algebraic variant):

Definition (*Search* LWE Problem (Regev))

Given $\mathbf{b} \in \mathbb{Z}_q^m$ and $A \in (\mathbb{Z}_q)^{m \times n}$, find an *unknown* vector $\mathbf{s} \in \mathbb{Z}_q^n$ such that

$$A\,\mathbf{s} + \mathbf{e} = \mathbf{b} \bmod q$$

where $\mathbf{e} \in \mathbb{Z}_q^m$ is *small* random error.

What are the parameters of an FHE scheme?

Almost all FHE schemes used today are based on the Learning with Errors (LWE) problem (or its algebraic variant):

Definition (*Search* LWE Problem (Regev))

Given $\mathbf{b} \in \mathbb{Z}_q^m$ and $A \in (\mathbb{Z}_q)^{m \times n}$, find an *unknown* vector $\mathbf{s} \in \mathbb{Z}_q^n$ such that

$$A\,\mathbf{s} + \mathbf{e} = \mathbf{b} \bmod q$$

where $\mathbf{e} \in \mathbb{Z}_q^m$ is *small* random error.

- $\mathbf{s}$ follows the distribution χ_s with standard deviation σ_s.
- $\mathbf{e}$ follows the distribution χ_e with standard deviation σ_e.
- Modulus q and the LWE dimension n.
- Security level λ.

Functioning of the scheme

Let m be the message.

Correct decryption occurs only if the error is below a given bound, which depends on q.

$$\|\text{error}\| \leq \text{Bound}(q)$$

Additional difficulty: error grows as homomorphic operations are performed.

Why choosing secure parameters is hard?

- To guarantee correctness , we need a large enough modulus q.

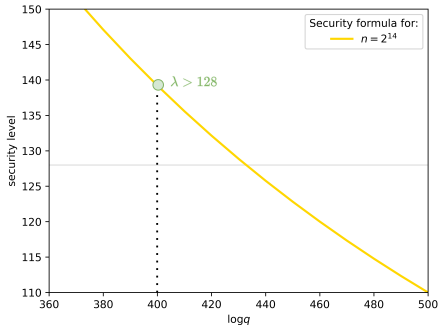

Why choosing secure parameters is hard?

- To guarantee correctness , we need a large enough modulus q.
- Security problem: larger q decreases the security.

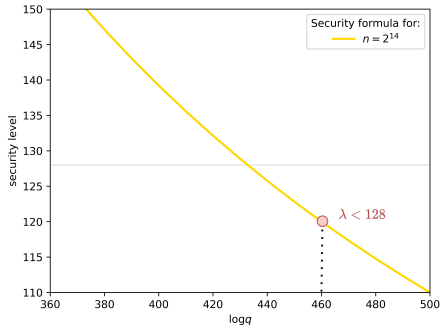

Why choosing secure parameters is hard?

- To guarantee correctness , we need a large enough modulus q.
- Security problem: larger q decreases the security.
- Efficiency problem: to increase the security level λ again, we need a larger dimension n.

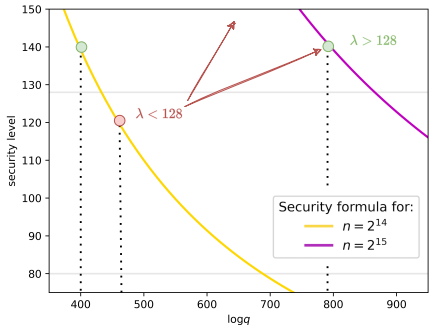

Current Approaches

To correctly select q, it is necessary to **precisely** estimate the error.

At present, **two** main **approaches**:

Worst-case approach

estimates the error by exploiting the norm properties

Average-case approach

models the coefficients of the error as random variables and studies their statistical properties

Average-case

Average-case approaches provide estimates that are **closer to experimental** observations
$\Rightarrow$ is generally regarded as more promising.

However, some **concerns** have recently been **raised**.

The problem is that *most existing* approaches model the coefficients of the error as:

- **Independent**.
- **Gaussian distributed**.

Average-case

Average-case approaches provide estimates that are **closer to experimental** observations
$\Rightarrow$ is generally regarded as more promising.

However, some **concerns** have recently been **raised**.

The problem is that *most existing* approaches model the coefficients of the error as:

- **Independent**.
- **Gaussian distributed**.

Is it really the case? We tried to answer for the BGV scheme.

The BGV Example

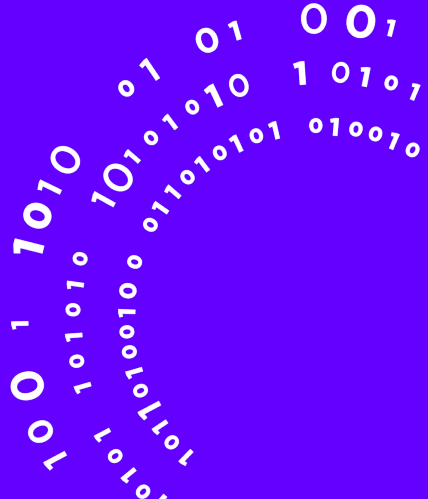

BGV

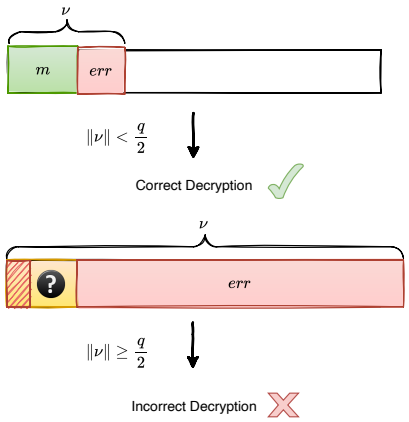

Leveled BGV scheme on RLWE

$m \in \mathbb{Z}_t[x]/\langle x^n+1\rangle \qquad \mathbf{c} \in (\mathbb{Z}_q[x]/\langle x^n+1\rangle)^2$

$\mathsf{sk} = s \qquad \mathsf{pk} = (b, a) \equiv (-a \cdot s + te, a)$

$s \leftarrow \chi_s \qquad a \leftarrow \mathcal{U}_q \qquad e \leftarrow \chi_e$

Enc(m)

$e \leftarrow \chi_e \quad u \leftarrow \chi_s \quad e_0, e_1 \leftarrow \chi_e$

$\mathbf{c} = (c_0, c_1) \equiv (b \cdot u + te_0 + m, a \cdot u + te_1) \mod q$

CRITICAL QUANTITY

$\nu = c_0 + c_1 \cdot s \equiv m + t(e \cdot u + e_1 \cdot s + e_0) \mod q$

Dec($\mathbf{c}$)

$\nu \equiv m \mod t$

BGV Error

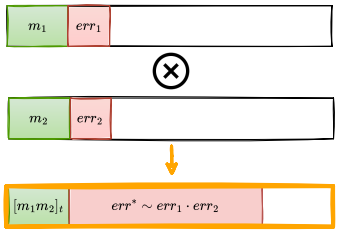

$$\nu_1 = m_1 + te_0 + tu \cdot e + te_1 \cdot s$$

$$\nu_2 = m_2 + te'_0 + tu' \cdot e + te'_1 \cdot s$$

$$\nu^* = (\dots) + (\dots)e + (\dots)s + (\dots)es + (\dots)e^2 + (\dots)s^2$$

THE CRITICAL QUANTITY IS ALWAYS OF THE FORM

$$\nu = \sum_\iota a_\iota s^\iota = \sum_\iota \sum_\mu b_\mu(\iota) e^\mu s^\iota$$

First Assumption: Independence

Are the errors mutually independent?

First Assumption: Independence

Are the errors mutually independent?

Short Answer: **NO!**

$\Longrightarrow$ Errors share powers of s and e

First Assumption: Independence

Are the errors mutually independent?

Short Answer: **NO!**

$\Longrightarrow$ Errors share powers of s and e

Theorem (Informal)

Let V, V' be the variances of the error coefficients of two ciphertexts. Then, the variance of the error coefficients obtained by a multiplication of these two ciphertexts is

$$V_{\text{mul}} \leq nV\,V' F_s F_e\,,$$

where F_s and F_e depend on the number of multiplications previously performed in the circuit.

First Assumption: Independence

Are the errors mutually independent?

Short Answer: **NO!**

$\Longrightarrow$ Errors share powers of s and e

Theorem (Biasioli, Marcolla, Murru, Urani)

Given $\nu_1 = \sum_\iota a_\iota s^\iota, \nu_2 = \sum_\iota a'_\iota s^\iota$, *then*

$$\mathsf{Var}((a_{\iota_1} s^{\iota_1} a'_{\iota_2} s^{\iota_2})|_i) \leq n\mathsf{Var}((a_{\iota_1} s^{\iota_1})|_i)\mathsf{Var}((a'_{\iota_2} s^{\iota_2})|_i) F(\iota_1, \iota_2) F(\mu_1, \mu_2),$$

where μ_1, μ_2 *represent, respectively, the highest power of* e *in* $a_{\iota_1}, a'_{\iota_2}$.

$$F(\alpha, \beta)^{[1]} = \frac{(\alpha + \beta)!}{\alpha!\, \beta!}$$

[1] This function extends the results in [1]

Experimental Validation

Key Idea

Let V_1, V_2 be the variances of two errors.

variance of the product $\neq$ product of variances

Dependencies introduce a **correction factor** ✗ nV_1V_2 ⟶ ✓ $nV_1V_2 \cdot \mathcal{C}_F$

	Encryption		1 Multiplication		6 Multiplications	
n	*our*	*exp*	*our*	*exp*	*our*	*exp*
2^{13}	48.76	48.76	95.68	95.65	95.71	95.25
2^{14}	49.76	49.76	98.68	97.62	98.71	97.63
2^{15}	50.76	50.72	101.68	101.62	101.71	101.59

Table 1: Values labeled *our* correspond to our estimates, and *exp* to experimental results.

Second Assumption: Gaussianity

Are the errors Gaussian distributed?

Second Assumption: Gaussianity

Are the errors Gaussian distributed?

Short Answer: **YES, but with some care!**

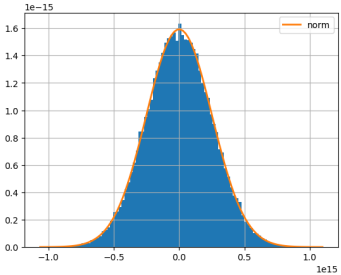

Error coefficients distribution after three multiplications *with* modulus switching: asymmetry 0.00 and kurtosis 2.99.

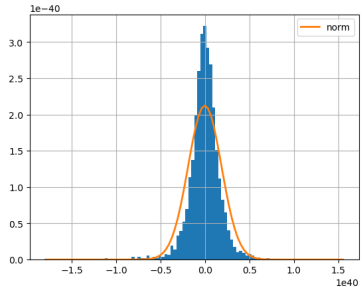

Error coefficients distribution after three multiplications *without* modulus switching: asymmetry -0.52 and kurtosis 14.38.

What is Modulus Switch?

Modulus switch consists in **reducing the error** by reducing the ciphertext modulus.

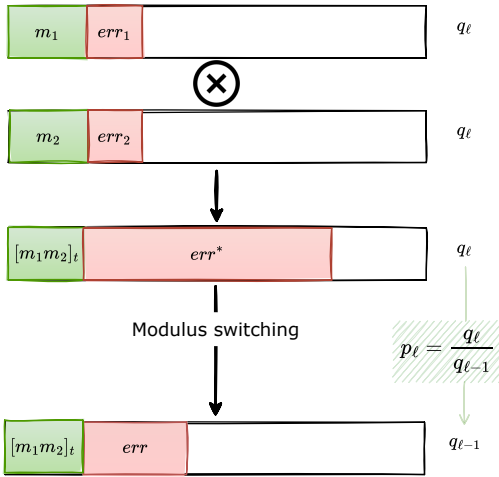

ERROR AFTER MODULUS SWITCH

$$\nu^* \longrightarrow \frac{\nu^*}{p_\ell} + \delta \cdot s$$

$s \rightarrow \mathcal{U}_t$

VARIANCE AFTER MODULUS SWITCH

$$V \longrightarrow \frac{V}{p_\ell^2} + V_{\text{ms}}$$

Our Moduli Choice

ERROR AFTER MODULUS SWITCH

$$\nu = \frac{\nu^*}{p_\ell} + \underbrace{\delta \cdot s}_{}$$

We proved it is **Gaussian** distributed

VARIANCE AFTER MODULUS SWITCH

$$\frac{V}{p_\ell^2} + V_{\text{ms}}$$

Thus, by choosing p_ℓ large enough $\nu \approx \delta \cdot s$

with coefficients variance $\approx V_{\text{ms}}$

Specifically, we required

$$p_\ell > \sqrt{\frac{V}{\alpha \cdot V_{\text{ms}}}}$$

where we choose $\alpha = 1/100$

Benefits

Choosing p_ℓ large enough  *Gaussian distribution* of the error coefficients

Benefits

Choosing p_ℓ large enough  *Gaussian distribution* of the error coefficients

Benefits

Choosing p_ℓ large enough

Gaussian distribution of the error coefficients

Variance no longer depends on the *multiplication depth*

Level Independence

With our parameter choice, by applying modulus switch before each multiplication, we have

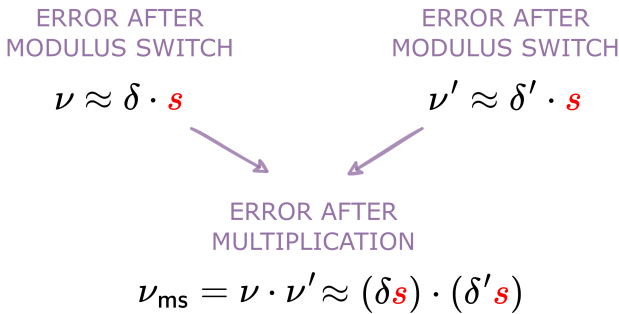

Level Independence

With our parameter choice, by applying modulus switch before each multiplication, we have

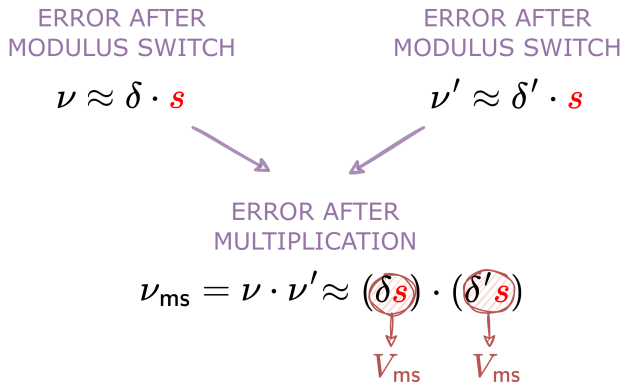

Level Independence

With our parameter choice, by applying modulus switch before each multiplication, we have

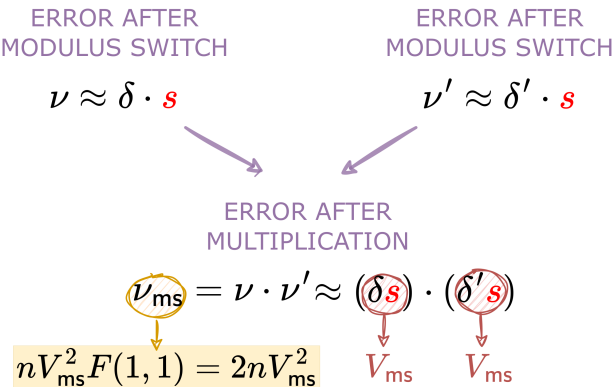

Does this overcome the benefits of our method?

At first sight: is the condition on p_ℓ too restrictive?

n	2^{12}	2^{13}	2^{14}	2^{15}
OpenFHE	147.3	151.8	156.3	161.6
our	121.0	124.5	128.0	131.5

Circuit of depth 3

n	2^{12}	2^{13}	2^{14}	2^{15}
OpenFHE	249.3	256.8	264.3	272.6
our	210.3	216.8	223.3	229.8

Circuit of depth 6

Comparison of $\log_2(q)$ for circuits of progressive multiplications.

Conclusions

Theorem (Biasioli, Marcolla , Murru, Urani)

Let V, V' be the variances of the error coefficients of two ciphertexts. Then, the variance of the error coefficients obtained by a multiplication of these two ciphertexts is

$$V_{\text{mul}} \leq nF_sF_eV\,V'$$

where F_s and F_e depend on the number of multiplications previously performed in the circuit.

NO ModSwitch then $V_{\text{mul}} \approx nF_sF_eVV'$

BUT distribution is not Gaussian

which gives a bound on q

SO we CANNOT infer $||\nu||_\infty \not\leq D\sqrt{2V_{\text{mul}}}$

HOWEVER If we find the correct distribution, **our variance estimation** is still **useful!**

Conclusions

Theorem (Biasioli, Marcolla , Murru, Urani)

Let V, V' be the variances of the error coefficients of two ciphertexts. Then, the variance of the error coefficients obtained by a multiplication of these two ciphertexts is

$$V_{\text{mul}} \leq n \cdot 2 \cdot 1 \cdot V\, V'$$

where F_s and F_e ***do not*** *depend on the number of multiplications previously performed in the circuit.*

⇨ WITH ModSwitch then $V_{\text{mul}} \approx 2nVV' \approx 2nV_{\text{ms}}^2$

⇨ AND distribution is Gaussian

SO we CAN infer $||\nu||_\infty \leq D\sqrt{2V_{\text{mul}}}$ → which gives a bound on q

ePrint 2025/2027

Thank you for your attention!

References

[1] M. Gao and H. Zheng. A Critique on Average-Case Noise Analysis in RLWE-Based Homomorphic Encryption. In *Proceedings of the 13th Workshop on Encrypted Computing & Applied Homomorphic Computing – WAHC 2025*, pages 26–37, 2025.