

Faster Amortized FHEW Bootstrapping Using Ring Automorphisms

Gabrielle De Micheli¹[0000-0002-2617-6878], Duhyeong Kim²[0000-0002-4766-3456], Daniele Micciancio¹[0000-0003-3323-9985], and Adam Suhl¹

¹ UC San Diego, USA

² Intel Labs, Hillsboro, OR, USA

Abstract. Amortized bootstrapping offers a way to simultaneously refresh many ciphertexts of a fully homomorphic encryption scheme, at a total cost comparable to that of refreshing a single ciphertext. An amortization method for FHEW-style cryptosystems was first proposed by (Micciancio and Sorrell, ICALP 2018), who showed that the amortized cost of bootstrapping n FHEW-style ciphertexts can be reduced from $\tilde{O}(n)$ basic cryptographic operations to just $\tilde{O}(n^\epsilon)$, for any constant $\epsilon > 0$. However, despite the promising asymptotic saving, the algorithm was rather impractical due to a large constant (exponential in $1/\epsilon$) hidden in the asymptotic notation. In this work, we propose an alternative amortized bootstrapping method with much smaller overhead, still achieving $O(n^\epsilon)$ asymptotic amortized cost, but with a hidden constant that is only linear in $1/\epsilon$, and with reduced noise growth. This is achieved following the general strategy of (Micciancio and Sorrell), but replacing their use of the Nussbaumer transform, with a much more practical Number Theoretic Transform, with multiplication by twiddle factors implemented using ring automorphisms. A key technical ingredient to do this is a new “scheme switching” technique proposed in this paper which may be of independent interest.

Keywords: Fully homomorphic encryption, Ring LWE, amortized bootstrapping.

1 Introduction

Fully Homomorphic Encryption (FHE) schemes support the evaluation of arbitrary programs on encrypted data. Since a first solution to the problem was proposed in [8], FHE has become both a central tool in the theory of cryptography, and an attractive cryptographic primitive to be used to secure privacy sensitive applications. Still, improving the efficiency of these schemes is a major obstacle to the use of FHE in practice, and a very active area of research.

All reasonably efficient currently known constructions of FHE are based on the “Ring Learning With Errors” (RingLWE) problem [19, 22]. There are two main approaches to design FHE schemes based on Ring LWE: the one pioneered by the BGV cryptosystem and its variants (e.g., see [4, 9, 10]) and the one put forward by the FHEW cryptosystem and follow up work (e.g., see [3, 5, 7].) In BGV, ring operations are directly used to implement (componentwise) addition and multiplication of ciphertexts encrypting *vectors* of values. The ability to simultaneously work on all the components of a vector (at the cost of a single cryptographic operation) makes these schemes very powerful. The downside is that they also require fairly large parameters, leading to stronger security assumptions (namely, the hardness of approximating lattice problems within superpolynomial factors), a very slow bootstrapping procedure, and complex programming model. By contrast, in FHEW, ciphertexts are simple LWE encryptions (which offer native support only for homomorphic addition,) while Ring LWE is used only internally, to implement a special “functional bootstrapping” procedure that, given an encryption of m , produces a (bootstrapped) encryption of $f(m)$, for a given function f . The combination of linearly homomorphic LWE addition and functional bootstrapping still allows to perform arbitrary computations: for example, as originally done in [7], one can represent bits $x, y \in \{0, 1\}$ as integers modulo 4, and then implement a (universal) NAND boolean gate as an addition followed by a (functional bootstrapping) rounding operation $\lfloor (x+y+2 \bmod 4)/2 \rfloor$. The FHEW approach has several attractive features: (1) since bootstrapping is performed after every operation, gates can be arbitrarily

composed, leading to a very simple and easy to use programming model; (2) since we only need to bootstrap basic LWE ciphertexts supporting a single homomorphic addition, the scheme can be instantiated with much smaller parameters; (3) in turn, this leads to weaker security assumptions (hardness of approximating lattice problems within polynomial factors), and substantially simpler and faster bootstrapping, orders of magnitude faster than BGV.

However, the lower speeds of BGV bootstrapping are largely compensated by its ability to encrypt and operate on many values (encrypted as a vector) at the same time, allowing, for example, to simultaneously bootstrap thousands of ciphertexts. This drastically reduces the *amortized* cost of BGV bootstrapping, making it preferable to FHEW in terms of overall performance in many settings.

In an effort to bridge the gap between the two approaches, a method to amortize FHEW bootstrapping was proposed in [21]. The suggested method consists in combining several (say n) FHEW/LWE input ciphertexts into a single RingLWE ciphertext, and then perform FHEW-style bootstrapping on a single RingLWE ciphertext. This results in a major asymptotic performance improvement, reducing the amortized cost of FHEW bootstrapping from $O(n)$ homomorphic multiplications to just $O(n^\epsilon)$, for any fixed constant $\epsilon > 0$. Unfortunately, the method of [21] is rather far from being practical, due in large part to a large constant $2^{O(1/\epsilon)}$ hidden in the asymptotic notation.

Challenges, Results and Techniques In this paper we propose a variant of the bootstrapping procedure of [21] with similar asymptotics, but substantially smaller multiplicative overhead. In particular, we reduce the amortized cost of FHEW bootstrapping from $2^{O(1/\epsilon)} \cdot n^\epsilon$ to just $(1/\epsilon) \cdot n^\epsilon$. In other words, we achieve a similar asymptotic cost $O(n^\epsilon)$ (for any constant $\epsilon > 0$), but with an exponentially smaller constant hidden in the asymptotic notation.

The main challenge faced by [21] was the use of RingGSW registers to implement the homomorphic Fourier transform required to bootstrap a RingLWE ciphertext. These registers, introduced in [7], encrypt messages in the exponent as X^m . This allows to implement homomorphic addition using some form of ciphertext multiplication $X^{m_0} \cdot X^{m_1} = X^{m_0+m_1}$, but other homomorphic operations required by FFT/NTT algorithms (like subtraction and constant multiplication by so-called “twiddle factors”) are much harder, seemingly requiring homomorphic division and exponentiation on ciphertexts. This is addressed in [21] by using the Nussbaumer transform, a variant of the FFT/NTT algorithm that does not require multiplication by twiddle factors. Unfortunately, the use of the Nussbaumer transform in [21] also introduces a $2^{O(1/\epsilon)}$ factor in the running time, making the algorithm impractical.

Methods to perform homomorphic multiplication in the exponent (i.e., exponentiation by a constant) are known, using automorphisms, and have been used in connection to the bootstrapping of FHEW-like cryptosystems [3, 15], but they only work for RingLWE ciphertexts, making them inapplicable to the RingGSW ciphertexts required by [21]. In this paper we introduce three technical innovations that allow to overcome these issues:

We introduce a new RingLWE-to-RingGSW “scheme switching” procedure, which allows us to transform RingLWE ciphertexts into equivalent RingGSW ones. The method is of independent interest and may find applications elsewhere. Note that a similar technique also appears in [14].

We design a new variant of the amortized FHEW bootstrapping of [21] that operates on RingLWE registers, rather than RingGSW. This allows us to implement multiplication by arbitrary twiddle factors using the automorphism techniques of [3, 15], and instantiate the amortized FHEW bootstrapping framework with a standard (homomorphic) FFT/NTT computation, which carries a much smaller overhead. Then, when RingGSW registers are required, we resort to our scheme switching procedure to convert RingLWE to RingGSW on the fly.

We replace the power-of-two cyclotomic rings [7, 15, 21] and circulant rings [3] used by previous FHEW bootstrapping algorithms, with prime cyclotomics. This requires a new error analysis for prime cyclotomics, which we describe in this paper. (Error analysis for power-of-two cyclotomic and circulant rings is comparatively much easier.) This speeds up and simplifies various steps of our bootstrapping procedure, e.g., by supporting a standard radix 2 FFT (as opposed to the radix 3 Nussbaumer transform of [21]), and completely bypassing the problem that automorphisms only exist for invertible exponents [15].

One important problem that still remains open is that of reducing the noise growth in amortized FHEW bootstrapping. Just as in previous work [21], the ciphertext noise of our bootstrapping procedure increases multiplicatively at every level of the FFT/NTT computation. In order to keep the RingLWE noise (and underlying lattice inapproximability factors) polynomial, this requires to limit the recursive depth of the FFT/NTT algorithms to a constant. This is the reason why both [21] and our work only achieve $O(n^\epsilon)$ amortized complexity, rather than the $O(\log n)$ one would expect from a full ($O(\log n)$ -depth) FFT algorithm. In practice, this limits the recursive depth to a small constant, typically just two levels or so. Further improving amortized FHEW bootstrapping, allowing the execution of a homomorphic FFT with $O(\log n)$ levels is left as an open problem.

Related and Concurrent Work: Ring automorphisms have been used in many other works aimed at improving the efficiency of lattice cryptography based on the RingLWE problem, most notably the evaluation of linear functions in HELib [13] and algebraic trace computations [2]. Our use of automorphisms is most closely related to [15], which recently used them to improve the performance of FHEW (sequential, non-amortized) bootstrapping. In a concurrent and independent work [11], an algorithm very similar to ours is presented. The algorithm achieves essentially the same results, improving the cost of amortized FHEW bootstrapping from $2^{O(1/\epsilon)} \cdot n^{1/\epsilon}$ to $(1/\epsilon) \cdot n^{1/\epsilon}$. The overall structure of the algorithm is very similar, using automorphisms to replace the Nussbaumer transform in [21] with a standard FFT. However, the algorithms differ in some technical details. For example, while [11] uses the circular rings [3], we use prime cyclotomics, which results in marginally smaller ciphertexts. Another difference is that while [11] extends the automorphism multiplication technique to work directly on RingGSW ciphertexts, we center our FFT algorithm around RingLWE registers (which are smaller than RingGSW by a factor 2) and convert them to RingGSW only when necessary using our new scheme switching technique. This allows us to exploit RLWE²-RGSW multiplications instead of RGSW-RGSW multiplications during the homomorphic inverse FFT, which gives a 2x performance improvement compared to RGSW key-switching as used in [11].

At Eurocrypt 2023, Liu and Wang introduced a new algebraic framework for batch homomorphic computation based on the tensoring of three rings [16]. Their new framework is also used in the context of bootstrapping algorithms [17] to improve the efficiency of the amortized bootstrapping algorithm following the Nussbaumer technique from [21], achieving an amortized bootstrapping cost of $\tilde{O}(1)$ FHE multiplications within a polynomial modulus. We compare our algorithm to theirs in Section 5.4.

More recently still, another line a work from Liu and Wang [18] (Asiacrypt 2023) obtained the asymptotic results of [16, 17] while also achieving concrete efficiency by exploiting both BFV and LWE ciphertexts. Asymptotically, [18] requires a super-polynomial modulus whereas our work considers only polynomial modulus. Considering a super-polynomial modulus would allow to increase the recursive depth in our work but would unlikely be effective in practice.

2 Preliminaries

2.1 Cyclotomic Rings and Embeddings

Given a positive integer N , the N^{th} cyclotomic polynomial is defined as $\Phi_N(X) = \prod_{i \in \mathbb{Z}_N^*} (X - \omega_N^i)$ for $\omega_N = e^{2\pi i/N} \in \mathbb{C}$ the complex N^{th} principal root of unity. The N^{th} cyclotomic ring is defined as $\mathcal{R}_N = \mathbb{Z}[X]/(\Phi_N(X))$. In this work, we will consider the d^{th} cyclotomic ring modulo q , for d a power-of-2, defined as $\mathcal{R}_d = \mathbb{Z}_q[X]/\Phi_d(X) \simeq \mathbb{Z}_q^{\phi(d)}$. Each element of this ring corresponds to a polynomial $\mathbf{a} \in \mathcal{R}_d$ of degree less than $\phi(d)$ and with coefficients taken modulo q . There exist various ways of representing a ring element. One can first map the polynomial $\mathbf{a}(X) = \sum_{i \leq \phi(d)} a_i \cdot X^i$ to its vector of coefficients $(a_1, a_2, \dots, a_{\phi(d)}) \in \mathbb{Z}_q^{\phi(d)}$. This is known as the coefficient embedding. The norm of any ring element then refers to the ℓ_2 norm of the corresponding vector in the coefficient embedding.

Another representation of a ring element is with its canonical embedding $\sigma : K \rightarrow \mathbb{C}^n$ which endows K , the d^{th} cyclotomic number field, with a geometry. Note that the ring of integers of K corresponds to the d^{th} cyclotomic ring $\mathbb{Z}[X]/\Phi_d(X)$. We know that K has exactly $\phi(d)$ ring homomorphisms, also called

embeddings, $\sigma_i : K \rightarrow \mathbb{C}$. The canonical embedding is then defined as the map $\sigma(a) = (\sigma_i(a))_{i \in \mathbb{Z}_d^*}$ for $a \in K$. The norm usually considered when using the canonical embedding is the ℓ_∞ norm $\|\sigma(a)\|_\infty = \max_i |\sigma_i(a)|$. More generally, for any $a \in K$ and any $p \in [1, \infty]$, the ℓ_p norm is defined as $\|a\|_p = \|\sigma(a)\|_p$. Since the σ_i are ring homomorphisms, we then have for any $a, b \in K$ the inequality $\|a \cdot b\|_p \leq \|a\|_\infty \cdot \|b\|_p$.

Working with prime cyclotomics, or more generally with non-power-of-two cyclotomics can be rather cumbersome, in particular, when considering the canonical embedding and not just the coefficient embedding. We know that any two embeddings are related to each other by a fixed linear transformation on \mathbb{R}^d . For power-of-2 cyclotomics, the transformation is even an isometry and thus the coefficient and canonical embeddings are equivalent up to a \sqrt{d} factor. In this work, we will be working with both \mathcal{R}_d , the d^{th} cyclotomic ring modulo q for which we will use the notation \mathcal{R}_{in} and the q^{th} cyclotomic modulo Q for a prime q and a positive integer $Q > 0$, which we will denote \mathcal{R}_{reg} . The latter is a prime cyclotomic ring where the two embeddings cannot be easily interchanged. This will affect the error growth analysis as we later explain in Section 3.2.

2.2 Encryption Schemes and Operations

We recall definitions and notations for the standard LWE encryption scheme used in the bootstrapping algorithm. We also extend our description to the ring version of LWE and introduce two related schemes, GadgetRLWE and RGSW, both used in our algorithm.

LWE: Consider some positive integers n and q . Let $\mathbf{sk} \leftarrow \chi$ be a secret key sampled from a distribution χ and $m \in \mathbb{Z}$ a message. The LWE encryption of the message m under the secret key \mathbf{sk} is given by

$$\text{LWE}_{q, \mathbf{sk}}(m) = [\mathbf{a}^T, b] \in \mathbb{Z}_q^{1 \times (n+1)},$$

where $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, $b = -\mathbf{a} \cdot \mathbf{sk} + e + m \in \mathbb{Z}_q$ and $e \leftarrow \chi'$ is the error, sampled from a distribution χ' , and ciphertexts are represented as *row* vectors.

RLWE: The ring version of LWE considers the ring \mathcal{R}_q . Let $\mathbf{sk} \leftarrow \chi$ be a secret key sampled from a distribution χ and $\mathbf{m} \in \mathcal{R}_q$ a message. The RLWE encryption of the message \mathbf{m} under the secret key \mathbf{sk} is given by

$$\text{RLWE}_{q, \mathbf{sk}}(\mathbf{m}) = [\mathbf{a}, \mathbf{b}] \in \mathcal{R}_q^{1 \times 2},$$

where $\mathbf{a} \leftarrow \mathcal{R}_q$ and $\mathbf{b} = -\mathbf{a} \cdot \mathbf{sk} + \mathbf{e} + \mathbf{m}$ and $e_i \leftarrow \chi'$ for each coefficient e_i of the error. If context is clear, we do not specify the modulus q or the secret key \mathbf{sk} .

Gadget RLWE or RLWE' Consider a gadget vector $\mathbf{v} = (v_0, v_1, \dots, v_{k-1})$. Gadget RLWE or equivalently referred to as RLWE' is expressed as a vector of RLWE ciphertexts of the form

$$\text{RLWE}'_{\mathbf{sk}}(\mathbf{m}) = (\text{RLWE}_{\mathbf{sk}}(v_0 \cdot \mathbf{m}), \text{RLWE}_{\mathbf{sk}}(v_1 \cdot \mathbf{m}), \dots, \text{RLWE}_{\mathbf{sk}}(v_{k-1} \cdot \mathbf{m})) \in \mathcal{R}_q^{k \times 2}$$

i.e., matrices with k rows, each representing a basic RLWE ciphertext. We remark that RLWE ciphertext can be regarded as a special case of RLWE' instantiated with a trivial gadget $\vec{v} = (1)$. So, anything we say about RLWE' applies to RLWE as well.

RingGSW Given a message $\mathbf{m} \in \mathcal{R}_q$, we define

$$\text{RGSW}_{\mathbf{sk}}(\mathbf{m}) = (\text{RLWE}'_{\mathbf{sk}}(\mathbf{sk} \cdot \mathbf{m}), \text{RLWE}'_{\mathbf{sk}}(\mathbf{m})) \in \mathcal{R}_q^{2k \times 2}.$$

We now summarize the operations that can be done with the different schemes presented above and focus in particular on the operations used in our algorithm. The main operation in our algorithm that serves as a building block for other operations is the scalar multiplication by arbitrary ring elements. In order to compute

this multiplication, one uses RLWE' with gadget vector $\mathbf{v} = (v_0, v_1, \dots, v_{k-1})$. The scalar multiplication is denoted as $\mathcal{R} \odot \text{RLWE}'$ and corresponds to $\odot : \mathcal{R} \times \text{RLWE}' \rightarrow \text{RLWE}$ defined as

$$\begin{aligned} \mathbf{t} \odot \text{RLWE}'_{\text{sk}}(\mathbf{m}) &:= \sum_{i=0}^{k-1} \mathbf{t}_i \cdot \text{RLWE}_{\text{sk}}(v_i \cdot \mathbf{m}) \\ &= \text{RLWE}_{\text{sk}} \left(\sum_{i=0}^{k-1} v_i \cdot \mathbf{t}_i \cdot \mathbf{m} \right) = \text{RLWE}_{\text{sk}}(\mathbf{t} \cdot \mathbf{m}) \end{aligned}$$

where $\sum_i v_i \mathbf{t}_i = \mathbf{t}$ is the gadget decomposition of \mathbf{t} into “short” vectors \mathbf{t}_i , for an appropriate notion of “short” depending on the gadget \mathbf{v} . Each operation performed with ciphertexts increases the error. When performing many of these operations, as in our bootstrapping algorithm, it is crucial to keep track of the error growth. More details will be given in Section 3.2. For now, we simply state that each error e_i in $\text{RLWE}_{\text{sk}}(v_i \cdot \mathbf{m})$, after the scalar multiplication, becomes $\sum_{i=0}^{k-1} \mathbf{t}_i \cdot \mathbf{e}_i$.

The RLWE and RLWE' schemes only support multiplication by constant values. In order to obtain multiplication by ciphertexts, we need to consider the RGSW scheme. Let us now consider the multiplication $\star : \text{RLWE} \times \text{RGSW} \rightarrow \text{RLWE}$ defined as

$$\begin{aligned} \text{RLWE}_{\text{sk}}(\mathbf{m}_1) \star \text{RGSW}_{\text{sk}}(\mathbf{m}_2) &:= \mathbf{a} \odot \text{RLWE}'_{\text{sk}}(\text{sk} \cdot \mathbf{m}_2) + \mathbf{b} \odot \text{RLWE}'_{\text{sk}}(\mathbf{m}_2) \\ &= \text{RLWE}_{\text{sk}}(\mathbf{a} \cdot \text{sk} \cdot \mathbf{m}_2 + \mathbf{b} \cdot \mathbf{m}_2) \\ &= \text{RLWE}_{\text{sk}}(\mathbf{m}_1 \cdot \mathbf{m}_2 + \mathbf{e}_1 \cdot \mathbf{m}_2) \end{aligned}$$

for $\text{RLWE}(\mathbf{m}_1) := (\mathbf{a}, \mathbf{b})$. The output of this multiplication is an RLWE ciphertext encrypting the message $\mathbf{m}_1 \cdot \mathbf{m}_2 + \mathbf{e}_1 \cdot \mathbf{m}_2$. The error thus additively increases by $\mathbf{e}_1 \cdot \mathbf{m}_2$. If the error term $\mathbf{e}_1 \cdot \mathbf{m}_2$ is sufficiently small, then this approximately results in an RLWE encryption of the product of the two messages. This multiplication can be extended to RLWE' ciphertext multiplication $\star' : \text{RLWE}' \times \text{RGSW} \rightarrow \text{RLWE}'$ defined as

$$\begin{aligned} &\text{RLWE}'_{\text{sk}}(\mathbf{m}_1) \star' \text{RGSW}_{\text{sk}}(\mathbf{m}_2) \\ &:= (\text{RLWE}_{\text{sk}}(v_0 \cdot \mathbf{m}_1) \star \text{RGSW}_{\text{sk}}(\mathbf{m}_2), \dots, \text{RLWE}_{\text{sk}}(v_{k-1} \cdot \mathbf{m}_1) \star \text{RGSW}_{\text{sk}}(\mathbf{m}_2)) \\ &\approx \text{RLWE}'_{\text{sk}}(\mathbf{m}_1 \cdot \mathbf{m}_2). \end{aligned}$$

Each component $\text{RLWE}_{\text{sk}}(v_i \cdot \mathbf{m}_1) \star \text{RGSW}_{\text{sk}}(\mathbf{m}_2)$ of the result has the same error growth as a \star operation in the RLWE \star RGSW case. In particular, it includes an error term $\mathbf{e}_{1,i} \cdot \mathbf{m}_2$ that requires the second message \mathbf{m}_2 to be small.

The \star' operation corresponds to k times the \star operations, and thus a total of $2k \odot$ operations.

2.3 Using Ring Automorphisms

Similarly as in [15], we use ring automorphisms to perform scalar multiplication with registers. Recall that an automorphism is a bijective maps from the ring \mathcal{R} to itself such that for a given $t \in \mathbb{Z}_q^*$, we have $\mathbf{a}(X) \mapsto \mathbf{a}(X^t)$.

Automorphism in RLWE and RLWE': Consider the following RLWE ciphertext $(\mathbf{a}(X), \mathbf{b}(X))$ which encrypts a given message $\mathbf{m}(X)$ under a certain key sk , *i.e.*, $(\mathbf{a}(X), \mathbf{b}(X)) = \text{RLWE}_{\text{sk}}(\mathbf{m}(X))$. We also consider a switching key $\mathbf{ak}_t = \text{RLWE}'_{\text{sk}}(\text{sk}(X^t))$, which is used to map ciphertexts $[\mathbf{a}, \mathbf{b}]$ from key $\text{sk}(X^t)$ to sk . Given an automorphism $\psi_t : \mathcal{R} \rightarrow \mathcal{R}$ such that $\mathbf{a}(X) \mapsto \mathbf{a}(X^t)$, we recall the procedure `Evalauto` given in [15]:

1. apply ψ_t to each of the RLWE components. One obtains $(\mathbf{a}(X^t), \mathbf{b}(X^t)) = \text{RLWE}_{\text{sk}(X^t)}(\mathbf{m}(X^t))$
2. apply a key switching function $[\mathbf{a}, \mathbf{b}] \mapsto \mathbf{a} \odot \mathbf{ak}_t + [\mathbf{0}, \mathbf{b}]$ to obtain a ciphertext $\text{RLWE}_{\text{sk}(X)}(\mathbf{m}(X^t))$

The same application can be done on RLWE' ciphertexts. The only difference comes during the second step where we require k key switching, one for each RLWE ciphertext. The only $\mathcal{R} \odot \text{RLWE}'$ operation comes from key switching. Hence, for automorphism on RLWE, we have a single $\mathcal{R} \odot \text{RLWE}'$ operation and when considering automorphisms on RLWE' we have $k \mathcal{R} \odot \text{RLWE}'$ operations, where k is the length of the gadget.

2.4 Homomorphic Operations on Registers

Following the FHEW framework, we use cryptographic registers that encrypt a \mathbb{Z}_q element “in the exponent”. In other words, a register storing $m \in \mathbb{Z}_q$ is an encryption of $X^m \in \mathcal{R}_{reg}$. In our algorithm, the encryption scheme will sometimes be RLWE’ and sometimes RGSW. Some operations require one scheme or the other. In order to perform some of these operations, we will need to scheme-switch from RLWE’ to RGSW. We describe our scheme-switching technique in Section 3.1. In our bootstrapping algorithm, we will primarily use three operations on registers, *i.e.*, either RLWE’ or RGSW ciphertexts. We have already mentioned these operations and recall them now:

- \star' : RLWE’ \times RGSW \rightarrow RLWE’ multiplications: this operation allows to multiply two ciphertexts, which in the exponent acts like an addition.
- Scheme-switching: this operation converts an RLWE’ register into an RGSW register.
- Automorphisms: this operation allows us to multiply the exponent of a RLWE’ ciphertext by some (known) value, and corresponds to multiplication by a constant.

Note that automorphisms can only operate on RLWE’ registers, not an RGSW ones. (This is because RGSW does not directly support the key switching operation required by the second step of the homomorphic automorphic application algorithm. See Section 3.1 for details.) On the other hand, multiplication requires one of the two registers to be in RGSW format. The scheme switching operation is used to combine the two operations, keeping all registers in RLWE’ form, and convert them to RGSW only when required for multiplication.

In order to analyse the performance and the correctness of our algorithm we will analyse these three operations in terms of number of $\mathcal{R} \odot$ RLWE’ operations needed to compute them and the related error growth (see Table 2).

2.5 Standard and Primitive (Inverse) FFT

We only mention in this paper some relevant facts about FFT algorithms that are useful for our algorithm. Note that when referring to FFT and related algorithms, we actually refer to the Number Theoretic Transform (NTT) algorithm.

An FFT algorithm can either evaluate a polynomial at all N^{th} roots of unity for a given N or only the primitive ones. The former case is referred to as a standard/cyclic FFT whereas the latter case is called a primitive/cyclotomic FFT. In the case of a standard FFT, the inverse direction reconstructs from these evaluations a polynomial mod $X^N - 1$. When multiplying two polynomials $a(x)$ and $z(x)$ modulo a cyclotomic polynomial, using a standard FFT (and its inverse) requires a “final reduction” step to take polynomials modulo $(X^N - 1)$ to polynomials modulo $\Phi_N(X)$. This “final reduction” increases the multiplicative depth of the circuit and, in our case, prevents some useful optimizations (namely, using RLWE instead of RLWE’ in the last FFT layer as we will explain in Section 4.2). We can avoid the final reduction step by using a primitive/cyclotomic FFT, which we recall only evaluates the polynomials at the primitive N^{th} roots of unity ω^i for $i \in \mathbb{Z}_N^*$. The inverse FFT then reconstructs from these evaluations a polynomial modulo the N^{th} cyclotomic polynomial $\Phi_N(X)$. We note however that, unlike with a standard FFT, the forward and inverse directions are not interchangeable. We focus the rest of the discussion on the case of power-of-two cyclotomics (which is the case we will use in this paper) where $N = d = 2^{\log_2 N}$, and $\phi(d) = d/2$. For the forward direction, let $0 \leq i < \phi(d) = d/2$, and ω be a primitive d^{th} root of unity. Then the Fourier coefficients $\hat{f} := (\hat{f}_0, \hat{f}_1, \dots, \hat{f}_{d/2})$ of a polynomial $f(X) = \sum_{i=0}^{d/2-1} f_i \cdot X^i \pmod{X^{d/2} + 1}$ are computed as

$$\hat{f}_i := \sum_{j=0}^{d/2-1} f_j \omega^{(2i+1)j},$$

and the inverse FFT of \hat{f} can be computed as

$$\hat{f}_\ell := \frac{2}{d} \cdot \sum_{i=0}^{d/2-1} \hat{f}_i \omega^{-(2i+1)\ell} = \frac{2}{d} \cdot \omega^{-\ell} \sum_{i=0}^{d/2-1} \hat{f}_i \omega^{-2i\ell}$$

for each $0 \leq \ell < d/2$ and output $\widehat{f}(X) = \sum_{i=0}^{d/2-1} \widehat{f}_i \cdot X^i$. It is easy to verify that these operations are inverses of each other, *i.e.*, $\widehat{\widehat{f}} = f$. The FFT also preserves both addition and multiplication, *i.e.*, $\widehat{f + g} = \widehat{f} + \widehat{g}$ and $\widehat{f \circ g} = \widehat{f} \cdot \widehat{g}$ where \circ denotes the component-wise multiplication of input vectors. Moreover, one notices that the inverse operation can be computed as a standard/cyclic length- $\phi(d)$ FFT (using ω^{-2} as the $\phi(d)^{th}$ root of unity) followed by a multiplication by a power of ω .

In this work, we mainly focus on homomorphic computation of the inverse FFT (while the forward FFT is done in cleartext), so the constant multiplication by $2/d$ becomes a (minor) computational overhead. We can easily remove this overhead by moving the constant from the inverse FFT to the forward FFT. If we move the constant $2/d$, we get

$$\text{FFT}(f)_i := \frac{2}{d} \cdot \widehat{f}_i = \frac{2}{d} \cdot \sum_{j=0}^{d/2-1} f_j \omega^{(2i+1)j},$$

$$\text{FFT}^{-1}(\text{FFT}(f))_\ell := \frac{d}{2} \cdot \widehat{\text{FFT}(f)}_\ell = \sum_{i=0}^{d/2-1} \text{FFT}(f)_i \omega^{-(2i+1)\ell}.$$

In this case, FFT^{-1} is still the inverse of FFT and addition is preserved in the same manner. However, note that there is a slight difference in multiplication: $\text{FFT}^{-1}(\text{FFT}(f) \circ \widehat{g}) = \text{FFT}^{-1}(\text{FFT}(f \cdot g)) = f \cdot g$.

In our algorithm, we will consider partial (primitive) FFT, denoted by PFT, where instead of reducing modulo $(X - \zeta)$ (*i.e.*, evaluating the polynomials at $X = \zeta$, we reduce modulo $(X^k - \zeta)$. In this reduction, an X^i term will not interact with an X^j term unless $i \equiv j \pmod k$. Hence, an equivalent description of a partial FFT is doing k FFTs in parallel, each with $1/k$ as many terms. More precisely, one FFT will operate on the terms which are 0 modulo k , one FFT on just the terms that are 1 modulo k , and so on. This also applies to the inverse direction. In our algorithm, we will thus divide by $\phi(d)/k$ and not $\phi(d)$.

2.6 Summary of Notations

We summarize the notations used throughout the paper in Table 1. For simplicity of exposition, this paper uses a standard power-of-B gadget $(1, B, B^2, \dots, B^{d_B-1})$. In practice, this can be replaced by a CRT gadget which typically supports more efficient implementation.

3 Novel Techniques

In this section, we introduce some novel techniques related to scheme switching and error analysis. We first introduce a new variant of scheme-switching. We then introduce an error analysis in the context of prime cyclotomics. Indeed, our algorithm will use a prime cyclotomic for the registers, whereas common FHE schemes use power-of-2 cyclotomic rings for which the error analysis differs.

3.1 RLWE' to RGSW Scheme Switching

When an automorphism is applied to a ciphertext, it modifies both the encrypted message and the encryption key. (This applies to RLWE, RLWE' and RGSW ciphertexts alike.) Therefore, in order to use automorphisms to operate homomorphically on ciphertexts, one needs a method to switch back to the original key. For RLWE and RLWE' ciphertexts, this is provided by a standard key switching operation as described in the previous section. However, for RGSW encryption, this does not quite work. The reason is that a RGSW encryption can be interpreted as a pair of RLWE' ciphertexts encrypting \mathbf{m} and $\mathbf{m} \cdot \mathbf{sk}$. The first component does not pose any problem, as it can be transformed using a standard RLWE' key switching operation. However, key switching cannot be directly applied to the second component, because it encrypts a *key-dependent* message $\mathbf{m} \cdot \mathbf{sk}$. So, RGSW key-switching would require not only to modify the encryption key, but also

Table 1. Summary of notations used in the paper

	Notation	Description
Modulus	q_{plain}	Ciphertext modulus for standard LWE.
	q	Prime ciphertext modulus for input RLWE ciphertext. Plaintext modulus for registers.
	Q	Ciphertext modulus used in RLWE' / RGSW registers.
Rings	\mathcal{R}_{in}	d th cyclotomic ring (mod q), $\frac{\mathbb{Z}_q[x]}{\Phi_d(x)} \simeq \mathbb{Z}_q^{\phi(d)}$.
	d	power-of-2 degree of \mathcal{R}_{in} .
	\mathcal{R}_{reg}	q th prime cyclotomic ring mod Q used by the registers.
FFT	k	degree at which we stop the partial FFT.
	$\phi(d)$	the number of Plain-LWE ciphertexts that are packed into an RLWE ciphertext; the number of coefficients in an \mathcal{R}_{in} element; the number of coefficients in the input polynomial of the FFT; the number of registers in any layer of the IFFT; the number of registers output by the IFFT.
	$N = \phi(d)/k$	the number of degree- $(k-1)$ polynomials output by the partial FFT.
	$\omega \in \mathbb{Z}_q$	a primitive (d/k) th root of unity in \mathbb{Z}_q for use in the FFT.
	r_i	radix for FFT layer.
	ℓ	number of FFT layers.
Secret keys	$\mathbf{s}_p \in \mathbb{Z}_{q_{plain}}^{n_{plain}+1}$	Plain LWE secret key.
	$z \in \mathcal{R}_{in}$	RLWE secret key for the input (packed) RLWE ciphertext.
	$s \in \mathcal{R}_{reg}$	RGSW secret key used for registers.
Gadget decomposition	B	Base for the powers-of-B gadget used in registers.
	d_B	$\lceil \log_B(Q) \rceil$, the length of the PowersOfB gadget.
Error variance	σ_{\odot}^2	The (expected) factor by which an $\mathcal{R}_{reg} \odot$ RLWE' operation scales up the error variance in a ciphertext.
	$\sigma_{\odot, RGSW}^2$	The resulting error variance of the \odot operation on each RLWE' component of RGSW(\mathbf{m}_2).
	$\sigma_{\odot, eval_key}^2$	The resulting error variance of the \odot operation on the evaluation key with error variance $\sigma_{eval_key}^2$.
	$\sigma_{\odot, aut_key}^2$	The resulting error variance of the \odot operation on the automorphism key $\text{RLWE}'_{\mathbf{sk}}(\psi(\mathbf{sk}))$.
	σ_{in}^2	The error variance of the input to an operation.

to change the message from $\mathbf{m} \cdot \mathbf{sk}$ to $\mathbf{m} \cdot \mathbf{sk}'$, where \mathbf{sk}' is the new key. For this reason, key switching (and homomorphic automorphism evaluation), is directly applicable only to RLWE and RLWE' ciphertexts. On the other hand, RGSW ciphertexts are required to perform homomorphic multiplications when both multiplicands are encrypted. We address this problem by providing a method to convert RLWE' ciphertexts to RGSW ones, which we call *scheme switching*. Let us now describe how this is done.

Since $\text{RGSW}_{\mathbf{sk}}(\mathbf{m}) = (\text{RLWE}'_{\mathbf{sk}}(\mathbf{sk} \cdot \mathbf{m}), \text{RLWE}'_{\mathbf{sk}}(\mathbf{m}))$ and we are given $\text{RLWE}'_{\mathbf{sk}}(\mathbf{m})$, we just need a way to compute $\text{RLWE}'_{\mathbf{sk}}(\mathbf{sk} \cdot \mathbf{m})$. To do so, we will use $\text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}^2)$ given as part of the evaluation key. We will operate in parallel on each of the $\text{RLWE}_{\mathbf{sk}}(v_i \cdot \mathbf{m})$ ciphertexts that make up the $\text{RLWE}'_{\mathbf{sk}}(\mathbf{m})$ ciphertext, lifting each $\text{RLWE}_{\mathbf{sk}}(v_i \cdot \mathbf{m})$ to $\text{RLWE}_{\mathbf{sk}}(v_i \cdot \mathbf{sk} \cdot \mathbf{m})$. More precisely, for each $\text{RLWE}_{\mathbf{sk}}(v_i \cdot \mathbf{m}) := (\mathbf{a}, \mathbf{b})$, we compute

$$\mathbf{a} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}^2) + (\mathbf{b}, 0).$$

By regarding $(\mathbf{b}, 0)$ as a noiseless RLWE encryption of $\mathbf{b} \cdot \mathbf{sk}$ under the secret key \mathbf{sk} , this computation gives $\text{RLWE}_{\mathbf{sk}}(\mathbf{a} \cdot \mathbf{sk}^2 + \mathbf{b} \cdot \mathbf{sk}) = \text{RLWE}_{\mathbf{sk}}((\mathbf{a} \cdot \mathbf{sk} + \mathbf{b}) \cdot \mathbf{sk}) = \text{RLWE}_{\mathbf{sk}}((v_i \cdot \mathbf{m} + \mathbf{e}) \cdot \mathbf{sk})$. Hence we do get $\text{RLWE}_{\mathbf{sk}}(v_i \cdot \mathbf{sk} \cdot \mathbf{m})$ as desired, but with an additional error $\mathbf{e} \cdot \mathbf{sk}$ scaled up by \mathbf{sk} from the input RLWE' ciphertext error \mathbf{e} . We will choose the secret key \mathbf{sk} with small norm (e.g., binary) so that this multiplicative error growth remains small. More details about the full error growth for this scheme switching will be given in Section 3.2.

When our scheme switching method is used in conjunction with key switching, it allows a small optimization. Say we are given a $\text{RLWE}'_{\mathbf{sk}'}(\mathbf{m})$, and we want to turn it into a RGSW encryption under \mathbf{sk} . This can be done in two steps, by first performing key-switching to $\text{RLWE}'_{\mathbf{sk}}(\mathbf{m})$, and then using the scheme switching key $\text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}^2)$ to compute $\text{RLWE}'_{\mathbf{sk}}(\mathbf{m} \cdot \mathbf{sk})$. The optimization consists in using a modified scheme switching key $\text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}' \cdot \mathbf{sk})$ to turn the input ciphertext (encrypted under \mathbf{sk}') directly into $\text{RLWE}'_{\mathbf{sk}}(\mathbf{m} \cdot \mathbf{sk})$, performing key switching $\mathbf{sk}' \rightarrow \mathbf{sk}$ and homomorphic multiplication by \mathbf{sk} at the same time. Notice that the running time is about the same as before as we still need another key switching $\mathbf{sk}' \rightarrow \mathbf{sk}$ to compute

the other component of the output RGSW ciphertext. However, combining key switching and multiplication in a single operation allows to slightly reduce the noise growth. To give a more modular presentation, we ignore this optimization in the description of our algorithm.

3.2 Error Growth in Prime Cyclotomics

Analysing the error growth in bootstrapping algorithms is crucial for the correctness of the scheme as it allows to set the proper modulus sizes and show the implementation can be run with concrete parameters. It is a standard practice in lattice cryptography to estimate the error growth during homomorphic operations under the heuristic assumption that the noise in ciphertexts behaves like independent gaussian (or subgaussian) random variables, with standard deviation that depends on the computation leading to the ciphertext. In order to fairly compare our algorithm to previous work, in this paper we use a similar technique and compute the total error estimation based on the error variance introduced by a single $\mathcal{R}_{reg} \odot \text{RLWE}'$ operation. In previous works, where a power-of-2 cyclotomic is being used, this value is equal to $\frac{1}{12}d_B q B^2 \sigma_{\text{input}}^2$ where B is the base for the power-of- B gadget, d_B is the length of the gadget and σ_{input}^2 is the error variance of the input RLWE' ciphertext considered. Because \mathcal{R}_{reg} is a prime cyclotomic, the analysis of this variance differs in our case as we do not directly have a bound on the ℓ_∞ norm in the canonical embedding of a ring element for which we know a bound on each coefficient. We thus propose the following theorem.

Theorem 1. *For an odd prime q and a positive integer Q , let \mathcal{R}_{reg} be the q^{th} cyclotomic ring modulo Q used for registers, and d_B be the length of the gadget decomposition. For an RLWE' ciphertext defined over \mathcal{R}_{reg} , the error variance of the result of a single $\mathcal{R}_{reg} \odot \text{RLWE}'$ operation is bounded by*

$$\sigma_{\odot}^2 \leq 2d_B q \sigma_r^2 \sigma_{\text{input}}^2,$$

where σ_{input}^2 is the error variance of the input RLWE' ciphertext, and σ_r^2 is the variance of the gadget decomposition of the input \mathcal{R}_{reg} ring element.

Proof. We model an element $r \in \mathcal{R}_{reg}$ as sampled uniformly at random — this is a reasonable model because in our algorithm the \mathcal{R}_{reg} elements always either come from a ciphertext (and hence are uniform) or are simply an integer constant (leading to even smaller error growth). The gadget decomposition of $r \in \mathcal{R}_{reg}$, denoted $\mathbf{G}^{-1}(r)$, then consists of d_B ring elements r_1, \dots, r_{d_B} (which we model as independently distributed). Note that we will use in our algorithm a balanced base- B digit decomposition but we leave the decomposition unspecified here for sake of generality. We model the error vector $\vec{e}_{\text{RLWE}'}$ = (e_1, \dots, e_{d_B}) where each component is the error of each RLWE ciphertext in the input RLWE' ciphertext, as independent random variables with variance σ_{input}^2 . The output error can then be computed as an inner product

$$e_{\text{output}} = \langle \mathbf{G}^{-1}(r), \vec{e}_{\text{RLWE}'} \rangle = \sum_{i=1}^{d_B} r_i \cdot e_i.$$

We will start by considering a single multiplication $v_i := r_i \cdot e_i \pmod{\Phi_q(X)}$. Recall that r_i and e_i are both ring elements of \mathcal{R}_{reg} , *i.e.*, polynomials of degree $q-2$ (as $\mathcal{R}_{reg} = \mathbb{Z}_Q[X]/(\Phi_q(X))$ and $\Phi_q(X) = 1 + X + \dots + X^{q-1}$). We want to compute the variance of each coefficient of

$$v_i = (r_{i,0} + r_{i,1}X + \dots + r_{i,q-2}X^{q-2}) \cdot (e_{i,0} + e_{i,1}X + \dots + e_{i,q-2}X^{q-2}) \pmod{\Phi_q(X)}.$$

For simplicity of notation in the formulas below, we will consider r_i and e_i to be polynomials of degree $q-1$ (instead of $q-2$) with leading coefficients 0, *i.e.*, the trivial terms $r_{i,q-1} = e_{i,q-1} := 0$. By computing $r_i \cdot e_i \pmod{X^q - 1}$ first and then taking the result modulo $\Phi_q(X)$, we can easily obtain the ℓ^{th} coefficient of v_i , which we denote $v_i^{(\ell)}$, for $0 \leq \ell \leq q-2$. First, note that the ℓ^{th} coefficient of $v_i' := r_i \cdot e_i \pmod{X^q - 1}$ is given by $v_i'^{(\ell)} = \sum_{j=0}^{q-1} r_{i,j} \cdot e_{i,\ell-j}$, where the subscripts of e are defined modulo q , *i.e.*, $e_{i,\ell-j} := e_{i,q+\ell-j}$ if $\ell < j$. Then, since $X^{q-1} = -X^{q-2} - \dots - 1 \pmod{\Phi_q(X)}$, the ℓ^{th} coefficient of v_i modulo $\Phi_q(X)$ is computed

as $v_i^{(\ell)} = v_i'^{(\ell)} \circ v_i'^{(q-1)} = \sum_{j=0}^{q-1} r_{i,j} \cdot (e_{i,\ell-j} - e_{i,q-j-1})$. Let $X_{i,j}^{(\ell)} := r_{i,j} \cdot (e_{i,\ell-j} - e_{i,q-j-1})$ for $0 \leq j \leq q-1$ and hence $v_i^{(\ell)} = \sum_{j=0}^{q-1} X_{i,j}^{(\ell)}$. Since $r_{i,q-1} = 0$ is a constant value, we trivially have that $\text{var}(X_{i,q-1}^{(\ell)}) = 0$. When $0 \leq j \leq q-2$, the variance of each $X_{i,j}^{(\ell)}$ equals to

$$\text{var}(X_{i,j}^{(\ell)}) = \text{var}(r_{i,j}) \cdot \text{var}(e_{i,\ell-j} - e_{i,q-j-1}) = \begin{cases} \sigma_r^2 \sigma_{\text{input}}^2 & \text{if } j = 0 \text{ or } \ell + 1 \\ 2\sigma_r^2 \sigma_{\text{input}}^2 & \text{else} \end{cases}.$$

The first variance corresponds to the case where $\text{var}(e_{i,\ell-j} - e_{i,q-j-1}) = \text{var}(e_{i,\ell-j})$ as $e_{i,q-j-1} = 0$ when $j = 0$ or when $\text{var}(e_{i,\ell-j} - e_{i,q-j-1}) = \text{var}(e_{i,q-j-1})$ as $\text{var}(e_{i,\ell-j}) = 0$ when $j = \ell + 1$. Since $\text{var}\left(\sum_{j=0}^{q-1} X_{i,j}^{(\ell)}\right) = \sum_{j=0}^{q-1} \text{var}(X_{i,j}^{(\ell)}) + 2 \sum_{0 \leq j < k < q} \text{cov}(X_{i,j}^{(\ell)}, X_{i,k}^{(\ell)})$, it now suffices to compute the covariance of each pair. We will first consider the special case where $k = j + \ell + 1$ as it is the only case where common terms appear between $X_{i,j}^{(\ell)}$ and $X_{i,k}^{(\ell)}$. Indeed, we have that for $k = j + \ell + 1$,

$$X_{i,k}^{(\ell)} = r_{i,j+\ell+1} \cdot (e_{i,-j-1} - e_{i,q-j-\ell-2}),$$

where $e_{i,-j-1} = e_{i,q-j-1}$ also appears in $X_{i,j}^{(\ell)}$. However, due to the distributive property of covariance, it holds that

$$\text{cov}(X_{i,j}^{(\ell)}, X_{i,k}^{(\ell)}) = -\text{cov}(r_{i,j} \cdot e_{i,q-j-1}, r_{i,j+\ell+1} \cdot e_{i,q-j-1}) = 0^3.$$

In all other cases we trivially have $\text{cov}(X_{i,j}^{(\ell)}, X_{i,k}^{(\ell)}) = 0$ since $X_{i,j}^{(\ell)}$ and $X_{i,k}^{(\ell)}$ are independent. Note that there exist two j indices ($j = 0, \ell + 1$) satisfying $\text{var}(X_{i,j}^{(\ell)}) = \sigma_r^2 \sigma_{\text{input}}^2$ when $0 \leq \ell < q-2$, while there exists only one such j index ($j = 0$) when $\ell = q-2$. As a result, we obtain the variance of $v_i^{(\ell)}$ as

$$\text{var}(v_i^{(\ell)}) = \sum_{j=0}^{q-1} \text{var}(X_{i,j}^{(\ell)}) = \begin{cases} (2q-4)\sigma_r^2 \sigma_{\text{input}}^2 & \text{if } 0 \leq \ell < q-2 \\ (2q-3)\sigma_r^2 \sigma_{\text{input}}^2 & \text{if } \ell = q-2 \end{cases}.$$

Finally, the variance of each coefficient of e_{output} denoted by σ_{\odot}^2 is bounded by $2d_B q \sigma_r^2 \sigma_{\text{input}}^2$. \square

Corollary 1. *For an odd prime q and a positive integer Q , let \mathcal{R}_{reg} be the q^{th} cyclotomic ring modulo Q used for registers, and d_B be the length of a balanced base- B gadget decomposition with uniform coefficients in $[-B/2, B/2)$. For an RLWE' ciphertext defined over \mathcal{R}_{reg} , the error variance of the result of a single $\mathcal{R}_{\text{reg}} \odot$ RLWE' operation is bounded by*

$$\sigma_{\odot}^2 \leq \frac{B^2}{6} d_B q \sigma_{\text{input}}^2,$$

where σ_{input}^2 is the error variance of the input RLWE' ciphertext.

Proof. If one considers a balanced base- B digit decomposition of $r \in \mathcal{R}_{\text{reg}}$ which consists of d_B ring elements r_1, \dots, r_{d_B} whose coefficients are each uniform in $[-B/2, B/2)$, then the variance of the gadget decomposition of the input \mathcal{R}_{reg} ring element satisfies $\sigma_r^2 = B^2/12$. By replacing this value in the upper bound for σ_{\odot}^2 given in Theorem 1, we get $\sigma_{\odot}^2 \leq \frac{B^2}{6} d_B q \sigma_{\text{input}}^2$. \square

Remark 1. Note that the variance σ_{\odot}^2 considered for error analysis in power-of-2 cyclotomic is $\sigma_{\odot}^2 = \frac{B^2}{12} d_B N \sigma_{\text{input}}^2$, (see [15, Section 4.2]), where N is a power of two and the $2N^{\text{th}}$ cyclotomic ring is considered. Interestingly, our analysis for prime cyclotomic rings only shows a difference by a factor 2.

³ In general, it holds that $\text{cov}(XY, XZ) = E(X^2)E(Y)E(Z) - E(X)^2 E(Y)E(Z)$ for any random variables X, Y and Z . Therefore, if $E(Y) = E(Z) = 0$, then $\text{cov}(XY, XZ) = 0$.

Error Growth in Previous Operations We now describe the error growth for the main operations used in our algorithm as a function of σ_{\odot}^2 .

RGSW \times RLWE' multiplication: Recall that a multiplication between $\text{RLWE}'_{\mathbf{sk}}(\mathbf{m}_1)$ and $\text{RGSW}_{\mathbf{sk}}(\mathbf{m}_2) = (\text{RLWE}'_{\mathbf{sk}}(\mathbf{m}_2), \text{RLWE}'_{\mathbf{sk}}(\mathbf{sk} \cdot \mathbf{m}_2))$ is computed as $\mathbf{a} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{sk} \cdot \mathbf{m}_2) + \mathbf{b} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{m}_2)$ for each RLWE component (\mathbf{a}, \mathbf{b}) of $\text{RLWE}'_{\mathbf{sk}}(\mathbf{m}_1)$. From this description, we easily see that two \odot computations are performed to which should be added the error coming from the RGSW ciphertext itself multiplicatively. Finally, as already mentioned when describing the operation, the error also additively increases by $\mathbf{e}_{\text{RLWE}'} \cdot \mathbf{m}_2$. Since \mathbf{m}_2 is a monomial, we simply add $\sigma_{\text{RLWE}'}^2$. Therefore, the total error variance is equal to $2\sigma_{\odot, \text{RGSW}}^2 + \sigma_{\text{RLWE}'}^2$ where $\sigma_{\odot, \text{RGSW}}^2$ denotes the resulting error variance of the \odot operation on each RLWE' component of $\text{RGSW}(\mathbf{m}_2)$.

RLWE'-to-RGSW Scheme Switching: Recall that the operation can be described as $(\mathbf{b}, 0) + \mathbf{a} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}^2)$ for each RLWE component (\mathbf{a}, \mathbf{b}) of the input RLWE' ciphertext. There are two sources of error. Firstly, an additive error growth comes from the \odot operation in $\mathbf{a} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}^2)$. Since $\text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}^2)$ is a fresh encryption that comes from the evaluation key, the error variance is relatively small. We thus have an additive error growth with variance $\sigma_{\odot, \text{eval_key}}^2$ which denotes the resulting error variance of the \odot operation on the evaluation key with error variance $\sigma_{\text{eval_key}}^2$.

Secondly, a multiplicative error growth comes from the fact that the existing error in the RLWE' ciphertext gets scaled by \mathbf{sk} . The secret key \mathbf{sk} is not a scalar but rather a ring element and recall we work in a prime cyclotomic. We know the error variance scales by a factor of no more than $\ell_1(\mathbf{sk})$, where the norm is with respect to the canonical embedding for \mathbf{sk} .

Combining these two sources of error under the assumption that each coefficient of \mathbf{sk} is binary/ternary, then the error variance of the output is $\ell_1(\mathbf{sk}) \cdot \sigma_{\text{RLWE}'}^2 + \sigma_{\odot, \text{eval_key}}^2$.

RLWE' Automorphism: Applying an automorphism ψ itself does not change the error. The following key-switching operation however introduces an additive error growth with variance $\sigma_{\odot, \text{aut_key}}^2$ which denotes the resulting error variance of the \odot operation on the automorphism key $\text{RLWE}'_{\mathbf{sk}}(\psi(\mathbf{sk}))$.

We summarize these error growth in Table 2.

Table 2. Summary of register operations with $\mathcal{R} \odot \text{RLWE}'$ opcount and error growth.

Operation	Computation (for each RLWE (\mathbf{a}, \mathbf{b}) of RLWE')	\odot ops	Error Variance
RLWE' \times RGSW	$\mathbf{a} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{s} \cdot \mathbf{m}_2) + \mathbf{b} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{m}_2)$	$2k$	$2\sigma_{\odot, \text{RGSW}}^2 + \sigma_{\text{RLWE}'}^2$
SchemeSwitch	$\mathbf{a} \odot \text{RLWE}'_{\mathbf{sk}}(\mathbf{sk}^2) + (\mathbf{b}, 0)$	k	$\ell_1(\mathbf{sk}) \cdot \sigma_{\text{RLWE}'}^2 + \sigma_{\odot, \text{eval_key}}^2$
Automorphism	$\psi(\mathbf{a}) \odot \text{RLWE}'_{\mathbf{sk}}(\psi(\mathbf{sk})) + (0, \psi(\mathbf{b}))$	k	$\sigma_{\text{RLWE}'}^2 + \sigma_{\odot, \text{aut_key}}^2$

4 Description of the Algorithm

The overall algorithm, at a high level, can be subdivided into various steps:

- Step 1: a *packing* step takes as input $\phi(d)$ LWE ciphertexts and “combines” them into a single RLWE ciphertext $(\mathbf{a}, \mathbf{b}) \in \mathcal{R}_{in} \times \mathcal{R}_{in}$.

- Step 2: a *homomorphic decryption* of the RLWE ciphertext consists in computing (an encryption of) the ring element $(\mathbf{a} \cdot \mathbf{z} + \mathbf{b}) \in \mathcal{R}_{in}$, homomorphically, given (as a bootstrapping key) an encryption of \mathbf{z} .
- Step 3: an *msbExtract* step recovers the $\phi(d)$ LWE ciphertexts with reduced noise.

Step 1 and Step 3, except for the use of different rings, are very similar to previous work [21]. We describe the 3 steps in detail with a particular emphasis on Step 2 which is the main novelty of this paper.

4.1 Packing

The very first step of bootstrapping procedure consists in taking a set of LWE ciphertexts and pack them into a single RLWE ciphertext. More precisely, the packing algorithm takes as input $\phi(d)$ LWE ciphertexts encrypting messages $m_i \in \mathbb{Z}$ as well as an RLWE' encryption $\text{RLWE}'(s_i)$ of each coefficients of the plain LWE secret key $\mathbf{s}_p = (s_0, s_1, \dots, s_{n_{plain}}) \in \mathbb{Z}_{q_{plain}}^{n_{plain}+1}$, in the d th cyclotomic ring with modulus q_{plain} and outputs $(\mathbf{a}, \mathbf{b}) \in \mathcal{R}_{in} \times \mathcal{R}_{in}$ encrypting the message $\mathbf{m}(X) = \sum_i m_i X^{i-1}$. The pseudo-code is given in Algorithm 1.

Algorithm 1 Ring packing

Input: $\phi(d)$ plain LWE ciphertexts $(\vec{a}_i, b_i) \in \mathbb{Z}_{q_{plain}}^{n_{plain}} \times \mathbb{Z}_{q_{plain}}$, $\text{RLWE}'(s_i)$

Output: RLWE ciphertext in \mathcal{R}_{in} .

for $0 \leq i < n_{plain}$ **do**

let $r_i = a_{0,i} + a_{1,i}X + a_{2,i}X^2 + \dots + a_{\phi(d)-1,i}X^{\phi(d)-1}$ in $(\mathcal{R}_{in})_{q_{plain}}$.

end for

$r' = (0, (b_0 + b_1X + b_2X^2 + \dots + b_{\phi(d)-1}X^{\phi(d)-1}))$

▷ (Noiseless RLWE' ciphertext)

$ct \leftarrow r' + \sum_{i=0}^{n_{plain}-1} r_i \odot \text{RLWE}'(s_i)$

return $\text{ModSwitch}_{q_{plain} \rightarrow q}(ct)$

For simplicity, we first built a ring ciphertext ct modulo q_{plain} (i.e., the original input modulus) and then switch the modulus to q . Alternatively, one can directly compute a ring ciphertext modulo q by using a packing key $\{\text{RLWE}'(s_i)\}_i$ already encrypted under modulus q . The packing key may also use a different gadget (e.g., the power-of-two gadget, instead of powers-of-B) than other ciphertexts used later in the algorithm.

Since this part of the algorithm is essentially identical to previous work [21], we omit these details, and move on to the second step.

4.2 Linear Step

This step of the algorithm takes as input a single RLWE ciphertext $(\mathbf{a}, \mathbf{b}) \in \mathcal{R}_{in}^2$ and outputs $\phi(d)$ RLWE ciphertexts, each encrypting a coefficient of $(\mathbf{a} \cdot \mathbf{z} + \mathbf{b}) \in \mathcal{R}_{in}$ (recall that an element of \mathcal{R}_{in} is a polynomial of degree $\phi(d)$). It can be further subdivided into two computations: a (homomorphic) polynomial multiplication between \mathbf{a} and (an encryption of) \mathbf{z} , where each coefficient of the polynomials (describing the key \mathbf{z} , all intermediate results, and the final ring element) is a distinct ciphertext, and the addition of the ring element \mathbf{b} . We now provide a detailed explanation of these computations along with a pseudo-code of the various steps of the algorithm.

An FFT-based Polynomial Multiplication. For this step, the algorithm uses a standard FFT-based method summarized in Figure 1. More precisely, we perform the following steps:

1. Compute a partial FFT of $\mathbf{a} \in \mathcal{R}_{in}$, i.e., $\text{PFT}(\mathbf{a})$ in cleartext form. Let $k - 1$ be the degree of the polynomials outputted by PFT. Note that a full (non-partial) FFT would have $k = 1$ as the algorithm recurses until the input polynomial is reduced modulo all $\phi(d)$ linear factors of $\Phi_d(X)$. When computing

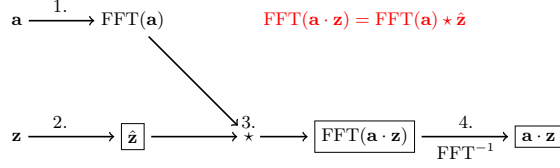


Fig. 1. High level description of the linear step of our algorithm. The notation \star refers to pointwise multiplication. The boxed information refers to encrypted data where homomorphic operations are required. Each step i . is described in detail in the paper.

PFT, the algorithm outputs $\phi(d)/k$ polynomials of degree $k - 1$ (and hence does not recurse all the way down to the linear factors). In other words, this corresponds to evaluating the CRT isomorphism

$$\frac{\mathbb{Z}_q[X]}{(X^{d/2} + 1)} \simeq \left(\frac{\mathbb{Z}_q[X]}{(X^k - \zeta_0)} \right) \times \cdots \times \left(\frac{\mathbb{Z}_q[X]}{(X^k - \zeta_{\phi(d)/k-1})} \right)$$

where the ζ_i are the solutions to $(\zeta^k)^{\phi(d)} = -1$, namely the primitive (d/k) -th roots of unity modulo q . This step thus outputs a list of $\phi(d)/k$ polynomials $\{\tilde{\mathbf{a}}_i\}_{0 \leq i < \phi(d)/k}$, where each $\tilde{\mathbf{a}}_i = \mathbf{a} \bmod (X^k - \zeta_i)$ is a polynomial with k coefficients. Note that this computation is done in the clear, and thus no homomorphic operations are needed.

Recall that when computing an inverse PFT, one must divide the polynomials $\tilde{\mathbf{a}}_i$ by $\phi(d)/k \pmod{q}$. In order to be able to compute this division in the clear rather than homomorphically, this step can be done now (Refer to Section 2.5). Hence the polynomials are updated to $\tilde{\mathbf{a}}_i \leftarrow \tilde{\mathbf{a}}_i / (\phi(d)/k) \pmod{q}$.

2. The evaluation key contains RGSW registers of $\text{PFT}(\mathbf{z})$. Similarly as before, let $\tilde{\mathbf{z}}_i = \mathbf{z} \bmod (X^k - \zeta_i)$, where each $\tilde{\mathbf{z}}_i$ is a polynomial with k coefficients. Let $\tilde{z}_i^{(j)}$ be the j^{th} coefficient of $\tilde{\mathbf{z}}_i$. Then the evaluation key contains the list of RGSW $(X^{\tilde{z}_i^{(j)}})$ for $0 \leq i < \phi(d)/k$ and $0 \leq j < k$.

3. We now want to homomorphically compute $\text{PFT}(\mathbf{a} \cdot \mathbf{z})$ from $\text{PFT}(\mathbf{a})$ and the RGSW registers of $\text{PFT}(\mathbf{z})$. Note that the polynomial multiplication in \mathcal{R}_{in} corresponds to component-wise multiplication in PFT representation, *i.e.*, $\text{PFT}(\mathbf{a} \cdot \mathbf{z}) = (\tilde{\mathbf{a}}_0 \cdot \tilde{\mathbf{z}}_0, \tilde{\mathbf{a}}_1 \cdot \tilde{\mathbf{z}}_1, \dots, \tilde{\mathbf{a}}_{\phi(d)/k-1} \cdot \tilde{\mathbf{z}}_{\phi(d)/k-1})$. For ease of notation, let us fix i (we drop the subscript i) and consider a single multiplication of $\tilde{\mathbf{a}} := \sum_{j=0}^{k-1} \tilde{a}_j X^j$ and $\tilde{\mathbf{z}} := \sum_{j=0}^{k-1} \tilde{z}_j X^j$ modulo $(X^k - \zeta)$. More precisely, we want to homomorphically compute

$$(\tilde{a}_0 + \tilde{a}_1 X + \cdots + \tilde{a}_{k-1} X^{k-1})(\tilde{z}_0 + \tilde{z}_1 X + \cdots + \tilde{z}_{k-1} X^{k-1}) \bmod (X^k - \zeta).$$

where each coefficient \tilde{z}_j is encrypted as an RGSW register.

Each coefficient of the resulting product can be computed as follows. For $j = 0, \dots, k - 1$, the j -th coefficient of $\mathbf{v} := \tilde{\mathbf{a}} \cdot \tilde{\mathbf{z}}$ is equal to

$$v_j = \tilde{z}_0 \tilde{a}_j + \tilde{z}_1 \tilde{a}_{j-1} + \cdots + \tilde{z}_{j-1} \tilde{a}_1 + \tilde{z}_j \tilde{a}_0 + \zeta (\tilde{z}_{j+1} \tilde{a}_{k-1} + \tilde{z}_{j+2} \tilde{a}_{k-2} + \cdots + \tilde{z}_{k-1} \tilde{a}_{j+1}),$$

which corresponds to the inner product taken between the vector of coefficients $\tilde{\mathbf{z}} = (\tilde{z}_0, \dots, \tilde{z}_{k-1})$ of the polynomial $\tilde{\mathbf{z}}$ and the new vector $\tilde{\mathbf{c}} = (\tilde{a}_j, \tilde{a}_{j-1}, \dots, \tilde{a}_0, \zeta \tilde{a}_{k-1}, \dots, \zeta \tilde{a}_{j+1})$. We emphasize again the fact that the coefficients of $\tilde{\mathbf{c}}$ are in the clear, whereas the coefficients of $\tilde{\mathbf{z}}$ are not. So, it is easy to multiply $\tilde{\mathbf{c}}$ by ζ .

Without loss of generality, let us assume all the coefficients c_i of $\tilde{\mathbf{c}}$ are nonzero and thus invertible. (Here we use the fact that q is a prime. So, all nonzero elements are invertible modulo q and multiplication (in the exponent) can be implemented using an automorphism of the prime cyclotomic ring.) Then we can compute the inner product in a telescoping manner as

$$v_j = ((\dots ((\tilde{z}_0 c_1^{-1} + \tilde{z}_1) c_1 c_2^{-1} + \tilde{z}_2) c_2 c_3^{-1} + \dots) c_{k-2} c_{k-1}^{-1} + \tilde{z}_{k-1}) c_{k-1}.$$

This will end up being the most efficient way to compute this inner product homomorphically.

Let us now explicit how one coefficient corresponding to a monomial X^j can be computed homomorphically (this computation will have to be repeated for all k coefficients of a single product as well as for all $\phi(d)/k$ pairs of $(\tilde{\mathbf{a}}_i, \tilde{\mathbf{z}}_i)$ polynomials).

- a) Let *accum* be an RLWE' register, initialized as RLWE' ($X^{\tilde{z}_0}$) from the evaluation key.
- b) For $j' \in [0, \dots, k-2]$, update *accum* as follows:
 - i. Apply the automorphism that sends X to $X^{c_j' c_{j'+1}^{-1}}$.
 - ii. Do an RLWE' \times RGSW multiplication with RGSW ($X^{\tilde{z}_{j'+1}}$) from the evaluation key.
- c) Finally apply the automorphism sending X to $X^{c_{k-1}}$, yielding RLWE' (X^{v_j}).

Since we repeat (a)-(c) for every coefficient of $\tilde{\mathbf{a}}_i \cdot \tilde{\mathbf{z}}_i$ for $1 \leq i \leq \phi(d)/k$, the output of this step consists of $\phi(d)$ RLWE' registers of the form

$$\left\{ \text{RLWE}' \left(X^{\mathbf{v}_i^{(j)}} \right) \right\}_{0 \leq i < \phi(d)/k, 0 \leq j < k},$$

where $\mathbf{v}_i^{(j)}$ denotes the j -th coefficient of $\mathbf{v}_i := \tilde{\mathbf{a}}_i \cdot \tilde{\mathbf{z}}_i \pmod{X^k - \zeta_i}$. This procedure is illustrated in Figure 2, and the corresponding pseudo-code for component-wise multiplication is given in Algorithm 2.

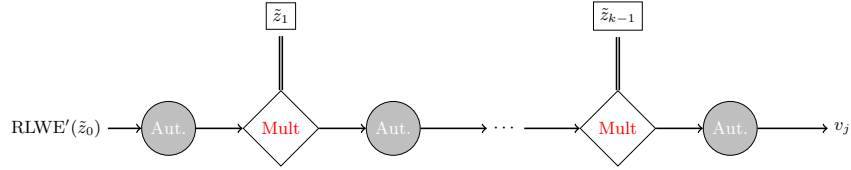


Fig. 2. Homomorphic computation of a X^j coefficient for pointwise multiplication. A single line corresponds to RLWE' ciphertexts and a double line to RGSW ciphertexts. Aut. stands for automorphisms and Mult. for multiplication. Boxed values are encrypted values.

4. We now have the encryption of $\text{PFT}(\mathbf{a} \cdot \mathbf{z})$. It remains to perform the inverse of PFT, denoted by PFT^{-1} , in order to recover the resulting polynomial product $\mathbf{a} \cdot \mathbf{z}$, more specifically RLWE encryptions of the coefficients of $\mathbf{a} \cdot \mathbf{z}$.

Recall from Section 2.5 that the inverse of a primitive FFT of length N (using a $2N$ th root of unity ω) can be computed by first taking a standard FFT of length N using ω^{-2} as the N th root of unity, then multiplying the i th term by ω^{-i} . Moreover, a partial FFT of length $\phi(d)$ that reduces modulo $(X^k - \zeta)$ is equivalent to k full FFTs of length $N = \phi(d)/k$ done in parallel, and the same remains true for the inverse (see section 2.5 for details about the equivalence). Hence, to homomorphically compute PFT^{-1} , we will

- a) Split the $\phi(d)$ registers output by the pointwise multiplication step into k groups of size N : each group corresponds to the coefficients of the monomial X^j for $0 \leq j \leq k-1$ of all $\phi(d)/k$ polynomials.
- b) Homomorphically perform a standard (not primitive) length- N FFT in the forward direction on each group of size N , using ω^{-2} as the N th root of unity. Overall, this step corresponds to computing k FFTs. We refer the reader to Section 2.5 for more details on the equivalence between partial FFT modulo $X^k - \zeta$ and k -parallel standard FFT.
- c) Multiply (homomorphically, via an automorphism) the i th output register in each group by ω^{-i} , for all i from 1 to $N-1$. More specifically, we apply the automorphism $X \mapsto X^{\omega^{-i}}$, corresponding to multiplication by ω^{-i} in the exponent, followed by a key switching operation.

The following algorithms provide pseudocodes for the above procedure to compute the homomorphic PFT^{-1} . More specifically, Algorithm 3 describes step (a) and then calls Algorithm 4 for each of the groups of

Algorithm 2 Pointwise multiplication between polynomials $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{z}}$.

1: **Input:** A set of degree- $(k - 1)$ polynomials $\{\tilde{\mathbf{a}}_i\}_{0 \leq i < N}$, $\{\text{RGSW}(X^{\tilde{z}_i^{(j)}})\}_{0 \leq i < N, 0 \leq j < k}$ for $N := \phi(d)/k$, ω : the $2N$ -th root of unity mod q
2: **Output:** $\phi(d)$ RLWE' ciphertexts
3: $\text{REG} \leftarrow [0, \dots, 0]$
4: **for all** $0 \leq i < N$ **do**
5: $\zeta \leftarrow \omega^{2i+1}$
6: Let $\tilde{\mathbf{a}} = (\tilde{a}_{i,0}, \tilde{a}_{i,1}, \dots, \tilde{a}_{i,k-1})$ $\triangleright \tilde{\mathbf{a}}_i := \sum_{j=0}^{k-1} \tilde{a}_{i,j} X^j$
7: **for all** $0 \leq j < k$ **do**
8: $\tilde{\mathbf{c}} \leftarrow (\tilde{a}_{i,j}, \tilde{a}_{i,j-1}, \dots, \tilde{a}_{i,0}, \zeta \tilde{a}_{i,k-1}, \zeta \tilde{a}_{i,k-2}, \dots, \zeta \tilde{a}_{i,j+1})$
9: $\text{accum} \leftarrow \text{RLWE}'(X^{\tilde{z}_i^{(0)}})$
10: **for** $j' \leftarrow 0, 1, \dots, k - 2$ **do**
11: $\text{accum} \leftarrow \text{EvalAut}(\text{accum}, c_{j'} c_{j'+1}^{-1})$
12: $\text{accum} \leftarrow \text{MulRGSW}(\text{RGSW}(\tilde{z}_i^{(j'+1)}), \text{accum})$
13: **end for**
14: $\text{accum} \leftarrow \text{EvalAut}(\text{accum}, c_{k-1})$ $\triangleright \text{accum} = \text{RGSW}(X^{(\tilde{\mathbf{a}}_i \cdot \tilde{\mathbf{z}}_i)^{(j)})}$
15: $\text{REG}[ik + j] \leftarrow \text{accum}$
16: **end for**
17: **end for**
18: **output** REG \triangleright Register of all coefficients of $\tilde{\mathbf{a}}_i \cdot \tilde{\mathbf{z}}_i$ for $0 \leq i < N$

Algorithm 3 IFFT stage of bootstrapping (BootstrapIFFT)

1: **Input:** a list of $\phi(d)$ registers REG , k , N , a list of radices $\{r_i\}_{0 \leq i < \ell}$, and ω
Require: ω a primitive $2N$ th root of unity mod q , $\prod_{0 \leq i < \ell} r_i = N$, and $kN = \phi(d) = \text{len}(\text{REG})$
2: **for all** $0 \leq j < k$ **do**
3: $\text{REG}[j, k + j, \dots, (N - 1)k + j] \leftarrow \text{N-IFFT}(\text{REG}[j, k + j, \dots, (N - 1)k + j], \{r_i\}, \omega)$
4: **end for**

Algorithm 4 Primitive length- N IFFT for a single group of size N (N-IFFT)

1: **Input:** List of RLWE' registers $\text{REG} = \{\text{RLWE}'(X^{(\tilde{\mathbf{a}}_i \cdot \tilde{\mathbf{z}}_i)^{(j)})}\}_{0 \leq i < N}$ for some fixed $0 \leq j < k$, list of radices $\{r_i\}_{0 \leq i < \ell}$, primitive $2N$ th root of unity ω modulo q .
2: **Output:** list of RLWE registers $\text{REG} = \{\text{RLWE}(X^{(\mathbf{a}_i \cdot \mathbf{z}_i)^{(j)})}\}_{0 \leq i < N}$.
3: **let** $\omega' = \omega^{-2}$
4: $\text{REG} \leftarrow \text{FFT}(\text{REG}, \{r_i\}_{0 \leq i < \ell}, \omega')$ \triangleright Step 4-(b)
5: **for** $i \leftarrow 1, \dots, N - 1$ **do**
6: $\text{REG}[i] \leftarrow \text{EvalAut}(\text{REG}[i], \omega^{-i})$ \triangleright Step 4-(c)
7: **end for**
8: **return** REG

registers. Algorithm 4 describes a primitive length- N (inverse) FFT for a single group of size N , consisting of a standard (cyclic) FFT (step (b)) as well as the multiplication by ω^{-i} (step (c)). It remains to describe more precisely what happens in the (cyclic) FFT call, line 4 of Algorithm 4.

Recall that FFT is a recursive algorithm that follows the structure of a remainder tree, see the procedure FFT given in Algorithm 6. We will now focus on what happens in a single layer of the FFT as described in the second procedure **FFT Layer** in Algorithm 6.

At a single layer: Let r_i be the radix used for the i -th FFT layer for $0 \leq i < \ell$. Then, for $R_i := \prod_{i \leq i' < \ell} r_{i'}$, the inputs to the i -th FFT layer are the coefficients of N/R_i polynomials modulo $(X^{R_i} - \omega'^j)$ (for varying values of j), each with R_i coefficients.⁴ Hence this corresponds to a total of N coefficients, *i.e.*, N registers. The outputs are the coefficients of N/R_{i+1} polynomials modulo $(X^{R_{i+1}} - \omega'^{j'})$ ranging over all j' such that $r_i \cdot j' \equiv j \pmod{N}$. Note that the total number of coefficients remains the same, *i.e.*, we still have N registers.

Let us now consider a single input polynomial (out of N/R_i), *i.e.*, one of the nodes in the remainder tree, and describe what computations are needed to produce the children nodes. This subroutine is described in Algorithm 5, called **FFT Subroutine** and is repeated for every node (meaning polynomial) of the layer, hence N/R_i times. We illustrate the reduction of this polynomial via an example to better describe the operations needed in this subroutine.

Algorithm 5 FFT Subroutine

```

1: Input: Radix  $r$  which divides  $R$  for  $R \mid N$ , index  $j$ , and RLWE' ciphertexts  $\{ct_i\}_{0 \leq i < R}$  storing coefficients of a
   single polynomial mod  $(x^R - \omega'^j)$ 
2: Output:  $r$  tuples each of which consists of index  $j'$  such that  $r \cdot j' \equiv j \pmod{N}$ , and  $R/r$  RLWE' ciphertexts
3: for all  $0 \leq i < R - R/r$  do
4:    $ct_i \leftarrow \text{SwitchToRGSW}(ct_i)$ 
5: end for
6: if this is the final FFT layer then
7:   for all  $R - R/r \leq i < R$  do
8:     let  $S = \frac{Q}{4}$  ▷ 4 is the plaintext modulus after bootstrapping
9:      $ct_i \leftarrow S \odot ct_i$  ▷  $ct_i$  is now RLWE instead of RLWE'
10:   end for
11: end if
12: let  $\{j'_0, \dots, j'_{r-1}\} =$  the set of all  $j'$ 's satisfying  $rj' \equiv j \pmod{N}$ 
13: for all  $0 \leq v \leq r - 1$  do
14:   let  $\zeta = \omega'^{j'_v}$ 
15:   for all  $0 \leq i < R/r$  do
16:      $accum[v][i] \leftarrow ct_{R - R/r + i}$ 
17:     for  $\kappa \leftarrow [2, 3, \dots, r]$  do
18:        $accum[v][i] \leftarrow \text{EvalAut}(accum, \zeta)$ 
19:        $accum[v][i] \leftarrow \text{MulRGSW}(ct_{R - \kappa \cdot R/r + i}, accum)$ 
20:     end for
21:   end for
22: end for
23: output  $r$  tuples  $(j'_v, accum[v])$  for  $0 \leq v < r$ 

```

Example 1. We describe in this example the reduction from an input polynomial to a single child node for the simple radix-2 FFT. Assume we have as input a polynomial of the form

$$g_0 + g_1X + g_2X^2 + g_3X^3 + g_4X^4 + g_5X^5 + g_6X^6 + g_7X^7$$

⁴ When $i = 0$, it starts with a single input polynomial modulo $X^N - \omega^0$.

Algorithm 6 Full radix- r standard FFT (FFT)

```
1: procedure FFT( $\{\text{REG}\}, \{r_i\}_{0 \leq i < \ell}, \omega'$ )
2:    $state \leftarrow \{(N, \text{REG})\}$  ▷ List of tuples
3:   for  $i$  in  $[0, 1, \dots, \ell - 1]$  do
4:      $state \leftarrow \text{FFT\_LAYER}(r_i, \omega', state)$ 
5:   end for
6:   return REG
7: end procedure

8: procedure FFT_LAYER( $r_i, \omega', \text{list of tuples}$ )
9:   Input:  $N/R_i$  tuples of the form  $(j, \{ct_{j,0}, \dots, ct_{j,R_i-1}\})$ , where each  $ct_{j,v}$  is an RLWE' ciphertext
10:    ▷  $ct_{j,v}$  represents the  $v$ -th coefficient of a polynomial mod  $(X^{R_i} - \omega'^j)$ 
11:   Output:  $N/R_{i+1}$  tuples of the form  $(j', \{ct_{j',0}, \dots, ct_{j',R_{i+1}-1}\})$ .
12:    ▷ Each input  $j$  has  $r_i$  corresponding outputs  $j'$  such that  $r_i \cdot j' \equiv j \pmod{N}$ .
13:    ▷  $ct_{j',v}$  represents the  $v$ -th coefficient of a polynomial mod  $(x^{R_{i+1}} - \omega'^{j'})$ 
14:   for all  $(j, \{ct_{j,0}, \dots, ct_{j,R_i-1}\})$  in input do
15:     FFT_SUBROUTINE( $\{ct_{j,0}, \dots, ct_{j,R_i-1}\}, r_i, j, \omega'$ )
16:   end for
17: end procedure
```

and we want to reduce it modulo $(X^2 - \zeta)$. Similarly as for pointwise multiplication, it is possible to compute the coefficient terms for each monomial X^j . In our example, we would have constant coefficient $g_0 + \zeta g_2 + \zeta^2 g_4 + \zeta^3 g_6$ and X coefficient $g_1 + \zeta g_3 + \zeta^2 g_5 + \zeta^3 g_7$. In the remainder tree, this operation would have to be repeated for r different values of ζ , in particular for this example, four different values.

The homomorphic circuit to perform this reduction is illustrated in Figure 3. We recall that the input coefficients g_i (both in the example and in Figure 3) correspond to registers, in particular RLWE' ciphertexts. The main operations needed for a reduction are scheme-switching for most of the coefficients, multiplication by a power of ζ , which can be done using automorphisms, and addition which corresponds to RLWE' \times RGSW multiplications.

Remark 2. The scheme switches at the beginning of the circuit convert RLWE' ciphertexts to RGSW ciphertexts for all coefficients except the last. As mentioned previously, the circuit for the same coefficients is performed for various values of ζ . For all these cases, the scheme-switching operations need only to be performed once (as the coefficients do not change) and thus the cost is amortized.

Details of this subroutine are given in Algorithm 5. Algorithm 6 provides the pseudocode for all $\phi(d)$ registers (**FFT Layer**) as well as the full FFT algorithm where all layers are considered (**FFT**).

One can note from Algorithm 5 that the case of the last layer of the FFT slightly differs (see line 6). Indeed, it is possible to optimize the running-time of the FFT algorithm by modifying the nature of the elements considered in the very last layer of the FFT. Indeed, one can notice that the outputs of the IFFT only need to be RLWE ciphertexts, not RLWE' ciphertexts. Hence, one can save operations by using RLWE registers instead of RLWE' registers when possible. While RLWE cannot be scheme-switched to RGSW without blowing up the error, we can modify the last IFFT layer to use RLWE instead of RLWE' for the registers that do not get scheme-switched (this corresponds to *accum* in Algorithm 5 or g_r in Figure 3). Concretely, each of the $\phi(d)/r$ non-scheme-switched RLWE' ciphertexts would be converted to RLWE by an $\mathcal{R} \odot \text{RLWE}'$ operation with \mathcal{R} element $\lceil Q/4 \rceil$, where 4 is the plaintext modulus the *msbExtract* stage expects. This concludes the description of the homomorphic computation of $\mathbf{a} \cdot \mathbf{z}$. The output of this multiplicative step is thus $\phi(d)$ RLWE registers, each encrypting a coefficient of $\mathbf{a} \cdot \mathbf{z} \in \mathcal{R}_{in}$.

Adding \mathbf{b} From the previous step, we have obtained registers encoding the coefficients of $\mathbf{a} \cdot \mathbf{z}$. We also have the polynomial \mathbf{b} in the clear. In order to obtain registers encoding the coefficients of $\mathbf{a} \cdot \mathbf{z} + \mathbf{b}$, we

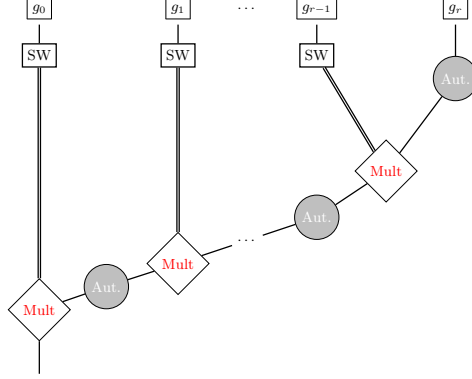


Fig. 3. One layer of FFT for a single input polynomial. A single line corresponds to RLWE ciphertexts and a double line to RGSW ciphertexts. SW stands for scheme-switching, Aut. for automorphisms and Mult for multiplication.

add \mathbf{b} via fixed rotations, *i.e.*, scaling the ciphertext by a monomial. Concretely, to add a coefficient b_i to a register RLWE($X^{\mathbf{a}\cdot\mathbf{z}}^i$), we simply scale the RLWE ciphertext by X^{b_i} resulting in RLWE($X^{\mathbf{a}\cdot\mathbf{z}+\mathbf{b}}^i$). Since X^{b_i} has norm 1, it does not increase the noise of the register. We thus now have $\phi(d)$ registers encoding the coefficients of $\mathbf{a} \cdot \mathbf{z} + \mathbf{b}$, as expected. This concludes the linear step of the algorithm.

4.3 msbExtract

The linear step outputs $\phi(d)$ RLWE registers each encrypting $\frac{Q}{4}X^{c_i}$, where each $c_i \in \mathbb{Z}_q$ is a noisy (unrounded) decryption of the i th input ciphertext. We now operate separately on each register (and drop the subscript i) in order to recover from each register a plain-LWE encryption of $f(c)$, for some function f that applies rounding and allows for computation. Output ciphertexts have plaintext modulus 4; to compute NAND gates, it suffices for $f(c)$ to be 1 for $c \in [-q/8, 3q/8]$ and 0 elsewhere (we refer to [20] for more details). Focusing on a single register $(\mathbf{a}, \mathbf{b}) \in \mathcal{R}_{reg}^2$, we have

$$\mathbf{b}(X) = -\mathbf{a}(X) \cdot \mathbf{s}(X) + \mathbf{e}(X) + \frac{Q}{4}\mathbf{m}(X) \pmod{Q, \Phi_q(X)}$$

where $\mathbf{m}(X) = X^c$. Looking at a single coefficient of these polynomials, the ring product $\mathbf{a}(X) \cdot \mathbf{s}(X)$ will become a vector inner product between the coefficients of \mathbf{s} and some permuted coefficients of \mathbf{a} . Because we use prime q , these polynomials have degree $\leq q-2$, and $X^{q-1} \equiv -1 - X - \dots - X^{q-2}$. Note that, as for error growth, the fact that we consider prime cyclotomics instead of power-of-2 cyclotomics slightly changes the setting. We get for a single coefficient

$$\begin{aligned} b_i &= -a_i s_0 - a_{i-1} s_1 - \dots - a_0 s_i - 0 \cdot s_{i+1} - a_{q-2} s_{i+2} - a_{q-3} s_{i+3} - \dots - a_{i+2} s_{q-2} \\ &\quad + a_{q-2} s_1 + a_{q-3} s_2 + \dots + a_2 s_{q-3} + a_1 s_{q-2} + \frac{Q}{4} m_i + e_i \end{aligned}$$

which can be re-written as

$$\begin{aligned} \frac{Q}{4} m_i + e_i &= b_i + \{(a_i, a_{i-1}, \dots, a_0, 0, a_{q-2}, a_{q-1}, \dots, a_{i+2}) \\ &\quad - (0, a_{q-2}, a_{q-1}, \dots, a_1)\} \cdot (s_0, s_1, \dots, s_{q-2}) \end{aligned}$$

For $0 \leq i \leq q-2$, letting \vec{a}_i denote the above vector $(a_i, a_{i-1} - a_{q-2}, \dots, a_{i+2} - a_1)$, we then have that (\vec{a}_i, b_i) is an LWE encryption with noise e_i under secret key $\mathbf{s}_p = (s_0, \dots, s_{q-2})$ of message m_i . Note that

m_i is 1 if $c = i$, -1 if $c = q - 1$, and 0 otherwise. To produce an encryption of $f(c)$, which should be 1 for $c \in (-q/8, 3q/8)$ and 0 elsewhere, we simply sum the relevant (\vec{a}_i, b_i) :

$$\sum_{i=\lceil 7q/8 \rceil}^{q-2} (\vec{a}_i, b_i) + \sum_{i=0}^{\lfloor 3q/8 \rfloor - 1} (\vec{a}_i, b_i)$$

taking care to ensure the number of summands is $3 \pmod 4$, so that when $c = q - 1$ the sum is $1 \pmod 4$ as desired. (When $q \equiv 1 \pmod 8$, this will be the case for the summation written above.)

This gives us an LWE encryption with plaintext modulus 4 and ciphertext modulus Q under a key $\mathbf{s}_p \in \mathbb{Z}^{q-1}$. To conclude bootstrapping, we can keyswitch back to the original plain LWE secret key, and modulus switch back down to the original (much smaller than Q) ciphertext modulus.

5 Analysis

To evaluate the performance of our algorithm, we analyse its running-time as well as the error growth. We will first show that our homomorphic decryption procedure takes no more than $\mathcal{O}((k + r \cdot \ell)\phi(d)d_B)$ homomorphic operations.

5.1 Counting Homomorphic Operations

We will evaluate the efficiency of our algorithm by first measuring the time complexity in terms of the number of $\mathcal{R} \odot \text{RLWE}'$ operations performed. We have already summarized in Section 3.2, Table 2 the number of $\mathcal{R} \odot \text{RLWE}'$ operations needed for the main operations used in our scheme: scheme switching, automorphisms (with key switching) and $\text{RGSW} \times \text{RLWE}'$. We now describe the number of $\mathcal{R} \odot \text{RLWE}'$ operations for the various steps of our algorithm.

Pointwise Multiplication Based on the description given in Section 4.2, we have the following analysis. For a single coefficient X^j in the computation of the inner product, our algorithm performed k automorphisms and $(k - 1)$ $\text{RGSW} \times \text{RLWE}'$ multiplications. Hence the number of $\mathcal{R} \odot \text{RLWE}'$ operations per register is $(3k - 2)d_B$. This computation needs to be repeated for all k coefficients of a single product and for $\phi(d)/k$ pairs of $(\vec{\mathbf{a}}_i, \vec{\mathbf{z}})$ polynomials. Thus the total number of $\mathcal{R} \odot \text{RLWE}'$ operations for the entire pointwise multiplication algorithm is $(3k - 2)\phi(d)d_B$.

Partial Inverse FFT Based on the description given in Section 4.2, considering a single register and a single radix- r layer of FFT, the algorithm computes $(r - 1)$ automorphisms, $(r - 1)$ $\text{RGSW} \times \text{RLWE}'$ multiplications but only amortized $(1 - 1/r)$ scheme switches as explained in Section 4.2. Thus the total number of operations per layer is

$$\left((r - 1) + 2(r - 1) + \left(1 - \frac{1}{r}\right) \right) \phi(d)d_B = \left(3r - 2 - \frac{1}{r} \right) \phi(d)d_B.$$

Last layer of IFFT optimization: Recall that the outputs of the IFFT only need to be RLWE ciphertexts and not RLWE' ciphertexts. This allowed us to optimize the cost of the last layer of the IFFT by using RLWE registers instead of RLWE' registers when possible. By using this modification, the multiplications and automorphisms in this layer will use a factor of d_B fewer operations. Thus the total number of operations for the last layer is only

$$\left((r - 1) + 2(r - 1) + \frac{1}{r} + \left(1 - \frac{1}{r}\right) d_B \right) \phi(d) = \left(3r - 3 + \frac{1}{r} + d_B \left(1 - \frac{1}{r}\right) \right) \phi(d).$$

After the last layer: Recall that the very last step of the homomorphic partial IFFT is to multiply the i^{th} output register in each group by ω^{-i} , via an automorphism that sends X to $X^{\omega^{-i}}$. Operation-wise, this corresponds to one automorphism per register. Since with the last-layer optimization the registers are RLWE instead of RLWE', this corresponds to $\phi(d)$ operations in total.

Adding \mathbf{b} Recall that to add b_i to the register $\text{RLWE}(X^{(\mathbf{a}\cdot\mathbf{z})^i})$, we simply scaled the RLWE ciphertext by X^{b_i} . No $\mathcal{R} \odot \text{RLWE}'$ operations are involved.

Table 3. Summary of $\mathcal{R} \odot \text{RLWE}'$ operation count.

Steps of the algorithm	$\mathcal{R} \odot \text{RLWE}'$ operations
Partial FFT of \mathbf{a}	–
Pointwise multiplication	$(3k - 2)\phi(d)d_B$
Partial IFFT (per layer)	$(3r - 2 - \frac{1}{r})\phi(d)d_B$
Last layer of IFFT	$(3r - 3 + \frac{1}{r} + d_B(1 - \frac{1}{r}))\phi(d)$
Last IFFT step	$\phi(d)$
Adding \mathbf{b}	–

5.2 Error Growth

We have already summarized in Table 2 the error growth coming from the scheme switching, automorphisms (with key switching) and $\text{RGSW} \times \text{RLWE}'$ operations. We now describe the error growth resulting from the various steps of our algorithm based on the error variance for each of these operations and the operation count described in the previous section.

Pointwise Multiplication Recall from the description given in Section 4.2 that the algorithm starts with an initial RLWE' ciphertext, denoted as $accum$, which is a “fresh” ciphertext from the evaluation key with error variance $\sigma_{eval_key}^2$. Each automorphism performed during pointwise multiplication adds $\sigma_{\odot, aut_key}^2$ error variance, and each multiplication with a fresh RGSW ciphertext adds $2\sigma_{\odot, \text{RGSW}}^2$ error variance. Hence, the error variance after pointwise multiplication is $(3k - 2)(\sigma_{\odot, aut_key}^2 + 2\sigma_{\odot, \text{RGSW}}^2) + \sigma_{eval_key}^2$.

Inverse FFT Again, recall that each automorphism adds $\sigma_{\odot, aut_key}^2$ error variance. The output of schemeswitching has $\sigma_{sw}^2 = \sigma_{\odot, eval_key}^2 + \ell_1(s)\sigma_{in}^2$ error variance. Let σ_{accum}^2 be initialized as σ_{in}^2 . Each automorphism and $\text{RGSW} \times \text{RLWE}'$ multiplication (performed a total amount of $r - 1$ times) updates the variance as $\sigma_{accum}^2 \leftarrow \sigma_{accum}^2 + \sigma_{\odot, aut_key}^2$ and $\sigma_{accum}^2 \leftarrow \sigma_{accum}^2 + 2\sigma_{\odot, sw}^2$ where $\sigma_{\odot, sw}^2 \leq \frac{B^2}{6}d_Bq\sigma_{sw}^2$. Hence, in total, the error variance after a radix- r layer becomes $\sigma_{in}^2 + (r-1)(\sigma_{\odot, aut_key}^2 + 2\sigma_{\odot, sw}^2)$ for $\sigma_{\odot, sw}^2 \leq \frac{B^2}{6}d_Bq\sigma_{sw}^2$.

After last layer Multiplying the i^{th} output register in each group by ω^{-i} with an automorphism increases (additively) the error variance by $\sigma_{\odot, aut_key}^2$.

Adding \mathbf{b} Scaling by a monomial does not increase the error. This result is summarized in Table 4.

Table 4. Summary of error growth.

Algorithms	Error growth
Partial FFT of \mathbf{a}	–
Pointwise multiplication	$(3k - 2)(\sigma_{\odot, aut_key}^2 + 2\sigma_{\odot, RGSW}^2) + \sigma_{eval_key}^2$.
Partial IFFT (per layer)	$\sigma_{in}^2 + (r - 1)(\sigma_{\odot, aut_key}^2 + 2(\sigma_{\odot, eval_key}^2 + \ell_1(s)\sigma_{in}^2))$
Last IFFT step	$\sigma_{\odot, aut_key}^2$
Adding \mathbf{b}	–

5.3 Asymptotic Analysis

Let $\lambda = O(n)$ be the security level considered. We study the performance of our algorithm as λ increases, *i.e.*, when n tends to infinity. Recall that the other parameters used in our algorithm are $d_B = \lceil \log_B Q \rceil = O(\log n)$, the number of layers ℓ (*i.e.*, the multiplicative depth) and $k = r = \phi(d)^{1/\ell}$ (k is the degree at which we stop the partial FFT and r is the radix for an FFT layer).

Theorem 2. *Let $\phi(d) = O(n)$ be the number of packed ciphertexts and $q, Q = \text{poly}(n)$ the moduli of the rings considered. The total cost of bootstrapping (non-amortized) then corresponds to $O(n^{1+\frac{1}{\ell}} \cdot \log n \cdot \ell)$ homomorphic operations (in terms of the number of $\mathcal{R} \odot \text{RLWE}'$ operations).*

Proof. The number of \odot operations in the pointwise multiplication step is $(3k - 2)\phi(d)d_B$ which asymptotically corresponds to $O(n^{1+\frac{1}{\ell}} \log n)$. Similarly, the inverse FFT requires $(3r - 2 - \frac{1}{r})\phi(d)d_B\ell$ operations (without including the last layer modification which asymptotically does not change the result) which asymptotically gives $O(n^{1+\frac{1}{\ell}} \cdot \log n \cdot \ell)$. \square

Corollary 2. *The amortized cost per message is $O(n^{\frac{1}{\ell}} \cdot \log n \cdot \ell)$ homomorphic operations (in terms of the number of $\mathcal{R} \odot \text{RLWE}'$ operations).*

Remark 3. Note that we can also reduce the total number of homomorphic operations in the case that we only need to pack less-than- $\phi(d)$ LWE ciphertexts, which we refer to sparse packing. Please refer to the full version of the paper [6] for the analysis on computational cost of the sparse packing case.

5.4 Comparison with Previous and Concurrent Work

Our algorithm can be compared to two lines of work: sequential bootstrapping algorithms such as FHEW/TFHE [5, 7] and the amortized bootstrapping algorithms of [21] and [17] (as our algorithm performs asymptotically better than [16], we focus on the comparison with the follow-up work [17]). The asymptotic running times of the algorithms are reported in Table 5, and show that our algorithm is asymptotically much faster than both [5, 7] (reducing the dependency on the main security parameter from linear $O(n)$ to just n^ϵ for arbitrary small $\epsilon = 1/\ell$), and [21] (reducing the dependency on ℓ from exponential 3^ℓ to linear $O(\ell)$.) While [17] is asymptotically faster, we now discuss why our algorithm is expected to outperform [17] for practical parameters.

Assume optimistically that the complexity of [17] is of the form $f(n, \ell) = 2^\ell \cdot \log(n)$, where ℓ is the recursive depth of the inverse FFT and n the ring dimension. The 2^ℓ term is inherited by the use of the Nussbaumer transform from [21], and $\log(n)$ is the “polylog” overhead. [17, section 7.3] sets $\ell = 5$, and, in fact, due to the use of the SIMD parallelization technique from [16], this is the smallest possible value of ℓ . On the other hand, the complexity of our algorithm is $g(n, \ell) = \ell \cdot n^{1/\ell}$, where ℓ can be set to any constant. To facilitate a more direct comparison to [17], consider setting $\ell = 5$ in our algorithm as well. Since noise growth of [17] and our algorithm depend in a similar way on the recursive depth ℓ , using the same value of ℓ should result in similar values of the ciphertext modulus Q and ring dimension n , for the same security level. So, for a fixed $\ell = 5$, basic operations in [17] and our work can be assumed to have roughly the same unit cost. The two algorithms can then be compared by checking when $f(n, 5) < g(n, 5)$. It is easy

to check that the crossover point $f(n, 5) = g(n, 5)$ is given by ring dimension $n = 2^{40}$, which is well beyond any conceivable instantiation of lattice cryptography. While this is only a rough comparison, it should be clear that for any reasonable value of the ring dimension n , our algorithm can be expected to outperform the (asymptotically faster) algorithm of [17]. In other words, [17] suffers from the very same limitations that made [21] completely impractical. This is precisely the problem addressed in our work, where the overhead of [21] is reduced from 2^ℓ to just ℓ , making amortization much closer to practicality. We have already mentioned the work of Liu and Wang [18] who achieved the same asymptotic complexity as [16, 17], but were able to make their bootstrapping scheme also efficient using BFV ciphertexts. As mentioned previously, their work differs by using a super-polynomial modulus. We refer to [18, Section 7] for detailed performance analysis, including timings.

Table 5. Comparing asymptotic cost of various bootstrapping algorithms in terms of homomorphic operations, where ℓ corresponds to the recursive depth in each algorithm. For uniformity with previous work, performance is expressed as the number of RGSW \times RGSW products, or equivalent operations. Alternatively, the number of basic $\mathcal{R} \odot$ RLWE' products can be obtained by multiplying these figures by $O(\log n)$, the length of the gadget vector.

Scheme	Total cost	Number of messages	Amortized cost
FHEW	$\tilde{O}(n)$	1	$\tilde{O}(n)$
TFHE	$O(n)$	1	$O(n)$
[21]	$\tilde{O}(3^\ell \cdot n^{1+1/\ell})$	$O(n)$	$\tilde{O}(3^\ell \cdot n^{1/\ell})$
[17]	$\tilde{O}(n)$	$O(n)$	$\tilde{O}(1)$
[18]	$\tilde{O}(n)$	$O(n)$	$\tilde{O}(1)$
our work	$O(\ell \cdot n^{1+1/\ell})$	$O(n)$	$O(\ell \cdot n^{1/\ell})$

Following previous work, performance in Table 5 is measured as the number of operations on cryptographic registers, and hides many important parameters that still have a big impact on the concrete performance of the algorithms in practice. In order to provide (still preliminary, but) more realistic estimates of the performance of the algorithm, we also evaluated the number of (integer) arithmetic operations required for concrete values of the parameters, with a target security level STD192. Next, we observe that since each NTT operation works on a q dimensional vector, it requires a number of arithmetic operation proportional to $q \log q$. So, in the before-last column of Table 6 we estimate the total number of arithmetic operation as $NTT \times q \log q$. We used [15, Table 1] for the number of \odot operations in FHEW-like schemes and recall that one can easily convert the number of \odot products to the number of NTTs as each \odot product requires $(d_B + 1)$ NTT operations. The reported number of \odot multiplication for [21] comes from a non-public report provided to us by the authors. The before-last column for [21] was obtained using the smallest values of q, d_B in the table as a conservative lower bound. Finally, the last column of Table 6 reports the key size (number of RLWE' ciphertexts). The parameters and numbers provided in Table 6 are tentative, preliminary estimates only meant to provide some intuition about the potential performance of our algorithm. Still, they are enough to draw some general conclusions: our algorithm clearly outperforms previous methods for amortized FHEW bootstrapping, improving [21] by two orders of magnitude. When comparing with sequential bootstrapping methods, our algorithm is practical enough to offer comparable (potentially better) performance. But the improvement is not yet sufficiently marked to make its implementation and use attractive in practice.

6 Conclusion and Future Work

We present a novel amortized bootstrapping algorithm with much smaller overhead than prior amortized work, in particular smaller than in the algorithm presented by Micciancio and Sorrell in [21]. We make use of ring automorphisms to perform the multiplication by twiddle factors and replace the Nussbaumer transform with the more practical Number Theoretic Transform.

In order to properly evaluate the practicality of our algorithm, and accurately compare it with previous work, we considered implementing it within some mainstream FHE library, like OpenFHE [1]. However, the

Table 6. Comparing cost of various bootstrapping algorithm in terms of \odot operations, number of NTTs performed and key size. The reported numbers for the amortized schemes are amortized over the number of packed ciphertexts. The parameter N is the ring dimension for registers (for us it is the same as q since q is prime in our algorithm).

FHEW-like schemes	n	q	N	$\log_2 Q$	B	d_B	\odot mult.	NTTs	NTTs	# keys
									$\times q \log q$	
FHEW-AP [7, 20]	1024	1024	2048	54	2^{27}	2	3968	11,904	10^8	126 976
TFHE-GINX [5]	1024	1024	2048	54	2^{27}	2	4096	12288	10^8	4096
FHEW-improved [15]	1024	1024	2048	37	2^{13}	3	3074	12,296	10^8	3073
[21] (amortized)	–	–	2187	37	–	–	47844	–	$> 10^9$	–
Ours (amortized)	1024	7681	7681	143	2^{71}	2	101	304	3×10^7	9728

preliminary estimates in Section 5.4 and our initial implementation effort suggest that both the algorithm and the support offered by state-of-the-art FHE libraries may still not be adequate to deliver concrete improvements in practice. In fact, implementing our algorithm within OpenFHE (or similar libraries) raises a number of technical challenges. Most FHE libraries (including OpenFHE) currently only support power-of-two cyclotomic rings, which are the most widely used rings in lattice cryptography. However, our algorithm makes essential use of prime cyclotomics. Some support of cyclotomics other than powers-of-two is offered by the HELib library [12], as well some undocumented functions within the OpenFHE codebase. However, in both cases, the implementation is based on Bluestein FFT, which for technical reasons is limited to ciphertext moduli of size at most 27-bit. As a result, the use of a 143-bit modulus (as in the example parameters in Table 6) would carry at least a $\times 6$ slowdown, compared to sequential bootstrapping methods which can be directly implemented using standard 64-bit arithmetics. Still, our theoretical estimates in Section 5.4 show that amortized FHEW bootstrapping has the potential of being practical, and that, with proper library support and further optimizations, it can offer a practical alternative to sequential bootstrapping algorithms. We hope our work will provide a motivation to extend OpenFHE and other libraries with optimized support for prime (or arbitrary) cyclotomics, and promote further investigation and improvement of amortized FHEW bootstrapping algorithm.

Acknowledgement. Research supported in part by the Swiss National Science Foundation Early Postdoc Mobility Fellowship, Intel Cryptographic Frontiers award, NSF Award CNS-1936703, and SAIT Global Research Cluster.

References

1. Al Badawi, A., Bates, J., Bergamaschi, F., Cousins, D.B., Erabelli, S., Genise, N., Halevi, S., Hunt, H., Kim, A., Lee, Y., et al.: Openfhe: Open-source fully homomorphic encryption library. In: Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography. pp. 53–63 (2022)
2. Alperin-Sheriff, J., Peikert, C.: Practical bootstrapping in quasilinear time. In: Canetti, R., Garay, J.A. (eds.) Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2013. Proceedings, Part I. Lecture Notes in Computer Science, vol. 8042, pp. 1–20. Springer (2013). , https://doi.org/10.1007/978-3-642-40041-4_1
3. Bonnoron, G., Ducas, L., Fillinger, M.: Large FHE gates from tensored homomorphic accumulator. In: Joux, A., Nitaj, A., Rachidi, T. (eds.) Progress in Cryptology - AFRICACRYPT 2018 - 10th International Conference on Cryptology in Africa, Marrakesh, Morocco, May 7–9, 2018, Proceedings. Lecture Notes in Computer Science, vol. 10831, pp. 217–251. Springer (2018). , https://doi.org/10.1007/978-3-319-89339-6_13
4. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. ACM Trans. Comput. Theory **6**(3), 13:1–13:36 (2014). , <https://doi.org/10.1145/2633600>
5. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: fast fully homomorphic encryption over the torus. Journal of Cryptology **33**(1), 34–91 (2020)
6. DeMicheli, G., Kim, D., Micciancio, D., Suhl, A.: Faster amortized fhe bootstrapping using ring automorphisms. Cryptology ePrint Archive, Paper 2023/112 (2023), <https://eprint.iacr.org/2023/112>, <https://eprint.iacr.org/2023/112>

7. Ducas, L., Micciancio, D.: FHEW: bootstrapping homomorphic encryption in less than a second. In: Annual international conference on the theory and applications of cryptographic techniques. pp. 617–640. Springer (2015)
8. Gentry, C.: A fully homomorphic encryption scheme. Stanford university (2009)
9. Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. In: Pointcheval, D., Johansson, T. (eds.) Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7237, pp. 465–482. Springer (2012). , https://doi.org/10.1007/978-3-642-29011-4_28
10. Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the AES circuit. In: Safavi-Naini, R., Canetti, R. (eds.) Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7417, pp. 850–867. Springer (2012). , https://doi.org/10.1007/978-3-642-32009-5_49
11. Guimarães, A., Pereira, H.V.L., van Leeuwen, B.: Amortized bootstrapping revisited: Simpler, asymptotically-faster, implemented. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology – ASIACRYPT 2023. pp. 3–35. Springer Nature Singapore, Singapore (2023)
12. Halevi, S., Shoup, V.: Algorithms in helib. In: Advances in Cryptology–CRYPTO 2014: 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I 34. pp. 554–571. Springer (2014)
13. Halevi, S., Shoup, V.: Faster homomorphic linear transformations in helib. In: Shacham, H., Boldyreva, A. (eds.) Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10991, pp. 93–120. Springer (2018). , https://doi.org/10.1007/978-3-319-96884-1_4
14. Kim, A., Deryabin, M., Eom, J., Choi, R., Lee, Y., Ghang, W., Yoo, D.: General bootstrapping approach for rlwe-based homomorphic encryption. IEEE Transactions on Computers (2023)
15. Lee, Y., Micciancio, D., Kim, A., Choi, R., Deryabin, M., Eom, J., Yoo, D.: Efficient FHEW bootstrapping with small evaluation keys, and applications to threshold homomorphic encryption. Cryptology ePrint Archive (2022)
16. Liu, F.H., Wang, H.: Batch bootstrapping i: a new framework for simd bootstrapping in polynomial modulus. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 321–352. Springer (2023)
17. Liu, F.H., Wang, H.: Batch bootstrapping ii: bootstrapping in polynomial modulus only requires $\tilde{o}(1)$ the multiplications in amortization. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 353–384. Springer (2023)
18. Liu, Z., Wang, Y.: Amortized functional bootstrapping in less than 7ms, with $\tilde{o}(1)$ polynomial multiplications. In: Advances in Cryptology – ASIACRYPT 2023. Springer-Verlag (2023)
19. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. J. ACM **60**(6), 43:1–43:35 (2013). , <https://doi.org/10.1145/2535925>
20. Micciancio, D., Polyakov, Y.: Bootstrapping in FHEW-like cryptosystems. In: Proceedings of the 9th on Workshop on Encrypted Computing & Applied Homomorphic Cryptography. pp. 17–28 (2021)
21. Micciancio, D., Sorrell, J.: Ring packing and amortized FHEW bootstrapping. In: Chatzigiannakis, I., Kaklamani, C., Marx, D., Sannella, D. (eds.) 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic. LIPIcs, vol. 107, pp. 100:1–100:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018). , <https://doi.org/10.4230/LIPIcs.ICALP.2018.100>
22. Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient public key encryption based on ideal lattices. In: Matsui, M. (ed.) Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5912, pp. 617–635. Springer (2009). , https://doi.org/10.1007/978-3-642-10366-7_36

A Sparse Packing based on MLWE

In Section 4.1, we only address the case where $\phi(d)$ LWE ciphertexts are packed into a single RLWE ciphertext and then bootstrapped in parallel. However, the ring dimension $\phi(d)$ is determined by the target security level, and we do not always need to bootstrap $\phi(d)$ ciphertexts at once in real-world applications. In this section, we introduce a simple variant of our amortized bootstrapping algorithm that requires an homomorphic inverse NTT of ring dimension smaller than $\phi(d)$ corresponding to a sparse packing case. Under the same security

level, the amortized bootstrapping for sparse packing shows *faster* performance in terms of latency, compared to the original amortized bootstrapping.

Sparse Packing. Let d_1 be a positive integer that divides $\phi(d)$ and $d_2 := \phi(d)/d_1 \geq 1$. We present an algorithm that packs d_1 LWE ciphertexts into a single RLWE ciphertext of ring dimension $\phi(d)$ in Algorithm 7.

Algorithm 7 Sparse Ring Packing

Input: d_1 plain LWE ciphertexts $(\vec{a}_i, b_i) \in \mathbb{Z}_{q_{\text{plain}}}^{n_{\text{plain}}} \times \mathbb{Z}_{q_{\text{plain}}}$, $\text{RLWE}'(s_i)$

Output: RLWE ciphertext in $\mathcal{R}_q^{1 \times 2}$.

for $0 \leq j < n_{\text{plain}}$ **do**

let $r_j = \sum_{i=0}^{d_1-1} a_{i,j} X^{d_2 i}$ in $\mathcal{R}_{q_{\text{plain}}}$.

end for

$r' = (0, \sum_{i=0}^{d_1-1} b_i X^{d_2 i})$

▷ (Noiseless RLWE' ciphertext)

$ct \leftarrow r' + \sum_{j=0}^{n_{\text{plain}}-1} r_j \odot \text{RLWE}'(s_j)$

return $\text{ModSwitch}_{q_{\text{plain}} \rightarrow q}(ct)$

Linear Step. The decryption of an RLWE ciphertext $[\mathbf{a}, \mathbf{b}] \in \mathcal{R}_q^{1 \times 2}$ of a sparsely-packed message \mathbf{m} under the secret key \mathbf{z} is expressed as $\mathbf{b} + \mathbf{a} \cdot \mathbf{z} \pmod{q} = \mathbf{m} + \mathbf{e}$ for some small polynomial $\mathbf{e} \in \mathcal{R}$. Since the message polynomial \mathbf{m} is of the form $\mathbf{m}(X) = \mathbf{m}'(X^{d_2})$ for some $\mathbf{m}' \in \mathcal{R}' := \mathbb{Z}[X]/(X^{d_1} + 1)$, we can apply the trace function from \mathcal{R}_q to \mathcal{R}'_q on the decryption formula. Then it holds that

$$\text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(\mathbf{b}) + \text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(\mathbf{a} \cdot \mathbf{z}) \pmod{q} = d_2 \cdot \mathbf{m}' + \text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(\mathbf{e}),$$

and it can be re-expressed in FFT form as

$$\text{FFT}'(\text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(\mathbf{b})) + \text{FFT}'(\text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(\mathbf{a} \cdot \mathbf{z})) \pmod{q} = d_2 \cdot \text{FFT}'(\mathbf{m}') + \text{FFT}'(\text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(\mathbf{e})),$$

where FFT' denoted the d_1 -dimensional FFT over \mathcal{R}'_q .

In general, $\text{FFT}(\text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(f))$ for $f \in R$ can be simply computed from $\text{FFT}(f)$ by taking block-wise summations. Indeed, for each d -th root of unity ω , the evaluation $\text{Tr}_{\mathcal{R}_q/\mathcal{R}'_q}(f)(\omega^{d_2}) = \sum_{\sigma \in \text{Gal}(\mathcal{R}_q/\mathcal{R}'_q)} f(\sigma(\omega))$ where ω^{d_2} is the $2d_1$ -th root of unity. Hence, the d -th roots of unity are partitioned into d_1 blocks, each of which are contained in the same coset of $\mathcal{R}_q/\mathcal{R}'_q$, and the summation is done over $f(\omega)$ for ω 's in the same block.

For $0 \leq j < d_1$, the summation corresponding to the j -th block can be expressed as $\sum_{i=0}^{d_2-1} u_{i,j} \cdot v_{i,j}$ for some $u_{i,j}, v_{i,j} \in \mathbb{Z}_q$, where $u_{i,j}$'s and $v_{i,j}$'s correspond to FFT evaluations on \mathbf{a} and \mathbf{z} , respectively. We compute the inner product as

$$\sum_{i=0}^{d_2-1} u_{i,j} \cdot v_{i,j} = u_{0,j} \cdot \left(v_{0,j} + \frac{u_{1,j}}{u_{0,j}} \cdot \left(\dots + \frac{u_{d_2-2,j}}{u_{d_2-3,j}} \cdot \left(v_{d_2-2,j} + \frac{u_{d_2-1,j}}{u_{d_2-2,j}} \cdot v_{d_2-1,j} \right) \right) \right).$$

Let us denote by $\vec{u}_i := (u_{i,j})_{0 \leq j < d_1}$ and $\vec{v}_i := (v_{i,j})_{0 \leq j < d_1}$ the d_1 -dimensional vectors over \mathbb{Z}_q . Starting from \vec{v}_{d_2-1} , we repeat 1) a point-wise multiplication with \vec{u}_i/\vec{u}_{i-1} and 2) taking RLWE'-RGSW multiplication with \vec{v}_{i-1} from $i = d_2 - 1$ to 1. The last step corresponds to a point-wise multiplication with \vec{u}_0 . Figure 4 describes a high-level procedure of the linear step of homomorphic RLWE decryption for sparse packing.

In order to compare the performance of this sparse-case algorithm, we can summarize the number of $\mathcal{R} \odot \text{RLWE}'$ operations in Table 7. The analysis closely follows what has been presented in the main body of the paper. To compare the total number of $\mathcal{R} \odot \text{RLWE}'$ with Table 3, we analysis the following ratio dividing

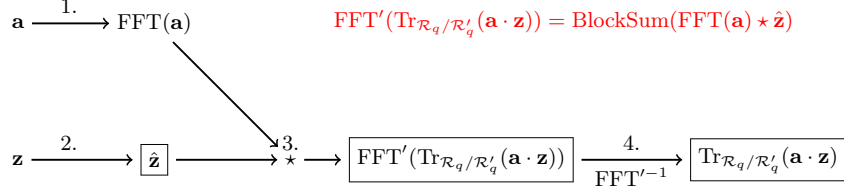


Fig. 4. High level description of the linear step of our sparse-case algorithm. The notation \star refers to pointwise multiplication. The boxed information refers to encrypted data where homomorphic operations are required.

Table 7. Summary of $\mathcal{R} \odot \text{RLWE}'$ operation count in sparse packing.

Steps of the algorithm	$\mathcal{R} \odot \text{RLWE}'$ operations
Partial FFT of \mathbf{a}	–
Mult & BlockSum	$(3k - 2)d_1 d_2 d_B + 2d_1(d_2 - 1)d_B$
Partial IFFT (per layer)	$(3r - 2 - \frac{1}{r})d_1 d_B$
Last layer of IFFT	$(3r - 3 + \frac{1}{r} + d_B(1 - \frac{1}{r}))d_1$
Last IFFT step	d_1
Adding \mathbf{b}	–

the total number of $\mathcal{R} \odot \text{RLWE}'$ operations in the sparse-case by the total number in the main algorithm:

$$\frac{3k\phi(d)d_B + (3r - 1 - \frac{2}{r})d_1 d_B + (3r - 2 + \frac{1}{r})d_1}{3k\phi(d)d_B + (3r - 3 - \frac{2}{r})\phi(d)d_B + (3r - 2 + \frac{1}{r})\phi(d)} \simeq \frac{3k + 3r/d_2}{3k + 3r}$$

Note that the ratio is close to $1/d_2$ if $3r/d_2$ is relatively larger than $3k$. Otherwise, the ratio is close to $k/(k + r)$. In both cases, this shows that the sparse case requires less operations, thus leads to a faster performance. The exact gain in the sparse case depends on the parameter values. Also, note that to obtain a complete analysis of this sparse algorithm, a thorough error growth analysis should also be done, which we leave for future work.