

Fast and Accurate: Efficient Full-Domain Functional Bootstrap and Digit Decomposition for Homomorphic Computation

Shihe Ma¹, Tairong Huang², Anyu Wang^{2,4,5}✉,
Qixian Zhou³ and Xiaoyun Wang^{2,4,5,6,7}

¹ Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing, China,
msh21@mails.tsinghua.edu.cn

² Institute for Advanced Study, BNRist, Tsinghua University, Beijing, China,
htr19@mails.tsinghua.edu.cn, anyuwang@tsinghua.edu.cn,
xiaoyunwang@tsinghua.edu.cn

³ Ant Group, qixian.zqx@antgroup.com

⁴ Zhongguancun Laboratory, Beijing, China

⁵ National Financial Cryptography Research Center, Beijing, China

⁶ Shandong Institute of Blockchain, Jinan, China

⁷ Key Laboratory of Cryptologic Technology and Information Security (Ministry of Education),
School of Cyber Science and Technology, Shandong University, Qingdao, China

Abstract. The functional bootstrap in FHEW/TFHE allows for fast table lookups on ciphertexts and is a powerful tool for privacy-preserving computations. However, the functional bootstrap suffers from two limitations: the negacyclic constraint of the lookup table (LUT) and the limited ability to evaluate large-precision LUTs. To overcome the first limitation, several full-domain functional bootstraps (FDFB) have been developed, enabling the evaluation of arbitrary LUTs. Meanwhile, algorithms based on homomorphic digit decomposition have been proposed to address the second limitation. Although these algorithms provide effective solutions, they are yet to be optimized. This paper presents four new FDFB algorithms and two new homomorphic decomposition algorithms that improve the state-of-the-art. Our FDFB algorithms reduce the output noise, thus allowing for more efficient and compact parameter selection. Across all parameter settings, our algorithms reduce the runtime by up to 39.2%. Our homomorphic decomposition algorithms also run at 2.0x and 1.5x the speed of prior algorithms. We have implemented and benchmarked all previous FDFB and homomorphic decomposition algorithms and our methods in OpenFHE.

Keywords: Homomorphic Encryption · TFHE · FHEW · Functional Bootstrap · FDFB · Homomorphic Decomposition

1 Introduction

Fully Homomorphic Encryption (FHE) is a powerful cryptographic tool that enables computation on encrypted data without requiring access to the decryption key. It has great potential for use in computing fields where data privacy is important, such as secure cloud computing [KSK⁺18, PKS⁺19, LATV12] and privacy-preserving machine learning [LKL⁺22, BMMP18, CJP21, LHH⁺21], as well as in the construction of cryptographic protocols such as private set intersection [CLR17, CHLR18, CMdG⁺21].

Since Gentry’s first construction of an FHE scheme utilizing the bootstrap technique [Gen09], various FHE schemes have been developed [FV12, BGV14, CKKS17, GSW13, DM15, CGGI20] and significant improvements have been made [LW23a, LW23b,

BIP⁺22, Klu22]. Among these FHE schemes, BGV/FV, CKKS and FHEW/TFHE have gained prominence recently because of their great efficiency. BGV/FV and CKKS have effective packing capabilities that allow for computations over vector data using *Single Instruction Multiple Data* (SIMD) instructions, making them ideal for simultaneously processing large arrays of numbers. However, these schemes are less efficient for evaluating deep circuits and inconvenient for evaluating non-polynomial functions. On the other hand, FHEW/TFHE utilize an efficient functional bootstrap (or programmable bootstrap) process that enables the evaluation of a *lookup table* (LUT) without additional cost, making these schemes ideal for evaluating boolean circuits and non-polynomial functions. Moreover, due to the switching method introduced in CHIMERA [BGGJ20] and later improved in PEGASUS [LHH⁺21], a CKKS ciphertext can be converted into multiple FHEW/TFHE ciphertexts to compute non-polynomial functions and then converted back to CKKS ciphertext for SIMD polynomial evaluation. This makes functional bootstrap a versatile tool for all FHE evaluation purposes.

Despite its strength, functional bootstrap still suffers from two limitations: (1) the evaluated LUT $f : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$ must be negacyclic such that $f(x + \frac{p}{2}) = -f(x)$ for all $x \in \mathbb{Z}_p$, preventing some LUTs from being evaluated directly; (2) the input plaintext modulus p is typically small due to efficiency constraints, limiting its ability to evaluate large precision LUTs. Numerous efforts have been made to address these two limitations. To circumvent the negacyclicity constraint, *Full Domain Functional Bootstrap* (FDFB) algorithms supporting arbitrary LUTs have been proposed. These FDFB algorithms can be categorized into Type-SelectMSB, Type-HalfRange and Type-Split. Type-SelectMSB selects between two negacyclic LUTs based on the *most significant bit* (MSB) of the encrypted message and is used in algorithms proposed by [CLOT21, KS22]. Type-HalfRange transforms the encrypted message to prevent it from exceeding $\frac{p}{2}$, thereby bypassing the negacyclic limitation. This method is adopted in algorithms proposed by [LMP22, YXS⁺21, GBA22]. Finally, Type-Split expresses an arbitrary LUT as the sum of a ‘pseudo-odd’ LUT and a ‘pseudo-even’ LUT, each of which can be evaluated using two functional bootstraps. This method is employed in the algorithm proposed by [CZB⁺22]. In addition to focusing on the construction of FDFB, a method for using FDFB to aid in evaluating CKKS ciphertexts is presented in [LY23]. To handle the evaluation of large-precision LUTs, Guimarães et al. [GBA21] propose tree-based and chaining methods to combine multiple functional bootstraps in TFHE. These two methods in [GBA21] assume that each ciphertext encrypts a digit of the original message. Therefore, when an input ciphertext has a large modulus, it must first be preprocessed with homomorphic decomposition before the methods can be applied. On the other hand, Liu et al. [LMP22] develop homomorphic digit decomposition algorithms and demonstrate how they can be used to evaluate large-precision sign functions. As a result, homomorphic decomposition is a crucial component in current techniques for evaluating large-precision LUTs.

In practice, functional bootstrap plays a critical role in many FHE applications, and thus its optimization is paramount for achieving high performance. Nevertheless, the efficiency of the FDFB and digit decomposition algorithms still requires further evaluation and optimization.

1.1 Our Contributions

This work presents new methods for optimizing the current FDFB and homomorphic decomposition algorithms. Our contributions can be summarized as follows.

(1) We present four novel FDFB algorithms: **FDFB-Compress**, **FDFB-CancelSign**, **FDFB-Select** and **FDFB-BFVMult (WoPPBS₁-Refine)**. **FDFB-Compress** improves Type-HalfRange to theoretical optimality, while the other three algorithms improve Type-SelectMSB but are suitable for different scenarios. In our experiments, we observe that our fastest algorithms demonstrate a significant speedup, ranging from 23.4% \sim 39.2%,

compared to the state-of-the-art results across various parameter settings.

(2) We present two new homomorphic decomposition algorithms **HomDecomp-Reduce** and **HomDecomp-FDFB**, whose running speed is 2x and 1.5x that of **HomFloor** and **HomFloorAlt** from [LMP22], respectively. Unlike **HomFloor**, our algorithms do not require the input ciphertext to have small noise. The speedup of our algorithms directly results in faster large-precision evaluations of functions such as sign, ReLU, max, ABS, etc.

(3) We provide a comprehensive theoretical noise analysis for our FDFB and homomorphic decomposition algorithms, as well as those developed by previous works. We have implemented and benchmarked all the algorithms in the OpenFHE library [BBB⁺22] to validate our results. Our implementation of all FDFB algorithms in a single library is a first-of-its-kind initiative, which provides standardized access to these algorithms.

1.2 Related Works

1.2.1 FDFB Algorithms

The current FDFB algorithms are summarized as follows.

WoP-PBS₁ [CLOT21] (Type-SelectMSB) introduces an extra MSB to the encrypted message by doubling the ciphertext modulus. The algorithm evaluates the LUT to obtain a ciphertext that possibly differs by a sign from the desired result. Then, it extracts the MSB using functional bootstrap and offsets the sign by invoking BFV multiplication. However, the rapid noise growth of BFV multiplication requires the algorithm to use inefficient parameters, thus degrading performance.

WoP-PBS₂ [CLOT21] (Type-SelectMSB) builds two sub-LUTs according to the MSB of the encrypted message. The algorithm evaluates both sub-LUTs to obtain two ciphertexts and extracts the MSB using functional bootstrap. Then BFV multiplication is invoked to select the correct ciphertext. Again, BFV multiplication still requires large parameters and degrades performance.

FDFB-KS [KS22] (Type-SelectMSB) builds two sub-LUTs similarly to **WoP-PBS₂**. The algorithm selects between the two sub-LUTs to obtain an encrypted LUT and then uses functional bootstrap to evaluate it. However, selecting the sub-LUTs requires multiple functional bootstraps and causes significant computational overhead.

EvalFunc [LMP22] (Type-HalfRange) introduces an extra MSB in a similar way to **WoP-PBS₁**. The algorithm extracts the MSB using functional bootstrap and cancels it to ensure that the message belongs to half of \mathbb{Z}_p . Then it can evaluate the LUT without being constrained by negacyclicity. We note that the **FullyFBS** of [YXS⁺21] and the **FDFB-C** of [GBA22] are essentially the same as **EvalFunc**.

Comp [CZB⁺22] (Type-Split) expresses an arbitrary LUT as the sum of a ‘pseudo-odd’ LUT and a ‘pseudo-even’ LUT. Then the algorithm evaluates each LUT using two functional bootstraps.

In [CIM19], Carpov et al. develop a multi-value bootstrap technique that allows several LUTs to be evaluated on the same input using a single functional bootstrap call. This technique can reduce the functional bootstraps required for **WoP-PBS₁**, **WoP-PBS₂** and **Comp** when the parameters support multi-value bootstrap.

1.2.2 Homomorphic Decomposition Algorithms

The current homomorphic decomposition algorithms are summarized as follows.

HomFloor [LMP22] uses two bootstraps to clear the lower bits of a large-precision message before modulus switching, which prevents the modulus switching noise from corrupting the higher digits. By iteratively applying these operations, a large-precision message can be decomposed into a vector of 4-bit digits. However, this algorithm does

Table 1: A summary of the intuition behind our algorithms and the improvement over previous methods.

Previous	Ours	Our Intuition/Improvement over Previous Works
EvalFunc [LMP22]	FDFB-Compress	Compress the coded message using a functional bootstrap and reduce the noise
WoP-PBS₁ [CLOT21] WoP-PBS₂ [CLOT21]	FDFB-CancelSign FDFB-Select	Replace BFV multiplication with LWE-to-RLWE packing and bootstrap
WoP-PBS₁ [CLOT21] WoP-PBS₂ [CLOT21]	WoPPBS₁-Refine FDFB-BFVMult	Use a refined noise analysis for BFV multiplication; use fewer BFV multiplications
HomFloor [LMP22] HomFloorAlt [LMP22]	HomDecomp-Reduce HomDecomp-FDFB	Reduce the range of the lower bits instead of clearing them and use fewer bootstraps

not apply to extracted CKKS ciphertexts because it requires a small noise in the input ciphertext.

HomFloorAlt [LMP22] uses three bootstraps to extract the digits of a large-precision message, allowing it to support the decomposition into 5-bit digits and decompose extracted CKKS ciphertexts.

1.3 Overview of Our Algorithms

We present the intuition behind our algorithm design and explain how it leads to better performance (see Table 1 for a summary). The key advantage of our algorithms is their reduced noise growth, which enables us to choose more compact LWE and RLWE parameters (such as decomposition bases in blind rotation and RLWE dimension) for a given plaintext modulus, resulting in shorter running time.

FDFB-Compress is a Type-HalfRange FDFB algorithm. Our key observation is that the LWE message must be in a coded (and thus redundant) form $\frac{q}{p}m' + e \in \mathbb{Z}_q$ to prevent decryption failures due to errors, where q is the ciphertext modulus. This enables us to design a compression function that can compress the coded LWE message into $[-\frac{q}{4}, \frac{q}{4} - 1]$ using one functional bootstrap. Then, we can perform another functional bootstrap on the compressed message to get the desired result. As a result, **FDFB-Compress** uses the same number of bootstraps as **EvalFunc** but reduces the error variance of the compressed message by half, resulting in a more compact parameter choice and better performance.

FDFB-CancelSign, **FDFB-Select** and **FDFB-BFVMult** (**WoPPBS₁-Refine**) are all Type-SelectMSB FDFB algorithms. The primary objective of **FDFB-CancelSign** and **FDFB-Select** is to replace the BFV multiplication in **WoP-PBS₁** and **WoP-PBS₂** with LWE-to-RLWE packing and an additional functional bootstrap. This approach prevents the multiplicative noise growth in BFV multiplication and instead achieves additive noise growth. As a result, although **FDFB-CancelSign** and **FDFB-Select** require an extra functional bootstrap compared to **WoP-PBS**, their slower noise growth allows for more compact parameter choices and better efficiency in most cases, according to our experiments. On the other hand, **WoPPBS₁-Refine** and **FDFB-BFVMult** are enhanced algorithms of **WoP-PBS₁** and **WoP-PBS₂**, respectively. They significantly reduce the error growth in **WoP-PBS₁** and **WoP-PBS₂** by roughly N times, where N is the RLWE dimension. This is achieved through a refined noise analysis of the BFV multiplication. Such an in-depth analysis allows for the choice of smaller bootstrapping

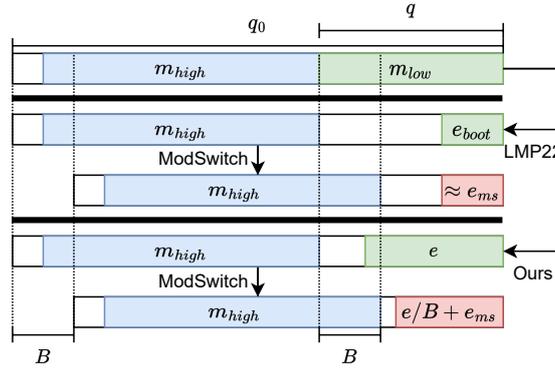


Figure 1: Comparison of our homomorphic digit decomposition approach and that of [LMP22]. The blue parts stand for higher bits, while the green and red parts stand for lower bits before and after modulus switching.

parameters, resulting in enhanced efficiency. Moreover, **FDFB-BFVMult** removes one BFV multiplication in **WoP-PBS₂** by combining two BFV multiplications with the sign bit into one multiplication, further reducing the noise growth by half.

The current homomorphic digit decomposition algorithms presented in [LMP22] extract digits by repeatedly clearing the lower bits m_{low} of the encrypted messages (leaving a small bootstrap error) and then modulus-switching it to a smaller modulus $\frac{q_0}{B}$. We observe that this goal can also be achieved by reducing the range of the lower bits instead of clearing them. In contrast to clearing the lower bits, reducing their range consumes fewer functional bootstraps. Still, it can reserve enough room to hold the modulus switching noise, thus preventing the higher digits from being destroyed by overflowed noise. Figure 1 illustrates a comparison of these two approaches. Following this idea, we design **HomDecomp-Reduce** and **HomDecomp-FDFB**, which run 2x and 1.5x faster compared to **HomFloor** and **HomFloorAlt** in our experiments.

2 Preliminaries

2.1 Notations

The ring of integers modulo q is denoted as $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$. Its elements are represented as integers in either $[0, q - 1]$ (positive form) or $[-\lfloor \frac{q}{2} \rfloor, \lfloor \frac{q-1}{2} \rfloor]$ (signed form). For an integer a , its positive form and signed form in \mathbb{Z}_q are denoted as $[a]_q^+$ and $[a]_q$, respectively.

For a power-of-2 N , the $2N$ -th cyclotomic ring is denoted as $R = \mathbb{Z}[X]/(X^N + 1)$, and its quotient ring is denoted as $R_q = R/qR$. Polynomials are represented using bold letters, e.g., \mathbf{a} . For a vector \vec{a} or a polynomial \mathbf{b} , we use a_i and \mathbf{b}_i respectively to denote \vec{a} 's i -th entry and \mathbf{b} 's coefficient of the X^i term. The coefficient vector of \mathbf{b} is denoted as $\vec{\mathbf{b}} = (\mathbf{b}_0, \mathbf{b}_1, \dots, \mathbf{b}_{N-1})$.

For a positive integer n , the set $\{0, 1, \dots, n - 1\}$ is denoted as $\llbracket n \rrbracket$. We use $a \leftarrow \chi$ to represent a random variable a sampled from the distribution χ , and $a \leftarrow S$ to indicate that a is uniformly sampled from the finite set S . We use $\mathcal{D}(\mathbb{Z}, \sigma)$ to denote the discrete Gaussian distribution of parameter σ over \mathbb{Z} . The infinity norm and 2-norm of a vector \vec{a} are denoted as $|\vec{a}|_\infty$ and $|\vec{a}|_2$ respectively. All logarithms are taken with a base of 2 unless otherwise stated.

2.2 FHEW/TFHE Encryption Schemes

2.2.1 LWE and RLWE Ciphertexts

Throughout this paper, we use lowercase q and n to denote the modulus and dimension of LWE instances, while uppercase Q and N are used for the RLWE modulus and dimension.

The LWE ciphertext encrypting an encoded message $m \in \mathbb{Z}_q$ is defined to be

$$\text{LWE}_{\vec{s},n,q}(m + e) = (-\langle \vec{a}, \vec{s} \rangle + m + e, \vec{a}) \in \mathbb{Z}_q^{n+1},$$

where $\vec{a} \leftarrow \mathbb{Z}_q^n$, $e \leftarrow \mathcal{D}(\mathbb{Z}, \sigma)$, and the secret vector $\vec{s} \leftarrow \{0, \pm 1\}^n$.

The RLWE ciphertext encrypting an encoded message $\mathbf{m} \in R_Q$ is defined to be

$$\text{RLWE}_{\mathbf{s},N,Q}(\mathbf{m} + \mathbf{e}) = (-\mathbf{a} \cdot \mathbf{s} + \mathbf{m} + \mathbf{e}, \mathbf{a}) \in R_Q^2,$$

where $\mathbf{a} \leftarrow R_Q$, $\mathbf{e}_i \leftarrow \mathcal{D}(\mathbb{Z}, \sigma)$, and the secret polynomial satisfies $\mathbf{s}_i \leftarrow \{\pm 1, 0\}$.

For simplicity, we may sometimes use the abbreviated notation $\text{LWE}_{\vec{s}}(m)$ and $\text{RLWE}_{\mathbf{s}}(\mathbf{m})$ (or $\text{LWE}(m)$ and $\text{RLWE}(\mathbf{m})$) to denote the LWE and RLWE ciphertexts respectively.

Messages in LWE and RLWE ciphertexts are typically encoded to prevent decryption failures caused by errors. For instance, in an RLWE ciphertext, \mathbf{m} is often an up-scaled version of the actual message $\mathbf{m}' \in R_p$, as given by $\mathbf{m} = \lfloor \frac{Q}{p} \mathbf{m}' \rfloor = \frac{Q}{p} \mathbf{m}' + \mathbf{e}_{rnd}$, where $p < Q$ is the plaintext modulus and \mathbf{e}_{rnd} accounts for the rounding errors. Then an RLWE ciphertext $(\mathbf{b}, \mathbf{a}) \in R_Q^2$ decrypts to $\lfloor \frac{p}{Q} (\mathbf{b} + \mathbf{a} \cdot \mathbf{s}) \rfloor = \lfloor \mathbf{m}' + \frac{p}{Q} (\mathbf{e} + \mathbf{e}_{rnd}) \rfloor$, which is equal to \mathbf{m}' modulo p as long as $|\frac{p}{Q} (\mathbf{e} + \mathbf{e}_{rnd})|_{\infty} < \frac{1}{2}$.

2.2.2 RLWE' and RGSW Ciphertexts

An RLWE' ciphertext is a vector of RLWE ciphertexts encrypting the same message at different scales, i.e.,

$$\text{RLWE}'_{\mathbf{s}}(\mathbf{m}) = (\text{RLWE}_{\mathbf{s}}(\mathbf{m}), \text{RLWE}_{\mathbf{s}}(\mathbf{m} \cdot B), \dots, \text{RLWE}_{\mathbf{s}}(\mathbf{m} \cdot B^{l-1})),$$

where $B \in \mathbb{Z}$ is the decomposition base and $l = \lceil \log_B Q \rceil$. For any $\mathbf{u} \in R_Q$, there is a decomposition $\mathbf{u} = \sum_{i=0}^{l-1} \mathbf{u}_i \cdot B^i$ such that \mathbf{u}_i 's coefficients are all in $[-\frac{B}{2}, \frac{B}{2}]$. Let $\text{Decomp}(\mathbf{u}) = (\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{l-1})$. Then the product $\odot : R_Q \times \text{RLWE}' \rightarrow \text{RLWE}$ can be defined as

$$\mathbf{u} \odot \text{RLWE}'_{\mathbf{s}}(\mathbf{m}) = \langle \text{Decomp}(\mathbf{u}), \text{RLWE}'_{\mathbf{s}}(\mathbf{m}) \rangle = \text{RLWE}_{\mathbf{s}}(\mathbf{u} \cdot \mathbf{m}).$$

The obtained RLWE ciphertext contains a noise much smaller than the regular $R_Q \times \text{RLWE}$ multiplication due to the small coefficients of \mathbf{u}_i 's. Besides, the LWE' ciphertext can be defined similarly, but we omit the details here.

An RGSW ciphertext is defined as

$$\text{RGSW}_{\mathbf{s}}(\mathbf{m}) = (\text{RLWE}'_{\mathbf{s}}(\mathbf{m}), \text{RLWE}'_{\mathbf{s}}(\mathbf{m} \cdot \mathbf{s})).$$

Then the external product $\diamond : \text{RLWE} \times \text{RGSW} \rightarrow \text{RLWE}$ between $(\mathbf{b}, \mathbf{a}) = \text{RLWE}_{\mathbf{s}}(\mathbf{u} + \mathbf{e})$ and $\text{RGSW}_{\mathbf{s}}(\mathbf{m})$ is defined as

$$(\mathbf{b}, \mathbf{a}) \diamond \text{RGSW}_{\mathbf{s}}(\mathbf{m}) = \mathbf{b} \odot \text{RLWE}'_{\mathbf{s}}(\mathbf{m}) + \mathbf{a} \odot \text{RLWE}'_{\mathbf{s}}(\mathbf{m} \cdot \mathbf{s}),$$

which is equal to $\text{RLWE}_{\mathbf{s}}((\mathbf{b} + \mathbf{a} \cdot \mathbf{s})\mathbf{m}) = \text{RLWE}_{\mathbf{s}}((\mathbf{u} + \mathbf{e})\mathbf{m})$.

2.3 Homomorphic Operators

We introduce some basic homomorphic operations that will be used in our constructions.

2.3.1 Mod Down/Up and Modulus Switching

Let $c = (b, \vec{a}) = \text{LWE}_{\vec{s}, n, q}(m + e)$ be an LWE ciphertext, and let q' be a positive modulus.

For $q' \mid q$, the ‘mod down’ is defined as

$$\text{ModDown}(c, q') = ([b]_{q'}, [\vec{a}]_{q'}) = \text{LWE}_{\vec{s}, n, q'}([m + e]_{q'}).$$

For $q \mid q'$, the ‘mod up’ is defined as

$$\text{ModUp}(c, q') = (b, \vec{a}) = \text{LWE}_{\vec{s}, n, q'}(m + e + vq),$$

where $v \in \mathbb{Z}_{q'/q}$.

For any modulus q' , the ‘modulus switching’ is defined as

$$\text{ModSwitch}(c, q') = (\lfloor \frac{q'}{q} b \rfloor, \lfloor \frac{q'}{q} \vec{a} \rfloor) = \text{LWE}_{\vec{s}, n, q'}(\frac{q'}{q}(m + e) + e_{ms}),$$

where e_{ms} is the noise modulus switching introduces. The three homomorphic operators described above can also be defined for RLWE ciphertexts similarly but are omitted for brevity.

2.3.2 Sample Extract

Given an RLWE ciphertext $c = (\mathbf{b}, \mathbf{a}) = \text{RLWE}_{\vec{s}, N, Q}(\mathbf{m} + \mathbf{e})$ and an index $i \in \llbracket N \rrbracket$, define

$$\text{SampleExtract}(c, i) = \text{LWE}_{\vec{s}, N, Q}(\mathbf{m}_i + \mathbf{e}_i),$$

which extracts the coefficient of the X^i term into an LWE ciphertext.

2.3.3 Key Switching

Given an LWE ciphertext $c = (b, \vec{a}) = \text{LWE}_{\vec{s}, n, q_{ks}}(m + e)$, a decomposition base B_{ks} and key switching keys $\text{ksk}_{i,j,k} = \text{LWE}_{\vec{s}', n', q'}(\lfloor \frac{q'}{q_{ks}} \vec{s}_i \cdot j \cdot B_{ks}^k \rfloor)$ for $i \in \llbracket n \rrbracket, j \in \llbracket B_{ks} \rrbracket$ and $k \in \llbracket \lceil \log_{B_{ks}}(q_{ks}) \rceil \rrbracket$, define

$$\text{KeySwitch}(c, \{\text{ksk}_{i,j,k}\}) = \text{LWE}_{\vec{s}', n', q'}(\lfloor \frac{q'}{q_{ks}}(m + e) \rfloor + e_{ks}),$$

where e_{ks} is the error key switching introduces.

Besides LWE-to-LWE key switching, it is possible to pack LWE ciphertexts into an RLWE ciphertext with similar techniques [GBA21, CZ22], which can be viewed as a specific instance of the public functional key switching method proposed in [CGGI20]. This homomorphic operator, denoted as $\text{PackingKS}(\text{LWE}(m), \{\text{ksk}_{i,j,k}\})$, is parameterized by a positive integer d and outputs $\text{RLWE}(m + mX + \dots + mX^{d-1})$. Its full definition is detailed in the full version of the paper.

2.3.4 Blind Rotation and Functional Bootstrap

Blind rotation is the key step in the bootstrap of FHEW/TFHE. Given an LWE ciphertext $c = \text{LWE}_{\vec{s}}(m + e)$ with modulus $q \mid 2N$, a polynomial $\text{TV} \in R_Q$ (often called the *test vector*) and blind rotation keys $\{\text{brk}_i^\pm\}$, define

$$\text{BlindRotate}(c, \text{TV}, \{\text{brk}_i^\pm\}) = \text{RLWE}_{\vec{s}'}(\text{TV} \cdot X^{-\frac{2N}{q}}(m + e) + \mathbf{e}_{acc}),$$

where \mathbf{e}_{acc} is the noise that blind rotation introduces. In other words, TV is rotated left by $\frac{2N}{q}(m + e)$. $\{\text{brk}_i^\pm\}$ are parameterized by the blind-rotation base B_g . A smaller B_g

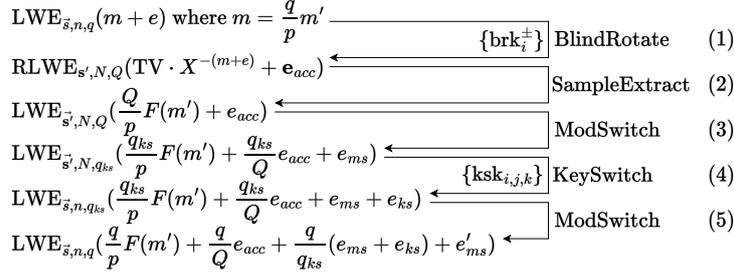


Figure 2: The five steps of FHEW/TFHE bootstrapping: (1) blind rotation of TV by the input ciphertext; (2) extracting the constant term of the rotated TV; (3) modulus switching to q_{ks} ; (4) key switching to the original secret key; (5) modulus switching to q . F is an LUT from \mathbb{Z}_p to \mathbb{Z}_p .

means longer running time and smaller e_{acc} . Since the inner structure of blind rotation is irrelevant to the focus of this paper, we omit the details about the use of $\{\text{brk}_i^{\pm}\}$. Interested readers can refer to [MP21] for more details. In this paper, we assume $q = 2N$ and omit the $\{\text{brk}_i^{\pm}\}$ in notations.

Note that the constant term of the rotated TV equals TV_{m+e} for $m+e \in [0, N-1]$, and equals $-\text{TV}_{[m+e]_N}^+$ for $m+e \in [N, 2N-1]$, then the blind rotation actually evaluates a negacyclic function $f: \mathbb{Z}_{2N} \rightarrow \mathbb{Z}_Q$ on $m+e$. To evaluate a negacyclic LUT $F: \mathbb{Z}_p \rightarrow \mathbb{Z}_p$ using blind rotation, the coefficients of TV are arranged in a redundant way to eliminate the error in input ciphertext. Specifically, by setting $\text{TV}_i = \lfloor \frac{Q}{p} F(\lfloor \frac{p}{q} i \rfloor) \rfloor$, the constant term of $\text{BlindRotate}(\text{LWE}_{\bar{s}}(\lfloor \frac{q}{p} m' + e \rfloor), \text{TV})$ is an encryption of $\lfloor \frac{Q}{p} F(m') \rfloor$.

The entire process of the functional bootstrap is illustrated in Figure 2. The noise introduced by the bootstrap process is denoted as e_{boot} . We use $\text{Boot}[f](c)$ to represent the result of performing functional bootstrap using function f on an LWE ciphertext c and use $\text{BootRaw}[f](c)$ to represent the freshly extracted LWE ciphertext after blind rotation (i.e., without any modulus switching or key switching). Notably, each TV uniquely corresponds to a negacyclic function f , so either TV or f can be used to parameterize the functional bootstrap. If the plaintext polynomial TV is replaced with an RLWE ciphertext c_{tv} , we denote the resulting output as $\text{Boot}[c_{tv}](c)$ or $\text{BootRaw}[c_{tv}](c)$.

2.3.5 Multi-Value Bootstrap

Multi-value bootstrap enables the evaluation of multiple LUTs on the same input LWE ciphertext with the cost of a single bootstrap [CIM19]. In this approach, the unscaled test vector is denoted as $\text{TV}' \in R_p$, and the goal is to compute $\lfloor \frac{Q}{p} \text{TV}' \rfloor X^{-(m+e)}$, where p is the plaintext modulus. To enable the computation of multiple LUTs, multi-value bootstrap decomposes $\lfloor \frac{Q}{p} \text{TV}' \rfloor$ approximately into $\text{TV}_0 \cdot \text{TV}_1$, where $\text{TV}_0 = \lfloor \frac{Q}{2p} \rfloor (1 + X + \dots + X^{N-1})$ is a constant polynomial, and $\text{TV}_1 = \text{TV}' - \text{TV}_0 \cdot X \in R_{2p}$ is LUT-specific. TV_0 is first multiplied by $X^{-(m+e)}$ using blind rotation, and the resulting RLWE ciphertext is multiplied by TV_1 , which also multiplies the output error variance by $|\text{TV}_1|_2^2 \leq p(p-1)^2$.

2.3.6 BFV Multiplication

Let p be the plaintext modulus. For two RLWE ciphertexts $c_i = \text{RLWE}_{s,N,Q}(\frac{Q}{p} \mathbf{m}'_i + \mathbf{e}_i)$ where $i = 0, 1$, define

$$\text{BFVMult}(c_0, c_1) = \text{RLWE}_{s,N,Q}(\frac{Q}{p} \mathbf{m}'_1 \mathbf{m}'_2 + \mathbf{e}_{mult}),$$

Table 2: Symbols used in our noise analysis.

Symbol	Meaning
σ^2	Encryption error variance
σ_{ms}^2	Modulus-switching error variance
σ_{ks}^2	Key-switching error variance
σ_{pk}^2	PackingKS error variance, $\sigma_{pk}^2 = \sigma_{ks}^2$
σ_{acc}^2	Blind-rotation error variance
σ_{com}^2	Variance of noise introduced in steps (3)~(5) in Figure 2
σ_{boot}^2	Bootstrap error variance
bnd, β	$\text{bnd} = \sqrt{2} \cdot \text{erfc}^{-1}(2^{-32}) \approx 6.338$, and $\beta = \text{bnd} \cdot \sigma_{boot}$ For $x \sim N(0, \sigma_{boot}^2)$, $ x < \beta$ with high probability
p	Plaintext modulus. p is an even number
q, n	LWE ciphertext modulus and dimension. q is a power of 2
Q, N	RLWE ciphertext modulus and dimension. $N = 2q$

where \mathbf{e}_{mult} is the noise of BFV multiplication. We note that re-linearization keys are required for BFV multiplication. See [KPZ21] for the detailed process.

2.4 Noise Introduced by the Operators

The variances of $e_{ms}, e_{ks}, e_{acc}, e_{boot}$ are denoted by $\sigma_{ms}^2, \sigma_{ks}^2, \sigma_{acc}^2, \sigma_{boot}^2$ respectively. Besides, recall that q_{ks} is the key switching modulus in blind rotation. B_{ks} and B_g are the decomposition bases for key switching and blind rotation, respectively. The values of these variances are listed in the following lemma, and the proof can be found in [MP21].

Lemma 1. *Let σ^2 be the variance of the encryption noise, and $d_g = \lceil \log_{B_g} Q \rceil$, $d_{ks} = \lceil \log_{B_{ks}}(q_{ks}) \rceil$. Then*

$$\sigma_{ms}^2(n) = \frac{n}{18} + \frac{1}{12},$$

$$\sigma_{ks}^2(n, q_{ks}, B_{ks}) = d_{ks} \left(1 - \frac{1}{B_{ks}}\right) n \left(\sigma^2 + \frac{1}{4}\right),$$

$$\sigma_{acc}^2(n, N, Q, B_g) = \frac{2d_g B_g^2 n N \sigma^2}{3},$$

$$\sigma_{boot}^2(n, N, Q, B_g, q_{ks}, B_{ks}) = \left(\frac{q}{q_{ks}}\right)^2 (\sigma_{ms}^2(N) + \sigma_{ks}^2(N, q_{ks}, B_{ks})) + \left(\frac{q}{Q}\right)^2 \sigma_{acc}^2(n, N, Q, B_g) + \sigma_{ms}^2(n).$$

PackingKS introduces the same amount of noise as KeySwitch. Besides, we denote $\sigma_{com}^2 = \left(\frac{q}{q_{ks}}\right)^2 (\sigma_{ms}^2 + \sigma_{ks}^2) + \sigma_{ms}^2$ as the variance of noise introduced by the last three steps in the functional bootstrap (Figure 2).

The literature generally assumes that error introduced by homomorphic operations follows a centered normal distribution. For a centered normal variable $x \sim N(0, \sigma^2)$, its range can be bounded by $\Pr[|x| > \text{bnd} \cdot \sigma] < 2^{-32}$, where $\text{bnd} = \sqrt{2} \cdot \text{erfc}^{-1}(2^{-32}) \approx 6.338$. We denote the bound of bootstrapping error as $\beta = \text{bnd} \cdot \sigma_{boot}$. Table 2 summarizes the symbols used in our noise analysis.

3 Improved FDFB Algorithms

This section introduces four new FDFB algorithms. We assume that the plaintext modulus p is a power of 2 for better presentation. Notably, changing p to any even number will not affect the correctness or efficiency of the algorithms presented because, as we will see later, the advantage of our algorithms comes from their slow noise growth, whose correctness is independent of the choice of p . We assume the ciphertext modulus $q = 2N$ is a power of 2 and view the message as an integer modulo q in the positive form. For an LWE ciphertext c encrypting $m = \frac{q}{p}m' + e$, we add $\frac{q}{2p}$ to c before performing any operations

to ensure that $e + \frac{q}{2p} \in [0, \frac{q}{p} - 1]$. This will simplify the understanding of homomorphic digit decomposition algorithms in Section 4 and is consistent with [LMP22]. To keep the description of the FDFB algorithms concise, we focus on input arguments like the LUT F and the input LWE ciphertext, omitting other arguments like the bootstrap key. In our noise analysis, we assume that the input ciphertext of the FDFB algorithms has an error variance of σ_{boot}^2 as in [LMP22]. The proof of correctness and noise analysis of the FDFB algorithms is provided in the full version of the paper.

3.1 FDFB-Compress

This algorithm employs the Type-HalfRange strategy. Specifically, it first compresses the coded message $\frac{q}{p}m' + e \in \mathbb{Z}_q$ into the range $[-\frac{q}{4}, \frac{q}{4} - 1]$ by evaluating the negacyclic function $f_C(x) : \mathbb{Z}_q \rightarrow \mathbb{Z}_q$ via a functional bootstrap, where

$$f_C(x) = \begin{cases} \frac{q}{2p}(\lfloor \frac{p}{q}x \rfloor + \frac{1}{2}) & x \in [0, \frac{q}{2} - 1] \\ -\frac{q}{2p}(\lfloor \frac{p}{q}x \rfloor - \frac{p}{2} + \frac{1}{2}) & x \in [\frac{q}{2}, q - 1] \end{cases}. \quad (1)$$

The design of f_C serves two purposes. Firstly, it maps messages encoding the same m' to the same value. Secondly, it ensures that the outputs of f_C for different m' s are at least $\frac{q}{2p}$ apart. $\frac{q}{2p}$ must be greater than 2β to prevent the bootstrapping noise from interfering with the compressed message. In other words, the plaintext modulus p is upper bounded by $p < \frac{q}{4\beta}$.

After compression, it is possible to bypass the negacyclicity constraint and evaluate an arbitrary LUT $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$ on the compressed message by using one functional bootstrap to compute $f_{eval} : \mathbb{Z}_q \rightarrow \mathbb{Z}_q$, which is defined as

$$f_{eval}(x) = \begin{cases} \lfloor \frac{q}{p}F(\lfloor \frac{2p}{q}x \rfloor) \rfloor & x \in [0, \frac{q}{4} - 1] \\ \lfloor \frac{q}{p}F(\lfloor \frac{2p}{q}(q - x) \rfloor + \frac{p}{2}) \rfloor & x \in [\frac{3q}{4}, q - 1] \\ -f_{eval}(x - \frac{q}{2}) & x \in [\frac{q}{4}, \frac{3q}{4} - 1] \end{cases}. \quad (2)$$

The algorithm for **FDFB-Compress** is fully described in Algorithm 1, with its parameter requirements and noise analysis provided in Theorem 1.

Algorithm 1: FDFB-Compress

input : Plaintext modulus p and an LUT $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$
input : An LWE ciphertext $(b, \vec{a}) = \text{LWE}_{\vec{s}, n, q}(\frac{q}{p}m' + e)$
output : An LWE ciphertext $\text{LWE}_{\vec{s}, n, q}(\frac{q}{p}F(m') + e_{boot})$
1 **ct** \leftarrow **Boot** $[f_C]((b + \frac{q}{2p}, \vec{a}))$
2 **return** **Boot** $[f_{eval}](\text{ct})$

Theorem 1. *Suppose $\beta < \frac{q}{4p}$ and $|e| < \frac{q}{2p}$, then **FDFB-Compress** $(F, \text{LWE}_{\vec{s}, n, q}(\frac{q}{p}m' + e)) = \text{LWE}_{\vec{s}, n, q}(\frac{q}{p}F(m') + e_{boot})$ and ct in line 1 of Algorithm 1 has an error variance of σ_{boot}^2 .*

3.2 FDFB-CancelSign

This algorithm employs the Type-SelectMSB strategy. Given $\text{LWE}_{\vec{s}, n, \frac{q}{2}}(\frac{q}{2p}m' + e)$, **FDFB-CancelSign** first executes **ModUp** to obtain a ciphertext $\text{LWE}_{\vec{s}, n, q}(\frac{q}{2}\text{MSB} + \frac{q}{2p}m' + e)$

and then performs a raw functional bootstrap to evaluate

$$f_{cs} = \begin{cases} \lfloor \frac{Q}{p} F(\lfloor \frac{2p}{q} x \rfloor) \rfloor & x \in [0, \frac{q}{2} - 1] \\ -f_{cs}(x - \frac{q}{2}) & x \in [\frac{q}{2}, q - 1] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_Q \quad (3)$$

and obtain a ciphertext encrypting $(-1)^{\text{MSB}} \lfloor \frac{Q}{p} F(m') \rfloor$. Finally, an LWE-to-RLWE packing key switching and another functional bootstrap cancel the extra $(-1)^{\text{MSB}}$ factor. The algorithm for **FDFB-CancelSign** is fully described in [Algorithm 2](#), and its parameter requirements and noise analysis are given in [Theorem 2](#).

Algorithm 2: FDFB-CancelSign

input : Plaintext modulus p and an LUT $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$
input : Base B_{pk} and modulus q_{pk} for PackingKS
input : $\{\text{ksk}'_{i,j,k}\}$, packing keys for PackingKS with $d = N$
input : An LWE ciphertext $(b, \vec{a}) = \text{LWE}_{\vec{s}, n, \frac{q}{2}}(\frac{q}{2p}m' + e)$
output : An LWE ciphertext $\text{LWE}_{\vec{s}, n, \frac{q}{2}}(\frac{q}{2p}F(m') + e')$

- 1 $\text{ct} \leftarrow \text{ModUp}((b + \frac{q}{4p}, \vec{a}), q)$
- 2 $\text{ct}_1 \leftarrow \text{BootRaw}[f_{cs}](\text{ct})$
- 3 $\text{ct}_{pk} \leftarrow \text{PackingKS}(\text{ct}_1, \{\text{ksk}'_{i,j,k}\})$
- 4 **return** $\text{Boot}[\text{ct}_{pk}](\text{ct})$

Theorem 2. *Suppose $|e| < \frac{q}{4p}$ and $|e'| < \frac{q}{4p}$, then $\mathbf{FDFB-CancelSign}(F, \text{LWE}_{\vec{s}, n, \frac{q}{2}}(\frac{q}{2p}m' + e)) = \text{LWE}_{\vec{s}, n, \frac{q}{2}}(\frac{q}{2p}F(m') + e')$. The output error e' has a variance of $(\frac{q}{Q})^2(2\sigma_{acc}^2 + \sigma_{pk}^2) + (\frac{q}{q_{pk}})^2\sigma_{ms}^2 + \sigma_{core}^2$. ciphertext.*

3.3 FDFB-Select

This algorithm employs the Type-SelectMSB strategy but does not perform the ModUp operation as in **FDFB-CancelSign**. In particular, let $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$ be an arbitrary LUT, let $\text{ct} = \text{LWE}_{\vec{s}, n, \frac{q}{2}}(\frac{q}{2p}m' + e)$ be a ciphertext encrypting m' , and let MSB be the most significant bit of m' . **FDFB-Select** first constructs two sub-LUTs from $\mathbb{Z}_{p/2}$ to \mathbb{Z}_p , which correspond to the LUT F with MSB = 0 or MSB = 1 respectively. These two sub-LUTs can be extended to $F_0, F_1 : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$ to fulfill the negacyclic constraint. i.e., $F_0(x) = F(x)$ and $F_1(x) = -F(x + p/2)$ for $x \in [0, p/2)$, $F_0(x) = -F(x - p/2)$ and $F_1(x) = F(x)$ for $x \in [p/2, p)$. F_0 and F_1 correspond to the functions in (4) and (5).

$$f_{pos} = \begin{cases} \lfloor \frac{Q}{p} F(\lfloor \frac{p}{q} x \rfloor) \rfloor & x \in [0, \frac{q}{2} - 1] \\ -f_{pos}(x - \frac{q}{2}) & x \in [\frac{q}{2}, q - 1] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_Q, \quad (4)$$

$$f_{neg} = \begin{cases} -f_{neg}(x + \frac{q}{2}) & x \in [0, \frac{q}{2} - 1] \\ \lfloor \frac{Q}{p} F(\lfloor \frac{p}{q} x \rfloor) \rfloor & x \in [\frac{q}{2}, q - 1] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_Q. \quad (5)$$

By evaluating these two functions on $\text{ct} + \frac{q}{2p}$ using a single functional bootstrap each, we can obtain two ciphertexts that encrypt $F_0(m')$ and $F_1(m')$, respectively. Additionally, we can obtain a ciphertext encrypting MSB by evaluating function (6) on $\text{ct} + \frac{q}{2p}$ using a single functional bootstrap.

$$f_{sgn} = \begin{cases} \frac{q}{8} & x \in [0, \frac{q}{2} - 1] \\ -\frac{q}{8} & x \in [\frac{q}{2}, q] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_Q. \quad (6)$$

Finally, we use the encryption of MSB to select $F_{\text{MSB}}(m')$ from $F_i(m')$ by a single functional bootstrap. The algorithm for **FDFB-Select** is fully described in [Algorithm 3](#), and its parameter requirements and noise analysis are given in [Theorem 3](#).

The first three functional bootstraps have the same input ciphertext ct , thus can be accomplished via a single multi-value bootstrap at the cost of increased noise growth. Therefore, when the parameter settings enable multi-value bootstrap, **FDFB-Select** needs only two functional bootstraps, otherwise it requires four functional bootstraps. In case multi-value bootstrap is unavailable, we develop a variant of **FDFB-Select**, called **FDFB-SelectAlt**, described in [Algorithm 4](#), which uses only three bootstraps. The parameter requirements and noise analysis of **FDFB-SelectAlt** are given in [Theorem 4](#).

Remark. We actually use an improved version of the base-aware LWE-to-RLWE packing proposed by [\[GBA21\]](#) to pack ct_{pos} and $-ct_{neg}$ into ct_{pk} . To pack $M|N$ messages $\text{LWE}(m_i)$ into $\text{RLWE}(\sum_{i=0}^{M-1} m_i(1+X+X^2+\dots+X^{\frac{N}{M}-1})X^{\frac{N}{M}i})$, [\[GBA21\]](#) generates M key switching keys, with each key corresponding to an index $i \in \llbracket M \rrbracket$. However, we observe that generating the key switching key for $i = 0$ is sufficient since the keys for $i \neq 0$ can be obtained by multiplying the key for $i = 0$ by $X^{\frac{N}{M}i}$. The storage cost of this optimized version of PackingKS is only $\frac{1}{M}$ that of [\[GBA21\]](#).

Algorithm 3: FDFB-Select

input : Plaintext modulus p and an LUT $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$
input : Base B_{pk} and modulus q_{pk} for PackingKS
input : $\{\text{ksk}'_{i,j,k}\}$, packing keys for PackingKS with $d = \frac{N}{2}$
input : An LWE ciphertext $(b, \vec{a}) = \text{LWE}_{\vec{s},n,q}(\frac{q}{p}m' + e)$
output : An LWE ciphertext $\text{LWE}_{\vec{s},n,q}(\frac{q}{p}F(m') + e')$

- 1 $ct \leftarrow (b + \frac{q}{2p}, \vec{a})$
- 2 $ct_{pos} \leftarrow \text{BootRaw}[f_{pos}](ct)$
- 3 $ct_{neg} \leftarrow \text{BootRaw}[f_{neg}](ct)$
- 4 $ct_{sgn} \leftarrow \text{Boot}[f_{sgn}](ct)$
- 5 $ct_{pk} \leftarrow \text{PackingKS}(ct_{pos}, \{\text{ksk}'_{i,j,k}\}) + \text{PackingKS}(-ct_{neg}, \{\text{ksk}'_{i,j,k}\}) \cdot X^{\frac{N}{2}}$
- 6 **return** $\text{Boot}[ct_{pk}](ct_{sgn})$

Theorem 3. Suppose $|e| < \frac{q}{2p}$, $\beta < \frac{q}{8}$ and $|e'| < \frac{q}{2p}$, then **FDFB-Select** $(F, \text{LWE}_{\vec{s},n,q}(\frac{q}{p}m' + e)) = \text{LWE}_{\vec{s},n,q}(\frac{q}{p}F(m') + e')$. The output error e' has a variance of $(\frac{q}{Q})^2(2\sigma_{acc}^2 + 2\sigma_{pk}^2) + (\frac{q}{q_{pk}})^2\sigma_{ms}^2 + \sigma_{com}^2$. Additionally, when multi-value bootstrap is employed, the variance becomes $(\frac{q}{Q})^2((p(p-1)^2 + 1)\sigma_{acc}^2 + 2\sigma_{pk}^2) + (\frac{q}{q_{pk}})^2\sigma_{ms}^2 + \sigma_{com}^2$.

Theorem 4. Suppose that $|e| < \frac{q}{2p}$ and $|e'| < \frac{q}{2p}$, then **FDFB-SelectAlt** $(F, \text{LWE}_{\vec{s},n,q}(\frac{q}{p}m' + e)) = \text{LWE}_{\vec{s},n,q}(\frac{q}{p}F(m') + e')$. The output error e' has a variance of $(\frac{q}{Q})^2(3\sigma_{acc}^2 + \sigma_{ks}^2) + (\frac{q}{q_{pk}})^2\sigma_{ms}^2 + \sigma_{com}^2$. Additionally, when multi-value bootstrap is employed, the variance becomes $(\frac{q}{Q})^2((6p(p-1)^2 + 1)\sigma_{acc}^2 + \sigma_{ks}^2) + (\frac{q}{q_{pk}})^2\sigma_{ms}^2 + \sigma_{com}^2$.

3.4 FDFB-BFVMult (WoPPBS₁-Refine)

This algorithm employs the Type-SelectMSB strategy but uses BFV multiplication to handle the MSB. It contains **WoPPBS₁-Refine** and **FDFB-BFVMult**.

Algorithm 4: FDFB-SelectAlt

input : Plaintext modulus p and an LUT $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$
input : Base B_{pk} and modulus q_{pk} for PackingKS
input : $\{\text{ksk}'_{i,j,k}\}$, packing keys for PackingKS with $d = N$
input : An LWE ciphertext $(b, \vec{a}) = \text{LWE}_{\vec{s},n,q}(\frac{q}{p}m' + e)$
output : An LWE ciphertext $\text{LWE}_{\vec{s},n,q}(\frac{q}{p}F(m') + e')$

- 1 $\text{ct} \leftarrow (b + \frac{q}{2p}, \vec{a})$
- 2 $\text{ct}_{diff} \leftarrow \text{BootRaw}[(f_{neg} - f_{pos})/2](\text{ct})$
- 3 $\text{ct}_{sum} \leftarrow \text{BootRaw}[(f_{neg} + f_{pos})/2](\text{ct})$
- 4 $\text{ct}_{pk} \leftarrow \text{PackingKS}(\text{ct}_{diff}, \{\text{ksk}'_{i,j,k}\})$
- 5 $\text{ct}_{res} \leftarrow \text{ct}_{sum} - \text{BootRaw}[\text{ct}_{pk}](\text{ct})$
- 6 $\text{ct}_{res} \leftarrow \text{KeySwitch}(\text{ModSwitch}(\text{ct}_{res}, q_{ks}), \{\text{ksk}_{i,j,k}\})$
- 7 **return** $\text{ModSwitch}(\text{ct}_{res}, q)$

The process of **WoPPBS₁-Refine** is identical to that of **WoP-PBS₁**, but it employs a much tighter noise analysis, as we will demonstrate later. It first obtains a ciphertext that encrypts $(-1)^{\text{MSB}} \lfloor \frac{Q}{p} F(m') \rfloor$ in the same way as **FDFB-CancelSign**. Then it evaluates the function (7) via a functional bootstrap to acquire the encryption of $\lfloor \frac{Q}{p} (-1)^{\text{MSB}} \rfloor$. Finally, it computes the product of the two LWE ciphertexts using LWE-to-RLWE packing and BFV multiplication. The algorithm is fully described in [Algorithm 5](#), and its parameter requirements and noise analysis are given in [Theorem 5](#).

$$f_{sgn1} = \begin{cases} \lfloor \frac{Q}{p} \rfloor & x \in [0, \frac{q}{2} - 1] \\ Q - \lfloor \frac{Q}{p} \rfloor & x \in [\frac{q}{2}, q - 1] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_Q \quad (7)$$

Algorithm 5: WoPPBS₁-Refine

input : Plaintext modulus p and an LUT $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$
input : Base B_{ks} and modulus q_{ks} for key switching
input : Base B_{pk} and modulus q_{pk} for PackingKS
input : $\{\text{ksk}_{i,j,k}\}$, key switching keys
input : $\{\text{ksk}'_{i,j,k}\}$, packing keys for PackingKS with $d = 1$
input : An LWE ciphertext $(b, \vec{a}) = \text{LWE}_{\vec{s},n,\frac{q}{2}}(\frac{q}{2p}m' + e)$
output : An LWE ciphertext $\text{LWE}_{\vec{s},n,\frac{q}{2}}(\frac{q}{2p}F(m') + e')$

- 1 $\text{ct} \leftarrow \text{ModUp}((b + \frac{q}{4p}, \vec{a}), q)$
- 2 $\text{ct}_0 \leftarrow \text{PackingKS}(\text{BootRaw}[f_{cs}](\text{ct}), \{\text{ksk}'_{i,j,k}\})$
- 3 $\text{ct}_{sgn} \leftarrow \text{PackingKS}(\text{BootRaw}[f_{sgn1}](\text{ct}), \{\text{ksk}'_{i,j,k}\})$
- 4 $\text{ct}_{prod} \leftarrow \text{SampleExtract}(\text{BFVMult}(\text{ct}_0, \text{ct}_{sgn}), 0)$
- 5 $\text{ct}_{res} \leftarrow \text{KeySwitch}(\text{ModSwitch}(\text{ct}_{prod}, q_{ks}), \{\text{ksk}_{i,j,k}\})$
- 6 **return** $\text{ModSwitch}(\text{ct}_{res}, \frac{q}{2})$

FDFB-BFVMult is an improved version of **WoP-PBS₂**. Unlike **WoP-PBS₂**, which requires the sign bit to be multiplied with both $f_{neg}(\text{ct})$ and $f_{pos}(\text{ct})$, **FDFB-BFVMult** only needs one BFV multiplication because the sign bit is multiplied with the fresh ciphertext $(f_{neg} - f_{pos})(\text{ct})$. Consequently, **FDFB-BFVMult** further halves the noise growth. Specifically, **FDFB-BFVMult** first constructs two LUTs F_0 and F_1 in the same way as **FDFB-Select**. Next, by using two functional bootstraps to evaluate f_{pos} and $f_{neg} - f_{pos}$ (defined in (4) and (5)), it obtains encryptions of $m_{pos} = \lfloor \frac{Q}{p} F_0(m') \rfloor$ and $m_{diff} = \lfloor \frac{Q}{p} (F_1 - F_0)(m') \rfloor$. Then it evaluates the function (8) via a functional bootstrap to

acquire the encryption of $m_{sgn} = \lfloor -\frac{Q}{2p}(-1)^{\text{MSB}} \rfloor + \lfloor \frac{Q}{2p} \rfloor \approx \lfloor \frac{Q}{p} \text{MSB} \rfloor$. Finally, it computes $\text{MSB} \cdot m_{diff} + m_{pos} \approx \lfloor \frac{Q}{p} F_{\text{MSB}}(m') \rfloor$ using LWE-to-RLWE packing and BFV multiplication. The algorithm is fully described in [Algorithm 6](#), and its parameter requirements and noise analysis are given in [Theorem 6](#).

Since the two bootstraps in **WoPPBS₁-Refine** (and the three bootstraps in **FDFB-BFVMult**) share the same input, they can be accelerated by employing a single multi-value bootstrap at the cost of increased noise growth.

$$f_{sgn2} = \begin{cases} Q - \lfloor \frac{Q}{2p} \rfloor & x \in [0, \frac{q}{2} - 1] \\ \lfloor \frac{Q}{2p} \rfloor & x \in [\frac{q}{2}, q - 1] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_Q \quad (8)$$

Algorithm 6: FDFB-BFVMult

input : Plaintext modulus p and an LUT $F : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$
input : Base B_{ks} and modulus q_{ks} for key switching
input : Base B_{pk} and modulus q_{pk} for PackingKS
input : $\{\text{ksk}_{i,j,k}\}$, key switching keys
input : $\{\text{ksk}'_{i,j,k}\}$, packing keys for PackingKS with $d = 1$
input : An LWE ciphertext $(b, \vec{a}) = \text{LWE}_{\vec{s}, n, q}(\frac{q}{p}m' + e)$
output : An LWE ciphertext $\text{LWE}_{\vec{s}, n, q}(\frac{q}{p}F(m') + e')$

- 1 $\text{ct} \leftarrow (b + \frac{q}{2p}, \vec{a})$
- 2 $\text{ct}_{pos} \leftarrow \text{BootRaw}[f_{pos}](\text{ct})$
- 3 $\text{ct}_{diff} \leftarrow \text{PackingKS}(\text{BootRaw}[f_{neg} - f_{pos}](\text{ct}), \{\text{ksk}'_{i,j,k}\})$
- 4 $\text{ct}_{sgn} \leftarrow \text{PackingKS}(\text{BootRaw}[f_{sgn2}](\text{ct}) + \lfloor \frac{Q}{2p} \rfloor, \{\text{ksk}'_{i,j,k}\})$
- 5 $\text{ct}_{prod} \leftarrow \text{SampleExtract}(\text{BFVMult}(\text{ct}_{diff}, \text{ct}_{sgn}), 0)$
- 6 $\text{ct}_{res} \leftarrow \text{ct}_{prod} + \text{ct}_{pos}$
- 7 $\text{ct}_{res} \leftarrow \text{KeySwitch}(\text{ModSwitch}(\text{ct}_{res}, q_{ks}), \{\text{ksk}_{i,j,k}\})$
- 8 **return** $\text{ModSwitch}(\text{ct}_{res}, q)$

Refined BFV Noise Analysis. Next, we provide a refined noise analysis for the BFV multiplication involved in **FDFB-BFVMult (WoPPBS₁-Refine)**. Our core observation is that in LWE-to-RLWE packing, only the constant term of the output polynomial message is assigned the value of the input LWE message, while the coefficients of non-constant terms are close to 0.

[Lemma 2](#) provides a noise analysis of this kind of BFV multiplication. We note that only the dominating term of the error variance is displayed in [Lemma 2](#) (as well as in [Theorem 5](#) and [Theorem 6](#)) due to the complexity of the full formula. Refer to the full version of the paper for the full formula and its proof.

In **FDFB-BFVMult (WoPPBS₁-Refine)**, each of the multiplicands for BFV multiplication is obtained by packing an LWE message with an error variance of σ_{acc}^2 into the constant term of an RLWE ciphertext. This means that the constant term of the encrypted polynomial has an error variance of $\sigma_{acc}^2 + \sigma_{ks}^2$, while the error variance of non-constant terms is σ_{ks}^2 . Note that σ_{acc}^2 and σ_{ks}^2 correspond to σ_i^2 and $\sigma_i^{2'}$ in [Lemma 2](#). In practice, σ_{acc}^2 is much larger than $N\sigma_{ks}^2$ and one of the packed LWE messages is a sign bit (i.e., in $\{0, \pm 1\}$). It then follows from [Lemma 2](#) that the output error variance is about $2p^2\sigma_{ms}^2\sigma_{acc}^2$.

On the other hand, for ordinary BFV multiplication where all terms have an error variance of $\sigma_{acc}^2 + \sigma_{ks}^2$, the output error variance is about $2Np^2\sigma_{ms}^2\sigma_{acc}^2$. This is because the dominating noise term becomes a polynomial-polynomial multiplication and introduces an extra factor N compared to scalar-polynomial multiplication (refer to the remark in the

full version of the paper for details). This means the noise growth is reduced by roughly N times compared to conventional BFV multiplication.

Lemma 2. *Let $c_i = (\mathbf{b}_i, \mathbf{a}_i) = \text{RLWE}_{s,N,Q}(\frac{Q}{p}m_i + e_i + \mathbf{e}_i)$ for $i = 0, 1$, where $e_i \sim N(0, \sigma_i^2)$, $\mathbf{e}_i \sim N(0, \sigma_i'^2)^N$, $\sigma_i^2 \gg N\sigma_i'^2$ and $m_0 \in \{0, \pm 1\}$. Then $\text{SampleExtract}(\text{BFVMult}(c_0, c_1), 0) = \frac{Q}{p}m_0m_1 + e$ and the variance of e is equal to $p^2\sigma_{ms}^2(\sigma_0^2 + \sigma_1^2)$ approximately¹.*

Theorem 5. *Suppose $|e| < \frac{q}{4p}$ and $|e'| < \frac{q}{4p}$, then $\text{WoPPBS}_1\text{-Refine}(F, \text{LWE}_{\bar{s},n,\frac{q}{2}}(\frac{q}{2p}m' + e)) = \text{LWE}_{\bar{s},n,\frac{q}{2}}(\frac{q}{2p}F(m') + e')$. The output error e' has a variance of $(\frac{q}{Q})^2\frac{N}{9}p^2\sigma_{acc}^2 + \sigma_{com}^2$ approximately. Additionally, when multi-value bootstrap is employed, the variance becomes $(\frac{q}{Q})^2\frac{N}{18}p^3(p-1)^2\sigma_{acc}^2 + \sigma_{com}^2$ approximately.*

Theorem 6. *Suppose $|e| < \frac{q}{2p}$ and $|e'| < \frac{q}{2p}$, then $\text{FDFB-BFVMult}(F, \text{LWE}_{\bar{s},n,q}(\frac{q}{p}m' + e)) = \text{LWE}_{\bar{s},n,q}(\frac{q}{p}F(m') + e')$. The output error e' has a variance of $(\frac{q}{Q})^2\frac{N}{9}p^2\sigma_{acc}^2 + \sigma_{com}^2$ approximately. Additionally, when multi-value bootstrap is employed, the variance becomes $(\frac{q}{Q})^2\frac{2N}{9}p^3(p-1)^2\sigma_{acc}^2 + \sigma_{com}^2$ approximately.*

4 Improved Homomorphic Digit Decomposition

This section presents two algorithms **HomDecomp-Reduce** and **HomDecomp-FDFB** to decompose an LWE ciphertext with a large modulus q_0 into multiple LWE ciphertexts with a smaller modulus q , each encrypting a digit of the original message. **HomDecomp-Reduce** creates buffer space for modulus switching noise by reducing the range of lower bits by half. It can handle digits of up to 4 bits and requires one bootstrap operation per decomposed digit. In contrast, **HomDecomp-FDFB** clears the lower bits approximately and can handle digits of up to 5 bits, but it requires two bootstrap operations per digit. We still assume $q = 2N$ as in the previous section. In our noise analysis, we assume that the input ciphertext of the decomposition algorithms has an error variance of σ_{boot}^2 as in [LMP22]. Proof of theorems is left to the full version of the paper due to space limit.

4.1 HomDecomp-Reduce

In **HomDecomp-Reduce**, the range of lower bits is first reduced by half using one bootstrap operation to accommodate the subsequent modulus switching noise. The reduction function $f_{red} : \mathbb{Z}_q \rightarrow \mathbb{Z}_{q_0}$ is defined in (9), with different input and output ranges.

$$f_{red} = \begin{cases} \frac{q}{4} & x \in [0, \frac{q}{2} - 1] \\ q_0 - \frac{q}{4} & x \in [\frac{q}{2}, q - 1] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_{q_0} \quad (9)$$

The complete algorithm is described in Algorithm 7. Its parameter requirements and noise analysis are given in Theorem 7.

Theorem 7. *If $\text{bnd}\sqrt{B^{-2}\sigma_{boot}^2 + \sigma_{ms}^2} < \frac{q}{4B}$, **HomDecomp-Reduce** outputs the decomposed digits correctly.*

¹Here, ‘approximately’ means that only the dominant term of the error variance is displayed, as the full formula is quite complex. For the full formula and an explanation of the approximation, please refer to the full version of the paper.

Algorithm 7: HomDecomp-Reduce

input : A base B for homomorphic decomposition
input : An LWE ciphertext $ct = \text{LWE}_{\vec{s}, n, q_0}(\frac{q_0}{p}m' + e)$
output : LWE ciphertexts $\{ct_i\}$ encrypting the digits of m'

```

1  $i \leftarrow 0$ 
2 while  $q_0 > q$  do
3    $ct_i \leftarrow \text{ModDown}(ct, q)$ 
4    $ct \leftarrow ct + (\frac{q_0}{2p}, \vec{0})$ 
5    $ct' \leftarrow \text{ModDown}(ct, q)$ 
6    $ct \leftarrow ct + \text{Boot}[f_{red}](ct') - (\frac{q}{2}, \vec{0})$ 
7    $ct \leftarrow \text{ModSwitch}(ct, \frac{q_0}{B})$ 
8    $i \leftarrow i + 1$ 
9  $ct_i \leftarrow ct$ 
10 return  $\{ct_i\}$ 

```

4.2 HomDecomp-FDFB

In **HomDecomp-FDFB**, we use **FDFB-Compress** to evaluate the continuous identity function $f_{id}(x) = x : \mathbb{Z}_q \rightarrow \mathbb{Z}_{q_0}$ (using zero extension), and the obtained result is used to approximately clear the lower bits in the input ciphertext. See [Algorithm 8](#) for a full description of **HomDecomp-FDFB** and [Theorem 9](#) for its parameter requirements and noise analysis.

Before beginning, we show how to evaluate a continuous function F' with **FDFB-Compress**, where the input and output scaling factors are Δ_{in} and Δ_{out} respectively. First, the compression function f_C in (1) is substituted with f'_C , which is defined in (10) and illustrated in [Figure 3](#).

$$f'_C = \begin{cases} \lfloor \frac{\frac{q}{2}-2\beta}{\frac{q}{2}-1}x + \beta \rfloor & x \in [0, \frac{q}{2} - 1] \\ q - f'_C(x - \frac{q}{2}) & x \in [\frac{q}{2}, q - 1] \end{cases} : \mathbb{Z}_q \rightarrow \mathbb{Z}_q \quad (10)$$

The strategy adopted to construct f'_C is called ‘ β -padding’, which creates a 2β distance between $f'_C(0)$ and $f'_C(\frac{q}{2})$ to separate the cases where the input is 0 and $\frac{q}{2}$. Otherwise, the bootstrapping error may intermix the two cases, making it impossible for f_{eval} to distinguish between them. As a result, when the input is positive and near 0, **FDFB-Compress** may yield an incorrect result $F'(-\frac{q}{2\Delta_{in}})$ instead of $F'(0)$. Also, $f'_C(\frac{q}{2} - 1)$ and $f'_C(q - 1)$ must be β away from $\frac{q}{4}$ and $\frac{3q}{4}$ respectively to ensure that the output message of f'_C always lies within half of \mathbb{Z}_q .

The modified version of f_{eval} in (2) (which we denote as f'_{eval}) is rather complicated. Intuitively f'_{eval} aims to recover the original input to f'_C , evaluate F' on the recovered input, and subsequently scales the result by Δ_{out} . As the evaluation of f'_C introduces a bootstrapping error, the input recovered by f'_{eval} also contains a bootstrapping error (multiplied by some constant), which means that the output error of **FDFB-Compress** depends on the Lipschitz constant of F' . The output error variance is given in [Theorem 8](#), and the proof can be found in the full version of the paper.

Theorem 8. *When evaluating a continuous function f with Lipschitz constant L , the output error variance of **FDFB-Compress** is $(\frac{L}{k_2\Delta_{in}})^2\sigma_{boot}^2 + \Delta_{out}^{-2}\sigma_{boot}^2$, where $k_2 = \frac{\frac{N}{2}-2\beta}{N-1} \approx \frac{1}{2}$.*

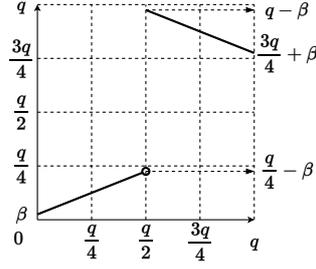


Figure 3: Compression function f'_C for continuous function evaluation.

HomDecomp-FDFB sets $\Delta_{in} = \Delta_{out} = 1$ and $F' = f_{id}$, which gives the following theorem.

Theorem 9. *Let e_f be the output error of **FDFB-Compress**, then its variance is $\sigma_f^2 = (1 + k_2^{-2})\sigma_{boot}^2$. If $\text{bnd}\sqrt{B^{-2}\sigma_f^2 + \sigma_{ms}^2} < \frac{q}{2B}$, **HomDecomp-FDFB** outputs the decomposed digits correctly.*

Algorithm 8: HomDecomp-FDFB

input : A base B for homomorphic decomposition
input : An LWE ciphertext $ct = \text{LWE}_{\vec{s}, n, q_0}(\frac{q_0}{p}m' + e)$
output : LWE ciphertexts $\{ct_i\}$ encrypting the digits of m'

- 1 $i \leftarrow 0$
- 2 **while** $q_0 > q$ **do**
- 3 $ct_i \leftarrow \text{ModDown}(ct, q)$
- 4 $ct \leftarrow ct + (\frac{q_0}{2p}, \vec{0})$
- 5 $ct' \leftarrow \text{ModDown}(ct, q)$
- 6 $ct \leftarrow ct - \text{FDFB-Compress}[f_{id}](ct')$
- 7 $ct \leftarrow \text{ModSwitch}(ct, \frac{q_0}{B})$
- 8 $i \leftarrow i + 1$
- 9 $ct_i \leftarrow ct$
- 10 **return** $\{ct_i\}$

5 Analysis and Comparison

This section analyzes the FDFB and the homomorphic decomposition algorithms, both previous ones and ours, concerning their noise growth and the number of required bootstraps.

5.1 Analysis of FDFB Algorithms

Table 3 presents the error variance ratio between our and previous FDFB algorithms and the number of bootstraps required. For Type-HalfRange FDFB algorithms (**FDFB-Compress** and **EvalFunc**), the coded message must first be compressed into half of \mathbb{Z}_q . Thus the error of the compressed message (e.g., the error in ct of line 1 of Algorithm 1) plays a major role in the selection of parameters. For Type-SelectMSB FDFB algorithms (other algorithms in Table 3), the output error plays a major role in the selection of parameters. The dominant term of the output error variance is the σ_{acc}^2 -term for most

Table 3: Comparison of previous and our FDFB algorithms regarding their noise growth and the number of bootstraps required.

Ours	Previous	Error Var Ratio (Ours/Prev)	Num of BTS (Prev→Ours)
FDFB-Compress	EvalFunc [LMP22]	1/2	2 → 2
FDFB-CancelSign	WoP-PBS₁ [CLOT21]	18/N ² p ²	2 → 2
WoPPBS₁-Refine		1/N	2 → 2
FDFB-Select	WoP-PBS₂ [CLOT21]	9/N ² p ²	3 → 4
FDFB-SelectAlt		27/2N ² p ²	3 → 3
FDFB-BFVMult		1/N	3 → 3
WoPPBS₁-Refine*	WoP-PBS₁* [CLOT21]	1/N	1 → 1
FDFB-Select*	WoP-PBS₂* [CLOT21]	9/2N ² p ²	1 → 2
FDFB-SelectAlt*		27/N ² p ²	1 → 2
FDFB-BFVMult*		1/N	1 → 1

* FDFB algorithms that use multi-value bootstrap.

algorithms (refer to the full version of the paper for the formula of the output error variance of all algorithms). Thus, in the table, the first row of the ratio column represents the ratio of the error variances of the compressed message. The remaining rows of the ratio column represent the ratios of the σ_{acc}^2 -terms of the output error variance. For **FDFB-CancelSign**, **FDFB-Select** and **FDFB-SelectAlt**, the ratios of the output error variance can be a small multiple of the displayed ones. For other algorithms, the output error variance ratios are very close to the displayed ones since the σ_{acc}^2 -term is dominant.

As stated earlier, the efficiency of an FDFB algorithm is not solely determined by the number of bootstraps it requires. The error variances also impact the compactness of parameters and thus affect the final efficiency. As shown in Table 3, the main advantage of our FDFB algorithms is their reduced noise growth. This allows for the selection of larger decomposition bases during blind rotation, resulting in a reduction in the decomposition dimension (denoted by l as described in Section 2.2.2). Since the number of NTTs required for a blind rotation is proportional to $(l + 1)$, our algorithms achieve better performance. To be more specific:

- **FDFB-Compress** reduces the error variance of the compressed message by half, resulting in a more relaxed parameter choice than **EvalFunc**.
- **FDFB-CancelSign**, **FDFB-Select**, **FDFB-SelectAlt** and their multi-value bootstrap variants use LWE-to-RLWE packing and blind rotation instead of BFV multiplication. This reduces the noise to $O(1/N^2p^2)$ that of **WoP-PBS**. Although our algorithms require an additional bootstrap to replace the BFV multiplication, we demonstrate in Section 6 that they are still faster than **WoP-PBS** in most cases due to their slower noise growth.
- **WoPPBS₁-Refine** and **FDFB-BFVMult** use significantly tighter noise analysis for BFV multiplication than **WoP-PBS₁** and **WoP-PBS₂**, reducing the noise growth to $1/N$ the original value.

The Optimality of FDFB-Compress. We observe that **FDFB-Compress** achieves optimality among Type-HalfRange algorithms. Recall that Type-HalfRange first uses functional bootstraps to transform the coded message $\frac{q}{p}m' + e \in \mathbb{Z}_q$ into $\phi(m') + \tilde{e} \in U \subseteq \mathbb{Z}_q$ and then evaluate the LUT with another functional bootstrap, where ϕ is an arbitrary map, U satisfies $U \cap (U + \frac{q}{2}) = \emptyset$ to bypass the negacyclic constraint, and \tilde{e} has a variance of at least $\sigma_{\tilde{e}}^2 \geq \sigma_{boot}^2$. Additionally, to ensure the correctness of evaluation, m' must be reconstructible from $\tilde{m} + \tilde{e}$, i.e., there is a map λ from U to \mathbb{Z}_p such that $\lambda(\phi(m') + \tilde{e}) = m'$

Table 4: Comparison of previous and our homomorphic decomposition algorithms.

Ours	HomDecomp-Reduce	HomDecomp-FDFB
Previous	HomFloor [LMP22]	HomFloorAlt [LMP22]
Number of BTS (Previous→Ours)	2 → 1	3 → 2
Constraints of Previous Methods	Cannot decompose extracted CKKS ciphertexts	$q > 8\sqrt{2}\beta$

for any $m' \in \mathbb{Z}_p$ and any $|\tilde{e}| < \text{bnd} \cdot \sigma_{\tilde{e}}$.

Thus, on the one hand, **FDFB-Compress** achieves the minimum number of bootstraps required for Type-HalfRange (i.e., 2). On the other hand, since $\phi(m') + \tilde{e} \in \lambda^{-1}(m')$, by the pigeonhole principle there exists an $m' \in \mathbb{Z}_p$ such that $|\lambda^{-1}(m')| \leq \frac{|\mathcal{U}|}{p} \leq \frac{q}{2p}$, implying $\frac{q}{2p} > 2 \cdot \text{bnd} \cdot \sigma_{\tilde{e}} \geq 2\beta$. This requires $\beta < \frac{q}{4p}$, which is also the *only* requirement for **FDFB-Compress**. This means that **FDFB-Compress** achieves the most compact parameter choice among Type-HalfRange algorithms, thus achieving optimality.

5.2 Analysis of Homomorphic Decomposition

Table 4 compares the number of bootstraps needed for previous and our homomorphic digit decomposition algorithms. Algorithms in the same row of the table share the same digit decomposition base B (i.e., their decomposed digits have the same plaintext modulus). According to the table, our algorithms need one less bootstrap than previous algorithms in [LMP22]. **HomFloor** requires that the input ciphertext encodes a discrete plaintext with small noise, which ensures a gap between two adjacent encoded messages to accommodate the noise introduced by subsequent bootstraps. Since an extracted CKKS ciphertext encodes messages continuously without any gaps, **HomFloor** cannot be applied to decompose it. Also, **HomFloorAlt** has an extra constraint for the ciphertext modulus. In contrast, our methods are free of these constraints, making them more flexible than previous methods. The full version of the paper provides a theoretical analysis of the noise growth and parameter choice.

6 Implementation

We implement all the FDFB algorithms and homomorphic decomposition algorithms, including both previous ones and ours, in OpenFHE [BBB⁺22] (commit id 745a492). We disable multi-threading, except during key generation. We build OpenFHE using the g++ compiler of version 12.2.1 with flag `WITH_NATIVEOPT=ON` (as the authors did in [LMP22]). The performance of algorithms is tested on a machine with Intel(R) Xeon(R) Gold 6248R CPU @ 3.00GHz and 125G of RAM, running Fedora Release 36.

Parameter Setting. We use two parameter sets in our LWE schemes, i.e., $\text{PARAM}_{\text{decomp}}$ and $\text{PARAM}_{\text{fast}}$, which have been verified to meet 128-bit security using lattice-estimator [APS15] (commit id 48fa49b). Table 5 presents the details of these parameter

Table 5: Parameter sets for LWE scheme and their use cases.

LWE Param Sets	n	q_{ks}	Use Cases
$\text{PARAM}_{\text{decomp}}$	1340	2^{35}	HomDecomp, Discrete FDFB
$\text{PARAM}_{\text{fast}}$	760	2^{20}	Discrete FDFB

Table 6: Running time of previous and our FDFB algorithms under four scenarios (A to D). Each running time is obtained by averaging over 100 tests and is measured in milliseconds (ms). For each scenario, the best algorithms from previous works and this paper are marked in blue and red, respectively. A ‘/’ indicates that the algorithm is unavailable in that scenario because the plaintext modulus p exceeds its parameter requirements.

Algorithm	PARAM _{decomp}		PARAM _{fast}	
	A: $p = 2^4$	B: $p = 2^5$	C: $p = 2^4$	D: $p = 2^5$
EvalFunc	/	/	598	/
WoP-PBS₁*	1160	/	682	/
WoP-PBS₂*	1200	1930	735	942
FDFB-KS	5360	6340	2940	3110
Comp*	1580	1760	897	985
FDFB-Compress	1050	/	598	/
FDFB-CancelSign	1060	/	611	/
FDFB-Select*	1260	1250	621	724
FDFB-SelectAlt*	1240	1250	717	718
WoPPBS₁-Refine*	777	/	458	/
FDFB-BFVMult*	785	1150	458	573

sets, and we briefly explain the selection criteria of q_{ks} below since n can be determined from q_{ks} . For PARAM_{decomp}, the maximum ciphertext modulus is set to 2^{35} such that the ciphertext to be digit-decomposed has a large modulus. This choice for q_{ks} is also consistent with [LMP22]. For PARAM_{fast}, we focus on FDFB algorithms for discrete LUTs. Thus q_{ks} can be set to a smaller value to accelerate FDFB. However, if q_{ks} is too small, it may lead to large key switching noise, corrupting the correctness of FDFB. Therefore, we set $q_{ks} = 2^{20}$ in PARAM_{fast}.

The performance of discrete LUT evaluation with FDFB variants is tested with the plaintext modulus set to 2^4 and 2^5 . To ensure fair comparisons, we have only recorded the best performance among the parameters for FDFB variants with multiple parameter choices (e.g., multi-value or not). In our experiments, the multi-value versions usually run faster than the non-multi-value ones. Thus, the multi-value versions of most algorithms are recorded.

Please refer to the full version of the paper for a complete list of the parameters used in the benchmarks.

Performance of FDFB Algorithms. Table 6 shows the running time of previous and our FDFB algorithms under four scenarios (two parameter sets \times two choices of p). We can draw the following conclusions from the benchmark data.

First, the experiment data validate our algorithms’ advantage over their predecessors, as suggested theoretically in Section 5. To be more specific:

- **FDFB-Compress** can support $p = 2^4$ in scenario A while **EvalFunc** cannot because the former benefits from a reduced error variance of the compressed message. In fact, **EvalFunc** would need to double the RLWE dimension N to support $p = 2^4$, which leads to worse efficiency.

- **FDFB-CancelSign** shows a speedup of 8.6%~10.4% compared to **WoP-PBS₁***, even though it requires one additional bootstrap and does not use multi-value bootstrap for acceleration. This is due to the slower noise growth of **FDFB-CancelSign**, which allows for the choice of a larger decomposition base B_g in blind rotation, resulting in improved performance. On the other hand, **FDFB-Select*** and **FDFB-SelectAlt*** have similar running time to **WoP-PBS₂*** in scenarios A & C and are 23.1%~35.2% faster than **WoP-PBS₂*** in scenarios B & D. This advantage grows with p , as a larger p results in less

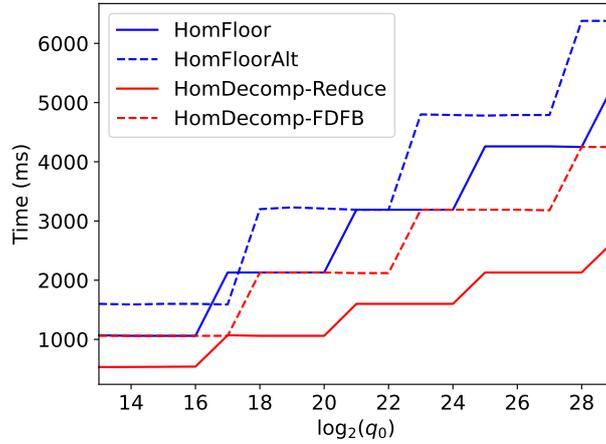


Figure 4: The "running time"- "input precision" graph for previous (blue) and our (red) homomorphic digit decomposition algorithms under $\text{PARAM}_{\text{decomp}}$.

tolerance for homomorphic noise and forces prior methods to use smaller B_g , degrading their performance.

- **WoPPBS₁-Refine*** is 32.8%~33.0% faster than **WoP-PBS₁*** and **FDFB-BFVMult*** is 37.7%~40.4% faster than **WoP-PBS₂***. Again, such performance improvement benefits from the choice of a larger B_g , which is possible due to the algorithms' reduced noise growth.

Second, when comparing the fastest algorithms from previous works and our algorithms, we observe a 23.4%~39.2% reduction in running time across all four scenarios (see Table 7). Among our algorithms, **FDFB-BFVMult*** is the fastest or very close to the fastest in all the scenarios. However, it does not render our other algorithms obsolete because (1) they support the addition of more bootstrapped ciphertexts since they have smaller output error than **FDFB-BFVMult** (**WoPPBS₁-Refine**); (2) they are useful for smaller RLWE dimensions, where BFV-based FDFB methods might be unavailable.

Performance of Homomorphic Digit Decomposition. Figure 4 illustrates the performance of different homomorphic decomposition algorithms (the raw data can be found in the full version of the paper). Data for $B = 2^4$ are drawn in solid lines, while data for $B = 2^5$ are drawn in dashed lines. For all choices of $\log_2(q_0)$, **HomDecomp-Reduce** runs roughly twice as fast as **HomFloor**, and **HomDecomp-FDFB** runs roughly at 1.5 times the speed of **HomFloorAlt**. Such speedup in homomorphic decomposition directly leads to speedup in the large-precision sign/ReLU/max/ABS evaluation, as they all require extracting the MSB of the input message.

Table 7: Performance improvement of our FDFB algorithms.

Scenario in Table 6	A	B	C	D
Best Running Time (Old, ms)	1160	1760	598	942
Best Running Time (New, ms)	777	1150	458	573
Reduction in Running Time	33.0%	34.7%	23.4%	39.2%

7 Conclusion

This paper develops four FDFB algorithms and two homomorphic decomposition algorithms. Our FDFB algorithms achieve a running time shorter than the best known results by up to 39.2%. Our homomorphic decomposition algorithms run 1.5x to 2x as fast as those presented in [LMP22], leading to speedup in large-precision ReLU, sign, max and ABS evaluation. We give a thorough theoretical noise analysis for FDFB and homomorphic decomposition algorithms, both in prior works and ours. We also implement all the algorithms in OpenFHE for a fair comparison between them.

Acknowledgments

This work is supported by the National Key R&D Program of China (2018YFA0704701, 2020YFA0309705), Shandong Key Research and Development Program (2020ZLYS09), the Major Scientific and Technological Innovation Project of Shandong, China (2019JZZY010133), the Major Program of Guangdong Basic and Applied Research (2019B030302008), and Tsinghua University Dushi Program.

References

- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
- [BBB⁺22] Ahmad Al Badawi, Jack Bates, Flavio Bergamaschi, David Bruce Cousins, Saroja Erabelli, Nicholas Genise, Shai Halevi, Hamish Hunt, Andrey Kim, Yongwoo Lee, Zeyu Liu, Daniele Micciancio, Ian Quah, Yuriy Polyakov, Saraswathy R.V., Kurt Rohloff, Jonathan Saylor, Dmitriy Suponitsky, Matthew Triplett, Vinod Vaikuntanathan, and Vincent Zucca. OpenFHE: Open-Source Fully Homomorphic Encryption Library. Cryptology ePrint Archive, Paper 2022/915, 2022. <https://eprint.iacr.org/2022/915>.
- [BGGJ20] Christina Boura, Nicolas Gama, Mariya Georgieva, and Dimitar Jetchev. CHIMERA: Combining Ring-LWE-based Fully Homomorphic Encryption Schemes. *Journal of Mathematical Cryptology*, 14(1):316–338, 2020.
- [BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) Fully Homomorphic Encryption without Bootstrapping. *ACM Trans. Comput. Theory*, 6(3), July 2014.
- [BIP⁺22] Charlotte Bonte, Ilia Iliashenko, Jeongeun Park, Hilder V. L. Pereira, and Nigel P. Smart. FINAL: Faster FHE Instantiated with NTRU and LWE. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022*, pages 188–215, Cham, 2022. Springer Nature Switzerland.
- [BMMP18] Florian Bourse, Michele Minelli, Matthias Minihold, and Pascal Paillier. Fast Homomorphic Evaluation of Deep Discretized Neural Networks. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018*, pages 483–512, Cham, 2018. Springer International Publishing.
- [CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast Fully Homomorphic Encryption Over the Torus. *Journal of Cryptology*, 33(1):34–91, 2020.

- [CHLR18] Hao Chen, Zhicong Huang, Kim Laine, and Peter Rindal. Labeled PSI from Fully Homomorphic Encryption with Malicious Security. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, page 1223–1237, New York, NY, USA, 2018. Association for Computing Machinery.
- [CIM19] Sergiu Carpov, Malika Izabachène, and Victor Mollimard. New Techniques for Multi-value Input Homomorphic Evaluation and Applications. In Mitsuru Matsui, editor, *Topics in Cryptology – CT-RSA 2019*, pages 106–126, Cham, 2019. Springer International Publishing.
- [CJP21] Ilaria Chillotti, Marc Joye, and Pascal Paillier. Programmable Bootstrapping Enables Efficient Homomorphic Inference of Deep Neural Networks. In Shlomi Dolev, Oded Margalit, Benny Pinkas, and Alexander Schwarzmann, editors, *Cyber Security Cryptography and Machine Learning*, pages 1–19, Cham, 2021. Springer International Publishing.
- [CKKS17] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. Homomorphic Encryption for Arithmetic of Approximate Numbers. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017*, pages 409–437, Cham, 2017. Springer International Publishing.
- [CLOT21] Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Improved Programmable Bootstrapping with Larger Precision and Efficient Arithmetic Circuits for TFHE. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021*, pages 670–699, Cham, 2021. Springer International Publishing.
- [CLR17] Hao Chen, Kim Laine, and Peter Rindal. Fast Private Set Intersection from Homomorphic Encryption. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, page 1243–1255, New York, NY, USA, 2017. Association for Computing Machinery.
- [CMdG⁺21] Kelong Cong, Radames Cruz Moreno, Mariana Botelho da Gama, Wei Dai, Ilia Iliashenko, Kim Laine, and Michael Rosenberg. Labeled PSI from Homomorphic Encryption with Reduced Computation and Communication. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, CCS '21*, page 1135–1150, New York, NY, USA, 2021. Association for Computing Machinery.
- [CZ22] Olive Chakraborty and Martin Zuber. Efficient and Accurate Homomorphic Comparisons. In *Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography, WAHC'22*, page 35–46, New York, NY, USA, 2022. Association for Computing Machinery.
- [CZB⁺22] Pierre-Emmanuel Clet, Martin Zuber, Aymen Boudguiga, Renaud Sirdey, and Cédric Gouy-Pailler. Putting up the swiss army knife of homomorphic calculations by means of TFHE functional bootstrapping. *Cryptology ePrint Archive*, Paper 2022/149, 2022. <https://eprint.iacr.org/2022/149>.
- [DM15] Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping Homomorphic Encryption in Less Than a Second. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015*, pages 617–640, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.

- [FV12] Junfeng Fan and Frederik Vercauteren. Somewhat Practical Fully Homomorphic Encryption. Cryptology ePrint Archive, Paper 2012/144, 2012. <https://eprint.iacr.org/2012/144>.
- [GBA21] Antonio Guimarães, Edson Borin, and Diego F. Aranha. Revisiting the functional bootstrap in TFHE. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021(2):229–253, February 2021.
- [GBA22] Antonio Guimarães, Edson Borin, and Diego F. Aranha. MOSFHET: Optimized Software for FHE over the Torus. Cryptology ePrint Archive, Paper 2022/515, 2022. <https://eprint.iacr.org/2022/515>.
- [Gen09] Craig Gentry. Fully Homomorphic Encryption Using Ideal Lattices. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 169–178, New York, NY, USA, 2009. Association for Computing Machinery.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013*, pages 75–92, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [Klu22] Kamil Klucznik. NTRU-v-Um: Secure Fully Homomorphic Encryption from NTRU with Small Modulus. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, CCS '22, page 1783–1797, New York, NY, USA, 2022. Association for Computing Machinery.
- [KPZ21] Andrey Kim, Yuriy Polyakov, and Vincent Zucca. Revisiting Homomorphic Encryption Schemes for Finite Fields. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021*, pages 608–639, Cham, 2021. Springer International Publishing.
- [KS22] Kamil Klucznik and Leonard Schild. FDFB: Full Domain Functional Bootstrapping Towards Practical Fully Homomorphic Encryption. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(1):501–537, Nov. 2022.
- [KSK⁺18] Andrey Kim, Yongsoo Song, Miran Kim, Keewoo Lee, and Jung Hee Cheon. Logistic regression model training based on the approximate homomorphic encryption. *BMC Medical Genomics*, 11(4):83, October 2018.
- [LATV12] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-Fly Multiparty Computation on the Cloud via Multikey Fully Homomorphic Encryption. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '12, page 1219–1234, New York, NY, USA, 2012. Association for Computing Machinery.
- [LHH⁺21] Wen-jie Lu, Zhicong Huang, Cheng Hong, Yiping Ma, and Hunter Qu. PEGASUS: Bridging Polynomial and Non-polynomial Evaluations in Homomorphic Encryption. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1057–1073, 2021.
- [LKL⁺22] Joon-Woo Lee, Hyungchul Kang, Yongwoo Lee, Woosuk Choi, Jieun Eom, Maxim Deryabin, Eunsang Lee, Junghyun Lee, Donghoon Yoo, Young-Sik Kim, and Jong-Seon No. Privacy-Preserving Machine Learning With Fully Homomorphic Encryption for Deep Neural Network. *IEEE Access*, 10:30039–30054, 2022.

- [LMP22] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. Large-Precision Homomorphic Sign Evaluation Using FHEW/TFHE Bootstrapping. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022*, pages 130–160, Cham, 2022. Springer Nature Switzerland.
- [LW23a] Feng-Hao Liu and Han Wang. Batch Bootstrapping I: A New Framework for SIMD Bootstrapping in Polynomial Modulus. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 321–352, Cham, 2023. Springer Nature Switzerland.
- [LW23b] Feng-Hao Liu and Han Wang. Batch Bootstrapping II: Bootstrapping in Polynomial Modulus only Requires $\tilde{O}(1)$ FHE Multiplications in Amortization. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 353–384, Cham, 2023. Springer Nature Switzerland.
- [LY23] Kang Hoon Lee and Ji Won Yoon. Discretization Error Reduction for High Precision Torus Fully Homomorphic Encryption. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *Public-Key Cryptography – PKC 2023*, pages 33–62, Cham, 2023. Springer Nature Switzerland.
- [MP21] Daniele Micciancio and Yuriy Polyakov. Bootstrapping in FHEW-like Cryptosystems. In *Proceedings of the 9th on Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, WAHC '21, page 17–28, New York, NY, USA, 2021. Association for Computing Machinery.
- [PKS⁺19] Sungjoon Park, Minsu Kim, Seokjun Seo, Seungwan Hong, Kyoohyung Han, Keewoo Lee, Jung Hee Cheon, and Sun Kim. A secure SNP panel scheme using homomorphically encrypted K-mers without SNP calling on the user side. *BMC Genomics*, 20(2):188, April 2019.
- [YXS⁺21] Zhaomin Yang, Xiang Xie, Huajie Shen, Shiyong Chen, and Jun Zhou. TOTA: Fully Homomorphic Encryption with Smaller Parameters and Stronger Security. Cryptology ePrint Archive, Paper 2021/1347, 2021. <https://eprint.iacr.org/2021/1347>.