

Rotate Once, Read Many Times: on the Output Noise of Multi-Value Bootstrapping

P. Chartier¹, M. Koskas² and M. Lemou³

¹ INRIA-IRMAR-University of Rennes, Rennes, France

² Ravel Technologies, Paris, France

³ CNRS-IRMAR-University of Rennes, Rennes, France

Abstract. In FHEW/TFHE, a programmable bootstrap evaluates an arbitrary function of the encrypted message, encoded as the *test polynomial* of a blind rotation. We work in a prime-power cyclotomic ring whose prime is the plaintext modulus p (a *design choice* that leaves the ring degree free as a security parameter) and compare the *Single-Value Mode* (SVM), one rotation per function, with the *Multi-Value Mode* (MVM), one *function-independent* rotation shared by all. Factoring the test polynomial as $v_f^* = \mathfrak{B}_f^* \cdot \mathfrak{w}$, the function carried by the dual module and the rotated factor by the torus, lets a single hypothesis (a centered accumulator error of *arbitrary* covariance G) cover both modes: the two variances are values of one G -form, and their ratio is the exact amplification. The f -independent factorizations are then parametrized by a non-zero ring element, the cofactor, and by the integer lift of the table: at a random lift the cofactor is chosen by a shortest-vector problem for an explicit quadratic form, then the lift by a closest-vector problem of rank $p - 1$. In the spherical model that form is a trace norm and two designs compete: the canonical cofactor, of noise cost $p(p + 1)/6$, and the pivot, which reads the integer table itself at noise cost $2p$, at a random lift; the optimized lift then favours the canonical cofactor on the generic tables. Under the other natural covariance, the block Laplacian, it is optimal for every p . Both cofactors are measured, in the key-noise regime, on a complete implementation, publicly available.

Keywords: Fully homomorphic encryption · Programmable bootstrapping · Look-up tables · FHEW/TFHE · Test polynomial · Output noise · Cyclotomic rings.

1 Introduction

Fully Homomorphic Encryption allows computations to be carried out on encrypted data, but every homomorphic operation increases the noise borne by a ciphertext, and beyond a certain threshold the ciphertext no longer decrypts. *Bootstrapping* [Gen09] allows a computation to proceed by re-evaluating the decryption map homomorphically and resetting that noise. In the FHEW/TFHE family [DM15,CGGI16,CGGI20] it is realized by a *blind rotation* followed by a *sample extraction*, and it can be made *programmable* [CGGI20,ABK⁺24]: the same operation that refreshes the noise simultaneously applies a function $f : \mathbb{T}_p \rightarrow \mathbb{T}_p$ to the encrypted message. The function is encoded once, as the *test polynomial* of the rotation module $\mathbb{T}[X]/\Phi_M$, and specified by a look-up table (LUT) rather than evaluated through an arithmetic circuit. This form is preferable whenever the computation must apply a fixed nonlinear map, such as a neural-network activation [CJP21], a comparison, or a polynomial evaluation.

A general cyclotomic ring, and no padding bit. We work throughout over a prime-power cyclotomic index $M = t^\alpha$ and, as a *design choice*, identify the ring prime t with the plaintext modulus p , so that $M = p^\alpha$; the general-index framework of [CKL26] permits $t \neq p$, but taking $t = p$ is what turns the test polynomial into a genuine look-up table, one coefficient block per plaintext value. This choice, rather than the power-of-two ring $X^N + 1$ of TFHE, lifts the *negacyclicity* (or

“padding-bit”) restriction $f(\mu + \frac{1}{2}) = -f(\mu)$, which a generic function does not satisfy and which costs one bit of message space. Over Φ_{p^α} the only constraint is the milder *cyclicity* $\sum_{k=0}^{p-1} f(k/p) = 0$ (Proposition 1), which *any* function satisfies after subtraction of a constant: a costless shift by a public constant, added back afterwards. The exponent α is then a free *security* parameter: the ring dimension $N = (p-1)p^{\alpha-1}$ is tuned to the target security level [APS15] (for example $\alpha = 5, 4, 3$ giving $N = 2500, 2058, 1210$ for $p = 5, 7, 11$), *independently* of the plaintext modulus p .

Two evaluation modes: SVM and MVM. The table may be read in two distinct ways, whose costs differ. Both read the final accumulator through the trace pairing of the cyclotomic field, $\langle P, Q \rangle = \text{Tr}(P\bar{Q})$ with $\bar{Q}(X) = Q(X^{-1})$, and differ only in the element they pair it with. Here $\Omega_0^* = (1 - X^N)/M$, the first vector of the basis dual to the monomials $(X^i)_{i < N}$, is invertible with inverse in the ring, and multiplication by it carries the ring onto its dual. In the *Single-Value Mode* (SVM), the function is *encoded* in the test polynomial v_f^* : the blind rotation rotates $v_f = (\Omega_0^*)^{-1}v_f^*$ by the phase and the *extraction* $\langle \bar{\Omega}_0^*, \cdot \rangle$ returns $f(\mu)$ – one blind rotation per function. In the *Multi-Value Mode* (MVM), one factorizes $v_f^* = V_f^* \cdot w$ with w *function-independent* and V_f^* in the dual module, performs the blind rotation (the dominant cost) *once* on w , and applies the function afterwards through the *f-dependent extraction* $\langle \bar{V}_f^*, \cdot \rangle$. Both modes thus rotate an element of $\mathcal{T}$ and differ only by the functional extracted from it, (SVM) being the degenerate case where the extracted functional is Ω_0^* and the rotated factor v_f itself. In (MVM), several functions of the *same* message are thereby evaluated for essentially the cost of one rotation.

The output-noise trade-off. The two modes differ in the noise they emit. The blind rotation combines the test polynomial with the bootstrapping keys through the noisy external product, and the *f*-dependent extraction of MVM then amplifies that error, whereas the plain trace of SVM does not. The same f moreover admits many function-independent factorizations $v_f^* = \mathfrak{B}_f^* \cdot \mathfrak{w}$; all are rotated at the same computational price, one blind rotation, but they do not emit the same output noise, since the extraction functional $\mathfrak{B}_f^*$ changes with the shared factor $\mathfrak{w}$. Which factorization emits the least noise, and how much precision MVM loses relative to SVM, is the question addressed in this paper.

Contributions. After situating the two modes in the literature (Section 2) and recalling the algebraic setting and the FHEW/TFHE bootstrap in the formulation of [CKL26] (Sections 3 and 4), we establish three results: the output-error estimate of *both* modes under a single hypothesis, an accumulator error of arbitrary covariance G (Assumption 1, Proposition 2); the structure of the *f*-independent factorizations (the shared factor is fixed by a non-zero ring element γ , the *cofactor*, and the extraction functional is then determined up to a lattice which, for the two designs studied, the integer lift of the table exhausts on the coordinates of the table) and their optimization, a problem (P) we do not solve but approach by two stages: a shortest-vector problem for the cofactor at a random lift, where the expected output variance is $\sigma_F^2 Q_G(\gamma)$, with σ_F^2 the variance of the table entries and Q_G an explicit quadratic form in the coordinates of γ built from G , and it cannot go below $(p-1)\sigma_F^2 \lambda_1^2$, where λ_1^2 (the squared first minimum of $G^{1/2}\mathbb{Z}^N$) is the least variance any integral extraction of the accumulator can have (Proposition 3), then a closest-vector problem of rank $p-1$ for the lift, which favours the canonical cofactor on the generic tables (Subsection 5.2); and the solution of the shortest-vector problem in the spherical model (Corollary 1), where the competition reduces to the canonical cofactor and the pivot, and under the block-Laplacian covariance, where the canonical cofactor is optimal for every p (Remark 4). These predictions, and the spherical covariance itself, are then measured on a complete implementation of the chain (Section 6), written in Julia and publicly available.¹

¹<https://github.com/CycloFHE/TFHE-Multi-Value-Bootstrapping->

2 Related work

Functional and programmable bootstrapping. The idea that a bootstrap can, at no additional cost, evaluate an arbitrary function encoded as a table originates with FHEW [DM15] and was made practical by TFHE [CGGI16, CGGI20], where it became known as *programmable* bootstrapping: the function is the test polynomial rotated by the blind rotation, and a sample extraction reads off the result. This is SVM in our terminology, and the baseline against which MVM should be measured. The padding bit of the negacyclic ring is lifted otherwise by the without-padding PBS of [CLOT21], which also bootstraps larger-precision integers by digit decomposition. Bootstrapping in the other FHE families proceeds otherwise (by homomorphic digit extraction in BGV [HS15], by polynomial approximation of the modular reduction in CKKS [CHK⁺18]) and carries no function with it.

Noise of the blind rotation. The variance of the extracted coefficient over $X^N + 1$ goes back to [DM15, CGGI16]; [KLKS26] derives it exactly, the independence of the gadget digits being deduced from the LWE assumption itself, and [dRDV25] removes the part of it due to the mean of the secret. They compute the *value* of one covariance entry; we take the whole covariance as hypothesis and study its *shape*, over a general cyclotomic ring, and what the extraction makes of it.

Multi-value bootstrapping (MVM) and the factorized test polynomial. MVM originates with the *multi-value* bootstrap of Carpov, Izabachène and Mollimard [CIM19]: the test polynomial is factorized into a common part, rotated once, and a per-function factor multiplied in afterwards. It is exactly the factorization $v_f^* = V_f^* \cdot w$ that this paper analyzes, except that here the per-function factor is a linear functional of the accumulator, incorporated directly into the extraction. Guimarães, Borin and Aranha [GBA21] study the *error propagation* of the shared rotation, lowering the growth of its output variance from quadratic to linear in the output base; their analysis is empirical and specific to the power-of-two ring, whereas Proposition 2 makes the dependence exact and general and Proposition 3 turns the choice among all admissible factorizations into a shortest-vector problem.

High-precision LUTs. When the table exceeds the capacity of a single ring, several bootstraps are combined: the tree-based and digit-decomposition constructions of [BBB⁺23], the homomorphic-floor and digit-removal line of [LMP22, CH18, DPdRV25], the LUT-aided bootstrap of [M⁺19] and the additive-noise Meta-PBS of [MHW⁺25]. These target the *size* of the evaluated table; our concern is logically prior to theirs – the noise emitted by a *single* bootstrap evaluation – so that the optimal MVM factorization studied here applies to any of these composite schemes.

3 Preliminaries

The algebraic objects used below (the cyclotomic field and its ring of integers, the dual module for the trace pairing, the canonical embedding and the dual bases) are classical; their systematic use in lattice cryptography is the Ring-LWE *toolkit* of Lyubashevsky, Peikert and Regev [LPR13a, LPR13b]. We follow throughout the notation and the torus formulation of [CKL26], which specializes them to FHEW/TFHE over a general cyclotomic index and to which we refer for detailed statements and proofs; this section collects, in condensed form, the objects and identities used in the sequel, and a reader familiar with [CKL26] may proceed directly to Section 4.

3.1 Cyclotomic rings, trace pairing and dual bases

Let Φ_M denote the M -th cyclotomic polynomial; for prime-power indices $M = t^\alpha$, t prime, $\alpha \geq 1$,

$$\Phi_{t^\alpha}(X) = \Phi_t(X^{t^{\alpha-1}}) = \sum_{k=0}^{t-1} X^{kt^{\alpha-1}}.$$

All sets of polynomials considered in this paper are embedded within the field

$$\mathcal{K} = \frac{\mathbb{Q}[X]}{(\Phi_M(X))} = \left\{ P \bmod \Phi_M, P \in \mathbb{Q}_{N-1}[X] \right\},$$

a Galois extension of $\mathbb{Q}$ of dimension $N = (t-1)t^{\alpha-1}$, whose ring of algebraic integers is $\mathcal{R} = \{P \bmod \Phi_M, P \in \mathbb{Z}_{N-1}[X]\}$. We write throughout $\mathfrak{m} = M/t$, so that $N = (t-1)\mathfrak{m}$ and $M-N = \mathfrak{m}$, and $\Omega_i = X^i \bmod \Phi_M$ for every integer i , the range $0 \leq i < N$ being the canonical basis.

The trace $\text{Tr} : \mathcal{K} \rightarrow \mathbb{Q}$, $P \mapsto \sum_{\gcd(d,M)=1} P(X^d) \bmod \Phi_M$, takes on a monomial the value [CKL26, Lemma 1]

$$\text{Tr}(X^b) = \begin{cases} N, & M \mid b, \\ -\mathfrak{m}, & \mathfrak{m} \mid b \text{ and } M \nmid b, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The trace induces the scalar product

$$\forall (P, Q) \in \mathcal{K}^2, \langle P, Q \rangle = \text{Tr}(P\bar{Q}) \in \mathbb{Q}, \bar{Q}(X) = Q(X^{M-1}) \bmod \Phi_M, \quad (2)$$

for which any basis $(B_i)_{0 \leq i < N}$ of $\mathcal{K}$ possesses a dual basis $(B_i^*)_{0 \leq i < N}$, characterized by $\langle B_i^*, B_j \rangle = \delta_{i,j}$. We write

$$\|P\|_{\text{Tr}}^2 := \langle P, P \rangle = \text{Tr}(P\bar{P}) = \sum_d |\tau_d(P)|^2 \geq 0, \quad (3)$$

the sum running over the N embeddings $\tau_d : X \mapsto e^{2i\pi d/M}$, $d \in (\mathbb{Z}/M\mathbb{Z})^\times$, of $\mathcal{K}$: a norm for which the involution $P \mapsto \bar{P}$ and the multiplication by any $\pm X^r$ are isometries. The dual basis $(\Omega_i^*)_{0 \leq i < N}$ of the canonical basis $(\Omega_i)_{0 \leq i < N} = (X^i)_{0 \leq i < N}$ has the explicit form [CKL26, Lemma 1]

$$\Omega_i^*(X) = \frac{1}{M} \left(X^i - X^{N+[i]_{\mathfrak{m}}} \right) \in \frac{1}{M} \mathcal{R}, \quad (4)$$

where $[i]_{\mathfrak{m}}$ denotes the integer $0 \leq j < \mathfrak{m}$ such that $i = j \bmod \mathfrak{m}$. By [CKL26, Proposition 3], $\Omega_0^*(X) = \frac{1-X^N}{M}$ is invertible modulo Φ_M with $(\Omega_0^*)^{-1} \in \mathcal{R}$; the $\mathbb{Z}$ -dual of $\mathcal{R}$ for the pairing (2) may accordingly be identified to

$$\mathcal{R}^\vee = \left\{ \sum_{i=0}^{N-1} \omega_i^* \Omega_i^* \bmod \Phi_M, (\omega_0^*, \dots, \omega_{N-1}^*) \in \mathbb{Z}^N \right\} = \Omega_0^* \mathcal{R}.$$

Finally, $\mathcal{T} = \mathcal{K}/\mathcal{R}$ and $\mathcal{T}^\vee = \mathcal{K}/\mathcal{R}^\vee$ are the $\mathbb{Z}[X]$ -modules of torus polynomials and of their duals, $\mathbb{T} = \mathbb{Q}/\mathbb{Z}$ being the torus.

3.2 Plaintext messages

Messages are either elements of the torus $\mathbb{T} = [-\frac{1}{2}, \frac{1}{2}) + \mathbb{Z}$ or polynomials. For an odd integer $p \geq 3$, the discrete torus is $\mathbb{T}_p = \frac{1}{p}\mathbb{Z}_p \subset \mathbb{T}$, with the structure of $(\mathbb{Z}_p, +, \times)$ carried by the symmetric representatives $i \bmod p \in \{-(p-1)/2, \dots, (p-1)/2\}$; from Section 4 on, p is moreover taken equal to the ring prime t , so that $M = p^\alpha$ (see Remark 2). Polynomial plaintexts are representatives of classes in $\mathcal{R}$ or in $\mathcal{T}_p = \frac{1}{p}\mathcal{R}/\mathcal{R} \subset \mathcal{T}$, the polynomials that can be encrypted, manipulated and decrypted exactly under standard parameter conditions.

3.3 LWE and RLWE cryptographic schemes

We recall the encryption schemes built upon [CGGI16, CGGI20], assuming the LWE problem on $\mathbb{Z}_q$ secure, assimilating $\mathbb{T} \equiv \frac{1}{q}\mathbb{Z}_q$, and denoting by $\mathbb{S} \subset \mathbb{Z}$ the finite set of key components. Given $\mathbf{s} \in \mathbb{S}^n$, the LWE-encryption of $\mu \in \mathbb{T}$ is

$$\text{LWE}_{\mathbf{s}}(\mu) = (\mathbf{a}, b) \in \mathbb{T}^{n+1}, \quad \mathbf{a} \xleftarrow{\$} \mathbb{T}^n, \quad e \leftarrow \mathcal{N}(0, \sigma^2), \quad b = \mathbf{s} \cdot \mathbf{a} + \mu + e,$$

decrypted as $\pi_p(b - \mathbf{s} \cdot \mathbf{a})$ with $\pi_p(\mu) = \frac{\lfloor p\mu \rfloor}{p}$, which recovers $\mu \in \mathbb{T}_p$ whenever $|e| < \frac{1}{2p}$. Polynomial messages are encrypted under the key $s(X) = \sum_{i < N} s_i X^i \in \mathcal{R}$, $(s_i)_{i < N} \in \mathbb{S}^N$, on the pair of dual bases $(\Omega_i)_{i < N}$ and $(\Omega_i^*)_{i < N}$: the RLWE*-encryption of $\mu^* \in \mathcal{T}^\vee$ is

$$c^* = \text{RLWE}_s^*(\mu^*) = (a^*, b^*) \in \mathcal{T}^\vee \times \mathcal{T}^\vee, \quad b^* = s \cdot a^* + \mu^* + e^* \bmod \mathcal{R}^\vee,$$

the coordinates of a^* on the dual basis being drawn uniformly on $\mathbb{T}$, and the error e^* Gaussian, spherical in the canonical embedding as in the Ring-LWE toolkit [LPR13a, LPR13b]: the $\tau_d(e^*)$ are centered, uncorrelated and of common variance σ^2 ; this is condition (32) below with $\sigma_c = \sigma$. The RLWE-encryption of $\mu \in \mathcal{T}$ is transported by Ω_0^* ,

$$c = \text{RLWE}_s(\mu) = (\Omega_0^*)^{-1} \text{RLWE}_s^*(\Omega_0^* \mu)$$

componentwise, so that $b = s \cdot a + \mu + e \bmod \mathcal{R}$ with error

$$e = (\Omega_0^*)^{-1} e^* \in \mathcal{T}.$$

Decryption is the coefficient-wise projection $\pi_p(b^{(*)} - s a^{(*)})$, and returns the message when every coefficient of the error is below $\frac{1}{2p}$ in absolute value. Finally, elements of $\mathcal{R}$ and $\mathcal{R}^\vee$ are encrypted, for integers $D \geq 2$ and $\ell \geq 1$, by the gadget and RGSW constructions of [CKL26, Definitions 14–17]: $\text{RGSW}_s^*(m^*)$ stacks the RLWE_s^* -encryptions of $-s m^*/D^t$ and of m^*/D^t , $1 \leq t \leq \ell$, and $\text{RGSW}_s(m)$, $m \in \mathcal{R}$, is defined identically with RLWE_s .

3.4 Homomorphic operations

The homomorphic counterpart of the modular product $\cdot$ of the $\mathbb{Z}[X]$ -module $\mathcal{T}$ is the external product $\boxdot$: if

$$C^* = \text{RGSW}_s^*(m^*) \text{ and } c = \text{RLWE}_s(\mu),$$

then

$$C^* \boxdot c$$

is an RLWE_s^* -encryption of $m^* \cdot \mu \in \mathcal{T}^\vee$, computed as the vector-matrix product $\text{dec}_{D,\ell}(c) C^*$, where $\text{dec}_{D,\ell}$ is the approximate base- D decomposition of the coordinates (on the monomial basis for $\mathcal{T}$, on the dual basis for $\mathcal{T}^\vee$) with balanced digits and a remainder at most $D^{-\ell}/2$ per coordinate; the dual products $C \boxdot c$ and $C \boxdot c^*$ are defined in the same way. The external product is the only noisy operation of the construction. Its error identity [CKL26, Proposition 6] sums the noise of the gadget product, the decomposition remainder and the propagated incoming noise; only the covariance of the accumulator error it produces enters the estimates below (Assumption 1), and we do not reproduce it.

4 The FHEW/TFHE bootstrapping procedure

Objective. Given an LWE encryption c of $\mu \in \mathbb{T}_p$ with error e , and a function $f : \mathbb{T}_p \rightarrow \mathbb{T}_p$, construct an LWE encryption of $f(\mu)$ with a new error smaller than e .

We follow [CKL26, CKLM24], whose derivations we do not repeat, and organize the procedure around three questions, one per subsection: which polynomial encodes f , how it is rotated homomorphically, and how much noise the rotation emits. How that noise is minimized by factorizing the encoding is the object of Section 5, the main contribution of this paper.

4.1 The bootstrap function and its test polynomial

We start from the following variation – through the shift parameter ϵ – of the bootstrap function of Ducas and Micciancio [DM15], later optimized by Chillotti et al. [CGGI16], as formulated in [CKL26, Definition 18]:

Definition 1. Let $v^*(X) \in \mathcal{T}^\vee$ be a polynomial and $\epsilon > 0$ a shift parameter. The bootstrap function Θ_{v^*} is defined as

$$\begin{aligned} \Theta_{v^*} : \mathbb{T} &\rightarrow \mathbb{T} \\ \mu &\mapsto \text{Tr} \left(X^{-\lfloor M(\mu+\epsilon) \rfloor} v^*(X) \right) = \langle X^{\lfloor M(\mu+\epsilon) \rfloor}, v^*(X) \rangle \bmod 1, \end{aligned}$$

where Tr denotes the trace operation. The polynomial v^* is referred to as the test polynomial.

The value is well defined in $\mathbb{T} = \mathbb{Q}/\mathbb{Z}$ since $v^* \in \mathcal{T}^\vee$ and $\text{Tr}(\mathcal{R} \cdot \mathcal{R}^\vee) \subset \mathbb{Z}$. Being constant on each interval $[\frac{i}{M} - \frac{1}{2M} - \epsilon, \frac{i}{M} + \frac{1}{2M} - \epsilon]$, the function Θ_{v^*} is entirely determined by the values $\Theta_{v^*}(\frac{i}{M})$, $0 \leq i \leq M-1$. The test polynomial of f should return $f(\mu)$ from any $\mu + e$ with $|e| < \frac{1}{2p}$ (the optimal correctness guarantee, e comprising in practice the incoming ciphertext error together with the modulo-switch error), that is, $\Theta_{v^*}(\mu) = f\left(\frac{\lfloor p\mu \rfloor}{p}\right)$ for all $\mu \in \mathbb{T}$, the projection π_p of Subsection 3.3. Existence and uniqueness of such a polynomial follow from the compatibility conditions of [CKL26, Proposition 9], which we quote in the form adapted to the shift ϵ :

Proposition 1 (adapted from Proposition 9 of [CKL26]). *Assume $p = t$, and define $\epsilon = \frac{1}{2p} - \frac{1}{2M}$. There exists a polynomial $v_f^* \in \mathcal{T}^\vee$ such that*

$$\forall \mu \in \mathbb{T}_p, \quad \Theta_{v_f^*}(\mu + e) = f(\mu), \quad (5)$$

for all errors $e \in \mathbb{T}$ with $|e| < \frac{1}{2p}$, if and only if the function f satisfies the compatibility condition

$$\sum_{k=0}^{p-1} f\left(\frac{k}{p}\right) = 0 \quad \text{in } \mathbb{T}. \quad (6)$$

The coordinates of v_f^* on the basis $(\Omega_i^*)_{0 \leq i < N}$ are then uniquely determined by:

$$v_{f, i\mathbf{m}+j}^* = f\left(\frac{i}{p}\right) \text{ in } \mathbb{T}, \quad i = 0, \dots, p-2, \quad j = 0, \dots, \mathbf{m}-1.$$

Sketch; see Section 7.1 of [CKL26] for the detailed argument. Note that $\langle X^i, v_f^* \rangle$ depends on i only modulo M . The shift ϵ is chosen so that

$$M(\mu + e + \epsilon) = i\mathbf{m} + Me + \frac{\mathbf{m}-1}{2}$$

at $\mu = i/p$: the admissible errors $|e| < \frac{1}{2p}$ round to the $\mathbf{m}$ exponents

$$i\mathbf{m}, \dots, i\mathbf{m} + \mathbf{m} - 1,$$

and to those only. These p blocks partition $\{0, \dots, M-1\}$, so that (5) is *equivalent* to the M conditions

$$\langle X^{i\mathbf{m}+j}, v_f^* \rangle = f(i/p), \quad 0 \leq i \leq p-1, \quad 0 \leq j < \mathbf{m}.$$

Duality turns the N of them with $i \leq p-2$ into

$$v_{f, i\mathbf{m}+j}^* = f(i/p),$$

prescribing each dual coordinate exactly once; whence the stated v_f^* , and its uniqueness. The last block has its exponents in $[N, M]$, where the defining relation $\Phi_M = 0$, multiplied by X^j and paired with v_f^* , gives

$$\langle X^{N+j}, v_f^* \rangle = - \sum_{k=0}^{p-2} v_{f,k\mathbf{m}+j}^*, \quad 0 \leq j < \mathbf{m} : \quad (7)$$

those $\mathbf{m}$ conditions constrain v_f^* no further and collapse to condition (6). $\square$

Definition 2. A map $f : \mathbb{T}_p \rightarrow \mathbb{T}_p$ is *admissible* if it satisfies the compatibility condition (6). The $p-1$ elementary admissible maps are the f_k , $0 \leq k \leq p-2$, defined by

$$f_k\left(\frac{i}{p}\right) = \frac{\delta_{i,k}}{p}, \quad 0 \leq i \leq p-2, \quad \text{whence} \quad f_k\left(\frac{p-1}{p}\right) = -\frac{1}{p}. \quad (8)$$

Remark 1. Condition (6) pins f down only modulo 1, whereas Section 5 works with an *integer* table. An *integer lift* of f is a choice of $F_i \in pf(i/p) + p\mathbb{Z}$ for $i \leq p-2$, the last entry being the integer form of (7):

$$F_{p-1} := - \sum_{i=0}^{p-2} F_i, \quad \text{so that} \quad \sum_{i=0}^{p-1} F_i = 0 \quad \text{in } \mathbb{Z}, \quad (9)$$

a lift of $pf(\frac{p-1}{p})$ by (6). That the sum vanishes in $\mathbb{Z}$ and not merely in $\mathbb{Z}_p$ is what Section 5 uses: the trace norm of the table $\sum_{i \leq p-1} F_i \Omega_{i\mathbf{m}}$ (the pivot design of Corollary 1) is $\mathbf{m}(p \sum_{i \leq p-1} F_i^2 - (\sum_{i \leq p-1} F_i)^2)$, and a sum vanishing only modulo p would leave its second term standing. The *canonical* lift takes $F_i = pf(i/p) \bmod p$ for $i \leq p-2$, leaving F_{p-1} as large as $(p-1)^2/2$. Two lifts differ by $p\mathbb{Z}^{p-1}$, hence share v_f^* , the bootstrap function and the output plaintext, but not the output *noise*, which depends on the lift through the factor V_f^* of Proposition 3(i); the lift is thus a parameter of the implementation, and Subsection 5.2 chooses it.

Figure 1 displays $\Theta_{v_f^*}$ for five maps at $p = 11$.

Remark 2. The proof extends to $p = t^\beta \mid M$, the compatibility condition becoming $\sum_{r < t} f(\mu + \frac{r}{t}) = 0$ for all $\mu \in \mathbb{T}_p$; an even modulus $p = 2^k$ then forces $t = 2$ (the negacyclic ring of [CGGI16], with $\Omega_i^* = X^i/2^{\alpha-1}$) and the condition becomes the padding-bit restriction $f(\mu + \frac{1}{2}) = -f(\mu)$ of Section 1.

4.2 Homomorphic implementation

We now convert Definition 1 into an algorithm operating on the ciphertext $(\mathbf{a}, b)$, which encrypts $\mu \in \mathbb{T}_p$ under the secret key $\mathbf{s}$, i.e. $\mu = b - \mathbf{s} \cdot \mathbf{a} - e$. Following the *collapsing* strategy of [CKLM24] (see also [CKL26, Section 7.4]), the rotation index is computed from the ciphertext as an integer

$$i \approx \lfloor M(\mu + e + \epsilon) \rfloor, \quad X^{-i} = \prod_{k=1}^{n/m} H_k(X), \quad H_k \in \mathcal{R}, \quad (10)$$

the rounding error being absorbed in e , where each H_k is a monomial determined by a block of m key components, $m \mid n$: the rotation of any $u \in \mathcal{T}$ is the result of n/m successive modular products $\mathcal{R} \times \mathcal{T}$ applied from the right to the left, $X^{-i} \cdot u = H_{n/m} \cdot (\dots (H_1 \cdot u) \dots)$, the collapse parameter m balancing the $|\mathbb{S}|^m$ keys per block against the number n/m of products. Since $\mathcal{R}^\vee = \Omega_0^* \mathcal{R}$, the test polynomial factors as $v_f^* = \Omega_0^* \cdot v_f$ with

$$v_f := (\Omega_0^*)^{-1} v_f^* \in \mathcal{T}, \quad (11)$$




Figure 1: The bootstrap function $\Theta_{v_f^*}$ of Definition 1 at $p = 11$, $\alpha = 3$ ($M = 1331$), for five maps $f : \mathbb{T}_p \rightarrow \mathbb{T}_p$. Each is split as $f = \hat{f} + \kappa$ with $\kappa = \frac{1}{p} \sum_{k < p} f(k/p)$ (see (6)): the test polynomial encodes $\hat{f}$, and the public constant κ is added back. The shaded band is the block of half-width $1/(2p)$ – the largest admissible incoming error.

and (SVM) rotates $u = v_f$, the pivot Ω_0^* being restored at extraction time. (MVM) rotates instead a factor of v_f^* shared by all functions, and applies the function at extraction time. What the rotation encrypts is an element of $\mathcal{T}$, and what the extraction returns is an element of $\mathbb{T}$, so that the factorization is naturally an identity in $\mathcal{T}^\vee$. The product it involves is well defined there: for $\mathfrak{B}^* \in \mathcal{R}^\vee$, $\mathfrak{w} \in \mathcal{K}$ and $r \in \mathcal{R}$, one has

$$\mathfrak{B}^*(\mathfrak{w} + r) - \mathfrak{B}^*\mathfrak{w} = \mathfrak{B}^*r \in \mathcal{R}^\vee,$$

$\mathcal{R}^\vee$ being an $\mathcal{R}$ -module, so that $(\mathfrak{B}^*, \mathfrak{w}) \mapsto \mathfrak{B}^* \cdot \mathfrak{w}$ descends from $\mathcal{R}^\vee \times \mathcal{K}$ to a bilinear map $\mathcal{R}^\vee \times \mathcal{T} \rightarrow \mathcal{T}^\vee$.

Definition 3. Let $(v_f^*)_f$ be the family of test polynomials defined in Proposition 1. An element $\mathfrak{w} \in \mathcal{T}$ is a *shared factor* if for every admissible f there exists $\mathfrak{B}_f^* \in \mathcal{R}^\vee$ with

$$\mathfrak{B}_f^* \cdot \mathfrak{w} = v_f^* \quad \text{in } \mathcal{T}^\vee. \quad (12)$$

A *factorization* is a pair $((\mathfrak{B}_f^*)_f, \mathfrak{w})$ of a shared factor and, for each admissible f , a *functional* $\mathfrak{B}_f^*$ satisfying (12).

Below, an element of $\mathcal{K}$ written for $\mathfrak{w}$ stands for its class in $\mathcal{T}$; an identity $\mathfrak{B}_f^* \mathfrak{w} = v_f^*$ between representatives in $\mathcal{K}$ implies (12) *a fortiori*. Given a factorization, (MVM) rotates $u = \mathfrak{w}$,

$$X^{-\iota} \cdot v_f^* = \mathfrak{B}_f^* \cdot (X^{-\iota} \cdot \mathfrak{w}) \quad \text{in } \mathcal{T}^\vee, \quad (13)$$

and extracts through $\mathfrak{B}_f^*$. Identity (12) is all that correctness requires:

Lemma 1. Let $\mathfrak{B}_f^* \in \mathcal{R}^\vee$ and $\mathfrak{w} \in \mathcal{T}$. The phase extracted from the rotated accumulator is correct, that is

$$\forall \iota \in \mathbb{Z}, \quad \text{Tr}(\mathfrak{B}_f^* X^{-\iota} \cdot \mathfrak{w}) = \text{Tr}(X^{-\iota} v_f^*) \bmod 1$$

if and only if $\mathfrak{B}_f^* \cdot \mathfrak{w} = v_f^*$ in $\mathcal{T}^\vee$.

Proof. Fix representatives in $\mathcal{K}$ of $\mathfrak{w}$ and v_f^* and put

$$D := \mathfrak{B}_f^* \mathfrak{w} - v_f^* \in \mathcal{K},$$

so that the difference of the two phases is $\text{Tr}(X^{-\iota} D) \bmod 1$. It vanishes for every $\iota \in \mathbb{Z}/M\mathbb{Z}$ if and only if

$$\text{Tr}(\mathcal{R} \cdot D) \subset \mathbb{Z},$$

the $(X^{-\iota})_\iota$ generating $\mathcal{R}$ over $\mathbb{Z}$; this is the definition of $D \in \mathcal{R}^\vee$, i.e. (12). $\square$

The functional of a factorization is thus determined by $\mathfrak{w}$ only up to the *admissible lattice* of $\mathfrak{w}$,

$$L(\mathfrak{w}) := \{\Delta \in \mathcal{R}^\vee : \Delta \cdot \mathfrak{w} \in \mathcal{R}^\vee\}, \quad \Lambda(\mathfrak{w}) := (\Omega_0^*)^{-1} L(\mathfrak{w}), \quad (14)$$

the second form,

$$\Lambda(\mathfrak{w}) = \{\delta \in \mathcal{R} : \delta \cdot \mathfrak{w} \in \mathcal{R}\},$$

being the primal one ($\Delta = \Omega_0^* \delta$, $\mathcal{R}^\vee = \Omega_0^* \mathcal{R}$). Indeed

$$(\mathfrak{B}_f' - \mathfrak{B}_f^*) \cdot \mathfrak{w} \in \mathcal{R}^\vee \iff \mathfrak{B}_f' - \mathfrak{B}_f^* \in L(\mathfrak{w}),$$

so that the functionals of $\mathfrak{w}$ for a given f form the coset $\mathfrak{B}_f^* + L(\mathfrak{w})$. The set $\Lambda(\mathfrak{w})$ is an ideal of $\mathcal{R}$, of full rank since it contains $d\mathcal{R}$ for any integer d with $d\mathfrak{w} \in \mathcal{R}$.

Remark 3 (the coset is linear in f). Condition (6) being one linear form on the p values $f(k/p) \in \mathbb{T}_p \simeq \mathbb{Z}_p$, the admissible maps form a free $\mathbb{Z}_p$ -module A of rank $p-1$; and $f \mapsto v_f^*$ is additive (Proposition 1). As $L(\mathfrak{w})$ is the kernel of $\Delta \mapsto \Delta \cdot \mathfrak{w}$, a morphism $\mathcal{R}^\vee \rightarrow \mathcal{T}^\vee$,

$$A \longrightarrow \mathcal{R}^\vee / L(\mathfrak{w}), \quad f \longmapsto \mathfrak{B}_f^* + L(\mathfrak{w}), \quad (15)$$

is well defined, injective and $\mathbb{Z}_p$ -linear. The functional itself is not as $\mathcal{R}^\vee \simeq \mathbb{Z}^N$.

(SVM) is the *degenerate* case $\mathfrak{B}_f^* = \Omega_0^*$, $\mathfrak{w} = v_f$ of (12), with a factor v_f that depends on f . Both modes therefore rotate an element of $\mathcal{T}$, and multiplying either accumulator by Ω_0^* turns it into an RLWE $_s^*$, so that a *single* hypothesis on e^* covers them both (Assumption 1). Each comprises two steps:

1. **Blind Rotate.** One first assembles the RGSW $_s$ -encryptions of the H_k from the bootstrapping keys, the RGSW $_s$ -encryptions of the indicators of the key blocks [CKL26, Section 7.4], and then runs the chain of external products

$$\text{RGSW}_s(H_{\frac{n}{m}}) \square (\dots (\text{RGSW}_s(H_1) \square c^{(0)}) \dots), \quad c^{(0)} = (0, v_f) \text{ or } (0, \mathfrak{w}), \quad (16)$$

on the trivial noise-free RLWE $_s$ -encryption of the f -dependent v_f of (11) in (SVM), of the f -independent $\mathfrak{w}$ in (MVM), giving an RLWE $_s$ -encryption of $X^{-\iota} \cdot v_f$ or $X^{-\iota} \cdot \mathfrak{w}$.

2. **Extract.** One then computes an LWE $_s$ -encryption of $\text{Tr}(X^{-\iota} \cdot v_f^*)$ under the key $\mathbf{s} = (s_0, \dots, s_{N-1})$ given by the coefficients of the polynomial key $s(X) = \sum_{j < N} s_j X^j$. Only the functional differs: $\langle \overline{\Omega_0^*}, \cdot \rangle$ in (SVM), which restores the pivot removed in (11),

$$\langle 1, X^{-\iota} \cdot v_f^* \rangle = \langle \overline{\Omega_0^*}, X^{-\iota} \cdot v_f \rangle, \quad (17)$$

and the f -dependent $\langle \overline{\mathfrak{B}_f^*}, \cdot \rangle$ in (MVM). It enters the formulas only through the M integers

$$\langle \overline{\mathfrak{B}_f^*}, \Omega_m \rangle = \text{Tr}(\mathfrak{B}_f^* X^m), \quad m \in \mathbb{Z}/M\mathbb{Z},$$

integral because $\mathfrak{B}_f^* \in \mathcal{R}^\vee$ and $\Omega_m \in \mathcal{R}$, and depending on m only modulo M . The phase of the output is

$$\langle \overline{\mathfrak{B}_f^*}, b - sa \rangle,$$

the functional applied to the phase $X^{-\iota}\mathbf{w} + e$ of the accumulator, which returns the phase of Lemma 1, up to the error. The error the extraction sees is likewise read through the pivot: with $e^* := \Omega_0^* e$, the two output errors are

$$\mathcal{E}_1 = \text{Tr}(e^*), \quad \mathcal{E}_2 = \text{Tr}((\Omega_0^*)^{-1} \mathfrak{B}_f^* e^*), \quad (18)$$

so that both are functionals of the *same* e^* , paired with elements of $\mathcal{R}$.

The marginal cost of one further function in (MVM) is thus one linear functional of the accumulator against a full blind rotation in (SVM). These integers $\text{Tr}(\mathfrak{B}_f^* X^m)$ depend on f alone (through the representative chosen for it in Section 5) and not on the ciphertext, so that they are computed once per function; the output is then $b' = \sum_k b_k \text{Tr}(\mathfrak{B}_f^* X^k)$ and $a'_j = \sum_k a_k \text{Tr}(\mathfrak{B}_f^* X^{j+k})$ by [CKL26, Proposition 7], one polynomial product against the ℓ products of each of the n/m external products of (16). This presupposes f *public*: a private f would force $\langle \mathfrak{B}_f^*, \cdot \rangle$ to be applied homomorphically. By contrast, in (SVM) only two of them are non-zero and the extraction is a loop of N index reads [CKL26, Algorithm 16].

4.3 Error estimates of the bootstrap output

The noise emitted by the chain (16) is obtained by iterating the error identity of the external product [CKL26, Proposition 6] along the n/m products. Its *value* plays no role below: every estimate of this paper depends on it through its covariance.

Assumption 1. *The error borne by the accumulator (16), in either mode, that is to say $e = b - s a - X^{-\iota} \cdot u$, with $u = \mathbf{w}$ in (MVM) and $u = v_f$ in (SVM), is read in the dual as $e^* = \Omega_0^* e \in \mathcal{T}^\vee$, whose coordinates on the dual basis $(\Omega_i^*)_{0 \leq i < N}$ are centered², with an arbitrary covariance matrix:*

$$e^* = \sum_{i=0}^{N-1} e_i^* \Omega_i^*, \quad \mathbb{E} e_i^* = 0, \quad G := \left(\text{Cov}(e_i^*, e_j^*) \right)_{0 \leq i, j < N} \succ 0. \quad (19)$$

For $x \in \mathbb{R}^N$ we write $\|x\|_G^2 := x^\top G x$, and we call

$$\lambda_1^2 := \min \{ \|x\|_G^2 : x \in \mathbb{Z}^N \setminus \{0\} \} > 0 \quad (20)$$

the squared *first minimum* of the lattice $G^{1/2} \mathbb{Z}^N$, λ_1 being in the standard notation the length of a shortest non-zero vector of that lattice. We further denote $\bar{u}_r$ the coordinates of $\bar{u}$:

$$\bar{U} = \sum_{r < N} \bar{u}_r \Omega_r, \quad \bar{u}_r = \langle \bar{U}, \Omega_r^* \rangle = \text{Tr}(U \Omega_r^*), \quad \bar{u} := (\bar{u}_r)_{r < N} \in \mathbb{Z}^N. \quad (21)$$

Lemma 2. *Let $U \in \mathcal{R}$, of conjugate coordinates $\bar{u} \in \mathbb{Z}^N$. Under Assumption 1 the extraction $\langle \bar{U}, e^* \rangle = \text{Tr}(U e^*)$ is centered, with*

$$\sigma^2(\langle \bar{U}, e^* \rangle) = \|\bar{u}\|_G^2 = \bar{u}^\top G \bar{u}, \quad (22)$$

and, if $U \neq 0$, $\|\bar{u}\|_G^2 \geq \lambda_1^2$.

Proof. Clearly, the coordinates e_r^* of $\langle \bar{U}, e^* \rangle = \sum_{r < N} e_r^* \langle \bar{U}, \Omega_r^* \rangle = \sum_{r < N} e_r^* \bar{u}_r$ are centered, of variance $\bar{u}^\top G \bar{u}$. The bound is (20), since $\bar{u} \in \mathbb{Z}^N \setminus \{0\}$. $\square$

The next proposition estimates the two output errors $\mathcal{E}_1$ (SVM) and $\mathcal{E}_2$ (MVM) of (18) for arbitrary $\mathbf{w} \in \mathcal{T}$ and $\mathfrak{B}_f^* \in \mathcal{R}^\vee$ satisfying (12).

²The output is centered only if the gadget digits are [KLKS26], which the balanced decomposition of Subsection 3.4 ensures.

Proposition 2 (adapted from Proposition 1 of [CKLM24]). *In the setting of Subsection 4.2 and under Assumption 1, the LWE-encryption errors of the bootstrap output are centered, and for every $(\mathfrak{B}_f^*, \mathfrak{w}) \in \mathcal{R}^\vee \times \mathcal{T}$ satisfying (12) their variances are*

$$\sigma^2(\mathcal{E}_2) = \|\bar{\mathfrak{b}}_f\|_G^2, \quad \sigma^2(\mathcal{E}_1) = G_{0,0}, \quad (23)$$

where $\mathfrak{B}_f := (\Omega_0^*)^{-1} \mathfrak{B}_f^* \in \mathcal{R}$ and $\bar{\mathfrak{b}}_f = (\langle \overline{\mathfrak{B}_f}, \Omega_r^* \rangle)_{r < N} \in \mathbb{Z}^N$. The amplification of (MVM) over (SVM) is therefore

$$\frac{\sigma(\mathcal{E}_2)}{\sigma(\mathcal{E}_1)} = \frac{\|\bar{\mathfrak{b}}_f\|_G}{\sqrt{G_{0,0}}}. \quad (24)$$

Over the coset $\mathfrak{B}_f^* + L(\mathfrak{w})$ of the functionals of $\mathfrak{w}$, $\sigma^2(\mathcal{E}_2)$ ranges over $\|\bar{\mathfrak{b}}_f + \bar{\delta}\|_G^2$, $\delta \in \Lambda(\mathfrak{w})$, and stays at or above λ_1^2 unless $v_f^* = 0$, that is $f = 0$.

Proof. The extraction applies

$$(\text{MVM}) : \langle \overline{\mathfrak{B}_f^*}, \cdot \rangle \quad \text{and} \quad (\text{SVM}) : \langle \overline{\Omega_0^*}, \cdot \rangle,$$

to the phase of the final accumulator of (16), hence to its error e . Moving the pivot from one factor to the other, $\text{Tr}(\mathfrak{B}_f^* e) = \text{Tr}(\mathfrak{B}_f e^*)$ and $\text{Tr}(\Omega_0^* e) = \text{Tr}(e^*)$: this is (18), the same e^* paired with $\mathfrak{B}_f$ and with 1, both in $\mathcal{R}$; this is the form in which Assumption 1 applies. We thus get (23), and (24) on dividing the two identities. Over the coset, $\mathfrak{B}_f$ becomes $\mathfrak{B}_f + \delta$ with $\delta \in \Lambda(\mathfrak{w})$; $\mathfrak{B}_f + \delta = 0$ would put $\mathfrak{B}_f^*$ in $L(\mathfrak{w})$ and give $v_f^* = \mathfrak{B}_f^* \cdot \mathfrak{w} = 0$ in $\mathcal{T}^\vee$, which is impossible unless $f = 0$. We conclude by Lemma 2. $\square$

5 Optimality of the decomposition for the MVM

The dependence of the (MVM) output error on $\mathfrak{B}_f^*$ raises the question of the best factorization. Proposition 3(ii) shows that the shared factors $\mathfrak{w}$ are the torus elements of which w is a non-zero multiple, $\gamma \cdot \mathfrak{w} = w$ with a cofactor $\gamma \in \mathcal{R} \setminus \{0\}$, the functionals of a given $\mathfrak{w}$ forming the coset $\gamma V_f^* + L(\mathfrak{w})$. Since $\mathfrak{w}$ is chosen once and the functional once per function, the problem is

$$(P) \quad \min_{\mathfrak{w}} \mathbb{E}_f \min_{\mathfrak{B}_f^* \in \gamma V_f^* + L(\mathfrak{w})} \sigma^2(\mathcal{E}_2),$$

over the shared factors (26), with f uniform among the maps satisfying (6); for a single f , (SVM) itself has amplification 1. Figure 2 summarises what follows.

The conjecture that $\gamma = 1$, $\mathfrak{w} = w$, solves (P) is out of reach along this route, and we proceed in two stages, which do not certify the optimum but approach it: (a) the cofactor first, at a random lift (Subsection 5.1), then (b) the integer lift of Remark 1 (Subsection 5.2). The *pivot* $M\Omega_0^* = 1 - X^N \in \mathcal{R}$ (the factor Ω_0^* of (11) up to the scalar M) is the natural candidate after $\gamma = 1$.

5.1 Stage (a): the cofactor, at a random lift

Proposition 3. *Let $f : \mathbb{T}_p \rightarrow \mathbb{T}_p$ satisfy condition (6), with $v_f^* \in \frac{1}{p}\mathcal{R}^\vee$ the representative of the test polynomial of Proposition 1 at a fixed integer lift $F \in \mathbb{Z}^{p-1}$ as in (9).*

(i) *With*

$$V_f^* = \sum_{i=0}^{p-2} F_i \Omega_{im}^* \in \mathcal{R}^\vee, \quad w = \frac{1}{p} \sum_{i=0}^{m-1} \Omega_i \in \frac{1}{p}\mathcal{R} \subset \mathcal{K}, \quad (25)$$

$((V_f^*)_f, w \bmod \mathcal{R})$ is a factorization, the identity $v_f^* = V_f^* \cdot w$ holding in fact in $\mathcal{K}$ for the representative v_f^* of the lift F .




Figure 2: (a) rests on Proposition 3(iii) and Corollary 1, (b) on Lemma 5.

(ii) An element $\mathfrak{w} \in \mathcal{T}$ is the shared factor of a factorization if and only if

$$\gamma \cdot \mathfrak{w} = w \quad \text{in } \mathcal{T} \quad \text{for some } \gamma \in \mathcal{R} \setminus \{0\}, \quad (26)$$

and the functionals of such a $\mathfrak{w}$ are then exactly

$$\mathfrak{B}_f^* \in \gamma V_f^* + L(\mathfrak{w}), \quad \text{that is} \quad \mathfrak{B}_f \in \gamma V_f + \Lambda(\mathfrak{w}), \quad (27)$$

where $V_f := (\Omega_0^*)^{-1} V_f^* \in \mathcal{R}$; this coset does not depend on the integer lift of f . For every $\gamma \in \mathcal{R} \setminus \{0\}$, the class in $\mathcal{T}$ of $w/\gamma \in \mathcal{K}$ (the quotient being taken in the field) satisfies (26); it is the shared factor retained below.

(iii) Take now the lift random, its coefficients $(F_i)_{0 \leq i \leq p-2}$ independent and centered with common variance σ_F^2 , and let $\Gamma = \Gamma(\gamma) \in \mathbb{Z}^{N \times (p-1)}$, attached to $\gamma \in \mathcal{R} \simeq \mathbb{Z}^N$, be the matrix

$$\Gamma_{r,k} = \left\langle \gamma (\Omega_0^*)^{-1} \Omega_{km}^*, \bar{\Omega}_r^* \right\rangle, \quad 0 \leq r < N, \quad 0 \leq k \leq p-2. \quad (28)$$

Then, for the representative $\mathfrak{B}_f^* = \gamma V_f^*$ of (27), the expected variance of the (MVM) output error is

$$\mathbb{E} \sigma^2(\mathcal{E}_2) = \sigma_F^2 Q_G(\gamma), \quad Q_G(\gamma) = \text{Trace}(\Gamma^\top G \Gamma) = \sum_{k=0}^{p-2} \|\Gamma_{\cdot, k}\|_G^2, \quad (29)$$

a positive definite quadratic form on $\mathcal{R} \simeq \mathbb{Z}^N$. The best of the representatives γV_f^* is therefore the *shortest-vector* solution

$$\min_{\mathfrak{w} = w/\gamma} \mathbb{E} \sigma^2(\mathcal{E}_2) = \sigma_F^2 \min_{\gamma \in \mathcal{R} \setminus \{0\}} Q_G(\gamma) \geq (p-1) \sigma_F^2 \lambda_1^2, \quad (30)$$

with equality if and only if the $p-1$ columns of Γ are minimal vectors of the lattice $(\mathbb{Z}^N, G)$. For a given f , the best element of the coset (27) has a variance at most that of the representative, and at least λ_1^2 unless $v_f^* = 0$.

Proof. (i) By Proposition 1 and $\Omega_{i\mathfrak{m}+j}^* = \Omega_{i\mathfrak{m}}^* \Omega_j$,

$$v_f^* = \left(\sum_{i \leq p-2} F_i \Omega_{i\mathfrak{m}}^* \right) \left(\frac{1}{p} \sum_{j < \mathfrak{m}} \Omega_j \right),$$

which is (25), the first factor lying in $\mathcal{R}^\vee$ and the second in $\frac{1}{p}\mathcal{R}$.

(ii) Let $\mathfrak{w} \in \mathcal{T}$ be the shared factor of a factorization $((\mathfrak{B}_f^*)_f, \mathfrak{w})$. For $f = f_0$, whose canonical lift is $(1, 0, \dots, 0)$, $V_{f_0}^* = \Omega_0^*$ and $v_{f_0}^* = \Omega_0^* w$ by (i); its functional $\mathfrak{B}_{f_0}^* \in \mathcal{R}^\vee = \Omega_0^* \mathcal{R}$ reads $\mathfrak{B}_{f_0}^* = \Omega_0^* \gamma$ with $\gamma \in \mathcal{R}$, and (12) gives

$$\Omega_0^* (\gamma \mathfrak{w} - w) \in \mathcal{R}^\vee = \Omega_0^* \mathcal{R},$$

that is $\gamma \mathfrak{w} - w \in \mathcal{R}$, since Ω_0^* is invertible in $\mathcal{K}$. As $w \notin \mathcal{R}$, $\gamma \neq 0$.

Conversely, let $\gamma \mathfrak{w} = w + r$ with $\gamma \in \mathcal{R} \setminus \{0\}$ and $r \in \mathcal{R}$. Then, by (i),

$$\gamma V_f^* \cdot \mathfrak{w} = V_f^* (w + r) = v_f^* + V_f^* r, \quad V_f^* r \in \mathcal{R}^\vee \cdot \mathcal{R} = \mathcal{R}^\vee,$$

so that $((\gamma V_f^*)_f, \mathfrak{w})$ is a factorization, and the functionals of $\mathfrak{w}$ form its coset $\gamma V_f^* + L(\mathfrak{w})$ (Subsection 4.2), which is (27) since $\mathcal{R}^\vee = \Omega_0^* \mathcal{R}$. Replacing the lift F by $F + pk$, $k \in \mathbb{Z}^{p-1}$, changes γV_f^* by

$$\Delta = p\gamma \sum_{i \leq p-2} k_i \Omega_{i\mathfrak{m}}^*, \quad \Delta \cdot \mathfrak{w} = \sum_{i \leq p-2} k_i \Omega_{i\mathfrak{m}}^* (pw + pr) \in \mathcal{R}^\vee,$$

as $pw = \sum_{j < \mathfrak{m}} \Omega_j \in \mathcal{R}$: $\Delta \in L(\mathfrak{w})$ and the coset is unchanged.

(iii) The k -th column of Γ is the vector of conjugate coordinates (21) of the primal $\mathfrak{B}_{f_k} = (\Omega_0^*)^{-1} \gamma \Omega_{k\mathfrak{m}}^*$ of $\mathfrak{B}_{f_k}^*$, an element of $\mathcal{R}$, hence integral and $\mathbb{Z}$ -linear in γ . Since $\mathfrak{B}_f^* = \sum_k F_k \gamma \Omega_{k\mathfrak{m}}^*$ by (25), the conjugate coordinates of $\mathfrak{B}_f$ give $\mathfrak{b}_f = \Gamma F$, and Proposition 2 with $\text{Cov}(F) = \sigma_F^2 I_{p-1}$ yields

$$\sigma^2(\mathcal{E}_2) = \|\Gamma F\|_G^2, \quad \mathbb{E} \sigma^2(\mathcal{E}_2) = \text{Trace}(\text{Cov}(F) \Gamma^\top G \Gamma) = \sigma_F^2 Q_G(\gamma). \quad (31)$$

Each column being non-zero by (21), Lemma 2 gives $\|\Gamma_{\cdot, k}\|_G^2 \geq \lambda_1^2$, whence (30); and $Q_G(\gamma) = 0$ forces $\Gamma_{\cdot, 0} = 0$, the conjugate coordinates of $\mathfrak{B}_{f_0} = \gamma$ itself, whence $\gamma = 0$. The last assertion is the coset statement of Proposition 2, applied to $\mathfrak{B}_f^* = \gamma V_f^*$. $\square$

Corollary 1 (the spherical model). *Suppose $e^* = \Omega_0^* e$ is spherical in the canonical embedding,*

$$\mathbb{E} \left[\tau_d(e^*) \overline{\tau_{d'}(e^*)} \right] = \sigma_c^2 \delta_{d, d'}, \quad d, d' \in (\mathbb{Z}/M\mathbb{Z})^\times. \quad (32)$$

Then, for the representative $\mathfrak{B}_f^* = \gamma V_f^*$ of (27),

$$\mathbb{E} \sigma^2(\mathcal{E}_2) = \frac{2p \sigma_F^2 \sigma_c^2}{M^2} \|(\Omega_0^*)^{-1} \gamma\|_{\text{Tr}}^2. \quad (33)$$

Writing $\gamma = \sum_{i < N} g_i \Omega_i$ and with $\mathcal{G} = \Psi \Delta^{-1} \Psi$ the integral matrix (38) of Lemma 3 below, so that the Gram matrix (36) is $M\mathcal{G} \otimes I_{\mathfrak{m}}$, the best representative γV_f^* is therefore given by a shortest-vector problem of rank $p-1$, whatever α :

$$\min_{\gamma \in \mathcal{R} \setminus \{0\}} \mathbb{E} \sigma^2(\mathcal{E}_2) = \frac{2p \sigma_F^2 \sigma_c^2}{M} \min_{c \in \mathbb{Z}^{p-1} \setminus \{0\}} c^\top \mathcal{G} c. \quad (34)$$

The two natural cofactors give

$$\begin{aligned} \gamma = 1 : \quad \mathfrak{B}_f^* &= V_f^*, & \mathfrak{w} &= w \in \frac{1}{p}\mathcal{R}, \\ \gamma = M\Omega_0^* : \quad \mathfrak{B}_f^* &= \Omega_0^* \sum_{i \leq p-1} F_i \Omega_{i\mathfrak{m}}, & \mathfrak{w} &= \frac{w}{M\Omega_0^*} \in \frac{1}{p^2}\mathcal{R}, \end{aligned} \quad (35)$$

of respective costs $\frac{p(p+1)}{6}$ and $2p$ in units of $\sigma_F^2 \sigma_c^2 N$, of ratio $(p+1)/12$.

Proof. Put $\beta := \frac{\gamma}{M\Omega_0^*}$, so that $(\Omega_0^*)^{-1}\mathfrak{B}_f^* = \beta MV_f^*$. By (18),

$$\mathcal{E}_2 = \langle \bar{\beta} \overline{MV_f^*}, e^* \rangle = \sum_{i=0}^{p-2} F_i \langle \bar{\beta} \overline{M\Omega_{im}^*}, e^* \rangle,$$

and, $\langle P, Q \rangle = \sum_d \tau_d(P) \overline{\tau_d(Q)}$ with $\mathcal{E}_2$ real,

$$\mathcal{E}_2 = \sum_{i=0}^{p-2} F_i \sum_{d \in (\mathbb{Z}/M\mathbb{Z})^\times} \tau_d(\beta) \tau_d(M\Omega_{im}^*) \tau_d(e^*).$$

Hence, by (32), $\mathbb{E}[\tau_d(e^*) \overline{\tau_{d'}(e^*)}] = \sigma_c^2 \delta_{d,d'}$ collapsing the double sum,

$$\sigma^2(\mathcal{E}_2) = \sigma_c^2 \sum_{i,j=0}^{p-2} F_i F_j \sum_d |\tau_d(\beta)|^2 \tau_d(M\Omega_{im}^*) \overline{\tau_d(M\Omega_{jm}^*)},$$

and finally, with $\mathbb{E}(F_i F_j) = \sigma_F^2 \delta_{i,j}$ and the *pointwise* identity

$$\sum_{i \leq p-2} |\tau_d(M\Omega_{im}^*)|^2 = 2p,$$

valid at every embedding,

$$\mathbb{E} \sigma^2(\mathcal{E}_2) = \sigma_c^2 \sigma_F^2 \sum_d |\tau_d(\beta)|^2 \sum_{i=0}^{p-2} |\tau_d(M\Omega_{im}^*)|^2 = \frac{2p \sigma_F^2 \sigma_c^2}{M^2} \|(\Omega_0^*)^{-1} \gamma\|_{\text{Tr}}^2.$$

Write $\gamma = \sum_{i < N} g_i \Omega_i$ with $i = k\mathfrak{m} + j$, $0 \leq k \leq p-2$ and $j < \mathfrak{m}$, and put $g^{(j)} := (g_{k\mathfrak{m}+j})_{k \leq p-2} \in \mathbb{Z}^{p-1}$. By bilinearity of the pairing,

$$\|(\Omega_0^*)^{-1} \gamma\|_{\text{Tr}}^2 = g^\top \text{Gr} \left(((\Omega_0^*)^{-1} \Omega_i)_{0 \leq i < N} \right) g,$$

and Lemma 3 evaluates that Gram matrix as $M(\mathcal{G} \otimes I_{\mathfrak{m}})$, so that

$$\|(\Omega_0^*)^{-1} \gamma\|_{\text{Tr}}^2 = M \sum_{j < \mathfrak{m}} (g^{(j)})^\top \mathcal{G} g^{(j)}.$$

As $\mathcal{G}$ is positive definite, this is minimised over $g \in \mathbb{Z}^N \setminus \{0\}$ by a single non-zero slice, which with (33) is (34). The values at $\gamma = 1$ and $\gamma = M\Omega_0^*$ are obtained from Proposition 3(ii). The two costs follow from (33) and $\|(\Omega_0^*)^{-1}\|_{\text{Tr}}^2 = M^2 \mathfrak{m}(p^2 - 1)/12$ and $\|M\|_{\text{Tr}}^2 = M^2 N$. $\square$

Lemma 3 (three Gram matrices). *Split the index as $i = k\mathfrak{m} + j$ with $0 \leq k \leq p-2$ and $0 \leq j < \mathfrak{m}$, the left Kronecker factor below carrying k and the right one j , and write $\Omega := (\Omega_i)_{i < N}$, $\eta\Omega := (\eta\Omega_i)_{i < N}$ for $\eta \in \mathcal{K}$. The Gram matrices, for the trace form, of the bases Ω , $\Omega_0^*\Omega$ and $(\Omega_0^*)^{-1}\Omega$ of $\mathcal{K}$ are*

$$\begin{aligned} \text{Gr}(\Omega) &= \Psi \otimes I_{\mathfrak{m}}, & \text{Gr}(\Omega_0^*\Omega) &= \frac{1}{M} \Delta \otimes I_{\mathfrak{m}}, \\ \text{Gr}((\Omega_0^*)^{-1}\Omega) &= M(\Psi \Delta^{-1} \Psi) \otimes I_{\mathfrak{m}}, \end{aligned} \tag{36}$$

with Δ the discrete Laplacian on $p-1$ points,

$$\Psi := \mathfrak{m}(p I_{p-1} - \mathbf{1}\mathbf{1}^\top), \quad \Delta := \text{tridiag}(-1, 2, -1), \tag{37}$$

and $\mathcal{G} := \Psi \Delta^{-1} \Psi$ is the integral matrix

$$\mathcal{G}_{k,k'} = p \mathfrak{m}^2 \left(\frac{p^2 - 1}{12} - \frac{c(p-c)}{2} \right), \quad c = (k - k') \bmod p. \tag{38}$$

The proof, a computation from (1), is given in Appendix A.

Below, $\|\beta\|_G^2$ denotes the G -norm of the conjugate coordinates (21) of $\beta \in \mathcal{R}$, which under (32) is $\sigma_c^2 \|\beta\|_{\text{Tr}}^2$; as β runs over $\mathcal{R}$ those coordinates run over all of $\mathbb{Z}^N$.

Lemma 4. *For every $\beta \in \mathcal{R} \setminus \{0\}$, $\|\beta\|_{\text{Tr}}^2 \geq N$, with equality if and only if $\beta = \pm \Omega_r$, $r \in \mathbb{Z}/M\mathbb{Z}$. Consequently*

$$\min_{\beta \in \mathcal{R} \setminus \{0\}} \|\beta\|_G^2 = \sigma_c^2 \mathfrak{m}(p-1) = \sigma_c^2 N = \sigma^2(\mathcal{E}_1) = \lambda_1^2, \quad (39)$$

the minimum being attained exactly at the $2M$ monomials $\pm \Omega_r$.

Proof. By the arithmetic-geometric inequality over the N embeddings,

$$\|\beta\|_{\text{Tr}}^2 = \sum_d |\tau_d(\beta)|^2 \geq N \left(\prod_d |\tau_d(\beta)|^2 \right)^{1/N} = N |N_{\mathcal{K}/\mathbb{Q}}(\beta)|^{2/N} \geq N,$$

the norm $N_{\mathcal{K}/\mathbb{Q}}(\beta)$ being a non-zero integer. Equality forces $|\tau_d(\beta)| = 1$ at every embedding, so that β is a root of unity by Kronecker's theorem, i.e. one of the $\pm X^r$. $\square$

By (39) the unit $\sigma_F^2 \sigma_c^2 N$ of (35) is $\sigma_F^2 \sigma^2(\mathcal{E}_1)$: at a random lift, the amplification grows linearly in p for the first design and as $\sqrt{p}$ for the second, the two costs crossing at $p = 11$ and the pivot winning beyond. The matrix $\mathcal{G}$ being integral, the shortest-vector problem (34) is solved exactly, in integer arithmetic, by the lattice enumeration of the JULIA package Hecke [FHHJ17], which gives $\min(p(p+1)/6, 2p)$ at every prime $p \leq 97$, with no better γ . The script that performs this enumeration is supplied as supplementary material.³

Remark 4. Neither Assumption 1 nor Proposition 3 requires (32): for any other measured (or assumed) covariance the cofactor of stage (a) is the shortest vector of Q_G . Under (32), $G = \sigma_c^2 \text{Gr}(\Omega)$ (see (44) below); the other natural choice is the block Laplacian of Lemma 3,

$$G = \sigma_r^2 \Delta \otimes I_{\mathfrak{m}} = \sigma_r^2 M \text{Gr}(\Omega_0^* \Omega), \quad Q_G(\gamma) = \frac{2\sigma_r^2}{\mathfrak{m}} \|\gamma\|_{\text{Tr}}^2, \quad (40)$$

for which, by Lemma 4, $\gamma = 1$ is optimal at every p , at cost $p-1$ in units $\sigma_F^2 \sigma^2(\mathcal{E}_1)$ against $2p$ for the pivot; this is the floor of (30), since $\lambda_1^2 = 2\sigma_r^2 = G_{0,0}$: the two natural covariances give opposite answers, and Subsection 6.4 decides between them.

5.2 Stage (b): optimizing the lift

Lemma 5 (the coset on the block coordinates). *Let $Y := X^{\mathfrak{m}}$ and $\mathbb{Z}[Y] \subset \mathcal{R}$ be the span of the Y^k , $0 \leq k \leq p-2$, which carry the primal representatives of the two designs, $V_f = \sum_{l \leq p-2} S_l Y^l$ with the tail sums $S_l = \sum_{i=l}^{p-2} F_i$, and $MV_f^* = \sum_{i \leq p-1} F_i Y^i$ with $F_{p-1} = -\sum_{i \leq p-2} F_i$. For $\delta \in \mathbb{Z}[Y]$,*

$$\delta \in \Lambda(w) \iff \delta \in p\mathbb{Z}[Y], \quad \delta \in \Lambda\left(\frac{w}{M\Omega_0^*}\right) \iff \begin{cases} \delta \in p\mathbb{Z}[Y] \\ \delta(1) \equiv 0 \pmod{p^2}, \end{cases} \quad (41)$$

and in both cases these are exactly the changes of the representative under the lift moves $F \mapsto F + pk$, $k \in \mathbb{Z}^{p-1}$, of Remark 1. Both equivalences hold within the block: $\Lambda(w)$ is of full rank N in $\mathcal{R}$, while $\mathbb{Z}[Y]$ is of rank $p-1$.

Proof. Write $\delta = \sum_{k \leq p-2} g_k Y^k$. Since $pw = \sum_{j < \mathfrak{m}} X^j$,

$$\delta \cdot pw = \sum_{k \leq p-2} \sum_{j < \mathfrak{m}} g_k X^{k\mathfrak{m}+j}$$

³It was co-developed by the authors with Claude (Anthropic).

Table 1: Amplification $\sigma(\mathcal{E}_2)/\sigma(\mathcal{E}_1)$ at $\alpha = 1$ for the five maps of Figure 4, at their canonical lift (rep.) and at the lift solving (42) (lift), for the two designs of Corollary 1.

map	$p = 7$				$p = 11$				$p = 13$			
	$\gamma = 1$		pivot		$\gamma = 1$		pivot		$\gamma = 1$		pivot	
	rep.	lift	rep.	lift	rep.	lift	rep.	lift	rep.	lift	rep.	lift
identity	5.72	4.04	5.72	5.72	16.32	8.52	11.00	11.00	24.32	11.87	14.04	14.04
sign	2.94	2.94	2.65	2.65	5.57	5.57	3.32	3.32	7.08	7.08	3.61	3.61
rectifier	7.55	3.87	7.64	5.07	22.26	7.46	17.55	9.50	33.29	12.01	24.01	13.49
square	3.42	3.42	5.72	5.72	26.03	9.84	24.60	12.05	10.61	10.61	10.61	10.61
comparison	3.65	3.65	3.74	3.74	6.93	6.93	5.74	5.74	8.77	8.77	6.75	6.75

has distinct exponents, all smaller than N : $\delta w \in \mathcal{R}$ if and only if $p \mid g_k$ for all k . For the pivot, $1 - X^N = -Y^{-1}(1 - Y)$ and $1 - Y = (1 - X)pw$ give

$$\delta \cdot \frac{w}{M\Omega_0^*} = \frac{\delta pw}{p(1 - X^N)} = -\frac{Y\delta}{p(1 - X)},$$

which lies in $\mathcal{R}$ if and only if $\delta = pc$ with $c \in \mathbb{Z}[Y]$ and

$$c \in (1 - X)\mathcal{R} \iff c(1) \equiv 0 \pmod{p},$$

since $P \mapsto P(1)$ identifies $\mathcal{R}/(1 - X)\mathcal{R}$ with $\mathbb{Z}/\Phi_M(1)\mathbb{Z} = \mathbb{Z}/p\mathbb{Z}$.

Finally, a lift move changes the tail sums by pT with $T_l = \sum_{i \geq l} k_i$, and $k \mapsto T$ is a bijection of $\mathbb{Z}^{p-1}$: the moves are all of $p\mathbb{Z}[Y]$. It changes the pivot's representative by $p \sum_{i \leq p-1} k_i Y^i$ with $k_{p-1} := -\sum_{i \leq p-2} k_i$, an element pc with $c(1) = 0$; conversely, if $c = \sum_{i \leq p-2} c_i Y^i$ has $c(1) = ps$, then $c = \sum_{i \leq p-1} k_i Y^i$ with $k_i = c_i - s$ for $i \leq p-2$ and $k_{p-1} = -s$, using $\sum_{i \leq p-1} Y^i = 0$, and $\sum_{i \leq p-1} k_i = 0$. $\square$

The cofactor being chosen, the element of the coset (27) remains free, and Proposition 2 makes its choice a closest-vector problem, one per function. By Lemma 5 the coset, on the block coordinates carrying the representative, is exhausted by the lifts of Remark 1, so that stage (b) is, for either design, the choice of the integer lift: with Γ as in (28),

$$\min_{k \in \mathbb{Z}^{p-1}} \|\Gamma(F + pk)\|_G^2 = \min_{k \in \mathbb{Z}^{p-1}} (F + pk)^\top \Gamma^\top G \Gamma (F + pk), \quad (42)$$

a closest-vector problem of rank $p-1$ in the lattice $(\mathbb{Z}^{p-1}, \Gamma^\top G \Gamma)$, solved exactly like the shortest-vector problem (34). Under (32), $\Gamma^\top G \Gamma$ is σ_c^2 times the Gram matrix, for the trace form, of the table basis $((\Omega_0^*)^{-1} \gamma \Omega_{k\mathbf{m}}^*)_{k \leq p-2}$, namely

$$\sigma_c^2 \mathbf{m}(p(\min(k, k') + 1) - (k+1)(k'+1)) \quad \text{and} \quad \sigma_c^2 p \mathbf{m}(\delta_{k,k'} + 1) \quad (43)$$

at $\gamma = 1$ and at the pivot respectively, whose traces are, in the units of (35), its two costs. Table 1 gives, for the five maps of Figure 4 at $\alpha = 1$, the amplification of the representative at its canonical lift and at the lift solving (42). The reduction never hurts, leaves the small tables untouched (sign, comparison), their canonical lift being already the closest vector, divides the amplification of the large tables by up to 3 at $\gamma = 1$ and 1.85 at the pivot, and reverses the comparison of the two designs on the identity, the rectifier and the square at $p = 11$, on the identity and the rectifier at $p = 13$, the square being a tie. For $\alpha \geq 2$ the lattice $\Lambda(w) = (1 - X)^{N-\mathbf{m}+1}\mathcal{R}$ exceeds $p\mathcal{R}$, but an exact closest-vector enumeration in it ($N = 20$ at $(p, \alpha) = (5, 2)$, $N = 42$ at $(7, 2)$) found for the five maps and both designs no better element than (42).

6 Numerical validation

Everything reported below is measured on a complete implementation (bootstrapping keys, blind rotation and extraction of an LWE ciphertext) of (SVM) and of (MVM) at each of its two designs:




Figure 3: Cost of the two designs of Corollary 1, in units of $\sigma_F^2 \sigma^2(\mathcal{E}_1)$, against the plaintext modulus p . Lines: the predicted $p(p+1)/6$ and $2p$ of (35), and the floor $p-1$ of (30). Points: $\mathbb{E} \sigma^2(\mathcal{E}_2)$ over 20 000 integer lifts drawn uniformly, so that $\sigma_F^2 = (p^2 - 1)/12$; error bars are two standard errors. Circles: $\alpha = 1$; crosses: $\alpha = 2$ at $p \leq 17$, $\alpha = 3$ at $p = 5$, $\alpha = 4$ at $p = 7$.

the accumulator error is the one the chain emits, not a postulated law, and the output error is the one a user of the scheme would observe.⁴

6.1 The reference implementation

The three variants are coded in the primal, on the monomial basis. The single-value mode rotates $v_f = (\Omega_0^*)^{-1} v_f^*$ and extracts through $\mathfrak{B}_f^* = \Omega_0^*$; the multi-value mode rotates $\mathfrak{w} = w/\gamma$ and extracts through $\mathfrak{B}_f^* = \gamma V_f^*$, at the two designs $\gamma = 1$ and $\gamma = M\Omega_0^*$ of Corollary 1.

The key material is common to the three variants. Every RLWE row is drawn as an RLWE* row (error spherical in the canonical embedding) and brought back by $(\Omega_0^*)^{-1}$, and the RGSW bootstrapping keys are assembled line by line from those rows.

The implementation carries the torus in double precision, a floating-point model of the integer torus $\frac{1}{q}\mathbb{Z}_q$ whose coarsest step, 2^{-54} , lies thirteen binary orders below every σ_{BK} used; the security instance of a setting is that of the integer torus so modelled, with $q = 2^{64}$.

Under (32) the covariance of the dual coordinates of e^* is σ_c^2 times the Gram matrix of $(\Omega_i)_{i < N}$ for the trace form: $e_i^* = \langle e^*, \Omega_i \rangle = \sum_d \tau_d(e^*) \overline{\tau_d(\Omega_i)}$, so that $\text{Cov}(e_i^*, e_j^*) = \sigma_c^2 \sum_d \tau_d(\Omega_i) \tau_d(\Omega_j) = \sigma_c^2 \langle \Omega_j, \Omega_i \rangle$, whence

$$G = \sigma_c^2 \mathfrak{m} (p I_{p-1} - \mathbb{1} \mathbb{1}^\top) \otimes I_{\mathfrak{m}} = \sigma_c^2 (M I_N - \mathfrak{m} \Pi), \quad \Pi_{r,r'} = \mathbb{1}_{\mathfrak{m} | r-r'}, \quad (44)$$

the second form being the first read on the flat index $r = k\mathfrak{m} + j$, and both output variances are trace norms, by the argument of the proof of Corollary 1,

$$\sigma^2(\mathcal{E}_1) = \sigma_c^2 N, \quad \sigma^2(\mathcal{E}_2) = \sigma_c^2 \|\mathfrak{B}_f\|_{\text{Tr}}^2, \quad \mathfrak{B}_f = \gamma V_f, \quad V_f = (\Omega_0^*)^{-1} V_f^* \in \mathcal{R}. \quad (45)$$

Remark 5. The statements of Sections 4 and 5 are set on $\mathbb{T} = \mathbb{Q}/\mathbb{Z}$, whereas the grid above is $\frac{1}{q}\mathbb{Z}/\mathbb{Z}$: they hold verbatim as soon as $p \mid q$, and the pivot alone requires $p^2 \mid q$. Multiplication by q identifies the q -torsion of $\mathcal{T}$ with $\mathcal{R}/q\mathcal{R}$ and that of $\mathcal{T}^\vee$ with $\mathcal{R}^\vee/q\mathcal{R}^\vee$, and $\mathbb{T}_p$ lies on the grid; the

⁴The Julia code behind this section was co-developed by the authors with Claude (Anthropic), which also produced the figures below from the recorded measurements; the design of the experiments and the reading of their results are the authors'.

objects to be rotated are integral there, $q v_f^* \in \mathcal{R}^\vee$ since $v_f^* \in \frac{1}{p} \mathcal{R}^\vee$ and $q \mathbf{w} \in \mathcal{R}$ for the canonical $\mathbf{w} = w \in \frac{1}{p} \mathcal{R}$, while the pivot, whose $\mathbf{w}$ lies in $\frac{1}{p^2} \mathcal{R}$ by (35), is what forces the stronger condition. Extraction and the error analysis descend to the quotient unchanged: the pairing $\mathcal{R}^\vee \times \mathcal{R} \rightarrow \mathbb{Z}$ is compatible with reduction modulo q , and Assumption 1, read on the integer vector $q e^*$ of covariance $q^2 G$, gives the same amplification (24) and the same costs (35), q^2 cancelling.

Remark 6. What the grid does change is the set of cofactors: over $\mathcal{K}$ the minimization of Proposition 3 runs over all of $\mathcal{R} \setminus \{0\}$, whereas on the grid it runs over the γ for which $q \mathbf{w} \in \mathcal{R}$. That subset still contains $\gamma = 1$ and, when the pivot is admissible, the rotated pivots $\pm X^r(1 - X^N)$, X^r being a unit, so that the optimum of Corollary 1 is unchanged. For $q = 2^{64}$, $\mathbb{T}_p$ is not on the grid; the same objects are then represented up to a rounding of w , deterministic and independent of γ , bounded by $\frac{p}{2} \sum_i |F_i|$ in units of $\mathbb{Z}_q$: more than forty binary orders below the decoding radius at every setting used.

6.2 The noise cost of the two designs

The first question is the one the chain plays no part in. Averaging (45) over the integer lifts F gives the two branches (35) of Corollary 1, and Figure 3 compares them with $\mathbb{E} \sigma^2(\mathcal{E}_2)$ estimated by forming $\mathfrak{B}_f = \gamma V_f$ exactly in $\mathcal{R}$. Over the forty configurations of that figure the ratio of the measured cost to the predicted one averages 0.999 with a dispersion of 0.6%, against a Monte-Carlo uncertainty of the same order. The three qualitative predictions are visible: the linear growth of the canonical design against the $\sqrt{p}$ growth of the pivot, their crossing at $p = 11$, and the fact that α (from $N = 2$ to $N = 2058$ over the points shown) leaves the amplification unchanged.

6.3 Amplification, map by map

Figure 4 reports the amplification the chain produces, at five settings from $p = 5$ to $p = 37$ and ring degrees 1210 to 2500. Over its fifty configurations the ratio of the measured amplification to the predicted one averages 1.009 with a dispersion of 1.0% and a worst case of 3.6%: the identity $\sigma^2(\mathcal{E}_2) = \sigma_c^2 \|\mathfrak{B}_f\|_{\text{Tr}}^2$ holds as an equality, not a bound. This corroborates the spherical model rather than establishing it. The key noise is set at each setting so that the largest of the ten amplifications leaves the output eight standard deviations inside the decoding radius $1/(2p)$; the margin attained is at least 9.4, and none of the 150 000 output errors, nor any of 3 000 complete bootstraps run as a check, falls outside the radius.

Two features are worth isolating. First, the amplification is strongly function-dependent within each design, by a factor of 13 at $p = 37$ between the sign (33.1) and the rectifier (429.1), so that the choice of design and the choice of table interact. Second, the comparison of the designs moves with f and with the lift, Subsection 5.2: at $p = 37$ the pivot divides the amplification by 4.8 on the identity and by 5.4 on the sign, but only by 1.9 on the square.

6.4 The covariance the blind rotation emits

Assumption 1 leaves G free, and Corollary 1 computes everything for the spherical one (44). Whether the chain actually produces that G is the one measurement that Proposition 2 does *not* shortcut.

Under fixed keys we draw K inputs, run the chain, and recover the accumulator error as the difference between the phase $\text{acc}_B - \text{acc}_A \cdot s$ and the noiseless accumulator $X^{-i} v_0$ at the rotation i of (10) the chain has realised; $\hat{G}$ is the empirical covariance of the N dual coordinates of $e^* = \Omega_0^* e$. The regime measured is the one that governs any realistic parameter set: $\sigma_{BK} > 0$ with the gadget precision ($D = 2^7$ as in Figure 4, $\ell = 5$, $n = N$) chosen so that only the key noise is seen.

Table 2 and Figure 5 report it. The diagonal comes out at 0.9904 on average, the band $r \equiv r'$ within 3.1% of $-1/(p-1)$ at worst and every remaining entry within 0.001 of zero (the three values (44) predicts). The relative Frobenius residual to (44), from 0.042 to 0.092, lies at all six settings




Figure 4: Amplification $\sigma(\mathcal{E}_2)/\sigma(\mathcal{E}_1)$ at five settings from $p = 5$ to $p = 37$. Bars: the predicted $\|\mathfrak{B}_f\|_{\text{Tr}}/\sqrt{N}$ of (45); dots: measured over 3 000 accumulator draws at $n = 512$, $\ell = 6$, $D = 2^7$. Lifts $F_i = p f(i/p)$, $i \leq p - 2$, completed by (9), with $i_p := i \bmod p$: identity i_p , sign $\text{sign}(i_p)$, rectifier $\max(i_p, 0)$, square $i^2 \bmod p$, comparison $\mathbb{1}_{i_p > 0}$.

Table 2: $\hat{G}/\hat{G}_{00}$ on the three orbits of (44), with the relative Frobenius residual $\|\hat{G}/\hat{G}_{00} - G/G_{00}\|_F/\|G/G_{00}\|_F$ to (44), to the white $\sigma^2 I_N$, and the sampling floor: the residual two independent halves of the sample show against each other. At $\alpha = 1$ the third orbit is empty, the band $r \equiv r' \pmod{\mathfrak{m}}$ being the whole off-diagonal.

p	α	N	K	$\hat{G}/\hat{G}_{00}$			$-\frac{1}{p-1}$	relative residual		
				diag.	$r \equiv r'$	other		(44)	white	floor
7	1	6	6 000	0.9680	-0.1615	—	-0.1667	0.042	0.364	0.075
11	1	10	6 000	1.0195	-0.1019	—	-0.1000	0.051	0.310	0.089
17	1	16	6 000	0.9834	-0.0615	—	-0.0625	0.059	0.246	0.135
5	2	20	6 000	0.9930	-0.2475	-0.0000	-0.2500	0.060	0.434	0.094
7	2	42	6 000	0.9748	-0.1625	-0.0001	-0.1667	0.092	0.376	0.168
3	4	54	6 000	1.0039	-0.5016	+0.0009	-0.5000	0.091	0.512	0.176

below the sampling floor, while the residual to the white $\sigma^2 I_N$ is 4.1 to 8.7 times larger: the emitted covariance is indistinguishable from (44) and clearly distinct from white.

The three orbit means are not, by themselves, the test: covariances that (44) does not describe share them, and it is the residual that separates them. This is exhibited by repeating the measurement at $\sigma_{BK} = 0$, where the keys are noiseless and only the decomposition remainder of the gadget decomposition (Subsection 3.4) remains. The three orbit means are then unchanged, but the residual to (44) is 0.75 to 1.5, an order of magnitude above the sampling floor. The spherical form is therefore a property of the key noise and not of the chain (one additional reason to keep σ_{BK} above the arithmetic floor of an implementation). The same remainder appears in [KLKS26, dRDV25] as a bias, cross-iteration correlations and a term in the mean of the secret.

That the band depends on p alone (see Figure 5(b)) is the covariance counterpart of the rank




Figure 5: The measured accumulator covariance, over the six settings of Table 2. (a) $\hat{G}/\hat{G}_{00}$ at $p=3$, $\alpha=4$ ($m=27$): the diagonal, the two bands $r \equiv r' \pmod{m}$ at $-\frac{1}{2}$, zero elsewhere. (b) that band against the $-1/(p-1)$ of (44) and the 0 of the white model. (c) every entry against the value (44) predicts for it: three clusters on the diagonal and nothing in between.

$p-1$ of Corollary 1; and the model has no free parameter beyond the scalar σ_c^2 , which cancels in every ratio of Section 5: Figure 4 may be read at one set of parameters and used at another.

What matters for Section 5 is finally not $\hat{G}$ but the quadratic forms it induces, reported to within 1% in Subsection 6.3: the costs of Corollary 1 hold for the covariance the implementation produces in the key-noise regime. Outside that regime, where at $\sigma_{BK}=0$ it is neither (44) nor (40), the factorization is recomputed from $\hat{G}$ as Remark 4 prescribes.

7 Conclusion

A complete implementation reproduces the predictions of Section 5 to within one percent and shows the accumulator covariance to be spherical in the key-noise regime. The joint optimization (P) is nevertheless left open; we conjecture that the canonical shared factor solves it. In practice one evaluates, for the tables at hand, the two branches at their optimized lifts and keeps the better: in every configuration we could search exhaustively this comes close to the optimum.

Acknowledgements

AI systems were used in the following parts of this work, on the authors' instructions and subject to their revision: the Julia script that solves the shortest-vector problem of Section 5.1 and the Julia implementation of Section 6 were co-developed with it, as were the scripts that plot the figures from the recorded measurements. Besides, it carried out L^AT_EX work and wording assistance on the manuscript. The design of the experiments and of the figures, and all definitions, statements and proofs, are the authors' own, who take full responsibility for the content.

A Proof of Lemma 3

Proof. For a basis $B = (B_i)_{i < N}$ of $\mathcal{K}$ and $\eta \in \mathcal{K} \setminus \{0\}$, write

$$\text{Gr}(B) := (\langle B_i, B_{i'} \rangle)_{i, i' < N},$$

B^* for the dual basis and $\eta B := (\eta B_i)_{i < N}$. Since

$$\langle \eta B_i, \bar{\eta}^{-1} B_{i'}^* \rangle = \langle B_i, B_{i'}^* \rangle,$$

the dual basis of ηB is $\bar{\eta}^{-1} B^*$, and the classical identity $\text{Gr}(B^*) = \text{Gr}(B)^{-1}$ gives

$$\text{Gr}(\eta^{-1} B) = \text{Gr}(B) \text{Gr}(\eta B)^{-1} \text{Gr}(B),$$

$\text{Gr}(\eta^{-1} B)$ and $\text{Gr}(\bar{\eta}^{-1} B)$ sharing the entries $\text{Tr}(B_i \bar{B}_{i'} / (\eta \bar{\eta}))$. We apply this to $\eta = \Omega_0^*$ and $B = \Omega := (\Omega_i)_{i < N}$. First, by (1),

$$\langle \Omega_i, \Omega_{i'} \rangle = \delta_{j,j'} (M \delta_{k,k'} - \mathbf{m}) \text{ i.e. } \text{Gr}(\Omega) = \Psi \otimes I_{\mathbf{m}}.$$

Next, $Y := X^{\mathbf{m}}$ being a primitive p -th root of unity with $X^N = Y^{-1}$,

$$M^2 \Omega_0^* \bar{\Omega}_0^* = (1 - Y^{-1})(1 - Y) = 2 - Y - Y^{-1},$$

so that

$$M^2 \langle \Omega_0^* \Omega_i, \Omega_0^* \Omega_{i'} \rangle = 2 \text{Tr}(X^{i-i'}) - \text{Tr}(X^{i-i'+\mathbf{m}}) - \text{Tr}(X^{i-i'-\mathbf{m}}),$$

which by (1) vanishes unless $j = j'$ and $|k - k'| \leq 1$, where it equals $2M$ and $-M$: thus

$$\text{Gr}(\Omega_0^* \Omega) = \frac{1}{M} \Delta \otimes I_{\mathbf{m}}.$$

For (38), compute the third Gram matrix directly: $(\Omega_0^*)^{-1} \Omega_i = M X^i / (1 - Y^{-1})$, so that, with $i = k\mathbf{m} + j$ and $i' = k'\mathbf{m} + j'$,

$$\begin{aligned} \langle (\Omega_0^*)^{-1} \Omega_i, (\Omega_0^*)^{-1} \Omega_{i'} \rangle &= M^2 \text{Tr} \left(\frac{X^{i-i'}}{2 - Y - Y^{-1}} \right) \\ &= \delta_{j,j'} M^2 \mathbf{m} g(k - k'), \\ g(c) &:= \text{Tr}_{\mathbb{Q}(Y)/\mathbb{Q}} \left(\frac{Y^c}{2 - Y - Y^{-1}} \right), \end{aligned}$$

since $\text{Tr}_{\mathbb{K}/\mathbb{Q}(Y)}(X^{j-j'}) = \mathbf{m} \delta_{j,j'}$ for $|j - j'| < \mathbf{m}$. The function g on $\mathbb{Z}/p\mathbb{Z}$ has zero mean and satisfies $(2 - Y - Y^{-1})g = \text{Tr}_{\mathbb{Q}(Y)/\mathbb{Q}}(Y^c) = p \delta_{c,0} - 1$, which determines it: $g(c) = \frac{p^2-1}{12} - \frac{c(p-c)}{2}$ for $0 \leq c < p$, as one checks on the two sides. Comparing with (36) gives $\mathcal{G}_{k,k'} = M \mathbf{m} g(k - k') = p \mathbf{m}^2 g(k - k')$. $\square$

References

- [ABK⁺24] S. Azogagh, Z. A. Birba, M.-O. Killijian, F. Larose-Gervais, and S. Gambs. RevoLUT: Rust efficient versatile oblivious look-up-tables. *Cryptology ePrint Archive*, Paper 2024/1935, 2024. <https://eprint.iacr.org/2024/1935>.
- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015. [doi:10.1515/jmc-2015-0016](https://doi.org/10.1515/jmc-2015-0016).
- [BBB⁺23] Loris Bergerat, Anas Boudi, Quentin Bourgerie, Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Parameter optimization and larger precision for (T)FHE. *Journal of Cryptology*, 36(3):28, 2023. [doi:10.1007/s00145-023-09463-5](https://doi.org/10.1007/s00145-023-09463-5).
- [CGGI16] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In *Advances in Cryptology – ASIACRYPT 2016*, pages 3–33. Springer, 2016. [doi:10.1007/978-3-662-53887-6_1](https://doi.org/10.1007/978-3-662-53887-6_1).
- [CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, 2020. [doi:10.1007/s00145-019-09319-x](https://doi.org/10.1007/s00145-019-09319-x).

- [CH18] Hao Chen and Kyoohyung Han. Homomorphic lower digits removal and improved FHE bootstrapping. In *Advances in Cryptology – EUROCRYPT 2018*, pages 315–337. Springer, 2018. doi:10.1007/978-3-319-78381-9_12.
- [CHK⁺18] Jung Hee Cheon, Kyoohyung Han, Andrey Kim, Miran Kim, and Yongsoo Song. Bootstrapping for approximate homomorphic encryption. In *Advances in Cryptology – EUROCRYPT 2018*, pages 360–384. Springer, 2018. doi:10.1007/978-3-319-78381-9_14.
- [CIM19] S. Carpov, M. Izabachène, and V. Mollimard. New techniques for multi-value input homomorphic evaluation and applications. In *Topics in Cryptology – CT-RSA 2019*, volume 11405 of *LNCS*. Springer, 2019. doi:10.1007/978-3-030-12612-4_6.
- [CJP21] Ilaria Chillotti, Marc Joye, and Pascal Paillier. Programmable bootstrapping enables efficient homomorphic inference of deep neural networks. In *Cyber Security Cryptography and Machine Learning (CSCML)*, pages 1–19. Springer, 2021. doi:10.1007/978-3-030-78086-9_1.
- [CKL26] Philippe Chartier, Michel Koskas, and Mohammed Lemou. Exploring general cyclotomic rings in torus-based fully homomorphic encryption. *IACR Communications in Cryptology*, 3(1), 2026. <https://cic.iacr.org/p/3/1/4>. doi:10.62056/ab01jb0kr.
- [CKLM24] Philippe Chartier, Michel Koskas, Mohammed Lemou, and Florian Méhats. Fully homomorphic encryption on large integers. Cryptology ePrint Archive, Paper 2024/155, 2024. <https://eprint.iacr.org/2024/155>.
- [CLOT21] Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for TFHE. In *Advances in Cryptology – ASIACRYPT 2021*, pages 670–699. Springer, 2021. doi:10.1007/978-3-030-92078-4_23.
- [DM15] Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In *Advances in Cryptology – EUROCRYPT 2015*, pages 617–640. Springer, 2015. doi:10.1007/978-3-662-46800-5_24.
- [DPdRV25] J.-P. D’Anvers, X. Pottier, T. de Ruijter, and I. Verbauwhede. Head start: Digit extraction in TFHE from MSB to LSB. Cryptology ePrint Archive, Paper 2025/2012, 2025. <https://eprint.iacr.org/2025/2012>.
- [dRDV25] Thomas de Ruijter, Jan-Pieter D’Anvers, and Ingrid Verbauwhede. Don’t be mean: Reducing approximation noise in TFHE through mean compensation. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2025. <https://eprint.iacr.org/2025/809>.
- [FHHJ17] Claus Fieker, William Hart, Tommy Hofmann, and Fredrik Johansson. Nemo/Hecke: Computer algebra and number theory packages for the Julia programming language. In *Proceedings of the 2017 ACM International Symposium on Symbolic and Algebraic Computation (ISSAC ’17)*. ACM, 2017. doi:10.1145/3087604.3087611.
- [GBA21] Antonio Guimarães, Edson Borin, and Diego F. Aranha. Revisiting the functional bootstrap in TFHE. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021(2):229–253, 2021. doi:10.46586/tches.v2021.i2.229-253.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 169–178, 2009. doi:10.1145/1536414.1536440.

- [HS15] Shai Halevi and Victor Shoup. Bootstrapping for HELib. In *Advances in Cryptology – EUROCRYPT 2015*, pages 641–670. Springer, 2015. doi:[10.1007/978-3-662-46800-5_25](https://doi.org/10.1007/978-3-662-46800-5_25).
- [KLKS26] Sin Kim, Seunghwan Lee, Dohyuk Kim, and Dong-Joon Shin. Exact error analysis for blind rotation in fully homomorphic encryption. Cryptology ePrint Archive, Paper 2026/269, 2026. <https://eprint.iacr.org/2026/269>.
- [LMP22] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. Large-precision homomorphic sign evaluation using FHEW/TFHE bootstrapping. In *Advances in Cryptology – ASIACRYPT 2022*, pages 130–160. Springer, 2022. doi:[10.1007/978-3-031-22966-4_5](https://doi.org/10.1007/978-3-031-22966-4_5).
- [LPR13a] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, 2013. doi:[10.1145/2535925](https://doi.org/10.1145/2535925).
- [LPR13b] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In *Advances in Cryptology – EUROCRYPT 2013*, pages 35–54. Springer, 2013. doi:[10.1007/978-3-642-38348-9_3](https://doi.org/10.1007/978-3-642-38348-9_3).
- [M⁺19] T. A. Melchor et al. LUT-aided bootstrapping. Cryptology ePrint Archive, Report 2019/458, 2019. <https://eprint.iacr.org/2019/458>.
- [MHW⁺25] Shihe Ma, Tairong Huang, Anyu Wang, Changtong Xu, Tao Wei, and Xiaoyun Wang. Meta-PBS: Compact high-precision programmable bootstrapping. Cryptology ePrint Archive, Paper 2025/2284, 2025. <https://eprint.iacr.org/2025/2284>.