Exact Error Bootstrapping and Comb Bootstrapping in FHEW/TFHE*

Philippe Chartier
Inria, IRMAR, Université de Rennes
philippe.chartier@inria.fr

Mohammed Lemou
CNRS, IRMAR, Université de Rennes
mohammed.lemou@univ-rennes.fr

Campus de Beaulieu, 263 avenue du Général Leclerc, 35042 Rennes Cedex, France

October 31, 2026

Abstract

In FHEW and TFHE, bootstrapping evaluates the decryption function via a blind rotation, whose resolution limits the message space. We propose instead to bootstrap the *error*. By rounding the ciphertext at a resolution that is a multiple of the message modulus, we preserve the message exactly. Reduction modulo a small integer then yields a noise-free encryption of the error, which a single blind rotation lifts exactly before it is subtracted. The message modulus is thus constrained solely by the noise and can be much larger than in standard TFHE. Since the error takes far fewer distinct values than the ring degree, a single blind rotation can also refresh several ciphertexts simultaneously. We call this procedure the *comb bootstrap*. Between refreshes, an integer remains encrypted in a single ciphertext, allowing additions modulo the message modulus and multiplications without programmable bootstrapping. Built on TFHE-rs, for 8-bit integers on a single thread, our refresh is 3.5 times faster than that of TFHE-rs, and 9.1 times faster with the comb bootstrap; multiplication is 6.5 times faster. In a residue number system, multiplication of integers modulo $\widehat{P} \geq 2^{32}$ is 19 times faster than in TFHE-rs. The error bootstrap serves exclusively to refresh ciphertexts; evaluating any other function still requires programmable bootstrapping.

Keywords: Fully homomorphic encryption · Bootstrapping · FHEW/TFHE · Blind rotation · Large plaintexts · Comb bootstrapping

1 Introduction

FHEW [26] and TFHE [22] refresh an LWE ciphertext by means of a blind rotation: the decryption function is evaluated homomorphically, thereby resetting the noise while simultaneously applying a function to the message. This is known as a programmable bootstrap (PBS). The price is precision: the phase is first rounded to the resolution $2N$ of a ring of degree N, and this rounding limits the message modulus $\widehat{p}$ (messages are represented by multiples of $1/\widehat{p}$ on the torus). The rounding error has standard deviation $\sqrt{V}/(2N)$, where $V \approx 13$ for the TFHE parameters of Appendix I (mask grouped with $m = 4$, recentered, and selected). To achieve a failure probability of 2^{-128}, it must remain a factor $u \approx 13$ below the decoding margin $1/(4\widehat{p})$:

$$\widehat{p} \leq \frac{N}{2u\sqrt{V}} \approx \frac{N}{95}. \tag{1}$$

Beyond a few bits, messages must therefore be split into blocks and refreshed one block at a time.

*The code and the logs of the measurements are available at https://github.com/CycloFHE/Comb-Bootstrap.

This paper removes this limitation by bootstrapping the *error* of an LWE ciphertext rather than its message. The ciphertext is rounded at resolution

$$R = \widehat{p}\,p,$$

which is an integer multiple of the message modulus. Thus, a message $\mu = i/\widehat{p}$ is mapped to $R\mu = pi$ and preserved exactly. Reduction modulo p then yields a *noise-free* LWE encryption of the error, represented by an integer in $\mathbb{Z}_p$, the group of exponents of the blind rotation. A single rotation, using a lookup table of p entries independently of $\widehat{p}$, lifts this error exactly. Subtracting the refreshed error ciphertext leaves a fresh encryption of μ (Figure 2). Kim, Seo, and Song [39] also refresh BFV ciphertexts through their error, but recover it only approximately, using several CKKS bootstraps that take tens of seconds per packed ciphertext. Our method achieves an exact lift with a single blind rotation. Two consequences follow:

(i) The message modulus is no longer constrained by the ring degree. The rounding error, with standard deviation $\sqrt{V}/R$, must satisfy the same decoding-margin constraint $1/(4\widehat{p})$ as in TFHE. It remains a factor u below this margin whenever

$$p \geq 4u\sqrt{V} \approx 190,$$

regardless of $\widehat{p}$. The message modulus is therefore limited by the noise, which is controlled by the keys, rather than by the ring degree. Under identical parameters, it can be much larger than in TFHE at the same ring degree. Moreover, achieving a failure probability of 2^{-128}, as required by the IND-CPA^{D} attacks of Cheon et al. [21] and Checri et al. [18], incurs no additional cost in ring degree.

(ii) The unused capacity of the ring can instead be exploited for batching. Each error occupies p of the $2N$ positions available in the rotation, allowing the TFHE ring to accommodate

$$\mathcal{B} = \frac{2N}{p}$$

errors, which coincides with TFHE's bound on $\widehat{p}$ when $p = 4u\sqrt{V}$. The errors are arranged in $\mathcal{B}$ columns, with coefficients X^i such that $i \equiv k \pmod{\mathcal{B}}$. These columns interleave like the teeth of a comb and are refreshed by a single blind rotation. We call this procedure the *comb bootstrap*, or CombBoot. For $N = 2048$, up to sixteen ciphertexts can share a rotation with messages of up to 8 bits. TFHE uses the rotation's capacity to encode the message; we use it to process multiple ciphertexts simultaneously. For instance, in a residue number system (RNS), eight residues of at most 5 bits each represent a 32-bit integer, and a single CombBoot refreshes all eight residues simultaneously[1].

Our first contribution is the error bootstrap for FHEW and TFHE (Section 4), together with its exact lift. We prove its correctness when the mask is grouped into blocks of m coordinates, reduce the rounding error through recentering and selection, and compare its message capacity with that of TFHE at the same ring degree. Our second contribution is the comb bootstrap (Section 5). The projections onto the columns are expressed as sums of automorphisms and computed using $\mathcal{B} - 1$ key switchings with a single decomposition (hoisting, as in Halevi and Shoup [34]). We derive a closed-form expression for the accumulator and analyze its noise. Appendix C extends the construction to rings of odd prime conductor, where the error window doubles. Our third contribution is an implementation in TFHE-rs [55], using its keys, FFT, and external product, and an experimental comparison with TFHE-rs on the same machine (Section 7).[2] At $N = 2048$,

[1] Additions use standard LWE arithmetic, while multiplication requires only a tensor product followed by relinearization.

[2] The product of Section 6, after a conversion to RLWE by a trace modulo an odd number, is not a new technique: it shows that additions and products need no programmable bootstrap, and its main noise term follows a Laplace law, whose exact tail sets its parameters.

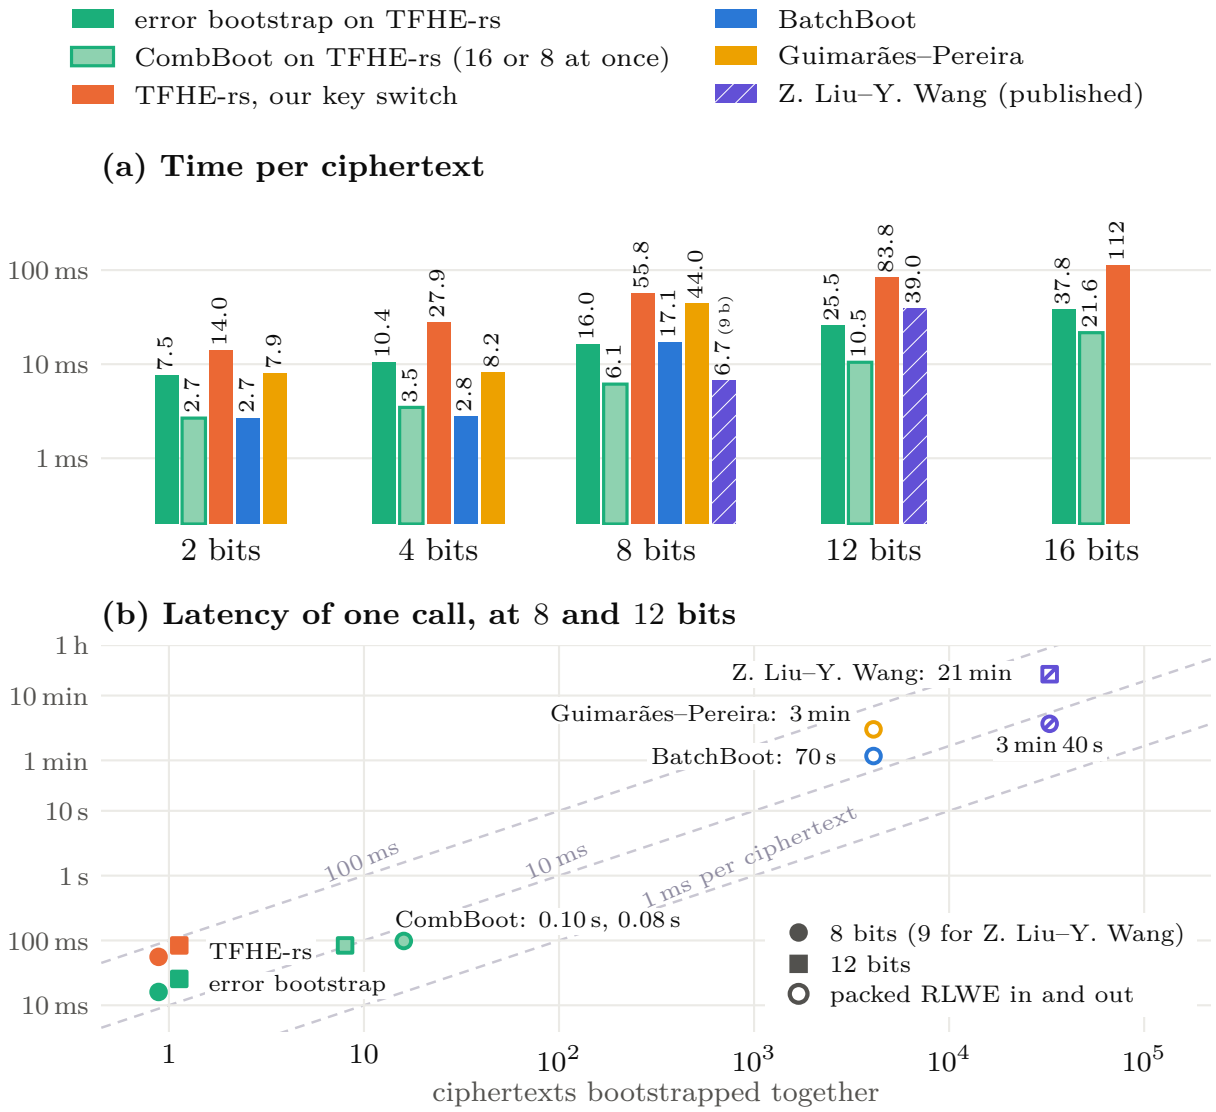

Figure 1: (a) Time per ciphertext at $N = 2048$, one thread, key switching included: error bootstrap and CombBoot (16 ciphertexts per rotation up to 8 bits, 8 beyond) built on TFHE-rs, TFHE-rs [55] with our key switch (one PBS per block), and the amortized bootstraps BatchBoot [44] and Guimarães–Pereira [32] with their code; hatched: published times of [48]. (b) Latency of one call against the number of ciphertexts bootstrapped together, at 8 bits (9 for [48]) and 12 bits; dashed: equal time per ciphertext; hollow: packed RLWE input and output (LWE packing not timed). Failure probabilities: Section 7; 2^{-30} for [48].

on one thread and at failure probability 2^{-128} after 1000 additions, an 8-bit refresh is 3.5 times faster than with TFHE-rs and the same key switch, 9.1 times when sixteen ciphertexts share a blind rotation, and a product, with the refreshes of both operands, 6.5 times; in a residue number system modulo $\widehat{P} \geq 2^{32}$, a refresh and a multiplication are 5.3 and 19 times faster than with the faster representation of TFHE-rs (Section 7). The amortized bootstraps reach their throughput only on batches of thousands of ciphertexts, refreshed together in 70 s to 21 min at 8 and 12 bits, whereas the error bootstrap refreshes one ciphertext in 16 ms at 8 bits and CombBoot sixteen in 0.1 s (Figure 1).

Section 2 reviews the related work and Section 3 fixes the notation. Sections 4, 5 and 6 present the error bootstrap, the comb bootstrap and the product, and Section 7 the measurements. The appendices give the general cyclotomic setting, the proofs, and further comparisons and measurements.

2 Related work

In FHEW and TFHE, the size of the messages is tied to the ring degree. Boolean gates use $N = 1024$, and TFHE-rs [55] uses $N = 2048$ for blocks of 4 bits; by (1), 8 bits in a single ring need $N = 2^{15}$. Several works go beyond this limit: they embed $\mathbb{Z}_q$ into the symmetric group [4], tensor two circulant rings [13], compute equality tests with traces [53], extend the domain of functional bootstraps [30, 40, 43, 51], use extended accumulators and digits [8], or iterate blind rotations on tables of low redundancy [50]. All of them pay with several blind rotations per ciphertext. The error bootstrap uses one, whose table has p entries whatever $\widehat{p}$.

Other works change the blind rotation itself. Automorphism-based rotations [11, 35, 42, 54] change how it is computed, not what it refreshes. Multi-value bootstraps [10, 14] evaluate several functions of one input in a single rotation. CombBoot does the converse: a single rotation refreshes several inputs.

In BGV, BFV and CKKS, several bootstraps already isolate the error, or borrow the blind rotation of FHEW. Digit removal [20, 27] deletes the error of BFV, alone in the low digits, inside the bootstrap. Kim, Seo and Song [39] take it out of the bootstrap and recover it approximately with CKKS (Section 1). META-BTS [5] bootstraps the error of a first bootstrap of CKKS to double its precision. Kim et al. [38] bootstrap BGV, BFV and CKKS with FHEW-type blind rotations, which frees BGV and BFV from the restrictions that digit extraction puts on the plaintext modulus.

Closer to us, in FHEW, the homomorphic floor of Liu, Micciancio and Polyakov [47] bootstraps a low part that holds message bits as well as the error, with two blind rotations, and Li, Liu and Ma [43] reduce a large error by blocks. Here, rounding at a multiple of $\widehat{p}$ leaves the error alone in the low part, and the reduced LWE ciphertext is noise-free: one rotation lifts it, with a table independent of $\widehat{p}$. This lift approximates nothing, hence the word *exact*; the output still carries the noise of the rotation.

Amortized bootstraps refresh many ciphertexts at once. Micciancio and Sorrell [52], Guimarães, Pereira and van Leeuwen [33], and De Micheli, Kim, Micciancio and Suhl [24] evaluate the LWE decryption of a packed vector by a transform whose butterflies are GSW products. F.-H. Liu and H. Wang [45, 46] reach $\tilde{O}(1)$ FHE multiplications per ciphertext in polynomial modulus, with traces in a ring whose conductor contains the LWE modulus [45]. Among the implemented ones, Z. Liu and Y. Wang [48] go through BFV, while Guimarães and Pereira [32] and BatchBoot [44] start from packed RLWE ciphertexts; all three evaluate an arbitrary function on thousands of ciphertexts. The bootstrap of CKKS also refreshes bits [6], small integers [7] and lookup tables [2]. Their throughput is high, and their latency is that of the whole batch. CombBoot batches at most sixteen ciphertexts at $N = 2048$.

We work in the framework of [15], which describes TFHE in general cyclotomic rings with the dual ring and the trace. Relative traces are a common tool: they underlie ring switching [28], the bootstrap of packed ciphertexts along towers of subrings [3], and the packing of LWE ciphertexts [19]. As in [45], the conductor of the comb bootstrap contains the modulus p of the rotated LWE ciphertexts, and traces separate the batch. The difference is that these ciphertexts encrypt errors, and that the ring is that of TFHE.

3 Notation

The torus is $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, and $\mathbb{T}_q = \frac{1}{q}\mathbb{Z}/\mathbb{Z}$ is its subgroup of order q, with $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$. For $x \in \mathbb{R}$, $\lfloor x \rceil$ is the nearest integer. The remainder of i modulo c is $i \bmod c \in \{0, \dots, c-1\}$, and j mods $p \in (-\frac{p}{2}, \frac{p}{2}]$ is the centered representative of j modulo p. The indicator $\mathbb{1}_{\mathcal{P}}$ is 1 if $\mathcal{P}$ holds and 0 otherwise. Vectors are in bold ($\mathbf{a}$, $\mathbf{s}$), and polynomials carry their variable ($s'(X)$, $a_i(X)$).

A message $\mu = i/\widehat{p} \in \mathbb{T}_{\widehat{p}}$ is encrypted under a key $\mathbf{s} \in \mathbb{S}^n$ as

$$(\mathbf{a}, b) \in \mathbb{T}^{n+1}, \qquad b = \mu + \mathbf{a} \cdot \mathbf{s} + e,$$

with a small noise $e \in \mathbb{T}$. The key is binary, $\mathbb{S} = \{0, 1\}$ and $K = 2$, or ternary, $\mathbb{S} = \{-1, 0, 1\}$ and $K = 3$.

The ring of TFHE is the cyclotomic ring of conductor $M = 2N$, a power of two; we use it with its torus,

$$\mathcal{R} = \mathbb{Z}[X]/(X^N + 1), \qquad \mathbb{T}_N[X] = \mathbb{T}[X]/(X^N + 1).$$

The coefficient of X^i in P is $[P]_i$, and X has order $2N$. Appendix C treats a general conductor, with the dual ring and the trace of [15].

An RLWE encryption of $\mu(X)$ under a ring key $s'(X) \in \mathcal{R}$ with coefficients in $\mathbb{S}$ is a pair (a, b) with $b = \mu + a\, s' + e$. Each coefficient of μ can be extracted as an LWE encryption under the vector $\mathbf{s}' \in \mathbb{S}^N$ of the coefficients of $s'(X)$. The gadget $\mathbf{g} = (B_g^{-1}, \ldots, B_g^{-\ell})$ is written (B_g, ℓ), and the decomposition $\mathrm{Dec}(a)$ has small entries, with $a \approx \mathrm{Dec}(a) \cdot \mathbf{g}$. The external product $\mathrm{Dec}(C) \cdot \mathrm{RGSW}(P)$ of an RLWE encryption C of μ with an RGSW encryption of P is an RLWE encryption of $P\mu$ (Appendix C.6.1 recalls $\mathrm{RGSW}(P)$). For $s \in \{0, 1\}$, the CMux

$$C \leftarrow C + \mathrm{Dec}(X^a C - C) \cdot \mathrm{RGSW}(s)$$

multiplies C by X^{as}.

4 Bootstrapping the error

Consider a message $\mu = \frac{i}{\widehat{p}} \in \mathbb{T}_{\widehat{p}}$ and an LWE encryption $(\mathbf{a}, b)$ of μ under the key $\mathbf{s} \in \mathbb{S}^n$, with noise e (Section 3). We first recall the TFHE bootstrap [22], then describe the bootstrap of the error and compare them.

4.1 TFHE bootstrapping with a grouped mask

As in [16, 36, 41], the mask is split into n/m groups of m consecutive coordinates (we assume for simplicity that m divides n),

$$\mathbf{a} = (\mathbf{a}_{[1]}, \ldots, \mathbf{a}_{[n/m]}), \qquad \mathbf{s} = (\mathbf{s}_{[1]}, \ldots, \mathbf{s}_{[n/m]}), \qquad \mathbf{a}_{[k]} \in \mathbb{T}^m,\ \mathbf{s}_{[k]} \in \mathbb{S}^m.$$

For a rounding resolution R (an integer), every group k and every pattern $\boldsymbol{\jmath} \in \mathbb{S}^m$, we define the grouped exponents and their rounding errors

$$\bar{a}_{k,\boldsymbol{\jmath}} = \left\lfloor R\, \mathbf{a}_{[k]} \cdot \boldsymbol{\jmath} \right\rceil, \qquad X_{k,\boldsymbol{\jmath}} = \bar{a}_{k,\boldsymbol{\jmath}} - R\, \mathbf{a}_{[k]} \cdot \boldsymbol{\jmath} \in [-\tfrac{1}{2}, \tfrac{1}{2}], \tag{2}$$

so that $X_{k,\mathbf{0}} = 0$, and we round the body separately:

$$\bar{b} = \lfloor Rb \rceil\,, \ \beta = \bar{b} - Rb.$$

Each group is rounded once, at the pattern $\mathbf{s}_{[k]}$ selected by the key; $m = 1$ is the usual coordinate-wise rounding. We gather the rounded exponents of the non-zero patterns into the *rounded mask*

$$\bar{\mathbf{a}} = (\bar{a}_{k,\boldsymbol{\jmath}})_{k,\boldsymbol{\jmath} \neq \mathbf{0}} \in \mathbb{Z}^D, \quad D = (K^m - 1)\frac{n}{m},$$

and the indicators $\delta_{\boldsymbol{\jmath},\mathbf{s}_{[k]}}$ (1 if $\boldsymbol{\jmath} = \mathbf{s}_{[k]}$, 0 otherwise) into the *extended key*

$$\boldsymbol{\delta} = (\delta_{\boldsymbol{\jmath},\mathbf{s}_{[k]}})_{k,\boldsymbol{\jmath} \neq \mathbf{0}} \in \{0, 1\}^D,$$

so that $\bar{\mathbf{a}} \cdot \boldsymbol{\delta} = \sum_k \bar{a}_{k,\mathbf{s}_{[k]}}$. Since $b - \sum_k \mathbf{a}_{[k]} \cdot \mathbf{s}_{[k]} = \mu + e$,

$$\bar{b} - \bar{\mathbf{a}} \cdot \boldsymbol{\delta} = R\mu + E_R \mod R, \qquad E_R = Re + \beta - \sum_{k=1}^{n/m} X_{k,\mathbf{s}_{[k]}} \in \mathbb{Z}, \tag{3}$$

as soon as $R\mu$ is an integer, i.e. $\widehat{p} \mid R$. Since E_R is an integer and $|\beta| \le \frac{1}{2}$, E_R is the integer closest to $Re - \sum_k X_{k,\mathbf{s}_{[k]}}$: the rounding of the body only rounds the rest and adds no error of its own.

TFHE works in the ring $\mathcal{R}$ of Section 3, of conductor $M = 2N$, and takes $R = 2N = M$. The accumulator is rotated by $X^{-(\bar{b}-\bar{\mathbf{a}}\cdot\boldsymbol{\delta})}$, the opposite of the phase (3), obtained as

$$X^{-\bar{b}} \prod_k H_k(X) \quad \text{with} \quad H_k(X) = \sum_{\boldsymbol{\jmath}} \delta_{\boldsymbol{\jmath},\mathbf{s}_{[k]}} X^{\bar{a}_{k,\boldsymbol{\jmath}}},$$

from RGSW encryptions of the indicators $\delta_{\boldsymbol{\jmath},\mathbf{s}_{[k]}}$, i.e. D bootstrapping keys (one per group and non-zero pattern). Since $H_k = 1 + \sum_{\boldsymbol{\jmath}\neq\mathbf{0}} \delta_{\boldsymbol{\jmath},\mathbf{s}_{[k]}}(X^{\bar{a}_{k,\boldsymbol{\jmath}}} - 1)$, each factor costs one decomposition in the *multi-bit form*

$$\mathrm{acc} \leftarrow \mathrm{acc} + \mathrm{Dec}(\mathrm{acc}) \cdot \sum_{\boldsymbol{\jmath}\neq\mathbf{0}} (X^{\bar{a}_{k,\boldsymbol{\jmath}}} - 1)\mathrm{RGSW}(\delta_{\boldsymbol{\jmath},\mathbf{s}_{[k]}}),$$

or, for $m = 1$, one CMux. Since $X^N = -1$, the test polynomial is negacyclic, and a message of $\mathbb{T}_{\widehat{p}}$ needs a padding bit: the bootstrap is correct when E_R ($R = 2N$) stays in a window of half-width $\frac{N}{2\widehat{p}}$, i.e. (up to half a unit of $\frac{1}{R}$, which we neglect in the statement but take into account in the computations)

$$\Big| e - \frac{1}{R} \sum_{k=1}^{n/m} X_{k,\mathbf{s}_{[k]}} \Big| < \frac{1}{4\widehat{p}}. \tag{4}$$

For $m = 1$ this is the usual TFHE condition. Attaching b to the first group, as in Proposition 5.2 of [16], slightly improves the bound.

4.2 The error bootstrap

We take $p = 2N$, so that X has order p in $\mathcal{R}$, and use the same grouping with the resolution $R = \widehat{p}\,p$. Since

$$R\mu = p\,i \in p\mathbb{Z},$$

the phase (3) holds with this R, and, dividing by R, the vector

$$\Big(\frac{\bar{\mathbf{a}}}{R}, \frac{\bar{b}}{R}\Big) \tag{5}$$

is a noise-free LWE encryption of $\mu + \frac{E_R}{R}$ under the extended key $\boldsymbol{\delta}$. Reducing (3) modulo p and dividing by p yields a noise-free encryption of the quantity $\frac{E_R}{p}$,

$$\frac{\bar{b} \bmod p}{p} - \frac{\bar{\mathbf{a}} \bmod p}{p} \cdot \boldsymbol{\delta} = \frac{E_R}{p} \mod 1. \tag{6}$$

The refresh applies to (6) the blind rotation of Section 4.1, with the exponents $\bar{b} \bmod p$ and $\bar{a}_{k,\boldsymbol{\jmath}} \bmod p$ and the same D keys, which gives an encryption of $\frac{E_R}{R}$; subtracting it from (5), brought under the output key $\mathbf{s}'$ (below), leaves a fresh encryption of μ. The table has p entries whatever $\widehat{p}$, and the division by R is built into the test polynomial

$$v(X) = \frac{1}{R} \sum_{j<N} \min(j, N-j)\, X^j :$$

rotated by X^{-E_R}, it has $\frac{E_R}{R}$ as constant coefficient whenever $|E_R| \le \frac{N}{2} = \frac{p}{4}$, the window left by the negacyclic restriction $X^N = -1$.

If p is an odd prime, the ring $\mathbb{Z}[X]/\Phi_p$ of conductor p and degree $N = p - 1$ has no negacyclic restriction, as with the general quotients of Joye and Walter [37]: the table is an arbitrary function on $\mathbb{Z}_p$, and the window doubles,

$$|E_R| \le \frac{p-1}{2},$$

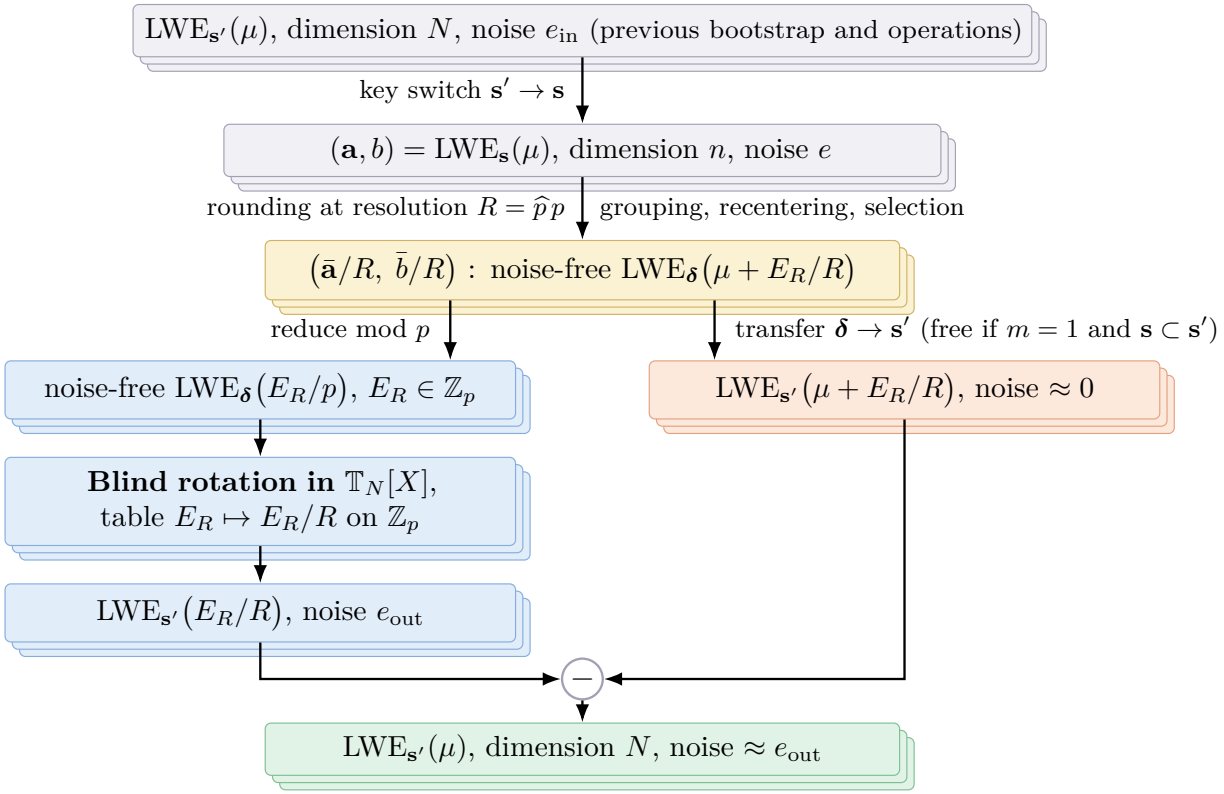

Figure 2: The error bootstrap ($p = 2N$).

without padding bit. This case needs the dual ring and the trace of [15], and a prime conductor assumed as secure as a power of two; Appendix C gives it.

The output of the rotation, extracted from its constant coefficient, is an LWE encryption of $\frac{E_R}{R}$ under $\mathbf{s}'$ (Section 3), whereas (5) is under $\boldsymbol{\delta}$: it is brought under $\mathbf{s}'$ before the subtraction. For $m = 1$ and a binary key, $\boldsymbol{\delta} = \mathbf{s}$, and this transfer is free with shared-randomness keys [9], where $\mathbf{s}$ is made of the first n coordinates of $\mathbf{s}'$: (5), padded with zeros, encrypts $\mu + \frac{E_R}{R}$ under $\mathbf{s}'$, and the key switch from $\mathbf{s}'$ to $\mathbf{s}$ before the bootstrap only involves the $N - n$ other coordinates. For $m \geq 2$, the rows of the bootstrapping keys bring (5) under $\mathbf{s}'$ with $D\ell$ products of a scalar by an RLWE ciphertext and a negligible noise; the coordinate-wise rounding of Section 4.3 avoids this transfer. The subtraction then leaves a fresh encryption of μ under $\mathbf{s}'$, of dimension N, with the noise of the output of the rotation (Figure 2).

Since (6) is noise-free and its exponents are exact, the refresh is correct as soon as E_R stays in the window $|E_R| \leq \frac{p}{4}$ above, that is, under (4) with $R = \widehat{p}\,p$: the rounding term is $\widehat{p}$ times smaller than in TFHE. Only the groups with $\mathbf{s}_{[k]} \neq \mathbf{0}$ contribute, so the rounding term is a sum of

$$n_* = \#\{k : \mathbf{s}_{[k]} \neq \mathbf{0}\} \approx \frac{n}{m}(1 - K^{-m})$$

terms for a key uniform in $\mathbb{S}^n$, instead of about $n(1 - \frac{1}{K})$ without grouping.

4.3 The noise before the bootstrap

For both bootstraps, condition (4) bounds the same quantity,

$$|\chi| < \lambda, \qquad \chi = e - \frac{T}{R}, \qquad T = \sum_{k=1}^{n/m} X_{k,\mathbf{s}_{[k]}}, \tag{7}$$

with $\lambda = \frac{1}{4\widehat{p}}$ in both cases, and $R = 2N$ for TFHE, $R = \widehat{p}\,p$ for the error bootstrap. It collects everything added since the last bootstraps,

$$\chi = e_{\mathrm{in}} + e_{\mathrm{KS}} - \frac{T}{R} :$$

the noise e_{in} of the input, of variance $\nu\,\sigma_{\text{out}}^2$ for a sum of ν fresh bootstrap outputs, the error e_{KS} of the key switch to $\mathbf{s}$, and the rounding term. We reduce the last two with known public techniques, detailed in Appendix B. The key switch is the stair key switch of [9] in its polynomial form, so that most coordinates are switched under keys of large dimension. The rounding term is recentered, as in the mean compensation of [25]: its public mean c over the key is rounded to $\bar{c}$ and added to $\bar{b}$. Its variance V over the key, also public, is brought below its 1% quantile V^* by selection, as in [12]; and the mask is rounded coordinate by coordinate, which makes the transfer of Section 4.2 free. Then χ has a mean below $\frac{1}{2R}$ and the variance

$$\sigma_\chi^2 = \sigma_e^2 + \frac{V^*}{R^2}, \qquad \sigma_e^2 = \nu\,\sigma_{\text{out}}^2 + \sigma_{\text{KS}}^2,$$

and, χ being close to Gaussian, the failure probability is about $\operatorname{erfc}(u/\sqrt{2})$, with

$$\frac{1}{u^2} = \frac{\sigma_e^2}{\lambda^2} + \frac{V^*}{(\lambda R)^2}. \tag{8}$$

A failure probability 2^{-128} requires $u \geq u_{128} \approx 13.11$; Appendix A gives the exact law in the tails. With $\nu = 1$, over 10^6 refreshes on TFHE-rs, χ has the predicted standard deviation within 3% and Gaussian tails down to probabilities of 10^{-6} (Appendix J). Since a combination $\sum_i c_i x_i$ of fresh outputs with public integers c_i has the variance $\sum_i c_i^2 \sigma_{\text{out}}^2$, the budget $\nu = 1000$ also covers multiplications by constants $|c| \leq 31$.

4.4 Comparison with TFHE

Both the error bootstrap and TFHE refresh one ciphertext at a time in the ring of TFHE, with almost the same parameters. Since $X^N = -1$, the test polynomial of TFHE has only N free coefficients: for the identity function, each of the $\widehat{p}$ values of μ gets $\Delta = N/\widehat{p}$ of them, and (4) reads, with $R = 2N$, $|2Ne - T| < \Delta/2$. The rounding term alone must fit in a box,

$$\Delta \gtrsim 2u\sqrt{V} \Longrightarrow \widehat{p} \lesssim N/(2u\sqrt{V}).$$

In the error bootstrap, μ occupies no coefficient: the same argument gives

$$p \gtrsim 4u\sqrt{V},$$

independent of $\widehat{p}$. With $u = u_{128}$, (8) gives

$$\widehat{p}_{\max}^{\text{TFHE}} \leq \frac{N}{2u\sqrt{V^*}}, \qquad \widehat{p}_{\max}^{\text{err}} \approx \frac{1}{4u\sigma_e}. \tag{9}$$

TFHE is capped by its rounding term, the error bootstrap by the noise only, the padding bit remaining in both. Between two bootstraps, with 1000 fresh outputs added and the parameters that maximize $\widehat{p}$ for TFHE at $N = 2048$ and 4096, the error bootstrap reaches a $\widehat{p}$ about five times larger. With parameters of its own, which increase n, $\widehat{p}_{\max}$ grows exponentially with N, far beyond the limit of TFHE. Appendix I gives the setting, the values, and the odd prime conductor, which gains another factor of two.

5 The comb bootstrap

The table of the error bootstrap only has to hold the p values of E_R, a few hundred, whereas the ring of TFHE offers $2N$ positions, a few thousand. Let $\mathcal{B}$ be a power of two and

$$p = \frac{2N}{\mathcal{B}},$$

so that $X^{\mathcal{B}}$ has order p. The rotations by powers of $X^{\mathcal{B}}$ act on $\mathcal{B}$ *columns*, the polynomials supported on the exponents $i \equiv k \pmod{\mathcal{B}}$, and one blind rotation refreshes $\mathcal{B}$ errors at once, one per column; for $\mathcal{B} = 1$, this is the error bootstrap of Section 4.2. Appendix C gives the construction in a general cyclotomic ring, including odd prime p, with the proofs of the general statements. The mask is grouped by m coordinates, as in Section 4.1.

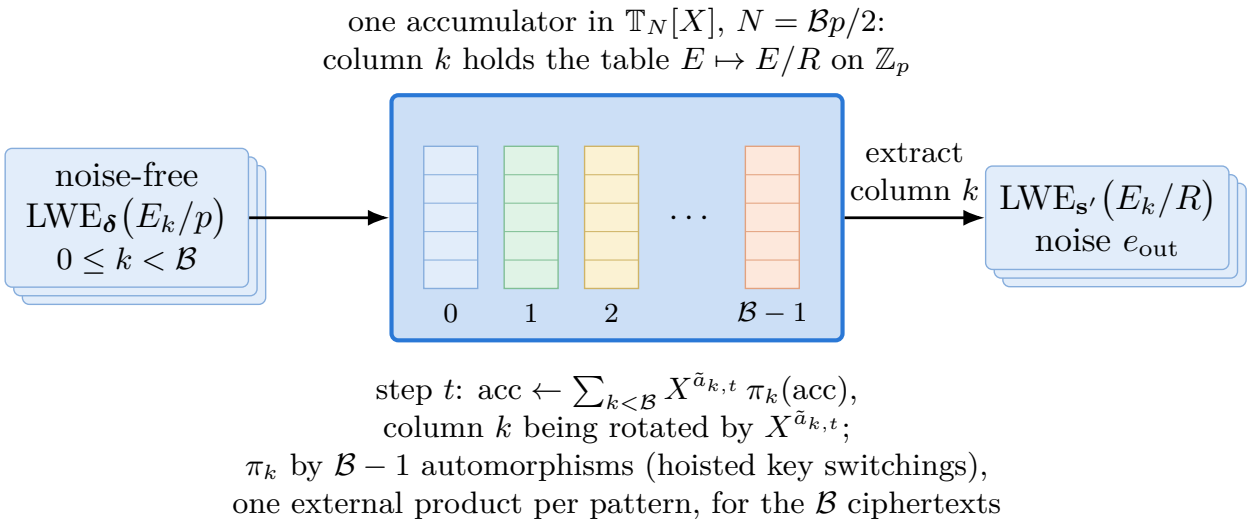

Figure 3: What the comb bootstrap changes in Figure 2: the $\mathcal{B}$ errors E_k share one blind rotation, in which each ciphertext owns a column of the accumulator. The key switch, the rounding, the transfer and the subtraction are those of Figure 2, for each ciphertext.

5.1 Setting

We use binary keys, which give the smallest p and the fewest external products per step (Appendix B), and the grouping of Section 4.1 with $m \in \{1, \dots, 4\}$. The groups are now indexed by t, and k indexes the ciphertexts: the D bootstrapping keys are the RGSW encryptions of the indicators $\delta_{\boldsymbol{\jmath},\mathbf{s}_{[t]}}$, $1 \le t \le \frac{n}{m}$, $\boldsymbol{\jmath} \neq \mathbf{0}$. The extended key and the transfer to $\mathbf{s}'$ are those of Section 4.2.

Let $\mu_0, \dots, \mu_{\mathcal{B}-1} \in \mathbb{T}_{\widehat{p}}$ be encrypted under the same key $\mathbf{s}$. Each ciphertext is rounded at the resolution $R = \widehat{p}\,p$ with the grouping (2), recentered and selected (Section 4.3) independently of the others: its body is replaced by $\bar{b}_k + \bar{c}_k$, and it is kept only if its rounding variance is at most V^*. Let $\bar{a}_{k,t,\boldsymbol{\jmath}}$ be the grouped exponents (2) of ciphertext k and $E_{R,k}$ its error (3). By (6), the reduction modulo p gives the noise-free LWE ciphertexts under $\boldsymbol{\delta}$

$$a_{k,t,\boldsymbol{\jmath}} = \frac{\bar{a}_{k,t,\boldsymbol{\jmath}} \bmod p}{p}, \qquad b_k = \frac{(\bar{b}_k + \bar{c}_k) \bmod p}{p} = \sum_{t=1}^{n/m} a_{k,t,\mathbf{s}_{[t]}} + \frac{E_k}{p},$$

$$E_k = E_{R,k} + \bar{c}_k,$$

with $a_{k,t,\mathbf{0}} = 0$. As in Section 4.2, the refresh is correct when $|E_k| \le \frac{p}{4}$, and we identify E_k with its representative in $\left(-\frac{p}{2}, \frac{p}{2}\right)$. We set

$$\tilde{a}_{k,t,\boldsymbol{\jmath}} = 2N\, a_{k,t,\boldsymbol{\jmath}} = \mathcal{B}\,(\bar{a}_{k,t,\boldsymbol{\jmath}} \bmod p),$$

a multiple of $\mathcal{B}$, so that $X^{\tilde{a}_{k,t,\boldsymbol{\jmath}}}$ is a power of $X^{\mathcal{B}}$.

Algorithm 1 below returns, for each k, an LWE encryption of E_k/R under $\mathbf{s}'$. It is subtracted from $(\bar{\mathbf{a}}_k/R, (\bar{b}_k + \bar{c}_k)/R)$, the noise-free encryption of $\mu_k + E_k/R$ under $\boldsymbol{\delta}$ given by (5), brought under $\mathbf{s}'$ (padded with zeros for $m = 1$, Section 4.2): this leaves a fresh encryption of μ_k under $\mathbf{s}'$. Only the blind rotation is shared by the $\mathcal{B}$ ciphertexts (Figure 3).

5.2 Columns

Column k, for $0 \le k < \mathcal{B}$, is the set of polynomials supported on the exponents $i \equiv k \pmod{\mathcal{B}}$. Let

$$U = \{1 + pr : \ 0 \le r < \mathcal{B}\} \subset (\mathbb{Z}/2N\mathbb{Z})^\times,$$

and, for $a \in U$, let $\tau_a(P) = P(X^a)$, an automorphism of $\mathcal{R}$ that fixes $X^{\mathcal{B}}$. The projection π_k on column k along the other columns keeps the coefficients of column k and cancels the others; for a real polynomial P,

$$\pi_k(P) = \frac{1}{\mathcal{B}} \sum_{a \in U} X^{k(1-a)}\, P\big(X^a\big).$$

Indeed, for $a = 1 + pr$,

$$X^{k(1-a)} \tau_a(X^i) = X^i\, (X^{p(i-k)})^r,$$

and since X^p is a primitive $\mathcal{B}$-th root of unity, the sum over r is $\mathcal{B}X^i$ if $i \equiv k \pmod{\mathcal{B}}$ and 0 otherwise. Thus

$$\sum_k \pi_k = \mathrm{id}, \quad \pi_k \pi_l = \delta_{kl} \pi_k,$$

and π_k commutes with the multiplication by any power of $X^{\mathcal{B}}$, which maps each column to itself. In a general cyclotomic ring, π_k is a relative trace (Appendix C.4).

5.3 The accumulator

Let

$$v(X) = \sum_{i<N} v_{\lfloor i/\mathcal{B} \rfloor}\, X^i, \qquad v_j = \frac{\min(j, \frac{p}{2} - j)}{R}, \quad 0 \le j < \frac{p}{2},$$

which repeats $\mathcal{B}$ times each coefficient of the test polynomial of Section 4.2, with p in place of $2N$.

Lemma 1. *For $0 \le k < \mathcal{B}$ and $c \in \mathbb{Z}$ with $|c \bmods p| \le \frac{p}{4}$,*

$$\big[X^{-\mathcal{B}c}\, \pi_k(v)\big]_k = \frac{c \bmods p}{R}.$$

Proof. Since $X^{\mathcal{B}p} = X^{2N} = 1$, both sides only depend on $c \bmod p$. We take $0 \le c < p$, so that the hypothesis reads $c \le \frac{p}{4}$ or $c \ge \frac{3p}{4}$. The projection π_k keeps the coefficients of v on column k:

$$\pi_k(v) = \sum_{j<p/2} v_j\, X^{\mathcal{B}j+k}, \qquad X^{-\mathcal{B}c}\, \pi_k(v) = \sum_{j<p/2} v_j\, X^{\mathcal{B}(j-c)+k}.$$

The term in X^k comes from the $j < \frac{p}{2}$ with $j \equiv c \pmod{\frac{p}{2}}$, that is $j = c$ or $j = c - \frac{p}{2}$, and $X^{\mathcal{B}p/2} = X^N = -1$ gives its sign:

$$c < \tfrac{p}{2}: \quad \big[X^{-\mathcal{B}c}\, \pi_k(v)\big]_k = v_c, \qquad\qquad c \ge \tfrac{p}{2}: \quad \big[X^{-\mathcal{B}c}\, \pi_k(v)\big]_k = -v_{c-p/2}.$$

With $v_j = \frac{\min(j, p/2-j)}{R}$, this is

$$c \le \tfrac{p}{4}: \quad \frac{c}{R}, \qquad\qquad c \ge \tfrac{3p}{4}: \quad -\frac{p-c}{R},$$

that is $\frac{c \bmods p}{R}$ in both cases. □

The accumulator starts from

$$\mathrm{acc}_0 = \sum_{k<\mathcal{B}} X^{-2Nb_k}\, \pi_k(v),$$

whose column k is that of v rotated by $-2Nb_k$, a multiple of $\mathcal{B}$ since $b_k \in \mathbb{T}_p$. For $0 \le t < \frac{n}{m}$,

$$\mathrm{acc}_{t+1} = \sum_{k<\mathcal{B}} X^{\tilde{a}_{k,t+1,\mathbf{s}_{[t+1]}}}\, \pi_k(\mathrm{acc}_t) = \mathrm{acc}_t + \sum_{\mathcal{J} \neq \mathbf{0}} \delta_{\mathcal{J},\mathbf{s}_{[t+1]}} \big(L^{\mathcal{J}}_{t+1}(\mathrm{acc}_t) - \mathrm{acc}_t\big),$$

with

$$L^{\boldsymbol{\jmath}}_{t+1}(P) = \sum_{k<\mathcal{B}} X^{\tilde{a}_{k,t+1,\boldsymbol{\jmath}}}\, \pi_k(P),$$

since $\tilde{a}_{k,t+1,\mathbf{0}} = 0$, $\sum_{\boldsymbol{\jmath}} \delta_{\boldsymbol{\jmath},\mathbf{s}_{[t+1]}} = 1$ and $\sum_k \pi_k = \mathrm{id}$: the key enters only through the indicators. For $m = 1$, $\mathrm{acc}_{t+1} = \mathrm{acc}_t + s_{t+1}\big(L^1_{t+1}(\mathrm{acc}_t) - \mathrm{acc}_t\big)$.

Theorem 1. *For $0 \le t \le \frac{n}{m}$ and $0 \le k < \mathcal{B}$,*

$$\pi_k(\mathrm{acc}_t) = X^{-2Nz_{k,t}}\, \pi_k(v), \qquad z_{k,t} = b_k - \sum_{w=1}^{t} a_{k,w,\mathbf{s}_{[w]}} \in \mathbb{T}_p.$$

In particular $z_{k,n/m} = E_k/p$ and, for $|E_k| \le \frac{p}{4}$,

$$\big[\mathrm{acc}_{n/m}\big]_k = \frac{E_k}{R} \bmod 1.$$

Proof. For $z \in \mathbb{T}_p$, $2Nz = \mathcal{B}\,(pz)$ is a multiple of $\mathcal{B}$, defined modulo $2N$: X^{-2Nz} is a power of $X^{\mathcal{B}}$, which commutes with every π_l (Section 5.2). Since $\pi_k\pi_l = \delta_{kl}\pi_k$, applying π_k to acc_0 or to acc_{t+1} keeps only the term $l = k$ of the sum. For $t = 0$, since $z_{k,0} = b_k$,

$$\pi_k(\mathrm{acc}_0) = \sum_{l<\mathcal{B}} X^{-2Nb_l}\, \pi_k\pi_l(v) = X^{-2Nz_{k,0}}\, \pi_k(v).$$

If the statement holds for t, then, since $\tilde{a}_{k,t+1,\mathbf{s}_{[t+1]}} = 2N\, a_{k,t+1,\mathbf{s}_{[t+1]}}$,

$$\begin{aligned}\pi_k(\mathrm{acc}_{t+1}) &= X^{2Na_{k,t+1,\mathbf{s}_{[t+1]}}}\, \pi_k(\mathrm{acc}_t)\\ &= X^{-2N(z_{k,t} - a_{k,t+1,\mathbf{s}_{[t+1]}})}\, \pi_k(v) = X^{-2Nz_{k,t+1}}\, \pi_k(v).\end{aligned}$$

Finally, by definition of b_k (Section 5.1), $z_{k,n/m} = E_k/p$, hence $2Nz_{k,n/m} = \mathcal{B}E_k$. The coefficient of X^k only depends on column k, and Lemma 1 with $c = E_k$ gives

$$\big[\mathrm{acc}_{n/m}\big]_k = \big[X^{-\mathcal{B}E_k}\, \pi_k(v)\big]_k = \frac{E_k}{R} \bmod 1.$$

□

This is Theorem 2 (Appendix C.5) in the ring of TFHE, whose accumulator in the dual torus is acc_t/N. Inserting the expression of π_k, for $\boldsymbol{\jmath} \neq \mathbf{0}$,

$$L^{\boldsymbol{\jmath}}_{t+1}(P) = \sum_{a\in U} A^{\boldsymbol{\jmath}}_{t+1,a}(X)\, P\big(X^a\big), \tag{10}$$

where, for $a = 1 + pr$,

$$\mathcal{B}A^{\boldsymbol{\jmath}}_{t+1,a} = \sum_{k<\mathcal{B}} X^{\tilde{a}_{k,t+1,\boldsymbol{\jmath}} - pkr}$$

is a public sum of $\mathcal{B}$ monomials. The automorphisms $P(X^a)$ do not depend on the pattern $\boldsymbol{\jmath}$: they are shared by the $2^m - 1$ maps $L^{\boldsymbol{\jmath}}_{t+1}$, which only differ by their public polynomials.

5.4 Homomorphic evaluation

The accumulator is encrypted as an RLWE ciphertext $C = (a, b)$ under $s'(X)$, stored in the FFT domain, where an automorphism is a free permutation of the evaluation points. An automorphism τ is evaluated by a key switching with the key K_τ, whose rows encrypt $\tau(s')\, B_g^{-1}, \ldots, \tau(s')\, B_g^{-\ell}$ under s' (Appendix C.6). For $a \in U$, τ_a permutes the coefficients up to sign and commutes with the decomposition, so that the $\mathcal{B} - 1$ automorphisms share one decomposition (hoisting [34]);

Algorithm 1 (comb bootstrap of the $\mathcal{B}$ LWE ciphertexts, mask grouped by m)

Input: the noise-free LWE ciphertexts of Section 5.1; the bootstrapping keys $\mathrm{RGSW}(\delta_{\jmath,\mathbf{s}_{[t]}})$, $1 \le t \le \frac{n}{m}$, $\jmath \neq \mathbf{0}$; the key-switching keys K_{τ_a}, $a \in U \setminus \{1\}$.

Output: LWE encryptions of E_k/R, $0 \le k < \mathcal{B}$, under $\mathbf{s}' = (s'_i)_{0 \le i < N}$.

1 $C \leftarrow (0, \mathrm{acc}_0)$, $\quad \mathrm{acc}_0 = \sum_{k=0}^{\mathcal{B}-1} X^{-2Nb_k}\, \pi_k(v)$ ▷ *trivial encryption of* acc_0

2 **for** $t = 0$ **to** $\frac{n}{m} - 1$ **do** ▷ *invariant: C encrypts* acc_t *under* s'

3 compute $\mathcal{B}\, A^{\jmath}_{t+1,a}$ in the FFT domain ($a \in U$, $\jmath \neq \mathbf{0}$) ▷ *public*

4 $(a', b') \leftarrow \big(\lfloor a/\mathcal{B} \rceil, \lfloor b/\mathcal{B} \rceil\big)$, $\quad C = (a, b)$ ▷ *division by* $\mathcal{B}$

5 $(d_1, \dots, d_\ell) \leftarrow \mathrm{Dec}(a')$; $\quad C_1 \leftarrow (a', b')$ ▷ *one decomposition*

6 **for** $a \in U \setminus \{1\}$ **do**

7 $C_a \leftarrow \big(0, \tau_a(b')\big) - \big(\tau_a(d_1), \dots, \tau_a(d_\ell)\big) \cdot \mathsf{K}_{\tau_a}$ ▷ *automorphisms*

8 **for** $\jmath \neq \mathbf{0}$ **do**

9 $C^{\jmath}_L \leftarrow \sum_{a \in U} \mathcal{B}\, A^{\jmath}_{t+1,a} \cdot C_a$ ▷ *public products*

10 $C \leftarrow C + \sum_{\jmath \neq \mathbf{0}} \mathrm{Dec}\big(C^{\jmath}_L - C\big) \cdot \mathrm{RGSW}\big(\delta_{\jmath,\mathbf{s}_{[t+1]}}\big)$ ▷ C *encrypts* acc_{t+1}

11 **for** $k = 0$ **to** $\mathcal{B} - 1$ **do** ▷ C *encrypts* $\mathrm{acc}_{n/m}$

12 $\mathrm{LWE}_k \leftarrow \mathrm{Extract}_k(C)$ ▷ *LWE encryption of* $[\mathrm{acc}_{n/m}]_k = E_k/R$

13 **return** $(\mathrm{LWE}_0, \dots, \mathrm{LWE}_{\mathcal{B}-1})$

giant steps do not pay for $\mathcal{B} \le 16$ (Appendix F). Finally, $\mathrm{Extract}_k(C)$ is the LWE encryption under $\mathbf{s}'$ of the coefficient k (Section 3).

The factor $\frac{1}{\mathcal{B}}$ of the projections is applied to the ciphertext before the automorphisms. For $a \in \mathbb{T}_N[X]$ with coefficients in $\mathbb{T}_q$, $q\, a_i$ being represented by an integer, we set

$$\left\lfloor \frac{a}{\mathcal{B}} \right\rceil = \sum_{i<N} \frac{1}{q} \left\lfloor \frac{q\, a_i}{\mathcal{B}} \right\rceil X^i :$$

for $q = 2^{64}$, a rounded shift by $\log_2 \mathcal{B}$ bits. The representative does not matter, and the division is exact for every q, although $\mathcal{B}$ is even (Lemma 4, Appendix G). The comments of Algorithm 1 hold by (10), Lemma 4 and Theorem 1. The public polynomials of line 3 are given in the FFT domain, for each pattern, by $p/2$ DFTs of length $\mathcal{B}$ (Appendix F).

An iteration has a part $\mathcal{C}_{\mathrm{sh}}$ shared by the patterns (decomposition and key switchings of the automorphisms) and a part $\mathcal{C}_{\mathrm{pat}}$ repeated for each of the $2^m - 1$ patterns (public products and external product), hence the cost per coordinate of the mask

$$\frac{\mathcal{C}_{\mathrm{sh}} + (2^m - 1)\, \mathcal{C}_{\mathrm{pat}}}{m} :$$

$m = 2$ is cheaper than $m = 1$ if $\mathcal{C}_{\mathrm{pat}} < \mathcal{C}_{\mathrm{sh}}$, and $m = 3$ cheaper than $m = 2$ if $5\,\mathcal{C}_{\mathrm{pat}} < \mathcal{C}_{\mathrm{sh}}$. Appendix J measures both parts, and Appendix F shows that grouping by m amounts to deferring the projections.

5.5 Noise

Appendix G fixes the usual heuristics and gives the mean squares of the fresh errors: V_{key} and V_{rnd} for the keys and the rounding of an external product, V_{aut} for a key switching and V_{div} for the division. Let n_* be the number of non-zero groups $\mathbf{s}_{[t]}$, as in Section 4.2.

Proposition 1. *The outputs of Algorithm 1 have errors of variance*

$$\sigma^2_{\text{out}} \le N\Big(\frac{n}{m}(2^m-1)V_{\text{key}} + n_*\big(V_{\text{rnd}} + (\mathcal{B}-1)\mathcal{B}V_{\text{aut}} + \mathcal{B}^2 V_{\text{div}}\big)\Big) + \frac{N}{q^2}.$$

This is Proposition 4 in the ring of TFHE (proof in Appendix H). The division contributes at most $n_*\mathcal{B}^2(1+N/2)/(12q^2)$, below 2^{-90} for $q = 2^{64}$, $\mathcal{B} \le 256$, $N \le 2^{13}$ and $n \le 2^{11}$.

Compared with $\mathcal{B}$ separate rotations, the only new term is that of the automorphisms, $Nn_*(\mathcal{B}-1)\mathcal{B}V_{\text{aut}}$, amplified by the public polynomials. Automorphism keys with more digits make it small: with single-digit bootstrapping keys $(2^{23}, 1)$, the gadget $(2^{14}, 2)$ brings the comb bootstrap within 0.3 bit of the separate rotations (Appendix G).

A tempting alternative bootstraps the $\mathcal{B}$ ciphertexts one by one in module LWE of rank $\mathcal{B}$ over $\mathbb{Z}[X]/(X^{p/2}+1)$, without automorphisms; but each ciphertext then pays a whole blind rotation: even the best of these rings ($\mathcal{B} = 4$) is 1.4 to 2.9 times slower per ciphertext than CombBoot (Appendix J).

6 Arithmetic modulo $\widehat{p}$ in one ciphertext

The error bootstrap evaluates no function, hence not the two squares by which TFHE multiplies,

$$xy = \frac{(x+y)^2 - (x-y)^2}{4}.$$

The standard tensor product and relinearization of Chillotti et al. [23] multiply without any programmable bootstrap; we give it only to show that additions and products need none, and a large $\widehat{p}$ makes it worthwhile. Let $(\mathbf{a}_i, b_i)$, $i = 1, 2$, be the fresh outputs of two distinct refreshes, with independent errors, of dimension N, under the key $\mathbf{s}'$ made of the coefficients of $s'(X)$, before the key switch. With

$$A_i(X) = a_{i,0} - \sum_{j=1}^{N-1} a_{i,j}\, X^{N-j}, \qquad [A_i\, s']_0 = \mathbf{a}_i \cdot \mathbf{s}',$$

(A_i, b_i) is an RLWE ciphertext whose phase

$$\psi_i = b_i - A_i s'$$

has $\mu_i + e_i$ as constant coefficient, the other coefficients being uniform. Left as they are, they would ruin the product: the phase is projected on its constant coefficient by $\frac{1}{N}\,\mathrm{Tr}$, where

$$\mathrm{Tr}(P) = \sum_{a \in (\mathbb{Z}/2N\mathbb{Z})^\times} P(X^a) = N\,[P]_0$$

is computed with $\log_2 N$ automorphisms [15, 19]. Modulo 2^{64}, N is not invertible; after a modulus switch to an odd q', whose rounding is negligible,

$$\mathrm{Tr}\big(N^{-1}\psi_i\big) = \psi_{i,0} = \mu_i + e_i \mod q',$$

and the noise of the input is not amplified. With

$$q' = 2^{64} - 1, \qquad N^{-1} = 2^{64-\log_2 N} \mod q',$$

the product by N^{-1} is a rotation of the 64 bits. The key switchings of the trace add errors that the following automorphisms accumulate: with a variance V_{KS} per key switching, the trace adds the variances

$$\text{(coefficient 0)}:\ V_0 = \frac{N^2-1}{3}\,V_{\text{KS}}, \qquad \text{(sum)}:\ V_{\text{tr}} = N(N-1)\,V_{\text{KS}},$$

since the key switching of the j-th step, $1 \le j \le \log_2 N$, passes through the remaining partial trace, of order $N/2^j$, which multiplies the constant coefficient and the squared norm by $N/2^j$. Lifting A_i and b_i to $[-\frac{1}{2}, \frac{1}{2})$, the cleaned phases are

$$\mu_i + E_i + I_i, \qquad I_i \in \mathbb{Z}[X],$$

where E_i is the error after the trace: e_i, of standard deviation σ_{out}, on the constant coefficient, plus the noise of the trace, and, since $\mu_i \in \frac{1}{\widehat{p}}\mathbb{Z}$,

$$\widehat{p}\,(\mu_1 + E_1 + I_1)(\mu_2 + E_2 + I_2) \equiv \widehat{p}\,\mu_1\mu_2 + \widehat{p}\,(E_1 I_2 + E_2 I_1) + \dots \mod 1,$$

where

$$\widehat{p}\,\mu_1\mu_2 = \frac{i_1 i_2 \bmod \widehat{p}}{\widehat{p}}.$$

After relinearization by a key for s'^2 and extraction, this is an encryption of the product modulo $\widehat{p}$. Since $\operatorname{Var} I_{i,j} \approx N/24$ for every j,

$$\sigma^2_{\text{prod}} \approx \widehat{p}^{\,2}\,\frac{N}{12}\left(\sigma^2_{\text{out}} + V_{\text{tr}}\right).$$

This variance does not control the tails. The main term of the constant coefficient, $\widehat{p}\,(E_{1,0}I_{2,0} + E_{2,0}I_{1,0})$, where $E_{i,0}$ has the variance $\sigma^2_{\text{out}} + V_0$, is a sum of two independent products of Gaussians, i.e. a Laplace law of scale

$$\gamma = \widehat{p}\,\sqrt{\frac{N}{24}\left(\sigma^2_{\text{out}} + V_0\right)},$$

whose tail decays like $e^{-x/\gamma}$: 2^{-128} requires 62.7 standard deviations instead of u_{128}. Since the trace puts the error on a single coefficient, the central limit theorem does not apply, and a Gaussian estimate would accept parameters that fail with probability 2^{-78} (Appendix J). We set the parameters with the exact law of the sum of this term and of the others, taken Gaussian (Appendix A), and likewise for a square, whose two operands are the same ciphertext: its main term is $2\widehat{p}\,E_{1,0}I_{1,0}$, of scale 2γ. Since both operands are refreshed just before it, products chain without bound on the multiplicative depth, a sum of products costing one refresh per term.

7 Implementation and comparison with TFHE-rs

7.1 Implementation and parameters

We implement the ring of TFHE at $N = 2048$ and $q = 2^{64}$ on the library TFHE-rs [55] (1.8.1, in Rust): the error bootstrap and CombBoot use its binary keys, its bootstrapping keys in the Fourier domain, its FFT, its CMux and its blind rotation, and add the polynomial key switch, the rounding with recentering and selection and, for CombBoot, the automorphism key switchings and the public products. The mask is rounded coordinate by coordinate, so that the transfer to $\mathbf{s}'$ is free. The error bootstrap takes $m = 2$ ($m = 1$ at 4 and 14 bits), and CombBoot $m = 1$. All the keys have the smallest noise allowed at 128 bits for $q = 2^{64}$ by the lattice estimator [1], $\sigma_{128}(2048) = 2^{-50.3}$ under the ring key.[3] Each configuration is the cheapest one of the noise model at failure probability 2^{-128} after 1000 additions, errors of the FFT included, with a margin of 0.2 bit on the standard deviations, and the measured noises agree with the predictions within 0.3 bit (Appendix J).

The times are measured on a MacBook Pro with an Apple M3 Max chip, on one thread: each is the minimum of 3 to 15 repetitions within a session, averaged over 10 sessions, with a standard deviation below 4% of the mean for the refreshes and the multiplications, and each ratio

[3] The same rule holds for every key: the bootstrapping, automorphism, trace and relinearization keys, under the ring key of degree 2048, have the noise $\sigma_{128}(2048)$; the keys of a step of the stair key switch to the dimension x, $\sigma_{128}(x)$ (Appendix B); and the 128 public encryptions of zero of the selection, under $\mathbf{s}$, $\sigma_{128}(n)$.

compares programs of the same session. All the times and ratios of this section are measured on TFHE-rs. A second implementation, in Julia, covers all the variants ($m \leq 4$, giant steps, odd prime conductors) and has an exact backend, without the error of the FFT: it validates the noise model, the errors given to the rotation end to end, and the independence of the outputs of one CombBoot (Appendix J). The code and the logs of the measurements are available at https://github.com/CycloFHE/Comb-Bootstrap.

7.2 Comparison with TFHE-rs

TFHE-rs computes in radix form on integers of $2h$ bits (`FheUint`$2h$), with h blocks of 2 bits of message and 2 bits of carry, at $N = 2048$: this is its answer to 8 bits, not 8-bit arithmetic in one ciphertext. Figure 4 compares it, with its default parameters at failure probability 2^{-128}, with the error bootstrap and CombBoot built on it, for

$$\widehat{p} = 4^h, \qquad h = 1, \ldots, 8.$$

A refresh of TFHE-rs is one key switch and one PBS per block, and runs our polynomial key switch, so that the comparison bears on the bootstraps. Its addition and its product keep its own key switch; ours is the product of Section 6, after refreshes of both operands. At 8 bits:

$\widehat{p} = 2^8$, times in ms	ours	TFHE-rs	ratio
refresh, one ciphertext	16.0	55.8	3.5
refresh, CombBoot (16 at once)	6.14	55.8	9.1
product, operands refreshed	64.7	422	6.5

From 2 to 16 bits, the refresh of the error bootstrap is 1.9 to 3.9 times faster than that of TFHE-rs, and that of CombBoot 5.2 to 9.1 times, which makes CombBoot 1.7 to 3.0 times faster per ciphertext than the error bootstrap alone. With the refreshes of both operands, the product is 3.3 to 6.5 times faster from 4 to 8 bits, and 1.6 times slower at 2 bits, where TFHE-rs multiplies a single block; beyond 8 bits, at $N = 2048$ in double precision, the error of the FFT and the noise of the keys of the trace exceed the window. An addition is an LWE addition, and 1000 of them fit between two refreshes; in TFHE-rs, the carry of a block holds a few additions, then two PBS per block propagate it.

These ratios are measured on one thread. On 8 threads, TFHE-rs refreshes the blocks of an integer in parallel, and its latency is that of the error bootstrap within 11%; the rotation of CombBoot, shared by the threads, takes 36 ms for sixteen 8-bit ciphertexts, three times less than on one (Appendix J).

The amortized bootstraps of BatchBoot [44] and of Guimarães and Pereira [32], run unchanged on the same laptop, evaluate a function at each bootstrap, on batches of 2048 to 8192 ciphertexts that take 5.5 s to 3 min, against 0.04 to 0.10 s for sixteen with CombBoot (Figure 1). Per ciphertext, BatchBoot is on par with CombBoot at 2 bits, faster at 4 bits and slower beyond, and the method of Guimarães and Pereira is slower throughout; both run at failure probabilities of 2^{-120} to 2^{-62} per coefficient, against 2^{-128} after 1000 additions here (Appendix J).

7.3 Integers of 32 and 16 bits

The message modulus enters CombBoot only through the multiplication by $\widehat{p}$ before the rounding and the values of the test polynomial, so that the $\mathcal{B}$ columns of one call can have different moduli: column k multiplies its ciphertext by its own $\widehat{p}_k$, any integer, and its test polynomial has $\widehat{p}_k$ in place of $\widehat{p}$. An integer is held by its residues modulo the $\mathcal{B}$ pairwise coprime integers of smallest maximum whose product exceeds its range,

$$(32 \text{ bits}) : \ \widehat{p}_k \in \{25, 23, 21, 19, 17, 16, 13, 11\}, \quad (16 \text{ bits}) : \ \widehat{p}_k \in \{19, 17, 16, 15\},$$

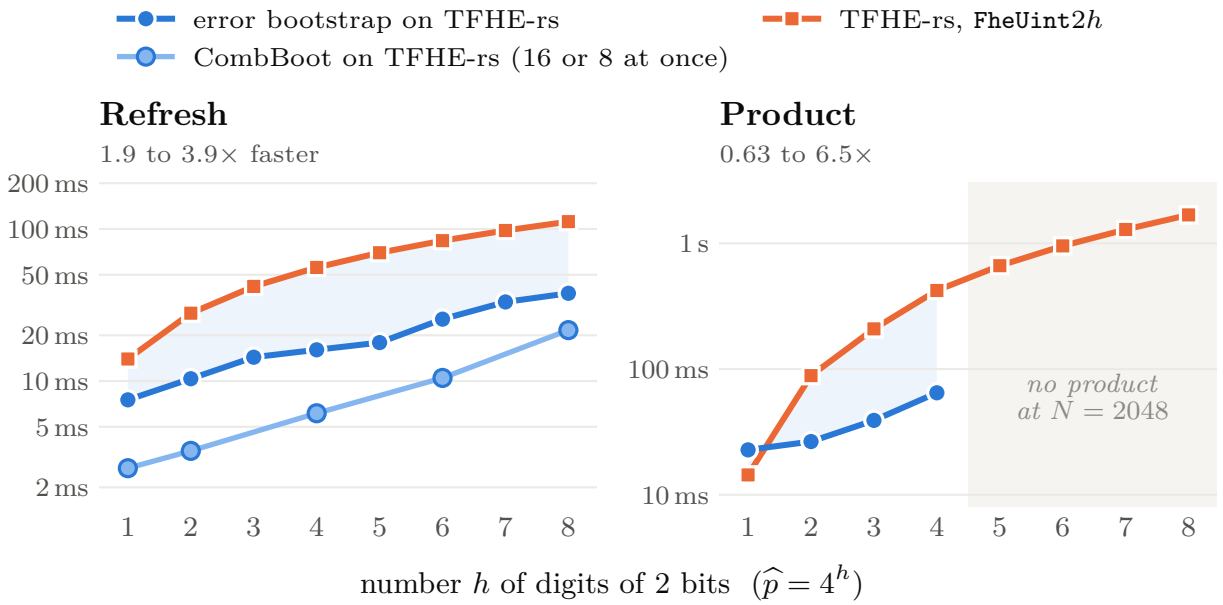

Figure 4: Error bootstrap and CombBoot (16 or 8 at once) built on TFHE-rs, against TFHE-rs on integers of $2h$ bits ($\widehat{p} = 4^h$, $N = 2048$): times in ms, one thread, key switching included; product: refreshes of both operands, plus the product; TFHE-rs: refresh with our key switch. Under each title: range of the ratio of the times of TFHE-rs and of the error bootstrap.

and one CombBoot refreshes them, with $\mathcal{B} = 8$ ($p = 512$) and $\mathcal{B} = 4$ ($p = 1024$). Additions and multiplications are computed residue by residue, the latter by the product of Section 6, whose operands come from two distinct calls of CombBoot (Appendix J). TFHE-rs [55] offers two representations: radix integers (`FheUint32`, `FheUint16`), and CRT integers [8, 55], whose residues of at most 5 bits sit in blocks with one bit of carry, in larger rings; both keep their own key switch.

Table 1: Integers of 32 and 16 bits, times in ms: residues refreshed by one CombBoot, against the radix and CRT integers of TFHE-rs; multiplication with the refreshes of both operands; ratio to the faster representation of TFHE-rs for each operation.

	32 bits		16 bits	
	refresh	multiplication	refresh	multiplication
residues, CombBoot	45.5	126.7	38.7	96.0
TFHE-rs, radix	242	6666	121	1681
TFHE-rs, CRT	816	2451	172	518
ratio	5.3	19.3	3.1	5.4

Table 1 gives the times. The radix integers refresh faster, their blocks being smaller, and the CRT integers multiply faster, with three PBS per block. Here one CombBoot refreshes all the residues, and the products need none: against the faster representation for each operation, a refresh is 5.3 times faster and a multiplication 19.3 times at 32 bits, and 3.1 and 5.4 times at 16 bits, where TFHE-rs uses smaller rings. The arithmetic is that of $\mathbb{Z}/\widehat{P}\mathbb{Z}$, $\widehat{P} = \prod_k \widehat{p}_k$: as for the CRT integers of TFHE-rs, the reduction modulo 2^{32}, comparisons and divisions require base extensions, hence programmable bootstraps.

8 Conclusion

Bootstrapping the error rather than the message removes the rounding term from the limit on $\widehat{p}$: with the same parameters, $\widehat{p}$ is about five times larger than in TFHE, ten times in a ring of odd

prime conductor, and the error bootstrap remains possible far beyond the limit of TFHE. Since its table only holds p values, $2N/p$ errors share one blind rotation. Built on TFHE-rs, an integer modulo $\widehat{p}$ stays in one ciphertext, refreshed several times faster than by TFHE-rs and multiplied without programmable bootstrap, and a residue number system extends this arithmetic to 32-bit integers. The error bootstrap is not tied to the blind rotation: the bootstrap of CKKS can also reduce the phases of the rounded ciphertexts modulo 1, which leaves their errors E_R/p and never the message, and thus refresh thousands of LWE ciphertexts at once, at a cost independent of $\widehat{p}$ but with the latency of the whole batch, as in Alexandru, Kim and Polyakov [2].

The error bootstrap only refreshes: any other function of μ still requires a programmable bootstrap or a decomposition in digits. The batch of the comb bootstrap is small, the product is limited to $\widehat{p} = 2^8$ at $N = 2048$ in double precision, and odd prime conductors are slower and assumed as secure as power-of-two rings of the same degree. The representation (the sign and the comparisons, which follow from the residues [17], the reduction modulo 2^{32}, a conversion to blocks), the grouped rounding for $m \geq 2$, complete applications, a comparison with leveled BFV and BGV, and other platforms are left to further work.

Use of AI systems

AI systems were used, on the authors' instructions and subject to their revision, to co-develop the implementations on TFHE-rs (Section 7) and in Julia (Appendix J) and the scripts that plot the figures from the recorded measurements, and for LaTeX 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; the authors take full responsibility for the content.

References

[1] M. R. Albrecht, R. Player, and S. Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.

[2] A. Alexandru, A. Kim, and Y. Polyakov. General functional bootstrapping using CKKS. In *CRYPTO*, 2025.

[3] J. Alperin-Sheriff and C. Peikert. Practical bootstrapping in quasilinear time. In *CRYPTO*, 2013.

[4] J. Alperin-Sheriff and C. Peikert. Faster bootstrapping with polynomial error. In *CRYPTO*, 2014.

[5] Y. Bae, J. H. Cheon, W. Cho, J. Kim, and T. Kim. META-BTS: bootstrapping precision beyond the limit. In *ACM CCS*, 2022.

[6] Y. Bae, J. H. Cheon, J. Kim, and D. Stehlé. Bootstrapping bits with CKKS. In *EUROCRYPT*, 2024.

[7] Y. Bae, J. Kim, D. Stehlé, and E. Suvanto. Bootstrapping small integers with CKKS. In *ASIACRYPT*, 2024.

[8] L. Bergerat, A. Boudi, Q. Bourgerie, I. Chillotti, D. Ligier, J.-B. Orfila, and S. Tap. Parameter optimization and larger precision for (T)FHE. *Journal of Cryptology*, 2023.

[9] L. Bergerat, I. Chillotti, D. Ligier, J.-B. Orfila, A. Roux-Langlois, and S. Tap. New secret keys for enhanced performance in (T)FHE. IACR ePrint 2023/979, 2023.

[10] O. Bernard, N. Carouge, M. Joye, J.-B. Orfila, and S. Tap. Low-noise multi-value bootstrapping via log-unit lattice search and LUT shifting. IACR ePrint 2026/2310, 2026.

[11] O. Bernard and M. Joye. Bootstrapping (T)FHE ciphertexts via automorphisms: closing the gap between binary and Gaussian keys. In *ASIACRYPT*, 2025.

[12] O. Bernard, M. Joye, N. P. Smart, and M. Walter. Drifting towards better error probabilities in fully homomorphic encryption schemes. In *EUROCRYPT*, 2025.

[13] G. Bonnoron, L. Ducas, and M. Fillinger. Large FHE gates from tensored homomorphic accumulator. In *AFRICACRYPT*, 2018.

[14] S. Carpov, M. Izabachène, and V. Mollimard. New techniques for multi-value input homomorphic evaluation and applications. In *CT-RSA*, 2019.

[15] P. Chartier, M. Koskas, and M. Lemou. Exploring general cyclotomic rings in torus-based fully homomorphic encryption. *IACR Communications in Cryptology*, 3(1), 2026.

[16] P. Chartier, M. Koskas, M. Lemou, and F. Méhats. Fully homomorphic encryption on large integers. IACR ePrint 2024/155, 2024.

[17] P. Chartier, M. Koskas, M. Lemou, and F. Méhats. Homomorphic sign evaluation with a RNS representation of integers. In *ASIACRYPT*, 2024.

[18] M. Checri, R. Sirdey, A. Boudguiga, and J.-P. Bultel. On the practical CPA^{D} security of "exact" and threshold FHE schemes and libraries. In *CRYPTO*, 2024.

[19] H. Chen, W. Dai, M. Kim, and Y. Song. Efficient homomorphic conversion between (ring) LWE ciphertexts. In *ACNS*, 2021.

[20] H. Chen and K. Han. Homomorphic lower digits removal and improved FHE bootstrapping. In *EUROCRYPT*, 2018.

[21] J. H. Cheon, H. Choe, A. Passelègue, D. Stehlé, and E. Suvanto. Attacks against the IND-CPA^{D} security of exact FHE schemes. In *ACM CCS*, 2024.

[22] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène. TFHE: fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, 2020.

[23] I. Chillotti, D. Ligier, J.-B. Orfila, and S. Tap. Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for TFHE. In *ASIACRYPT*, 2021.

[24] G. De Micheli, D. Kim, D. Micciancio, and A. Suhl. Faster amortized FHEW bootstrapping using ring automorphisms. In *PKC*, 2024.

[25] T. de Ruijter, J.-P. D'Anvers, and I. Verbauwhede. Don't be mean: reducing approximation noise in TFHE through mean compensation. *IACR TCHES*, 2026:82–104, 2026.

[26] L. Ducas and D. Micciancio. FHEW: bootstrapping homomorphic encryption in less than a second. In *EUROCRYPT*, 2015.

[27] R. Geelen and F. Vercauteren. Bootstrapping for BGV and BFV revisited. *Journal of Cryptology*, 2023.

[28] C. Gentry, S. Halevi, C. Peikert, and N. P. Smart. Ring switching in BGV-style homomorphic encryption. In *SCN*, 2012.

[29] J. Gil-Pelaez. Note on the inversion theorem. *Biometrika*, 38(3–4):481–482, 1951.

[30] A. Guimarães, E. Borin, and D. F. Aranha. Revisiting the functional bootstrap in TFHE. In *IACR TCHES*, 2021.

[31] A. Guimarães, E. Borin, and D. F. Aranha. MOSFHET: optimized software for FHE over the torus. *Journal of Cryptographic Engineering*, 14:559–576, 2024.

[32] A. Guimarães and H. V. L. Pereira. Fast amortized bootstrapping with small keys and polynomial noise overhead. In *ACM CCS*, 2025.

[33] A. Guimarães, H. V. L. Pereira, and B. van Leeuwen. Amortized bootstrapping revisited: simpler, asymptotically-faster, implemented. In *ASIACRYPT*, 2023.

[34] S. Halevi and V. Shoup. Faster homomorphic linear transformations in HElib. In *CRYPTO*, 2018.

[35] I. Hwang, S. Min, and Y. Song. Carousel: fully homomorphic encryption with bootstrapping over automorphism group. In *ASIACRYPT*, 2025.

[36] M. Joye and P. Paillier. Blind rotation in fully homomorphic encryption with extended keys. In *CSCML*, 2022.

[37] M. Joye and M. Walter. Liberating TFHE: programmable bootstrapping with general quotient polynomials. In *WAHC*, 2022.

[38] A. Kim, M. Deryabin, J. Eom, R. Choi, Y. Lee, W. Ghang, and D. Yoo. General bootstrapping approach for RLWE-based homomorphic encryption. IACR ePrint 2021/691, 2021.

[39] J. Kim, J. Seo, and Y. Song. Simpler and faster BFV bootstrapping for arbitrary plaintext modulus from CKKS. In *ACM CCS*, 2024.

[40] K. Kluczniak and L. Schild. FDFB: full domain functional bootstrapping towards practical fully homomorphic encryption. *IACR TCHES*, 2023(1):501–537, 2023.

[41] C. Lee, S. Min, J. Seo, and Y. Song. Faster TFHE bootstrapping with block binary keys. In *ACM ASIACCS*, 2023.

[42] Y. Lee, D. Micciancio, A. Kim, R. Choi, M. Deryabin, J. Eom, and D. Yoo. Efficient FHEW bootstrapping with small evaluation keys, and applications to threshold homomorphic encryption. In *EUROCRYPT*, 2023.

[43] H. Li, D. Liu, and G. Ma. New strategies for bootstrapping large-error ciphertext in large-precision FHEW/TFHE cryptosystem. IACR ePrint 2024/1672, 2024.

[44] Z. Li, H. Wang, Y. Zhao, L. Li, Z. Wang, J. He, C. Wei, Y. Yan, and L. Guo. BatchBoot: fast batched bootstrapping for TFHE scheme and practical applications. In *USENIX Security*, 2026.

[45] F.-H. Liu and H. Wang. Batch bootstrapping I: a new framework for SIMD bootstrapping in polynomial modulus. In *EUROCRYPT*, 2023.

[46] F.-H. Liu and H. Wang. Batch bootstrapping II: bootstrapping in polynomial modulus only requires $\tilde{O}(1)$ FHE multiplications in amortization. In *EUROCRYPT*, 2023.

[47] Z. Liu, D. Micciancio, and Y. Polyakov. Large-precision homomorphic sign evaluation using FHEW/TFHE bootstrapping. In *ASIACRYPT*, 2022.

[48] Z. Liu and Y. Wang. Amortized functional bootstrapping in less than 7 ms, with $\tilde{O}(1)$ polynomial multiplications. In *ASIACRYPT*, 2023.

[49] R. Lugannani and S. Rice. Saddle point approximation for the distribution of the sum of independent random variables. *Advances in Applied Probability*, 12(2):475–490, 1980.

[50] S. Ma, T. Huang, A. Wang, C. Xu, T. Wei, and X. Wang. Meta-PBS: compact high-precision programmable bootstrapping. IACR ePrint 2025/2284, 2025.

[51] S. Ma, T. Huang, A. Wang, Q. Zhou, and X. Wang. Fast and accurate: efficient full-domain functional bootstrap and digit decomposition for homomorphic computation. *IACR TCHES*, 2024(1), 2024.

[52] D. Micciancio and J. Sorrell. Ring packing and amortized FHEW bootstrapping. In *ICALP*, 2018.

[53] H. Xia, F.-H. Liu, and H. Wang. More efficient functional bootstrapping for general functions in polynomial modulus. In *TCC*, 2024.

[54] B. Xiang, J. Zhang, Y. Deng, Y. Dai, and D. Feng. Fast blind rotation for bootstrapping FHEs. In *CRYPTO*, 2023.

[55] Zama. TFHE-rs: a pure Rust implementation of the TFHE scheme for Boolean and integer arithmetics over encrypted data. https://github.com/zama-ai/tfhe-rs, 2022.

A Probability that condition (4) holds

Set

$$(p \text{ odd}): \quad \lambda = \frac{1}{2\widehat{p}}, \qquad (p = 2^\alpha): \quad \lambda = \frac{1}{4\widehat{p}},$$

and

$$\chi = e - \frac{1}{R}\sum_k X_{k,\mathbf{s}_{[k]}},$$

so that (4), and (11) for p odd, read $|\chi| < \lambda$. Without recentering or selection, assume that the masks are uniform and independent of e, and that $e \sim \mathcal{N}(0, \sigma_e^2)$. Conditionally on the key, the

$$n_* = \#\{k : \mathbf{s}_{[k]} \neq \mathbf{0}\}$$

nonzero terms $X_{k,\mathbf{s}_{[k]}}$ are then independent and uniform on $[-\frac{1}{2}, \frac{1}{2}]$, since distinct blocks involve disjoint coordinates of the mask.

Proposition 2 (Exact law)**.** *With $a = \frac{1}{2R}$, χ has the cumulant generating function*

$$\Psi(t) = \frac{\sigma_e^2 t^2}{2} + n_* \log \frac{\sinh at}{at},$$

and, for p odd, since $\log \frac{\sinh x}{x} \leq \frac{x^2}{6}$,

$$\mathbb{P}(|\chi| \geq \lambda) \leq 2e^{-u^2/2}, \qquad u^{-2} = (2\widehat{p}\,\sigma_e)^2 + \frac{n_*}{3p^2}.$$

As in Proposition 5.2 of [16], with the symmetric bounds $\pm\lambda$, the law of χ follows from Ψ by the Gil-Pelaez inversion [29], and small failure probabilities by the Lugannani–Rice formula [49]. Compared with this exact law, checked by Monte Carlo, the Gaussian estimate $\operatorname{erfc}(u/\sqrt{2})$ of Section 4.3 is conservative: for $100 \leq n_* \leq 700$, it overestimates the smallest admissible p by at most 4.4% at failure probability 2^{-64} and 10% at 2^{-128}, the most for $\sigma_e = 0$ and $n_* = 100$. Without recentering and with $\sigma_e = 0$, $|\chi| \leq \frac{n_*}{2R}$, so that the refresh never fails as soon as (p odd): $n_* < p$, ($p = 2^\alpha$): $n_* < \frac{p}{2}$.

Tail of the product. For a Laplace term of scale γ and an independent Gaussian term of variance σ^2 (Section 6), the sum exceeds x in absolute value with probability

$$2Q\Big(\frac{x}{\sigma}\Big) + e^{\sigma^2/2\gamma^2}\Big(e^{-x/\gamma}Q\Big(\frac{\sigma}{\gamma} - \frac{x}{\sigma}\Big) - e^{x/\gamma}Q\Big(\frac{\sigma}{\gamma} + \frac{x}{\sigma}\Big)\Big),$$

with Q the Gaussian tail.

B Reducing the noise before the bootstrap

Key switch. The keys $\mathbf{s} = \mathbf{s}'|_{x_0} \subset \cdots \subset \mathbf{s}'|_{x_J} = \mathbf{s}'$ are nested prefixes of $\mathbf{s}'$ (shared-randomness keys [9]). Cut into blocks of d coordinates, the LWE ciphertext of dimension N is a GLWE ciphertext of rank N/d over $R_d = \mathbb{Z}[X]/(X^d+1)$, and each step of the stair is a GLWE key switch over R_d that lowers the rank, with module-LWE keys of noise $\sigma_{128}(x_i)$, whose security is taken as that of LWE in dimension x_i. With the gadget $(B_{\mathrm{KS}}, \ell_{\mathrm{KS}})$ and $\rho = \mathbb{E}[s_j^2]$, a step from y to x adds

$$(y-x)\Big(\ell_{\mathrm{KS}}\,\frac{B_{\mathrm{KS}}^2+2}{12}\,\sigma_{128}(x)^2 + \rho\,\frac{B_{\mathrm{KS}}^{-2\ell_{\mathrm{KS}}}}{12}\Big).$$

The noise $v_{\mathrm{KS}}(x)$ of one coordinate switched with keys of dimension x decreases by about $2^{-0.05}$ per dimension, so that, with $d = 16$ and a step every d coordinates,

$$\sigma_{\mathrm{KS}}^2 \approx \frac{d\,v_{\mathrm{KS}}(n)}{1-2^{-0.05d}} = 38\,v_{\mathrm{KS}}(n),$$

against $(N-n)\,v_{\mathrm{KS}}(n)$ in a single step ($1200\,v_{\mathrm{KS}}(n)$ at $N = 2048$, $n = 848$). The digits must have zero mean (a balanced digit $\frac{B_{\mathrm{KS}}}{2}$ is replaced by $-\frac{B_{\mathrm{KS}}}{2}$, with a carry, with probability $\frac{1}{2}$); otherwise the noise of the keys becomes a bias that depends on the key, up to 1.3 bit above the prediction. At $N = 2048$, on TFHE-rs, the key switches take 0.13 to 0.40 ms for the configurations of Figure 4 (0.9 to 2.3% of a refresh, keys of 1.8 to 6.4 MB), and 0.11 to 0.33 ms for those of Table 3 (1.5 to 4.0% of the time per ciphertext of CombBoot, keys of 1.6 to 5.2 MB).

Rounding. The errors $X_{k,\boldsymbol{\jmath}}$ are public, whereas the key is not: for a key uniform in $\mathbb{S}^n$, the terms $X_{k,\mathbf{s}_{[k]}}$ are independent, with the public means and variances

$$c_k = K^{-m}\sum_{\boldsymbol{\jmath}\in\mathbb{S}^m} X_{k,\boldsymbol{\jmath}}, \qquad v_k = K^{-m}\sum_{\boldsymbol{\jmath}\in\mathbb{S}^m} X_{k,\boldsymbol{\jmath}}^2 - c_k^2,$$

and T has the mean $c = \sum_k c_k$ and the variance $V = \sum_k v_k$. *Recentering* adds $\bar{c} = \lfloor c \rceil$ to $\bar{b}$ before the rotation (and subtracts $\bar{c}/R$ after it, for the error bootstrap), which leaves the mean $c - \bar{c} \in [-\frac{1}{2}, \frac{1}{2}]$. *Selection* adds public encryptions of zero, which change the mask but not μ, until $V \le V^*$, at the cost of the noise of an encryption of zero in e_{in}; V is close to Gaussian, and, with about 100 trials on average,

$$V^* = \frac{n}{m}\,\mathbb{E}[v_k] - 2.33\sqrt{\frac{n}{m}\,\mathrm{Var}\,v_k}.$$

Rounding coordinate by coordinate, $\bar{a}_i = \lfloor R\,a_i \rceil$, with the sums $\bar{\mathbf{a}}_{[k]} \cdot \boldsymbol{\jmath}$ as exponents of the patterns, keeps the law of $m = 1$ for every m: the rounded ciphertext $(\bar{\mathbf{a}}, \bar{b})/R$ is under $\mathbf{s}$ itself, and the transfer of Section 4.2 is free. For a binary key, it gives $\mathbb{E}[v_k] = \frac{1}{48}$ and $\mathrm{Var}\,v_k = \frac{1}{2880}$ per coordinate, against $\mathbb{E}[v_k] \approx \frac{(1-2^{-m})^2}{12}$ per group of m with the grouped rounding ($\frac{3}{64}$ for $m = 2$); for a ternary key, $c_k = 0$, recentering brings nothing, and $\mathbb{E}[v_k] = \frac{1-3^{-m}}{12}$, so that grouping helps. Without recentering, V is replaced by $V + c^2$, of average $\frac{n}{m}\cdot\frac{1-K^{-m}}{12}$. For the comb bootstrap (Section 5.1), binary keys are the most favorable case for the smallest p: after recentering, $\mathbb{E}[v_k] = \frac{1}{48}$ for a binary key, against $\frac{1}{18}$ for a ternary key, for which recentering brings nothing, so that the smallest p is about 1.6 times smaller. Grouping hardly changes it ($m = 4$ lowers it by about 6%); m is chosen for the cost (Section 7). Moreover, a binary key needs $2^m - 1$ external products per step of the rotation, against $3^m - 1$ for a ternary key.

C The comb bootstrap in a general cyclotomic ring

This appendix gives the error bootstrap and the comb bootstrap for a general conductor, with the duality of [15], and states the results proved in Appendices E to H. In the ring of TFHE (case (B) below, with $M_0 = \mathcal{B}$), they reduce to Sections 4.2 and 5, whose accumulator is N times the one below.

C.1 Cyclotomic rings and the error bootstrap

Cyclotomic rings. For a conductor M, Φ_M is the M-th cyclotomic polynomial, of degree $N = \varphi(M)$ ($\Phi_M = X^N + 1$ for $M = 2N$ a power of two), and

$$\mathcal{K} = \mathbb{Q}[X]/\Phi_M, \qquad \mathcal{R} = \mathbb{Z}[X]/\Phi_M, \qquad \mathcal{T} = \mathcal{K}/\mathcal{R} \ \subset \ \mathbb{T}[X]/\Phi_M.$$

As in [15], the trace inner product defines the dual of $\mathcal{R}$ and its torus:

$$\langle P,\, Q\rangle = \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}}\big(P(X)\, Q(X^{-1})\big), \qquad \mathcal{R}^\vee = \big\{Q \in \mathcal{K} :\ \langle Q,\, \mathcal{R}\rangle \subset \mathbb{Z}\big\},$$

$$\mathcal{T}^\vee = \mathcal{K}/\mathcal{R}^\vee,$$

with $\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}}(P) = \sum_{a\in(\mathbb{Z}/M\mathbb{Z})^\times} P(X^a)$, and $\mathcal{R}^\vee = \frac{1}{N}\mathcal{R}$ for M a power of two (Appendix D.2).

Let p be an odd prime or a power of two $p = 2^\alpha$, and $N = \varphi(p)$, i.e. $N = p-1$ or $N = p/2$. The error bootstrap of Section 4.2 applies with the blind rotation in $\mathbb{T}[X]/\Phi_p$, of degree N. If p is an odd prime, X has odd order p and the table is an arbitrary function on $\mathbb{Z}_p$: there is no negacyclic restriction, and the correction

$$E_R \mapsto \frac{E_R}{R} \ \text{ is exact for } \ |E_R| \le \frac{p-1}{2}.$$

If $p = 2^\alpha$, then $\Phi_p = X^N + 1$ is the ring of TFHE and $X^N = -1$, so the table must satisfy $f(j+N) = -f(j)$ for $j \in \mathbb{Z}_p$. The correction

$$E_R \mapsto \frac{E_R}{R} \ \text{ is exact for } \ |E_R| \le \frac{p}{4}.$$

The table is carried by the test polynomial $v^* \in \mathcal{T}^\vee$ of the dual torus [15], which after rotation by X^{-E_R} is read by $\langle X^{-E_R} v^*,\, 1\rangle = \langle v^*,\, X^{E_R}\rangle$, so the coordinates of v^* are the values of the table. For p an odd prime,

$$v^* = \sum_{j=0}^{p-1} \frac{j \bmod p}{p\,R}\, X^j, \qquad \langle v^*,\, X^j\rangle = \frac{j \bmod p}{R} \quad (j \in \mathbb{Z}_p);$$

for $p = 2^\alpha$,

$$v^* = \frac{1}{NR} \sum_{j=0}^{N-1} \min(j, N-j)\, X^j, \qquad \langle v^*,\, X^j\rangle = \frac{j}{R} \quad \Big(|j| \le \frac{N}{2}\Big),$$

the other coordinates following from $\langle v^*,\, X^{j+N}\rangle = -\langle v^*,\, X^j\rangle$. In this case $\mathcal{R}^\vee = \frac{1}{N}\mathcal{R}$, and Nv^* is an ordinary test polynomial with coefficients in $\mathbb{T}$. The refresh is correct when (4) holds for $p = 2^\alpha$, and, for p odd, when

$$\Big| e - \frac{1}{R} \sum_{k=1}^{n/m} X_{k,\mathbf{s}_{[k]}} \Big| < \frac{1}{2\widehat{p}}: \tag{11}$$

the noise before the bootstrap is that of Section 4.3, with $\lambda = \frac{1}{2\widehat{p}}$ for p odd. The output (a, b) of the rotation, after extraction of the coordinate $\langle\cdot,\, 1\rangle$, is an LWE encryption of $\frac{E_R}{R}$ under the key $\mathbf{s}' = (s'_i)_{0\le i<N}$: since $\langle a\, s',\, 1\rangle = \sum_{i<N} s'_i \langle a,\, X^{-i}\rangle$, the extracted mask is $\big(\langle a,\, X^{-i}\rangle\big)_{0\le i<N}$ and the body is $\langle b,\, 1\rangle$. For $p = 2^\alpha$, this is the usual extraction of the constant coefficient of $N \cdot (a, b)$.

C.2 Algebraic setting

Let M_0 be a power of 2 and $M = M_0 p$, with

$$\text{(A)}: \quad p > 2 \text{ prime},\ M_0 \ge 4, \qquad\qquad \text{(B)}: \quad p = 2^\alpha,\ \alpha \ge 2,\ M_0 \ge 2.$$

Then

$$\text{(A)}: \quad \Phi_M(X) = \frac{\Phi_p(X^{M_0})}{\Phi_p(X^{M_0/2})}, \qquad \text{(B)}: \quad \Phi_M(X) = \Phi_p(X^{M_0}) = X^{M/2} + 1,$$

and the degree of Φ_M is $N = \mathcal{B}\,\varphi(p)$, with

$$\text{(A)}: \quad \mathcal{B} = \frac{M_0}{2},\ \varphi(p) = p - 1, \qquad\qquad \text{(B)}: \quad \mathcal{B} = M_0,\ \varphi(p) = \frac{p}{2}.$$

$\mathcal{B} \geq 2$ is the number of columns (Section C.4). In case (A) with $M_0 = 2$, or in case (B) with $M_0 = 1$, there is a single column and the construction reduces to the error bootstrap of Section 4.2.

Recall that $\mathcal{K} = \mathbb{Q}[X]/\Phi_M(X) \simeq \mathbb{Q}(\zeta_M)$. Since $\Phi_p(X^{M_0}) = 0$ in $\mathcal{K}$, the subfield generated by X^{M_0} is

$$\mathbb{Q}(X^{M_0}) = \mathbb{Q}[Y]/\Phi_p(Y), \qquad Y = X^{M_0},$$

of degree $\varphi(p)$ over $\mathbb{Q}$ (in case (B), $\Phi_p(Y) = Y^{p/2} + 1$), and $[\mathcal{K} : \mathbb{Q}(X^{M_0})] = \mathcal{B}$.

Definition 1. The *full* trace of $P \in \mathcal{K}$ is

$$\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}}(P) = \sum_{a \in (\mathbb{Z}/M\mathbb{Z})^\times} P(X^a) \in \mathbb{Q},$$

and the *relative* traces of $P \in \mathcal{K}$ and $Q \in \mathbb{Q}(X^{M_0})$ are

$$\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(P) = \sum_{a \in U} P(X^a) \in \mathbb{Q}(X^{M_0}),$$

and

$$\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}(Q) = \sum_{b \in (\mathbb{Z}/p\mathbb{Z})^\times} Q(X^{M_0 b}) \in \mathbb{Q},$$

where

$$U = \big\{a \in (\mathbb{Z}/M\mathbb{Z})^\times : a \equiv 1 \mod p\big\}, \qquad |U| = \mathcal{B}.$$

Then

$$\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}} = \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}} \circ \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}.$$

In case (A), $a \mapsto a \bmod M_0$ is a bijection from U onto $(\mathbb{Z}/M_0\mathbb{Z})^\times$; in case (B),

$$U = \{1 + pr : r \in \mathbb{Z}_{M_0}\}.$$

The trace inner product reads $\langle P, Q\rangle = \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}}(P\bar{Q})$, with $\bar{Q}(X) = Q(X^{-1})$; the norm $\|P\|_{\mathrm{Tr}} = \langle P, P\rangle^{1/2}$ is such that $\|P\|_{\mathrm{Tr}}^2$ is the sum of $|P|^2$ over the complex embeddings of $\mathcal{K}$.

Inputs. The keys and the inputs are those of Section 5.1, and the refresh is correct when $|E_k| \leq E_{\max}$, with

$$\text{(A)}: \quad E_{\max} = \frac{p-1}{2}, \qquad\qquad \text{(B)}: \quad E_{\max} = \frac{p}{4},$$

which holds except with probability $2^{-\kappa}$ as soon as (Section 4.3)

$$\text{(A)}: \quad (2\widehat{p}\,\sigma_e)^2 + \frac{4V^*}{p^2} \leq u_\kappa^{-2}, \qquad\qquad \text{(B)}: \quad (4\widehat{p}\,\sigma_e)^2 + \frac{16V^*}{p^2} \leq u_\kappa^{-2},$$

with $\mathrm{erfc}(u_\kappa/\sqrt{2}) = 2^{-\kappa}$. We then identify E_k with its representative in $\left(-\frac{p}{2}, \frac{p}{2}\right)$, and we set

$$\tilde{a}_{k,t,\jmath} = M\, a_{k,t,\jmath} \in M_0\mathbb{Z}_M,$$

so that $X^{\tilde{a}_{k,t,\jmath}} = (X^{M_0})^{p\, a_{k,t,\jmath}} \in \mathbb{Q}(X^{M_0})$ is well defined, since $X^M = 1$.

C.3 Dual bases

The duality of [15] is refined along the tower $\mathcal{K} \supset \mathbb{Q}(X^{M_0}) \supset \mathbb{Q}$ (Appendix D). Let

$$\text{(A)}: \quad \varpi_k = X^{kp}, \qquad\qquad \text{(B)}: \quad \varpi_k = X^k, \qquad\qquad 0 \le k < \mathcal{B},$$

and $\omega_j = X^{jM_0}$, $0 \le j < \varphi(p)$. The products $\Omega_{j,k} = \varpi_k\,\omega_j$, for (j,k) in $I = \{0,\dots,\varphi(p)-1\} \times \{0,\dots,\mathcal{B}-1\}$, form a $\mathbb{Z}$-basis of $\mathcal{R}$, and their dual basis for the trace inner product is $\Omega^*_{j,k} = \varpi^*_k\,\omega^*_j$, with $\varpi^*_k = \varpi_k/\mathcal{B}$ and

$$\text{(A)}: \quad \omega^*_j = \frac{1}{p}\big(X^{jM_0} - X^{(p-1)M_0}\big), \qquad\qquad \text{(B)}: \quad \omega^*_j = \frac{2}{p}\,X^{jM_0}$$

(Propositions 5 and 6). The dual of $\mathcal{R}$ and its torus are

$$\mathcal{R}^\vee = \bigoplus_{(j,k)\in I} \mathbb{Z}\,\Omega^*_{j,k} = \Omega^*_{0,0}\,\mathcal{R}, \qquad \mathcal{T}^\vee = \mathcal{K}/\mathcal{R}^\vee,$$

and $P \in \mathcal{T}^\vee$ has the coordinates $\langle P,\,\Omega_{j,k}\rangle \bmod 1$ on the basis $(\Omega^*_{j,k})$. In case (B), $\Omega^*_{j,k} = \frac{1}{N}\,X^{jM_0+k}$ and $\mathcal{R}^\vee = \frac{1}{N}\mathcal{R}$.

C.4 Column projections

Column k is $\varpi^*_k\,\mathbb{Q}(X^{M_0}) = \varpi_k\,\mathbb{Q}(X^{M_0})$: in case (B), the polynomials supported on the exponents $i \equiv k \bmod \mathcal{B}$. Then $\mathcal{K} = \bigoplus_{k<\mathcal{B}} \varpi^*_k\,\mathbb{Q}(X^{M_0})$, and the projection on column k along the other columns is

$$\pi_k(P) = \varpi^*_k\,\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}\big(\bar{\varpi}_k\,P\big), \qquad \sum_{k=0}^{\mathcal{B}-1} \pi_k = \mathrm{id}, \qquad \pi_k\pi_l = \delta_{kl}\,\pi_k.$$

We say that P is *supported on column k* if $P = \pi_k(P)$.

C.5 The polynomial v^* and the accumulator iteration

C.5.1 The polynomial v^*

Definition 2. Let $v^* \in \mathcal{T}^\vee$ be the class of

$$v^*(X) = \sum_{(j,k)\in I} v^*_j\,\Omega^*_{j,k}(X),$$

with, for $0 \le j < \varphi(p)$,

$$\text{(A)}: \quad v^*_j = \frac{j \bmod p}{p\,\widehat{p}}, \qquad\qquad \text{(B)}: \quad v^*_j = \frac{\min\big(j,\,\frac{p}{2}-j\big)}{p\,\widehat{p}},$$

where $\widehat{p}$ is a nonzero integer.

Closed form. Since $\Omega^*_{j,k} = \varpi^*_k\,\omega^*_j$,

$$v^* = W\,\mathcal{V}^*, \qquad W = \sum_{k=0}^{\mathcal{B}-1} \varpi^*_k = \frac{1}{\mathcal{B}}\sum_{k=0}^{\mathcal{B}-1} \varpi_k, \qquad \mathcal{V}^* = \sum_{j<\varphi(p)} v^*_j\,\omega^*_j \in \mathbb{Q}[X^{M_0}],$$

given explicitly in Appendix F; in case (B), v^* repeats $\mathcal{B}$ times each coefficient of the test polynomial of Section 4.2.

Column decomposition. For $0 \le k < \mathcal{B}$, let

$$v^{*,k} = \sum_{j<\varphi(p)} v_j^*\,\Omega_{j,k}^* = \varpi_k^*\,\mathcal{V}^*, \quad \text{so that} \quad v^* = \sum_{k=0}^{\mathcal{B}-1} v^{*,k} \quad \text{and} \quad \pi_k(v^*) = v^{*,k}.$$

The table. For $c \in \mathbb{Z}_p$, let

$$\vartheta(c) = \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\mathcal{V}^* X^{-cM_0}\big).$$

Lemma 2. *In case (A), $\vartheta(c) = \frac{c \bmod p}{p\,\widehat{p}}$ for all c. In case (B), $\vartheta(c) = v_c^*$ for $0 \le c < \frac{p}{2}$ and $\vartheta(c + \frac{p}{2}) = -\vartheta(c)$. In both cases,*

$$\vartheta(c) = \frac{c \bmod p}{p\,\widehat{p}} \qquad \textit{for} \quad |c \bmod p| \le E_{\max}.$$

C.5.2 The accumulator

Initial accumulator. We set

$$\mathrm{acc}_0 = \sum_{k=0}^{\mathcal{B}-1} X^{-Mb_k}\,\pi_k(v^*) = \Big(\sum_{k=0}^{\mathcal{B}-1} X^{-Mb_k}\,\varpi_k^*\Big)\,\mathcal{V}^* \;\in\; \mathcal{T}^\vee :$$

column k of v^* is rotated by $-Mb_k$, where b_k is the body of the k-th ciphertext.

By Proposition 7 (Appendix F), $X^{M_0 c}\,v^{*,k}$ is supported on column k for $c \in \mathbb{Z}$, $\pi_k(\mathrm{acc}_0) = X^{-Mb_k}\,v^{*,k}$ and $\langle \mathrm{acc}_0,\, \varpi_k \rangle = \vartheta(p\,b_k)$.

Iteration. For $0 \le t < \frac{n}{m}$, set

$$\mathrm{acc}_{t+1} = \sum_{k=0}^{\mathcal{B}-1} X^{\tilde{a}_{k,t+1,\mathbf{s}_{[t+1]}}}\,\pi_k(\mathrm{acc}_t) = \mathrm{acc}_t + \sum_{\jmath \neq \mathbf{0}} \delta_{\jmath,\mathbf{s}_{[t+1]}}\big(L_{t+1}^{\jmath}(\mathrm{acc}_t) - \mathrm{acc}_t\big),$$

with

$$L_{t+1}^{\jmath}(P) = \sum_{k=0}^{\mathcal{B}-1} X^{\tilde{a}_{k,t+1,\jmath}}\,\pi_k(P),$$

since $\tilde{a}_{k,t+1,\mathbf{0}} = 0$, $\sum_{\jmath} \delta_{\jmath,\mathbf{s}_{[t+1]}} = 1$ and $\sum_k \pi_k = \mathrm{id}$: the key enters only through the indicators. For $m = 1$, with $L_{t+1} = L_{t+1}^{1}$,

$$\mathrm{acc}_{t+1} = \mathrm{acc}_t + s_{t+1}\big(L_{t+1}(\mathrm{acc}_t) - \mathrm{acc}_t\big).$$

Note that acc_0 is obtained from v^* in the same way, with the rotations X^{-Mb_k}.

Theorem 2 (closed form)**.** *For every $0 \le t \le \frac{n}{m}$ and $0 \le k < \mathcal{B}$,*

$$\pi_k(\mathrm{acc}_t) = X^{-Mz_{k,t}}\,v^{*,k}, \qquad \langle \mathrm{acc}_t,\, \varpi_k \rangle = \vartheta(p\,z_{k,t}),$$

with

$$z_{k,t} = b_k - \sum_{w=1}^{t} a_{k,w,\mathbf{s}_{[w]}} \in \mathbb{T}_p.$$

In particular $z_{k,n/m} = E_k/p$: the final accumulator is

$$\mathrm{acc}_{n/m} = \sum_{k=0}^{\mathcal{B}-1} X^{-M_0 E_k}\,v^{*,k},$$

and, for $|E_k| \le E_{\max}$,

$$\langle \mathrm{acc}_{n/m},\, \varpi_k \rangle = \frac{E_k}{p\,\widehat{p}} = \frac{E_k}{R} \bmod 1.$$

C.5.3 Simplified form of the iteration

For $a \in U$, ϖ_k being a monomial, $\bar{\varpi}_k(X^a) = \varpi_k^{-a}$, hence

$$\varpi_k^*\, \bar{\varpi}_k(X^a) = \frac{\varpi_k^{1-a}}{\mathcal{B}}.$$

Proposition 3. *For* $P \in \mathcal{K}$ *and* $0 \le k < \mathcal{B}$,

$$\pi_k(P) = \frac{1}{\mathcal{B}} \sum_{a \in U} \varpi_k^{1-a}\, P\big(X^a\big).$$

Consequently, for $0 \le t < \frac{n}{m}$ *and* $\boldsymbol{\jmath} \neq \mathbf{0}$,

$$L_{t+1}^{\boldsymbol{\jmath}}(P) = \sum_{a \in U} A_{t+1,a}^{\boldsymbol{\jmath}}(X)\, P\big(X^a\big), \qquad A_{t+1,a}^{\boldsymbol{\jmath}} = \frac{1}{\mathcal{B}} \sum_{k=0}^{\mathcal{B}-1} X^{\tilde{a}_{k,t+1,\boldsymbol{\jmath}}}\, \varpi_k^{1-a},$$

and $\mathcal{B} A_{t+1,a}^{\boldsymbol{\jmath}} \in \mathcal{R}$ *is a sum of* $\mathcal{B}$ *monomials. In case (B), with* $a = 1 + pr$, $\mathcal{B} A_{t+1,a}^{\boldsymbol{\jmath}} = \sum_{k<\mathcal{B}} X^{\tilde{a}_{k,t+1,\boldsymbol{\jmath}} - pkr}$.

The automorphisms $P(X^a)$ do not depend on the pattern $\boldsymbol{\jmath}$: they are shared by the $2^m - 1$ maps $L_{t+1}^{\boldsymbol{\jmath}}$, which only differ by their public polynomials.

C.6 Homomorphic evaluation

Algorithm 1 applies in the general case with $\mathrm{acc}_0 = \sum_k X^{-Mb_k}\, \pi_k(v^*)$, the automorphisms τ_a, $a \in U$, and $\mathrm{Extract}_{i_k}$, $X^{i_k} = \varpi_k$, which gives an LWE encryption of $\langle \mathrm{acc}_{n/m}, \varpi_k \rangle = E_k/R$ by Theorem 2. The division by $\mathcal{B}$ is defined on the coordinates: the ciphertext coordinates in the basis $(\Omega_{j,k}^*)$ lie in $\mathbb{T}_q$, for any modulus q, for instance $q = 2^{64}$. For $a = \sum_{(j,k)\in I} a_{j,k}\, \Omega_{j,k}^*$, with $q\, a_{j,k}$ represented by an integer, we set

$$\left\lfloor \frac{a}{\mathcal{B}} \right\rceil = \sum_{(j,k)\in I} \frac{1}{q} \left\lfloor \frac{q\, a_{j,k}}{\mathcal{B}} \right\rceil \Omega_{j,k}^* :$$

for $q = 2^{64}$, a rounded shift by $\log_2 \mathcal{B}$ bits. The representative does not matter (Lemma 4).

C.6.1 Ciphertexts and elementary operations

Ciphertexts. We use the RLWE and RGSW encryptions of Section 3, with the gadget $\mathbf{g} = (g_1, \ldots, g_\ell) = \Omega_{0,0}^*\, (B_g^{-1}, \ldots, B_g^{-\ell})$ (in case (B), $\Omega_{0,0}^* = \frac{1}{N}$), under the ring key $s'(X) \in \mathcal{R}$ with binary coefficients, over $\mathcal{T}^\vee$: a pair $C = (a, b) \in (\mathcal{T}^\vee)^2$ encrypts $\mu \in \mathcal{T}^\vee$ if $b - a\, s' = \mu + \varepsilon$ with a small error ε. Since $\mathcal{R}^\vee = \Omega_{0,0}^* \mathcal{R}$, $\mathrm{Dec}(a) \in \mathcal{R}^\ell$ is given by the balanced digits in base B_g of the coordinates of $a/\Omega_{0,0}^* \in \mathcal{T}$ on $(\Omega_{j,k})$ (Appendix D.2), so that $a \approx \mathrm{Dec}(a) \cdot \mathbf{g}$, and $\mathrm{Dec}(C) = \big(\mathrm{Dec}(a), \mathrm{Dec}(b)\big) \in \mathcal{R}^{2\ell}$. The 2ℓ rows of $\mathrm{RGSW}(P) = Z + P\,(I_2 \otimes \mathbf{g}^\top) \in (\mathcal{T}^\vee)^{2\ell \times 2}$ are RLWE encryptions under s' of $-s' P\, g_i$ and $P\, g_i$, $1 \le i \le \ell$.

FFT domain. Ring elements are stored in the FFT domain, where an automorphism of $\mathcal{K}$ is a permutation of the evaluation points and costs nothing. A decomposition requires the coefficient domain: one inverse FFT and ℓ forward FFTs.

Elementary operations. Let $P \in \mathcal{R}$ be a public polynomial (its FFT is precomputed) and $C = (a, b)$ an RLWE encryption of μ under s': (a) the public product $P \cdot C = (Pa, Pb)$ encrypts $P\mu$; (b) the external product $\mathrm{Dec}(C) \cdot \mathrm{RGSW}(P)$ encrypts $P\mu$; (c) for an automorphism τ of $\mathcal{K}$, $\mathrm{KS}_\tau(C) = \big(0, \tau(b)\big) - \mathrm{Dec}\big(\tau(a)\big) \cdot \mathsf{K}_\tau$ encrypts $\tau(\mu)$ under s', the rows of $\mathsf{K}_\tau \in (\mathcal{T}^\vee)^{\ell \times 2}$ encrypting $\tau(s')g_1, \ldots, \tau(s')g_\ell$; (d) an automorphism $X \mapsto X^u$ with $u \in U$ fixes $\Omega^*_{0,0} \in \mathbb{Q}(X^{M_0})$ and permutes the $\Omega_{j,k}$ up to sign, hence commutes with the decomposition: automorphisms $\tau_1, \ldots, \tau_r$ of this form share $\mathrm{Dec}(a)$ for the same C (hoisting); (e) $\mathrm{Extract}_k(C) = \big((\langle a,\, X^{k-i}\rangle)_{0 \le i < N}, \langle b,\, X^k\rangle\big)$ is an LWE encryption of $\langle \mu,\, X^k\rangle$ under $\mathbf{s}'$ (Section 4.2). Their costs are:

Operation	Encrypts	Decompositions	Products
(a) $P \cdot C$	$P\mu$	0	2
(b) $\mathrm{Dec}(C) \cdot \mathrm{RGSW}(P)$	$P\mu$	2	4ℓ
(c) $\mathrm{KS}_\tau(C)$	$\tau(\mu)$	1	2ℓ
(d) $\mathrm{KS}_{\tau_1}(C), \ldots, \mathrm{KS}_{\tau_r}(C)$	$\tau_j(\mu)$	1	$2\ell r$
(e) $\mathrm{Extract}_k(C)$	$\langle \mu,\, X^k\rangle$ (LWE)	0	0

C.6.2 Automorphisms

For $a \in U$, let τ_a be the automorphism $\tau_a(P)(X) = P(X^a)$ of $\mathcal{K}$. The $\mathcal{B} - 1$ automorphisms $\tau_a(\mathrm{acc}) = \mathrm{acc}(X^a)$, $a \neq 1$, are obtained from one decomposition (item (d) above), with one key-switching key per a. They do not depend on the pattern $\boldsymbol{\jmath}$: for $m \ge 2$ they are computed once for the $2^m - 1$ maps $L^{\boldsymbol{\jmath}}$, and only the public products and the external products are repeated. The polynomials $\mathcal{B}A_a \in \mathcal{R}$ are public: they are computed in the clear from the LWE masks.

Giant steps [34] do not pay for $\mathcal{B} \le 16$ or $m \ge 2$ (Appendix F).

C.7 Noise

Throughout, a ciphertext $C = (a, b) \in (\mathcal{T}^\vee)^2$ with $a, b \in \frac{1}{q}\mathcal{R}^\vee$ has *phase* $b - a\,s' \in \frac{1}{q}\mathcal{R}^\vee/\mathcal{R}^\vee$; it encrypts μ with error ε if $b - a\,s' = \mu + \varepsilon$. The plaintexts acc_t have coordinates in $\frac{1}{p\widehat{p}}\mathbb{Z}$, so that ε contains the rounding of acc_t to $\frac{1}{q}\mathcal{R}^\vee$. With the norm and the mean squares V_{key}, V_{rnd}, V_{aut} and V_{div} of Appendix G, and n_* the number of active groups (Section 5.5):

Proposition 4. *Let ε_t be the error of C after t iterations of Algorithm 1. Then $\|\varepsilon_0\| \le 2N/(qM)$ and*

$$\mathbb{E}\|\varepsilon_{n/m}\|^2 \le \|\varepsilon_0\|^2 + \frac{n}{m}\,(2^m - 1)\,V_{\mathrm{key}} + n_*\Big(V_{\mathrm{rnd}} + (\mathcal{B} - 1)\,\mathcal{B}\,V_{\mathrm{aut}} + \mathcal{B}^2\,V_{\mathrm{div}}\Big).$$

For $m = 1$, the right-hand side is at most $\|\varepsilon_0\|^2 + n\,V_{\mathrm{ext}} + n_1\big((\mathcal{B}-1)\mathcal{B}\,V_{\mathrm{aut}} + \mathcal{B}^2 V_{\mathrm{div}}\big)$. The output LWE_k has error $\langle \varepsilon_{n/m},\, \varpi_k\rangle$, of variance $N\,\mathbb{E}\|\varepsilon_{n/m}\|^2$: this is the output noise e_{out} of the refreshed ciphertexts (Figure 3).

D Duality in the tower

D.1 Relative duality

The duality of [15] is refined along the tower $\mathcal{K} \supset \mathbb{Q}(X^{M_0}) \supset \mathbb{Q}$. Let

$$(\mathrm{A}): \quad \varpi_k = X^{kp}, \qquad\qquad (\mathrm{B}): \quad \varpi_k = X^k, \qquad\qquad 0 \le k < \mathcal{B},$$

and

$$\omega_j = X^{jM_0}, \quad 0 \le j < \varphi(p).$$

Proposition 5. *The field $\mathcal{K}$ admits the tower $\mathcal{K} \supset \mathbb{Q}(X^{M_0}) \supset \mathbb{Q}$ with relative bases:*

(a) $(\varpi_k)_{k<\mathcal{B}}$ *is a* $\mathbb{Q}(X^{M_0})$*-basis of* $\mathcal{K}$*, with relative dual basis* $\varpi_k^* = \frac{\varpi_k}{\mathcal{B}}$*, i.e.*

$$0 \le k, l < \mathcal{B}, \quad \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}\big(\varpi_k^*\,\bar{\varpi}_l\big) = \delta_{kl},$$

and every $P \in \mathcal{K}$ *can be written uniquely as*

$$P = \sum_{k=0}^{\mathcal{B}-1} \varpi_k^*\, P_k, \qquad P_k = \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}\big(\bar{\varpi}_k\, P\big) \in \mathbb{Q}(X^{M_0}).$$

(b) $(\omega_j)_{j<\varphi(p)}$ *is a* $\mathbb{Q}$*-basis of* $\mathbb{Q}(X^{M_0})$*, with dual basis for* $\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}$

$$(A): \quad \omega_j^* = \frac{1}{p}\big(X^{jM_0} - X^{(p-1)M_0}\big), \quad (B): \quad \omega_j^* = \frac{2}{p}\,X^{jM_0},$$

that is to say

$$\forall 0 \le j, j' < \varphi(p), \quad \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}(\omega_j^*\,\bar{\omega}_{j'}) = \delta_{jj'}.$$

Moreover, for $0 \le j, j' < \varphi(p)$*,*

$$\begin{aligned}(A): &\quad \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\omega_j^*\,X^{-(p-1)M_0}\big) = -1,\\ (B): &\quad \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\omega_j^*\,X^{-(j'+\frac{p}{2})M_0}\big) = -\delta_{jj'},\end{aligned}$$

the first identity not depending on j'*.*

The proof is given in Appendix E. The data of the two cases are:

	p	$\mathcal{B}$	$\varphi(p)$	N	ϖ_k	ϖ_k^*	$E_{\max}$
(A)	odd prime	$M_0/2$	$p-1$	$M_0(p-1)/2$	X^{kp}	$X^{kp}/\mathcal{B}$	$\frac{p-1}{2}$
(B)	2^α	M_0	$p/2$	$M_0 p/2$	X^k	$X^k/\mathcal{B}$	$\frac{p}{4}$

The products of the two dual bases are

$$\begin{aligned}(\mathrm{A}): &\quad \Omega_{j,k}^* = \varpi_k^*\,\omega_j^* = \frac{1}{p\mathcal{B}}\,X^{kp}\big(X^{jM_0} - X^{(p-1)M_0}\big),\\ (\mathrm{B}): &\quad \Omega_{j,k}^* = \varpi_k^*\,\omega_j^* = \frac{1}{N}\,X^{jM_0+k}.\end{aligned}$$

D.2 Duality in $\mathcal{K}$

Let

$$I = \{0, \dots, \varphi(p)-1\} \times \{0, \dots, \mathcal{B}-1\}, \qquad |I| = N,$$

and, for $(j,k) \in I$,

$$\Omega_{j,k} = \varpi_k\,\omega_j, \qquad \Omega_{j,k}^* = \varpi_k^*\,\omega_j^*.$$

Proposition 6. $(\Omega_{j,k})_{(j,k)\in I}$ *is a basis of* $\mathcal{K}$ *with dual basis* $(\Omega_{j,k}^*)_{(j,k)\in I}$*:*

$$\forall\,(j,k),(j',k') \in I, \quad \big\langle \Omega_{j,k}^*,\, \Omega_{j',k'} \big\rangle = \delta_{jj'}\,\delta_{kk'},$$

and, for $(j,k) \in I$*,* $0 \le k' < \mathcal{B}$ *and* $j' \in \mathbb{Z}$*,*

$$\Big\langle \Omega_{j,k}^*,\, \varpi_{k'} X^{j'M_0} \Big\rangle = \delta_{kk'}\ \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\omega_j^*\,X^{-j'M_0}\big),$$

given by Proposition 5(b).

The ring $\mathcal{R}$ and its dual $\mathcal{R}^\vee$. The ring of integers of $\mathcal{K}$ is

$$\mathcal{R} = \mathbb{Z}[X]/\Phi_M(X) = \bigoplus_{(j,k)\in I} \mathbb{Z}\,\Omega_{j,k},$$

with, in case (A), $\mathcal{R} = \mathbb{Z}[X^p, X^{M_0}]$ and, in case (B), $(\Omega_{j,k}) = (X^i)_{0\le i<N}$. Its dual with respect to the trace inner product is

$$\mathcal{R}^\vee = \big\{\, Q \in \mathcal{K} : \langle Q,\, P\rangle \in \mathbb{Z} \text{ for all } P \in \mathcal{R} \,\big\} = \bigoplus_{(j,k)\in I} \mathbb{Z}\,\Omega^*_{j,k} = \Omega^*_{0,0}\,\mathcal{R},$$

with $M\,\mathcal{R}^\vee \subset \mathcal{R}$, and, in case (B), $\mathcal{R}^\vee = \frac{1}{N}\mathcal{R}$.

The tori $\mathcal{T}$ and $\mathcal{T}^\vee$. Recalling that $\mathcal{T} = \mathcal{K}/\mathcal{R}$, we define

$$\mathcal{T}^\vee = \mathcal{K}/\mathcal{R}^\vee.$$

Since $\langle \mathcal{R}^\vee,\, \mathcal{R}\rangle \subset \mathbb{Z}$, the inner product induces well-defined $\mathbb{Z}$-bilinear pairings

$$\mathcal{R}^\vee \times \mathcal{T} \longrightarrow \mathbb{Q}/\mathbb{Z} \subset \mathbb{T}, \qquad \mathcal{T}^\vee \times \mathcal{R} \longrightarrow \mathbb{Q}/\mathbb{Z} \subset \mathbb{T},$$

and every $P \in \mathcal{T}$ and every $P \in \mathcal{T}^\vee$ can be written

$$\begin{aligned} (\mathcal{T}):&\quad P = \sum_{(j,k)\in I} P_{j,k}\,\Omega_{j,k}, \quad P_{j,k} = \big\langle P,\, \Omega^*_{j,k}\big\rangle \bmod 1,\\ (\mathcal{T}^\vee):&\quad P = \sum_{(j,k)\in I} P_{j,k}\,\Omega^*_{j,k}, \quad P_{j,k} = \langle P,\, \Omega_{j,k}\rangle \bmod 1; \end{aligned}$$

hence $\mathcal{T} \simeq (\mathbb{Q}/\mathbb{Z})^N \simeq \mathcal{T}^\vee$ as groups.

E Proof of Proposition 5

Lemma 3. *Let $m' = r^s$ with r prime and $s \ge 1$, and ζ a primitive m'-th root of unity. For $i \in \mathbb{Z}$,*

$$\mathrm{Tr}_{\mathbb{Q}(\zeta)/\mathbb{Q}}(\zeta^i) = m'\,\mathbb{1}_{m'|i} - \frac{m'}{r}\,\mathbb{1}_{\frac{m'}{r}|i}.$$

Proof. The sum of ζ^{ic} over $c \in (\mathbb{Z}/m'\mathbb{Z})^\times$ is the sum over $c \in \mathbb{Z}/m'\mathbb{Z}$, which is $m'\mathbb{1}_{m'|i}$, minus the sum over $c = rc'$, $c' \in \mathbb{Z}/\frac{m'}{r}\mathbb{Z}$, which is $\frac{m'}{r}\mathbb{1}_{\frac{m'}{r}|i}$. □

Proof of (a). In case (B),

$$U = \{1 + pr : r \in \mathbb{Z}_{M_0}\}$$

and X^p has order M_0 in $\mathcal{K}$, so that, for $i \in \mathbb{Z}$,

$$\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(X^i) = \sum_{r=0}^{M_0-1} X^{i(1+pr)} = X^i \sum_{r=0}^{M_0-1} \big(X^{ip}\big)^r = M_0\,X^i\,\mathbb{1}_{M_0|i}.$$

For $0 \le k, l < \mathcal{B} = M_0$, $|k-l| < M_0$, hence

$$\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(\varpi_k\bar{\varpi}_l) = \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(X^{k-l}) = \mathcal{B}\,\delta_{kl}.$$

In case (A), $\varpi_k\bar{\varpi}_l = X^{(k-l)p}$ belongs to

$$\mathbb{Q}(X^p) = \mathbb{Q}[Z]/\Phi_{M_0}(Z), \qquad Z = X^p.$$

For $P \in \mathbb{Q}(X^p)$, $P(X^a)$ only depends on $a \bmod M_0$, and $a \mapsto a \bmod M_0$ is a bijection from U onto $(\mathbb{Z}/M_0\mathbb{Z})^\times$, so that $\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(P) = \mathrm{Tr}_{\mathbb{Q}(X^p)/\mathbb{Q}}(P)$. Lemma 3 with $m' = M_0$ gives

$$\mathrm{Tr}_{\mathbb{Q}(X^p)/\mathbb{Q}}\big(X^{ip}\big) = M_0\, \mathbb{1}_{M_0|i} - \frac{M_0}{2}\, \mathbb{1}_{\frac{M_0}{2}|i},$$

and, for $0 \le k, l < \mathcal{B} = \frac{M_0}{2}$,

$$\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(\varpi_k \bar{\varpi}_l) = \Big(M_0 - \frac{M_0}{2}\Big)\delta_{kl} = \mathcal{B}\,\delta_{kl}.$$

In both cases, the matrix

$$\big(\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(\varpi_k \bar{\varpi}_l)\big)_{k,l} = \mathcal{B}\, I_{\mathcal{B}}$$

is invertible, so that the $\mathcal{B}$ elements ϖ_k are linearly independent over $\mathbb{Q}(X^{M_0})$; since $[\mathcal{K} : \mathbb{Q}(X^{M_0})] = \mathcal{B}$, they form a basis, with relative dual basis $\varpi_k/\mathcal{B}$. If

$$P = \sum_k \varpi_k^* P_k,$$

applying the $\mathbb{Q}(X^{M_0})$-linear map $\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(\bar{\varpi}_l\, \cdot)$ gives

$$P_l = \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(\bar{\varpi}_l P).$$

Proof of (b). X^{M_0} is a primitive p-th root of unity, and Lemma 3 with $m' = p$ gives, for $j \in \mathbb{Z}$,

$$\begin{aligned}
&(\mathrm{A}): \quad \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(X^{jM_0}\big) = p\, \mathbb{1}_{p|j} - 1,\\
&(\mathrm{B}): \quad \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(X^{jM_0}\big) = p\, \mathbb{1}_{p|j} - \frac{p}{2}\, \mathbb{1}_{\frac{p}{2}|j} = \frac{p}{2}\, X^{jM_0}\, \mathbb{1}_{\frac{p}{2}|j},
\end{aligned}$$

since, in case (B), $X^{\frac{p}{2}M_0} = -1$.

In case (A), for $0 \le j \le p-2$ and $0 \le j' \le p-1$,

$$\begin{aligned}
\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\omega_j^*\, X^{-j'M_0}\big) &= \tfrac{1}{p}\Big(\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(X^{(j-j')M_0}\big)\\
&\qquad - \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(X^{(p-1-j')M_0}\big)\Big)\\
&= \mathbb{1}_{p|j-j'} - \mathbb{1}_{p|p-1-j'},
\end{aligned}$$

which is

$$(j' \le p-2): \quad \delta_{jj'}, \qquad\qquad (j' = p-1): \quad -1.$$

In case (B), for $0 \le j, j' < \frac{p}{2}$, $|j - j'| < \frac{p}{2}$ and

$$\begin{aligned}
\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\omega_j^*\, X^{-j'M_0}\big) &= \frac{2}{p}\, \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(X^{(j-j')M_0}\big)\\
&= X^{(j-j')M_0}\, \mathbb{1}_{\frac{p}{2}|j-j'} = \delta_{jj'},
\end{aligned}$$

and the last identity follows from

$$X^{-(j'+\frac{p}{2})M_0} = -X^{-j'M_0}.$$

In both cases, the

$$\varphi(p) = [\mathbb{Q}(X^{M_0}) : \mathbb{Q}]$$

elements ω_j are linearly independent by duality, hence form a basis with dual basis (ω_j^*). $\square$

F The test polynomial, the accumulator and its evaluation

With the notation of Section C.5, $v^* = W\,\mathcal{V}^*$ with

$$\text{(A)}: \quad \mathcal{V}^* = \sum_{j=0}^{p-1} \frac{j \bmod p}{p^2\,\widehat{p}}\, X^{jM_0},$$

$$\text{(B)}: \quad \mathcal{V}^* = \sum_{j<p/2} \frac{2\min\bigl(j, \frac{p}{2}-j\bigr)}{p^2\,\widehat{p}}\, X^{jM_0},$$

the expression in case (A) following from $\sum_{j=0}^{p-1}(j \bmod p) = 0$. In case (B),

$$v^* = \frac{1}{NR}\sum_{i<N} \min\Bigl(\Bigl\lfloor\frac{i}{\mathcal{B}}\Bigr\rfloor, \frac{p}{2} - \Bigl\lfloor\frac{i}{\mathcal{B}}\Bigr\rfloor\Bigr) X^i :$$

each coefficient of the test polynomial of Section 4.2 is repeated $\mathcal{B}$ times.

Example 1. $M = 12$, case (A), $p = 3$, $M_0 = 4$: $\mathcal{B} = 2$, $N = 4$, $(v_0^*, v_1^*) = (0, \frac{1}{3\widehat{p}})$ and

$$v^* = \frac{X^3 + 2X^2 - 2X - 1}{18\,\widehat{p}}.$$

$M = 16$, case (B), $p = 8$, $M_0 = 2$: $\mathcal{B} = 2$, $N = 8$, $(v_j^*)_{j<4} = \frac{1}{8\widehat{p}}(0, 1, 2, 1)$ and

$$v^* = \frac{X^2 + X^3 + 2X^4 + 2X^5 + X^6 + X^7}{64\,\widehat{p}}.$$

The initial accumulator of Section C.5 satisfies:

Proposition 7.

(a) $\mathrm{acc}_0 = w_0\,\mathcal{V}^*$, *with* $w_0 = \sum_{k=0}^{\mathcal{B}-1} X^{-Mb_k}\,\varpi_k^* \in \frac{1}{\mathcal{B}}\mathcal{R}$.

(b) For $c \in \mathbb{Z}$, $0 \le k, l < \mathcal{B}$ *and* $0 \le j < \varphi(p)$,

$$\Bigl\langle X^{M_0 c}\,v^{*,k},\, \Omega_{j,l}\Bigr\rangle = \delta_{k,l}\,\vartheta(j-c). \tag{12}$$

In particular $X^{M_0 c}\,v^{*,k}$ *is supported on column* k, *and* $\pi_k(\mathrm{acc}_0) = X^{-Mb_k}\,v^{*,k}$.

(c) For $0 \le k < \mathcal{B}$,

$$\langle \mathrm{acc}_0,\, \varpi_k\rangle = \vartheta(p\,b_k),$$

where $p\,b_k \in \mathbb{Z}_p$.

Public polynomials. At a root ζ of Φ_M, with (A): $\varpi_1 = X^p$ and (B): $\varpi_1 = X$,

$$\mathcal{B}A_{t+1,a}^{\jmath}(\zeta) = S^{\jmath}\bigl(\zeta^{M_0},\, \varpi_1(\zeta)^{1-a}\bigr), \qquad S^{\jmath}(\theta, y) = \sum_{k<\mathcal{B}} \theta^{p\,a_{k,t+1,\jmath}}\,y^k,$$

where θ is a primitive p-th root of unity and $y^{\mathcal{B}} = 1$ (in case (A), $1-a$ is even; in case (B), $1 - a \in p\mathbb{Z}$). For each pattern, the $\varphi(p)$ DFTs of length $\mathcal{B}$ of the sequences $(\theta^{p\,a_{k,t+1,\jmath}})_{k<\mathcal{B}}$ give all the values of line 3.

Deferred projections. Grouping by m (Section 5.4) amounts to deferring the projections: π_k commutes with the monomials $X^{M_0 c}$, so that, over r groups,

$$\mathrm{acc}_{t+r} = \sum_{k=0}^{\mathcal{B}-1} X^{\sum_{w=t+1}^{t+r} \tilde{a}_{k,w,\mathbf{s}_{[w]}}} \, \pi_k(\mathrm{acc}_t),$$

but the exponent depends on the key bits of the r groups, and its expansion on the indicators is the grouping by rm: one set of automorphisms for rm coordinates, and $2^{rm} - 1$ patterns.

Giant steps. Halevi and Shoup [34] also split U into g baby steps and h giant steps, $gh \geq \mathcal{B}$, which brings the number of key switchings down to about $2\sqrt{\mathcal{B}}$; but each giant step needs its own decomposition, repeated for each pattern. For $\mathcal{B} \leq 16$, or as soon as $m \geq 2$, giant steps do not pay (Appendix J), and we do not use them; they could pay for much larger batches.

G Noise of the comb bootstrap

This appendix completes Section 5.5, with its notation.

G.1 Exactness of the division

Lemma 4. *L maps $\mathcal{R}^\vee$ and $\frac{1}{q}\mathcal{R}^\vee$ into themselves. Let $x \in \frac{1}{q}\mathcal{R}^\vee$, $y \in \mathcal{R}^\vee$ and $f \in \mathcal{K}$ with $x' = \frac{x+y}{\mathcal{B}} + f \in \frac{1}{q}\mathcal{R}^\vee$. Then*

$$\sum_{a \in U} \mathcal{B} A_a(X)\, x'\big(X^a\big) \equiv L(x) + \mathcal{B}\, L(f) \pmod{\mathcal{R}^\vee}.$$

Consequently, in Algorithm 1, let $C = (a, b)$ have phase $\mathrm{acc}_t + \varepsilon_t$. The representatives used in line 4 differ from a and b by $y_a, y_b \in \mathcal{R}^\vee$, and the roundings f_a, f_b have coordinates in $[-\frac{1}{2q}, \frac{1}{2q}]$ on the basis $(\Omega^*_{j,k})$. Since $s' \in \mathcal{R}$, C_1 has phase

$$\frac{\mathrm{acc}_t + \varepsilon_t + y}{\mathcal{B}} + f^{\mathrm{div}}, \qquad y = y_b - y_a s' \in \mathcal{R}^\vee, \qquad f^{\mathrm{div}} = f_b - f_a s',$$

and, by Lemma 4, $C_L^{\mathcal{J}}$ has phase $L^{\mathcal{J}}_{t+1}(\mathrm{acc}_t) + L^{\mathcal{J}}_{t+1}(\varepsilon_t) + \mathcal{B}\, L^{\mathcal{J}}_{t+1}(f^{\mathrm{div}}) + E^{\mathrm{aut},\mathcal{J}}$, where $E^{\mathrm{aut},\mathcal{J}}$ collects the key-switching errors of lines 7–9. Division by $\mathcal{B}$ thus works for every q, in particular $q = 2^{64}$, although $\mathcal{B}$ is even.

G.2 Norm and heuristics

For $P \in \mathcal{K}$, let

$$\|P\|^2 = \frac{1}{N} \langle P,\, P\rangle = \frac{1}{N} \operatorname{Tr}_{\mathcal{K}/\mathbb{Q}}\big(P\, \bar{P}\big).$$

For $a \in (\mathbb{Z}/M\mathbb{Z})^\times$ and $i \in \mathbb{Z}$, $\|P(X^a)\| = \|P\|$ (the trace is invariant under $X \mapsto X^a$) and $\|X^i P\| = \|P\|$; hence $\|QP\| \leq \xi \|P\|$ if $Q \in \mathcal{R}$ is a sum of ξ monomials with coefficients ± 1. By Cauchy–Schwarz and $\langle X^i,\, X^i\rangle = N$, $|\langle P,\, X^i\rangle| \leq N \|P\|$. As usual, errors are modelled as random elements of $\mathcal{K}$ whose values at the roots of Φ_M are independent with the same variance. Then, for independent random P and Q,

$$\mathbb{E}\|P + Q\|^2 = \mathbb{E}\|P\|^2 + \mathbb{E}\|Q\|^2, \qquad \mathbb{E}\|PQ\|^2 = \mathbb{E}\|P\|^2\, \mathbb{E}\|Q\|^2,$$

$$\mathbb{E}\,\langle P,\, X^i\rangle^2 = N\, \mathbb{E}\|P\|^2,$$

and $\langle P,\, X^i\rangle$ is Gaussian-like. For $Q \in \mathcal{R}$ a sum of ξ monomials with coefficients ± 1 and pseudo-random exponents, $\mathbb{E}\|QP\|^2 = \xi\, \mathbb{E}\|P\|^2$. A polynomial with N independent coefficients of variance ς^2 in the basis $(X^i)_{i<N}$ has $\mathbb{E}\|\cdot\|^2 = N\varsigma^2$.

G.3 Fresh errors

Let σ^2 be the variance of the coordinates on the basis $(\Omega^*_{j,k})$ of the errors of the RLWE encryptions in K_τ and RGSW(δ) (in case (B), the coefficients of the torus polynomials of TFHE), and $\rho = \mathbb{E}[s_i'^2] = \frac{1}{2}$ for the coefficients of the binary key s'. The digits d_i of a decomposition have coordinates on $(\Omega_{j,k})$ uniform in $[-\frac{B_g}{2}, \frac{B_g}{2}]$, and the rounding $r = a - \mathrm{Dec}(a)\cdot\mathbf{g}$ has coordinates on $(\Omega^*_{0,0}\Omega_{j,k})$ in $[-\frac{{B_g}^{-\ell}}{2}, \frac{{B_g}^{-\ell}}{2}]$. Since $\omega_j\bar{\omega}_j = 1$, the elements $\Omega^*_{j,k}$ and $\Omega^*_{0,0}\Omega_{j,k}$ all have $\|\cdot\|^2 = \frac{2}{Np\mathcal{B}}$ (Proposition 5): N independent centered coordinates of variance ς^2 on either basis give $\mathbb{E}\|\cdot\|^2 = \eta\,\varsigma^2/N$, where

$$\eta = \frac{2N}{p\mathcal{B}}: \qquad (\mathrm{A}): \quad \eta = 2\Big(1 - \frac{1}{p}\Big), \qquad\qquad (\mathrm{B}): \quad \eta = 1.$$

(a) *Public product* by $Q \in \mathcal{R}$ with ξ monomials: the error ε becomes $Q\varepsilon$, $\mathbb{E}\|Q\varepsilon\|^2 = \xi\,\mathbb{E}\|\varepsilon\|^2$.

(b) *External product* $\mathrm{Dec}(C)\cdot\mathrm{RGSW}(\delta)$, $\delta \in \{0,1\}$: the error is $\delta\big(\varepsilon + r_a s' - r_b\big) + \sum_{i=1}^{2\ell} d_i e_i$, of mean square $\delta\big(\mathbb{E}\|\varepsilon\|^2 + V_{\mathrm{rnd}}\big) + V_{\mathrm{key}} \le \delta\,\mathbb{E}\|\varepsilon\|^2 + V_{\mathrm{ext}}$ with

$$V_{\mathrm{key}} = \frac{\eta}{12}\,2\ell({B_g}^2 + 2)\sigma^2, \quad V_{\mathrm{rnd}} = \frac{\eta}{12}\,{B_g}^{-2\ell}\big(\rho + \tfrac{1}{N}\big), \quad V_{\mathrm{ext}} = V_{\mathrm{key}} + V_{\mathrm{rnd}}.$$

(c) *Key switching* $\mathrm{KS}_\tau(C)$: the error is $\tau(\varepsilon) + \tau(rs') - \sum_{i=1}^{\ell}\tau(d_i)e_i$, of mean square $\mathbb{E}\|\varepsilon\|^2 + V_{\mathrm{aut}}$ with

$$V_{\mathrm{aut}} = \frac{\eta}{12}\Big(\ell({B_g}^2 + 2)\sigma^2 + {B_g}^{-2\ell}\rho\Big);$$

hoisting does not change this.

(d) *Extraction* $\mathrm{Extract}_i(C)$: the LWE error is $\langle\varepsilon,\, X^i\rangle$, of variance $N\,\mathbb{E}\|\varepsilon\|^2$.

With the same B_g, ℓ and σ for both key types, $V_{\mathrm{aut}} \approx V_{\mathrm{ext}}/2$.

Division. By Proposition 5, $\Big\langle \Omega^*_{j,k},\, \Omega^*_{j,k}\Big\rangle = \frac{1}{\mathcal{B}}\,\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}(\omega^*_j\bar{\omega}^*_j) = \frac{2}{p\mathcal{B}}$ in both cases, so that the error $f^{\mathrm{div}} = f_b - f_a s'$ of line 4 has

$$\mathbb{E}\|f^{\mathrm{div}}\|^2 \le V_{\mathrm{div}} = \frac{1 + \rho N}{6\,p\mathcal{B}\,q^2}.$$

G.4 Rotation of the error

Lemma 5. *Let $\Lambda = \sum_{k<\mathcal{B}} \lambda_k\,\pi_k$ with $\lambda_k \in \{X^{M_0 c} : c \in \mathbb{Z}\}$ (for instance L, the identity, and their compositions). Then $\|\Lambda(P)\| = \|P\|$ for every $P \in \mathcal{K}$.*

G.5 Comparison with separate blind rotations

The $\mathcal{B}$ LWE ciphertexts can also be rotated one by one, in the same ring and with the same keys. For $m = 1$, a step on an accumulator is the CMux

$$\mathrm{acc} \leftarrow \mathrm{acc} + \mathrm{Dec}\big(X^{\tilde{a}}\mathrm{acc} - \mathrm{acc}\big)\cdot\mathrm{RGSW}(s_t),$$

which maps the error ε to $X^{s_t\tilde{a}}\varepsilon + f^{\mathrm{ext}}$: after n steps, $\mathbb{E}\|\varepsilon_n\|^2 \approx n\,V_{\mathrm{key}} + n_1 V_{\mathrm{rnd}} \le n\,V_{\mathrm{ext}}$. For $m \ge 2$, each accumulator is multiplied by the combined key

$$\Gamma = \sum_{\jmath \neq \mathbf{0}} \big(X^{\tilde{a}_\jmath} - 1\big)\,\mathrm{RGSW}\big(\delta_{\jmath,\mathbf{s}_{[t]}}\big), \qquad \mathrm{acc} \leftarrow \mathrm{acc} + \mathrm{Dec}(\mathrm{acc})\cdot\Gamma$$

(multi-bit form), with $\tilde{a}_{\mathcal{J}}$ the grouped exponent of the accumulator: the fresh errors are multiplied by $X^{\tilde{a}_{\mathcal{J}}}-1$. For $\tilde{a}\in M_0\mathbb{Z}$, $\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}}(X^{\tilde{a}})=\mathcal{B}\,\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}(X^{\tilde{a}})$, hence, for $X^{\tilde{a}}\neq 1$,

$$\text{(A)}:\quad \|X^{\tilde{a}}-1\|^2=2\Big(1+\frac{1}{p-1}\Big),\qquad \text{(B)}:\quad \|X^{\tilde{a}}-1\|^2=2,\ \text{or 4 if } X^{\tilde{a}}=-1,$$

with the same mean $2(1+\frac{1}{p-1})$ for a uniform $\tilde{a}$, as for the conductor p [15], and

$$\mathbb{E}\|\varepsilon_{n/m}\|^2\approx 2\Big(\frac{n}{m}\,(2^m-1)\,V_{\mathrm{key}}+n_*V_{\mathrm{rnd}}\Big),$$

twice the corresponding terms of the comb bootstrap (Proposition 1). For the comb bootstrap with $m=1$, $n_1=\frac{n}{2}$ and, with the same gadget and noise for both key types, $V_{\mathrm{aut}}\approx V_{\mathrm{ext}}/2$,

$$\mathbb{E}\|\varepsilon_n\|^2\approx n\,V_{\mathrm{ext}}\Big(1+\tfrac{1}{4}(\mathcal{B}-1)\mathcal{B}\Big),$$

about 60 times the variance of the CMux for $\mathcal{B}=16$, through the key-switching error of the automorphisms amplified by the public polynomials. Automorphism keys with a smaller base and more digits bring it down; with single-digit bootstrapping keys $(2^{23},1)$, as at 2 bits in Table 3, the gadget $(2^{14},2)$ already brings the comb bootstrap within 0.3 bit of the output noise of the separate rotations.

H Proofs of the comb bootstrap

Proof of Proposition 6. $(\Omega_{j,k})$ is a $\mathbb{Q}$-basis of $\mathcal{K}$ by Proposition 5. Since $\omega_j^*X^{-j'M_0}$ lies in $\mathbb{Q}(X^{M_0})$, the transitivity of the trace gives

$$\begin{aligned}\Big\langle\Omega_{j,k}^*,\,\varpi_{k'}X^{j'M_0}\Big\rangle&=\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\Big(\omega_j^*\,X^{-j'M_0}\,\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}\big(\varpi_k^*\,\bar{\varpi}_{k'}\big)\Big)\\&=\delta_{kk'}\,\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\omega_j^*\,X^{-j'M_0}\big),\end{aligned}$$

which is $\delta_{kk'}\delta_{jj'}$ for $0\le j'<\varphi(p)$.

Proof of Lemma 2. Take $0\le c<p$. By Proposition 5(b), for $c<\varphi(p)$,

$$\vartheta(c)=\sum_j v_j^*\,\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}(\omega_j^*X^{-cM_0})=v_c^*.$$

In case (A), for $c=p-1$, $\vartheta(c)=-\sum_{j\le p-2}v_j^*=-\frac{1}{p\widehat{p}}=\frac{(p-1)\bmod p}{p\widehat{p}}$. In case (B), for $c\ge\frac{p}{2}$, $\vartheta(c)=-v_{c-p/2}^*$; then $\vartheta(c)=\frac{c}{p\widehat{p}}$ for $c\le\frac{p}{4}$, and for $c\ge\frac{3p}{4}$, $\vartheta(c)=-\frac{p-c}{p\widehat{p}}=\frac{c-p}{p\widehat{p}}$.

Proof of Proposition 7. (a) is immediate. (b) By Proposition 6 with $j'=j-c$,

$$\begin{aligned}\Big\langle X^{M_0c}v^{*,k},\,\varpi_lX^{jM_0}\Big\rangle&=\sum_{j''}v_{j''}^*\,\Big\langle\Omega_{j'',k}^*,\,\varpi_lX^{(j-c)M_0}\Big\rangle\\&=\delta_{k,l}\,\mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(\mathcal{V}^*X^{-(j-c)M_0}\big)=\delta_{k,l}\,\vartheta(j-c).\end{aligned}$$

The coordinates of $X^{M_0c}v^{*,k}$ on the columns $l\neq k$ vanish, and $-Mb_k\in M_0\mathbb{Z}$. (c) By (b) with $c=-pb_k$, $j=0$ and $\varpi_k=\Omega_{0,k}$.

Proof of Theorem 2. For $t=0$, $z_{k,0}=b_k$ and the statements are those of Proposition 7. Assume them for some $t<\frac{n}{m}$. By Proposition 7(b), the term $X^{-Mz_{k,t}}v^{*,k}$ is supported on column k, so $\pi_k(\mathrm{acc}_t)=X^{-Mz_{k,t}}v^{*,k}$. Since $z_{k,t+1}=z_{k,t}-a_{k,t+1,\mathbf{s}_{[t+1]}}$,

$$\mathrm{acc}_{t+1}=\sum_{k=0}^{\mathcal{B}-1}X^{Ma_{k,t+1,\mathbf{s}_{[t+1]}}}\,X^{-Mz_{k,t}}\,v^{*,k}=\sum_{k=0}^{\mathcal{B}-1}X^{-Mz_{k,t+1}}\,v^{*,k}.$$

The coordinates follow from (12) with $c=-pz_{k,t+1}$ and $j=0$. Finally, $b_k=\sum_t a_{k,t,\mathbf{s}_{[t]}}+E_k/p$ gives $z_{k,n/m}=E_k/p$, and Lemma 2 gives $\vartheta(E_k)=E_k/R$.

Proof of Proposition 3. Expand $\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(\bar{\varpi}_k P) = \sum_{a\in U}(\bar{\varpi}_k P)(X^a)$ in the definition of π_k, insert it in $L^{\mathcal{J}}_{t+1}(P) = \sum_k X^{\tilde{a}_{k,t+1,\mathcal{J}}}\pi_k(P)$ and exchange the sums over k and a.

Proof of Lemma 4. $\Omega^*_{j,l}$ is supported on column l and $X^{\tilde{a}_k} \in \mathcal{R}$; since $\mathcal{R}^\vee$ is an $\mathcal{R}$-module, $L = \sum_k X^{\tilde{a}_k}\pi_k$ maps $\mathcal{R}^\vee$, hence $\frac{1}{q}\mathcal{R}^\vee$, into itself. By Proposition 3,

$$\sum_{a\in U} \mathcal{B}A_a(X)\, x'(X^a) = \mathcal{B}\, L(x') = L(x) + L(y) + \mathcal{B}\, L(f),$$

and $L(y) \in \mathcal{R}^\vee$.

Proof of Lemma 5. Write $P = \sum_k \varpi^*_k P_k$ with $P_k \in \mathbb{Q}(X^{M_0})$, so that $\Lambda(P) = \sum_k \varpi^*_k \lambda_k P_k$. Since $\varpi^*_l = \varpi_l/\mathcal{B}$, $\mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}(\varpi^*_k\bar{\varpi}^*_l) = \frac{\delta_{kl}}{\mathcal{B}}$, and by transitivity

$$\begin{aligned}\langle P,\, P\rangle &= \sum_{k,l} \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\Big(P_k\bar{P}_l\ \mathrm{Tr}_{\mathcal{K}/\mathbb{Q}(X^{M_0})}\big(\varpi^*_k\bar{\varpi}^*_l\big)\Big)\\ &= \tfrac{1}{\mathcal{B}}\sum_k \mathrm{Tr}_{\mathbb{Q}(X^{M_0})/\mathbb{Q}}\big(P_k\bar{P}_k\big),\end{aligned}$$

which is unchanged when P_k is replaced by $\lambda_k P_k$, since $\lambda_k\bar{\lambda}_k = 1$.

Proof of Propositions 4 and 1. Line 1: ε_0 is the rounding of acc_0 to $\frac{1}{q}\mathcal{R}^\vee$, with coordinates in $[-\frac{1}{2q}, \frac{1}{2q}]$ on the N elements $\Omega^*_{j,k}$, each of norm at most $4/M$. Fix t. By hoisting, C_a ($a \in U$, $a \neq 1$) has phase $\tau_a(\psi_1) + f_a$, where ψ_1 is the phase of C_1 and f_a a fresh key-switching error, and line 9 gives $C^{\mathcal{J}}_L$ the phase

$$\sum_{a\in U} \mathcal{B}A^{\mathcal{J}}_{t+1,a}\,\tau_a(\psi_1) + \sum_{a\in U,\ a\neq 1} \mathcal{B}A^{\mathcal{J}}_{t+1,a}\, f_a.$$

By Lemma 4 the first term is $L^{\mathcal{J}}_{t+1}(\mathrm{acc}_t) + L^{\mathcal{J}}_{t+1}(\varepsilon_t) + \mathcal{B}\,L^{\mathcal{J}}_{t+1}(f^{\mathrm{div}})$ modulo $\mathcal{R}^\vee$, and $\mathbb{E}\|\mathcal{B}\,L^{\mathcal{J}}_{t+1}(f^{\mathrm{div}})\|^2 \le \mathcal{B}^2 V_{\mathrm{div}}$ by Lemma 5. The second term has $\mathcal{B}-1$ terms, each of mean square $\mathcal{B}\,V_{\mathrm{aut}}$. Hence $C^{\mathcal{J}}_L - C$ has error $L^{\mathcal{J}}_{t+1}(\varepsilon_t) - \varepsilon_t + F^{\mathrm{aut},\mathcal{J}}$ with $\mathbb{E}\|F^{\mathrm{aut},\mathcal{J}}\|^2 \le (\mathcal{B}-1)\mathcal{B}\,V_{\mathrm{aut}} + \mathcal{B}^2 V_{\mathrm{div}}$. In line 10, only the pattern $\mathcal{J} = \mathbf{s}_{[t+1]}$ has $\delta_{\mathcal{J},\mathbf{s}_{[t+1]}} = 1$, and item (b) of Appendix G.3 gives

$$\varepsilon_{t+1} = \Lambda_t(\varepsilon_t) + \mathbb{1}_{\mathbf{s}_{[t+1]}\neq\mathbf{0}}\left(F^{\mathrm{aut},\mathbf{s}_{[t+1]}} + f^{\mathrm{rnd}}\right) + \sum_{\mathcal{J}\neq\mathbf{0}} f^{\mathrm{key},\mathcal{J}},$$

$$\Lambda_t = \begin{cases} L^{\mathbf{s}_{[t+1]}}_{t+1} & \text{if } \mathbf{s}_{[t+1]} \neq \mathbf{0},\\ \mathrm{id} & \text{if } \mathbf{s}_{[t+1]} = \mathbf{0},\end{cases}$$

with fresh errors f^{rnd} of mean square V_{rnd} and $f^{\mathrm{key},\mathcal{J}}$ of mean square V_{key}. By induction,

$$\varepsilon_{n/m} = \Lambda_{n/m-1}\cdots\Lambda_0(\varepsilon_0) + \sum_{t<n/m} \Lambda_{n/m-1}\cdots\Lambda_{t+1}(F_t)$$

with F_t the fresh error of iteration t; every product of Λ's preserves the norm (Lemma 5), and the terms are independent. Extraction is item (d) of Appendix G.3, and the phase of LWE_k is $E_k/R + \langle\varepsilon_{n/m},\, \varpi_k\rangle$, of variance $N\,\mathbb{E}\|\varepsilon_{n/m}\|^2$. In the ring of TFHE, $\|\varepsilon_0\| \le 2N/(qM) = 1/q$, which gives Proposition 1.

I Further comparisons with TFHE

This appendix completes Section 4.4. Besides the ring of TFHE, with $p = 2N$ (case (B)), the error bootstrap can use the ring of conductor p, of degree $\varphi(p) = p - 1 \le N$, with p the largest such prime (case (A), Appendix C.1).

Setting. Between two bootstraps, $\nu = 1000$ fresh bootstrap outputs are added, key-switched to $\mathbf{s}$ and bootstrapped, with failure probability at most 2^{-128}, the level commonly adopted against the IND-CPA^{D} attacks of Cheon et al. [21] and Checri et al. [18], based on decryption failures. This probability is that of the noise model, on average over the keys; Table 2 assumes exact transforms (NTT), whereas the configurations of Section 7 add the error of the FFT as an independent Gaussian term; it is not a proof of IND-CPA^{D} security. The keys are binary or ternary, the mask is grouped by $m \le 4$, recentered and selected, every key has the smallest noise σ_{128} allowed at 128 bits for $q = 2^{64}$ by the lattice estimator [1] (below), and the key switch is the polynomial key switch of Section 4.3 ($d = 16$). The ring of conductor p is assumed as secure as a power-of-two ring of the same degree. In case (A), the error bootstrap has the full window $\frac{1}{2\widehat{p}}$; in case (B), the window $\frac{1}{4\widehat{p}}$ of TFHE.

Case (A). With $u = u_{128}$ and $\lambda = \frac{1}{2\widehat{p}}$, (8) gives

$$(2\widehat{p}\,\sigma_e)^2 + \frac{4V^*}{p^2} \le \frac{1}{u^2}, \qquad \widehat{p}^{\,\mathrm{err}}_{\max} = \frac{\sqrt{1 - 4u^2V^*/p^2}}{2u\sigma_e} \approx \frac{1}{2u\sigma_e},$$

twice the bound of case (B) given by (9), since there is no padding bit; the counting argument gives $p \gtrsim 2u\sqrt{V}$.

Largest message modulus. Table 2 maximizes $\widehat{p}$ for TFHE over all the parameters, with $\ell_{\mathrm{BR}} \le 4$ or $\ell_{\mathrm{BR}} \le 3$, and evaluates the error bootstrap with the same parameters or with its own; TFHE and case (B) need N a power of two, case (A) a prime p. With the same parameters, at $N = 2048$ and 4096, the error bootstrap reaches a $\widehat{p}$ 9.9 to 11.1 times larger in case (A) and 5.0 to 5.6 times larger in case (B), limited by the key-switching noise of these parameters ($n = 752$ and 784, $\sigma_e \approx 2^{-12.6}$ and $2^{-13.4}$). With its own, which increase n, $\widehat{p}_{\max}$ grows exponentially with N, until the ring noise reaches the floor (210 to 280 times larger at $N = 4096$ than at $N = 2048$), case (B) staying within a factor 1.5 of case (A). Beyond the limit of TFHE, only the error bootstrap remains.

Noise model. Every key has the smallest noise allowed at 128 bits for $q = 2^{64}$, given by the lattice estimator [1]:

$$\begin{aligned} &\text{(binary keys)}: && \log_2 \sigma_{128}(N) = \max\big(-0.02591\,N + 2.77,\ -62\big),\\ &\text{(ternary keys)}: && \log_2 \sigma_{128}(N) = \max\big(-0.02600\,N + 1.64,\ -62\big). \end{aligned}$$

With a gadget $(B_{\mathrm{BR}}, \ell_{\mathrm{BR}})$ for the rotation and the density $\rho = \mathbb{E}[s_i^2]$ of the key ($\frac{1}{2}$ binary, $\frac{2}{3}$ ternary), the output of the TFHE bootstrap has the variance

$$\sigma^2_{\mathrm{out}} = \frac{n}{m}\Big((K^m - 1)\,2\ell_{\mathrm{BR}}N\,\frac{B_{\mathrm{BR}}^2 + 2}{12}\,\sigma_{128}(N)^2 + (1 + \rho N)\,\frac{B_{\mathrm{BR}}^{-2\ell_{\mathrm{BR}}}}{12}\Big),$$

doubled for the error bootstrap in case (A) [15] and unchanged in case (B). This is the variance of a rotation by one CMux per group and non-zero pattern, the rounding term being counted for every group (an upper bound); the combined key of Appendix G, with a single decomposition per group, at most doubles it. The key switch is the polynomial key switch of Appendix B ($d = 16$), with the chain of least noise, and the transfer of (5) to $\mathbf{s}'$ is free for $m = 1$ or with the coordinate-wise rounding, and adds a negligible noise otherwise (Section 4.2).

J Further measurements

This appendix completes Section 7, with the same setting; the logs give the remaining details.

Table 2: Largest message modulus, 1000 additions, failure 2^{-128}, polynomial key switch: TFHE (key, m, n), and the error bootstrap with the same parameters and with its own, in case (A) (conductor p) and in case (B) ($p = 2N$).

		TFHE		error bootstrap (A)			error bootstrap (B)		
N	p		(key, m, n)	same	own	(key, m, n)	same	own	(key, m, n)
$\ell_{\mathrm{BR}} \leq 4$									
2048	2039	21.2	bin, 4, 752	236	$1.1 \cdot 10^6$	ter, 1, 1296	118	$8.9 \cdot 10^5$	ter, 1, 1328
4096	4093	41.3	bin, 4, 784	410	$3.1 \cdot 10^8$	bin, 1, 1680	205	$2.2 \cdot 10^8$	bin, 1, 1696
$\ell_{\mathrm{BR}} \leq 3$									
2048	2039	21.2	bin, 4, 752	236	$1.9 \cdot 10^5$	ter, 1, 1200	118	$1.5 \cdot 10^5$	ter, 1, 1232
4096	4093	41.3	bin, 4, 784	410	$4.5 \cdot 10^7$	bin, 2, 1568	205	$3.2 \cdot 10^7$	bin, 2, 1584

Implementations. On TFHE-rs, the rotation of the error bootstrap is the CMux of TFHE-rs ($m = 1$) or the multi-bit form of Appendix G ($m = 2$), 1.12 to 1.21 times faster, whichever is faster for the configuration; CombBoot keeps $m = 1$, its public products making $m = 2$ slower. The Julia implementation runs on FFTW, with an exact backend (products modulo 2^{64}) for the noise; its times are within a factor 1.5 of those on TFHE-rs (within 20% for CombBoot). Each configuration is the cheapest of the noise model, the error of the FFT being an independent Gaussian term of variance $cB_{\mathrm{BR}}^2 \ell_{\mathrm{BR}}$ per external product or per pair of coordinates, with c calibrated on each FFT (2^{-87} and $2^{-84.5}$ on TFHE-rs, 2^{-85} with FFTW); this heuristic is checked on the standard deviation and in the tails down to 10^{-6} (below), not at 2^{-128}. The selection takes about 100 trials. For the refresh, TFHE-rs runs our polynomial key switch in place of its own (0.05 instead of 1.27 ms per block), which lowers its time by 7 to 8%; it uses $n = 918$ and $(2^{23}, 1)$.

Parameters of CombBoot. CombBoot takes $p = 256$ and $\mathcal{B} = 16$ up to 8 bits, and $p = 512$ and $\mathcal{B} = 8$ beyond, where the rounding term leaves no room at $p = 256$; at $p = 128$, it alone exceeds the budget unless $n \leq 323$. From $\widehat{p} = 16$, the automorphism keys need more digits: each iteration adds the rounding terms of $(\mathcal{B} - 1)\mathcal{B}$ key switchings per active coordinate. The mask of C_1, divided by $\mathcal{B}$, has its top $\log_2 \mathcal{B}$ bits equal to zero, and the shifted gadget

$$g_j = \frac{B_{\mathrm{aut}}^{-j}}{\mathcal{B}}, \qquad 1 \leq j \leq \ell_{\mathrm{aut}},$$

gains them: the rounding term is divided by $\mathcal{B}^2$, and at 8 bits the keys go from five digits to three (17% faster). Its digits represent the mask modulo $\frac{1}{\mathcal{B}}$, which is replaced by the representative that they reconstruct (Lemma 4). Table 3 gives the configurations, run with $m = 2$ in Julia and $m = 1$ on TFHE-rs; the error bootstrap alone uses those of Figure 4. The keys of CombBoot take 0.06 to 0.54 GiB, against 0.15 GiB for the error bootstrap at 8 bits. CombBoot is 2.1 to 3.4 times faster than the error bootstrap in Julia, and 1.8 to 3.0 times on TFHE-rs (1.7 to 3.0 in Figure 4). In module LWE of rank 4 over $\mathbb{Z}[X]/(X^{512} + 1)$, the fastest of the ranks 2, 4 and 8 (Section 5.5), a refresh takes 5.9 to 31.4 ms from 2 to 16 bits, against 2.7 to 21.7 ms per ciphertext for CombBoot in the same session.

Noise. On 500 refreshes per configuration of Figure 4, the noise of the error bootstrap agrees with the prediction within 0.22 bit, and on 1600 outputs per size that of CombBoot within 0.31 bit; the margins are 0.19 to 0.36 bit. In Julia, the noises of Table 3 exceed the predictions of Proposition 1 by 0.01 bit at most. End to end, on 2000 ciphertexts per size at the noise of $\nu = 1000$ fresh outputs, the errors E_k given to the rotation have the predicted standard deviation within 0.02 bit, and reach at most 16 and 35 against the windows $\frac{p}{4} = 64$ and 128. The $\mathcal{B}$ outputs of one

Table 3: CombBoot at $N = 2048$: size of all the keys in GiB ($m = 2$, Fourier domain), gadgets of the rotation (BR) and of the automorphisms (aut.; *: shifted), time per ciphertext in ms (key switch and selection included), and margin in bits at failure 2^{-128} with the measured output noise.

					Julia ($m = 2$)		TFHE-rs		margin	
bits	p ($\mathcal{B}$)	n	gadgets (BR / aut.)	keys	CombBoot	alone	CombBoot	alone	Julia	TFHE-rs
2	256 (16)	608	$(2^{23}, 1)$ / $(2^{14}, 2)$	0.06	2.58	6.14	2.67	7.59	0.29	0.27
4	256 (16)	688	$(2^{14}, 2)$ / $(2^{14}, 2)^*$	0.13	3.47	11.85	3.47	10.33	0.21	0.22
8	256 (16)	896	$(2^{11}, 3)$ / $(2^{11}, 3)^*$	0.25	5.72	14.92	6.14	16.13	0.19	0.19
12	512 (8)	986	$(2^{11}, 3)$ / $(2^{11}, 3)^*$	0.28	10.42	24.93	10.48	25.70	0.26	0.19
16	512 (8)	1158	$(2^8, 5)$ / $(2^6, 6)^*$	0.54	18.19	37.53	21.62	37.94	0.33	0.29

Table 4: Tail of the error with $\nu = 1$, 10^6 refreshes on TFHE-rs: standard deviation against the model, kurtosis, and number of values beyond t measured standard deviations, out of 10^6.

			values beyond $t\sigma$				
	$\sigma/\sigma_{\text{model}}$	kurtosis	$t = 3$	3.5	4	4.5	5
Gaussian law	1	3	2700	465	63.3	6.8	0.6
E_R, $\widehat{p} = 4$	1.030	2.994	2623	428	70	3	0
E_R, $\widehat{p} = 2^8$	1.005	3.000	2674	446	59	7	1
e_{out}, $\widehat{p} = 4$	0.967	2.995	2638	446	62	5	0
e_{out}, $\widehat{p} = 2^8$	1.023	3.006	2770	467	61	7	1

CombBoot add up as independent outputs: over 300 calls per configuration, the correlations of their errors have mean at most 0.01 in absolute value, and the variance of the sum of two of them is the sum of their variances within 1% on average.

Tail of the error. The failure probability rests on a Gaussian law of χ (Section 4.3). We test it with $\nu = 1$, where no sum of outputs makes it Gaussian: on TFHE-rs, 10^6 fresh encryptions per configuration of Figure 4, at $\widehat{p} = 4$ and 2^8, are refreshed, key-switched and rounded as for the next bootstrap (Table 4). The shapes are Gaussian within the Poisson fluctuations, and the standard deviations agree with the model within 3.5%, inside the margins of the configurations: this checks the law down to probabilities of 10^{-6}, not at 13 standard deviations.

Product. The trace and the tensor product are computed with exact products modulo q' and 2^{64}. On 200 products per configuration of Figure 4, on TFHE-rs, the measured noises of the trace and of the product agree with Section 6 within 0.2 bit (within 0.3 bit in Julia), and the refresh that follows the product keeps margins of 0.09 to 0.22 bit under the exact law. The bootstrap before a product needs more digits than a refresh ($(2^8, 5)$ instead of $(2^{14}, 2)$ at $\widehat{p} = 2^8$); a Gaussian estimate of the tails would accept $(2^9, 4)$, which fails with probability 2^{-78}. The product itself takes 1.15 to 1.52 ms, against 10.8 to 31.6 ms for the key switch and the bootstrap that precede it for each operand. Beyond $\widehat{p} = 2^8$, the error of the FFT is too large: exact transforms (NTT) remain to be implemented.

Other measurements. At 2 bits, where both use $(2^{23}, 1)$, a PBS of TFHE-rs alone takes 13.9 ms ($n = 918$), against 7.4 ms for our rotation ($n = 570$, $m = 2$). An addition of TFHE-rs propagates the carries with two PBS per block (119 ms at 8 bits), against 0.3 to 0.4 μs for an LWE addition. On 8 threads, in one session, the latencies are within 11%: from 2 to 16 bits, 5.7 to 18.3 ms for the error bootstrap ($m = 1$), against 6.0 to 16.5 ms for TFHE-rs with its multi-bit parameters; the rotation of CombBoot, shared by the threads, takes 19.5 to 51.4 ms per call

(35.7 ms at 8 bits), 2.4 to 3.4 times less than on one thread. In Julia ($p = 256$, $\mathcal{B} = 16$), an iteration costs about 43 μs shared by the patterns and 28 μs per pattern: $m = 2$ is optimal, as predicted in Section 5.4, and giant steps only pay for $\mathcal{B} = 32$ and $m = 1$. In rings of odd prime conductor (Julia, $\widehat{p} = 2^8$), the error bootstrap and CombBoot correct the error exactly over the whole window $\frac{p-1}{2}$, allow a $\widehat{p}$ twice as large as case (B) with the same keys, and are 2.2 to 2.4 and 1.35 to 1.42 times slower. For the integers of Section 7.3, the messages $r/\widehat{p}_k$ lie on $2^{-64}\mathbb{Z}$ up to an error below 2^{-64}, which the noise absorbs; the noises of CombBoot are within 0.06 bit of the predictions, the refresh after a product of operands from two distinct calls keeps a margin of 0.20 bit, and the decrypted residues give $x + y$, $(x + y)\, y$ and their refreshes modulo $\widehat{P}$. The CRT integers of TFHE-rs use the blocks $\{16, 15, 13, 11, 7\}$ (`MESSAGE_4_CARRY_1`, $N = 4096$) at 16 bits and $\{19, 23, 25, 27, 29, 31, 32\}$ (`MESSAGE_5_CARRY_1`, $N = 8192$) at 32 bits.

Amortized bootstraps. BatchBoot [44] and the method of Guimarães and Pereira [32] are written with the C library MOSFHET [31], whose default build requires AVX-512 and VAES; we run their code, unchanged, with its portable build (FFT in C), one thread, one parameter set per process. Both bootstrap a packed RLWE ciphertext, BatchBoot (in its code) to an RLWE ciphertext under another key, Guimarães and Pereira back to the input key. BatchBoot uses the parameter sets of [32] with 2048 and 4096 messages. BatchBoot is faster here than on its server with AVX-512 (3.6, 3.6, 8.4 and 25.1 ms); for Guimarães and Pereira, the 8.4 ms at 2 bits for 2048 messages match the 8.1 ms that BatchBoot reports for them, against 1.5 ms in their paper. Their failure probabilities, from the estimator of [32], are per coefficient, for a bootstrap that directly follows the previous one: 2^{-120}, 2^{-94}, 2^{-64} and 2^{-62} for the sets of BatchBoot, and 2^{-167} to 2^{-62} for all the sets of Guimarães and Pereira (2^{-62} and 2^{-72} at 8 bits). With this estimator, four additions between two bootstraps raise those of these sets to 2^{-27} to 2^{-51}, and ten to 2^{-12} to 2^{-36}; ours are 2^{-128} per ciphertext after 1000 additions.