

BatchBoot: Fast Batched Bootstrapping for TFHE scheme and Practical Applications

Zhihao Li¹, Hongyu Wang², Yuan Zhao¹, Lichun Li¹, Zhiwei Wang^{3*}

Jiaxing He¹, Changzheng Wei¹, Ying Yan¹, and Lifeng Guo²

¹Ant Digital Technologies, Ant Group ²Shanxi University

³State Key Laboratory of Cyberspace Security Defense, Institute of Information Engineering, CAS

{lzh458070, caesar.zy, lichun.llc, jiaxing.hjx, changzheng.wcz, fuying.yy}@antgroup.com

{wanghongyu, lfguo}@sxu.edu.cn, wangzhiwei@iie.ac.cn

Abstract

Torus-based Fully Homomorphic Encryption (TFHE) is distinguished by its unique bootstrapping mechanism, which enables arbitrary computation while refreshing the noise budget. However, this mechanism exhibits limited scalability since it can handle only a *single* encrypted message at a time. To address this, recent studies have proposed *batched* bootstrapping schemes that allow TFHE to process ciphertexts in parallel, thereby achieving promising amortization benefits. Despite these advances, this emerging direction remains underexplored, leaving ample room for further investigation.

In this paper, we present BATCHBOOT, an efficient batched bootstrapping framework for TFHE that enables amortized processing of encrypted messages. Specifically, our work makes three key contributions. First, we redesign the core submodule, i.e., homomorphic polynomial multiplication, to substantially reduce the reliance on expensive FFT operations. Second, we propose a sparsity-aware message packing strategy that flexibly supports varying packing scales. Third, we extend functional bootstrapping to circuit bootstrapping, thereby greatly enhancing the expressiveness of supported functions. Together, these contributions enable BATCHBOOT to deliver a 2.4× speedup over the state-of-the-art batched scheme (Guimarães et al., CCS’25) and a 43.8× improvement over the non-batched TFHE-rs implementation.

At the application level, we highlight the versatility of BatchBoot through two practical use cases. First, we present the first TFHE-based PSI protocol under the unbalanced setting, which achieves a 294× reduction in communication cost and a 4.1× speedup compared to the best BFV-based solution (PEPSI, USENIX Security’24). Second, we design an 8-bit FHE instruction set based on the BATCHBOOT that delivers up to a 5.4× speedup over the existing results (Wang et al., CCS’25).

*Zhiwei Wang is the corresponding author.

1 Introduction

Fully Homomorphic Encryption (FHE) is a cornerstone of privacy-preserving computation, enabling arbitrary operations directly on encrypted data without compromising confidentiality [31]. Over the past decade, two primary classes of FHE schemes have emerged, reflecting different trade-offs between functionality and efficiency. Word-wise schemes, such as BGV [10], BFV [30], and CKKS [19], pack multiple plaintexts into a single ciphertext to enable SIMD-style execution. However, their reliance on high-degree polynomial approximations for non-linear functions (e.g., comparisons) incurs prohibitive computational costs and accuracy loss. In contrast, bit-wise schemes such as FHEW [28, 44] and TFHE [20, 22] are optimized for Boolean and circuit-based computation. By embedding logic gates into the bootstrapping process, these schemes natively support expressive, unbounded non-linear operations with far greater efficiency, positioning TFHE as the preferred choice for general-purpose workloads.

The core strength of TFHE lies in functional bootstrapping (FBS) [9], also known as programmable bootstrapping (PBS) [23], which evaluates arbitrary functions via look-up tables (LUTs) while refreshing noise to sustain long computation chains. For example, TFHE-rs [12] uses FBS to implement complex logic and relational operations. After that, TFHE incorporates multi-valued bootstrapping [13, 23] for computing multiple functions in one call and circuit bootstrapping (CBS) [21, 40, 50] for multi-input functions and higher-precision operations. These advances broaden TFHE’s expressiveness and practical applicability.

Despite these innovations, TFHE remains fundamentally limited by the lack of batching and message-packing mechanisms, creating a major bottleneck for high-throughput and vectorized workloads. Some work [27, 34, 45] explores *batch-friendly* algorithms for TFHE schemes that amortize bootstrapping costs across multiple messages, but practical adop-

tion is constrained by large parameter sizes and rapid noise growth. A recent breakthrough, AmortBoot [33], introduced an amortization framework and implementation that improved both computational efficiency and key size, setting a new benchmark for memory efficiency, scalability, and real-world applicability in large-scale homomorphic computation.

1.1 Problems and Challenges

Our investigation reveals that existing batched (or amortized) designs continue to exhibit several critical limitations, which remain pressing challenges for the field. The following critique assumes familiarity with TFHE building blocks and batched bootstrapping. Unfamiliar readers may reference these foundations in Sec. 2 and Sec. 3 before proceeding.

❶ Excessive FFT invocations leads to poor performance.

To achieve batched bootstrapping, the prevailing approach leverages the sparsity of the RLWE secret key, attaining a complexity of $O(hn \log n)$, i.e., $O(h \log n)$ per message, where n and h ($h < n$) denote the dimension and Hamming weight of the secret key, respectively. Although this asymptotic improvement represents a significant algorithmic advance, it still depends on FFT-intensive subroutines in the core construction. These subroutines not only limit the realized performance gains but also constitute a fundamental bottleneck to scalability, underscoring the need to rethink algorithmic design to advance the performance frontier.

❷ Sparse message packing remains unexplored.

Existing techniques primarily assume a *full-packing* regime, where all n ciphertext slots are simultaneously populated. In practice, however, diverse applications demand varying degrees of sparsity in packing, which cannot be efficiently accommodated under this assumption. This rigidity stems from the inherent structure of RLWE decryption, where the computational complexity scales with the ciphertext dimension n rather than the actual number of active slots. The lack of sparsity-aware optimizations thus reveals a fundamental mismatch between algorithmic design and workload characteristics.

❸ Existing schemes compromise on functional generality.

Existing schemes are limited to functional bootstrapping, which supports only single-input LUTs with 8-bit precision. In contrast, for more variants, such as multi-input LUT evaluation, the circuit bootstrapping needs to be introduced to support these functions, rather than merely adjusting the encoding. Extending batched bootstrapping to incorporate these features is not only technically challenging, requiring careful parameter-level optimizations, but also critical for realizing the full practical potential of TFHE-based systems.

1.2 Our contributions

To address these challenges, we propose BATCHBOOT, a novel batched bootstrapping framework that significantly en-

hances both efficiency and functionality. Specifically, we reduce the overhead of expensive FFTs by restructuring the batched bootstrapping subroutines. Additionally, we introduce an efficient sparse packing method derived from the MLWE decryption paradigm. Finally, we present the first batch circuit bootstrapping extension that broadens bootstrapping functionality. The effectiveness of BATCHBOOT is validated through empirical deployment on PSI and instruction set architectures. Our contributions are as follows:

Efficient batched bootstrapping through significant FFT reduction.

To address issue ❶, BATCHBOOT employs two key strategies that substantially reduce the number of costly FFT operations. First, we develop a multi-bit CMux variant for batched bootstrapping, effectively reducing both the number of iterations and the FFTs, compared to the single-bit CMux method used in AmortBoot [33]. Second, we observe that automorphism operations can be post-processed after performing the FFT and redesign the polynomial multiplication pipeline for homomorphic evaluation. This enables FFT computation sharing between automorphism and CMux, thereby removing the additional FFT burden traditionally associated with automorphism. As a result, our approach reduces FFT operations, delivering a 2.4× speedup over the optimal amortization scheme (AmortBoot) and a 43.8× improvement over the non-batched TFHE-rs implementation.

Sparsity-aware message packing leveraging the MLWE decryption paradigm.

To overcome issue ❷, BATCHBOOT incorporates an MLWE-based decryption paradigm to enable sparsity-aware optimization. In the sparse packing setting, where an RLWE ciphertext encodes only $\underline{n} < n$ messages, we apply RLWE-to-MLWE extraction to derive a lower-dimensional MLWE ciphertext. This transformation reduces the polynomial degree from n to \underline{n} while preserving the secret key’s sparsity, thereby maintaining the efficiency benefits of sparse structures. Consequently, the computational complexity drops from $O(hn \log n)$ to $O(h\underline{n} \log \underline{n})$, yielding performance improvements that scale linearly with the packing sparsity compared to the full-packing setting.

Extended functional expressiveness through batched circuit bootstrapping.

We address issue ❸ by extending BATCHBOOT from functional bootstrapping to circuit bootstrapping to deal with multi-input LUT evaluation, resulting in BATCHCBOOT. Our design builds upon the circuit bootstrapping paradigm [40]. Specifically, the architecture proposed in [40] splits the circuit bootstrapping process into two subroutines: the ciphertext conversion from LWE to RGSW, and LUT evaluation by performing external operations within a tree-like structure. We note that in the first phase, the BATCHBOOT can be utilized to efficiently perform batch ciphertext conversions from RLWE to RGSWs, thereby enhancing performance. To the best of our knowledge, this is the first work that supports batch circuit bootstrapping. Benchmark results on 12-bit input/output LUTs demonstrate speedups ranging

from 4.9× to 134.3× over the non-batched method in [50].

Application I: Communication-friendly private set intersection (PSI) protocol. We present the first TFHE-based PSI protocol designed to minimize communication overhead. Leveraging BATCHBOOT, our design supports ciphertext transmission with an exceptionally small modulus (as low as 12 bits), whereas BFV-based schemes typically require larger FHE parameters, such as 100–200-bit moduli. This results in at least a two-order-of-magnitude reduction in communication cost relative to existing FHE-based PSI solutions. Moreover, in certain unbalanced settings, our protocol surpasses the state-of-the-art solution [43]; for example, it achieves a 4.1× speedup on datasets where the client holds 2,048 elements and the server has 2^{28} elements.

Application II: High-performance 8-bit instruction set implementation. We further develop an optimized 8-bit FHE instruction set encompassing fundamental operations such as AND, ADD, MUL, and MAX. Leveraging the BATCHBOOT algorithm, we enable the evaluation of multiple distinct instruction sets within a single bootstrapping operation, incorporating customized designs tailored to the unique characteristics of each set. Compared to prior solutions limited to a single instruction set [49, 50], our approach delivers 3.3×–6.5× amortized performance improvements.

2 Preliminary

We first present some definitions of symbols in Appendix B.1 and then introduce some necessary background for FHE schemes.

2.1 Ciphertexts Types

GLWE ciphertext [4]. Given the secret key $\mathbf{S} = (s_0, \dots, s_{k-1}) \in R_N^k$, the message $m \in R_{N,p}$ can be encrypted into GLWE ciphertext as $(\mathbf{a}_0, \dots, \mathbf{a}_{k-1}, \mathbf{b}) \in R_{N,Q}^k$, where $\mathbf{b} = \sum_{i=0}^{k-1} \mathbf{a}_i \cdot s_i + \Delta m + \epsilon \pmod{Q}$, where $\Delta = \lfloor \frac{Q}{p} \rfloor$, $(\mathbf{a}_0, \dots, \mathbf{a}_{k-1})$ are uniformly sampled from $R_{N,Q}$, and the error ϵ is sampled from the Gaussian Distribution χ . Here p, Q are the plaintext and ciphertext modulus, k and N are the dimension and degree of the GLWE, respectively, and we define $\tilde{m} = \Delta m + \epsilon$.

To avoid confusion, a GLWE ciphertext where $N = 1$ is termed an *LWE* ciphertext [47], denoted $(\mathbf{a}, b) \in \text{LWE}_{\mathbf{S}, p/Q}^k(m)$. When $k = 1$ and $N > 1$, it is called the ring *LWE* (*RLWE*) ciphertext [42], that is $(\mathbf{a}, \mathbf{b}) \in \text{RLWE}_{\mathbf{S}, p/Q}^N(m)$. A GLWE with $k > 1$ and $N > 1$ is referred to as *MLWE* $\text{MLWE}_{\mathbf{S}, p/Q}^{k,N}(m)$. These ciphertexts support homomorphic addition and scalar multiplication. In addition, we may omit some parameters for notation simplicity.

Gadget RLWE ciphertext [44]. Given the gadget vector $\mathbf{g} = (B_0, \dots, B_{d-1})$, the gadget RLWE ciphertext is defined as

$$\text{RLWE}'(m) = (\text{RLWE}(B_0 \cdot m), \dots, \text{RLWE}(B_{d-1} \cdot m)).$$

RGSW ciphertext [28]. RGSW encryption of the message m is defined as $\text{RGSW}(m) = (\text{RLWE}'(-s \cdot m), \text{RLWE}'(m))$, where $\text{half}(\text{RGSW}(m))$ denotes ciphertext $\text{RLWE}'(m)$.

2.2 Building Blocks

Gadget Decomposition (Decom). For a modulus Q , and a base B , the gadget vector is defined as $\mathbf{g} = (B_0, \dots, B_{d-1}) = \left(\left\lfloor \frac{Q}{B^d} \right\rfloor, \left\lfloor \frac{Q}{B^d} \right\rfloor \cdot B, \dots, \left\lfloor \frac{Q}{B^d} \right\rfloor \cdot B^{d-1} \right)$. Then, given a polynomial $\mathbf{a} \in R_Q$, the approximate decomposition can be used to output $(\mathbf{a}_0, \dots, \mathbf{a}_{d-1})$ such that $\mathbf{a} = \sum_{i=0}^{d-1} \mathbf{a}_i \cdot \left\lfloor \frac{Q}{B^{d-i}} \right\rfloor$. Note that the approximate gadget decomposition introduces an approximate error ϵ , i.e., $\langle \mathbf{g}^{-1}(\mathbf{a}), \mathbf{g} \rangle \equiv \mathbf{a} + \epsilon$, where $\|\epsilon\| \leq \frac{1}{2} \left\lfloor \frac{Q}{B^d} \right\rfloor$.

Gadget Product (GP). Given a polynomial $t \in R_Q$, let $(t_0, \dots, t_{d-1}) = \text{Decom}(t)$ is the result of the gadget decomposition with respect to the gadget vector $\mathbf{g} = (B_0, \dots, B_{d-1})$. The Gadget product \odot is defined by:

$$t \odot \text{RLWE}'(m) := \sum_{i=0}^{d-1} t_i \cdot \text{RLWE}(B_i \cdot m) = \text{RLWE}(t \cdot m).$$

External Product (EP). Given ciphertexts $\text{RLWE}(m_1) = (\mathbf{a}, \mathbf{b})$ and $\text{RGSW}(m_2)$, the external product is defined by:

$$\begin{aligned} & \text{RLWE}(m_1) \boxtimes \text{RGSW}(m_2) \\ &= \mathbf{a} \odot \text{RLWE}'(-s \cdot m_2) + \mathbf{b} \odot \text{RLWE}'(m_2) \\ &= \text{RLWE}(m_1 \cdot m_2 + \epsilon_1 \cdot m_2). \end{aligned}$$

The external product outputs encryption of the product of m_1 and m_2 . Note that the gadget product and external products involve polynomial operations, which can be accelerated with the Fast Fourier Transform (FFT).

CMux Gate (CMux). The CMux gate takes as inputs two RLWE ciphertexts c_0, c_1 and a RGSW ciphertext $C \in \text{RGSW}(m)$, and CMux operation is defined as

$$\text{CMux}(c_0, c_1, C) = c_0 + (c_1 - c_0) \boxtimes C = \begin{cases} c_0, & \text{if } m = 0; \\ c_1, & \text{else } m = 1, \end{cases}$$

where m satisfies $m \in \{0, 1\}$. We present the noise growth for $V_{\text{GP}}, V_{\text{EP}}$, and V_{CMux} in Appendix B.2.

Sample Extraction (SE). Given the ciphertext $c = (\mathbf{a}, \mathbf{b}) \in \text{RLWE}_{\mathbf{S}, Q}^N(m)$, the sample extraction operation can extract N LWE ciphertexts associated with the coefficient terms. For example, for the first term, we can define $\text{SampleExtract}(ct) = (a_0, -a_{N-1}, \dots, -a_1, b_0) \in \text{LWE}_{\text{Coef}(\mathbf{S}), Q}^N(\text{Coef}(m)[0])$.

RLWE-to-MLWE Extraction (RtM.Ext). Given an RLWE ciphertext $(\mathbf{a}, \mathbf{b}) \in \text{RLWE}_{\mathbf{S}, p/Q}^N(m)$, where $m = \sum_{i=0}^{n-1} m_i \cdot X^{i \cdot \frac{N}{n}}$. Let $m' = \sum_{i=0}^{n-1} m_i \cdot X^i$, we can extract the MLWE ciphertext as $(\mathbf{a}'_0, \dots, \mathbf{a}'_{k-1}, \mathbf{b}') \in \text{MLWE}_{\mathbf{S}, p/Q}^{k,n}(m')$, where $k = \frac{N}{n}$, $(\mathbf{a}'_0, \dots, \mathbf{a}'_{k-1}) = \text{RingSwitch}_{k,N}(\mathbf{a})$, $\mathbf{S} = (s'_0, \dots, s'_{k-1}) =$

RingSwitch(\mathfrak{s}), and $\mathfrak{b}' = \text{RingSwitch}(\mathfrak{b})_{k,N}[0]$. We show the details of RingSwitch in Appendix B.3.

Modulus Switching (MS). Given an RLWE ciphertext $c \in \text{RLWE}_{\mathfrak{s},p/q}^N(\mathfrak{m})$, the modulus switching $\text{ModSwitch}_{Q \rightarrow q}(c)$ can change the ciphertext modulus from Q to q as $\text{RLWE}_{\mathfrak{s},p/q}^N(\mathfrak{m})$ as shown in Appendix B.3.

Key Switching (KS). It can convert a ciphertext encrypted under the secret key \mathfrak{s}_1 into a ciphertext encrypted under the new secret key \mathfrak{s}_2 with the key switching key. In addition, the RLWE key switching variant scheme switching SchemSwitch [27, 40, 51] can embed the secret key \mathfrak{s} into the message. We show the RLWE key switching (RLWE.KeySwitch) and scheme switching steps [22] in Appendix B.3.

Homomorphic Automorphism (HomoAuto). There are N automorphisms as $\tau_j : \mathfrak{m}(X) \rightarrow \mathfrak{m}(X^j)$ for $j \in \mathbb{Z}_{2N}^*$ for $\mathbb{Z}_Q[X]/(X^N + 1)$. Given ciphertext $c = (\mathfrak{a}, \mathfrak{b}) \in \text{RLWE}_{\mathfrak{s},Q}^N(\mathfrak{m})$, and the RLWE key switching key $\text{Atk}_{\tau_j} \in \text{RLWE}'_{\mathfrak{s}}(\mathfrak{s}(X^j))$, the Homomorphic automorphism $\text{HomoAuto}_{\tau_j}(\text{ct}, \text{Atk}_{\tau_j})$ is divided into two steps: first performs $\text{Auto} : \tau_j(c) = (\mathfrak{a}(X^j), \mathfrak{b}(X^j)) \in \text{RLWE}_{\tau_j(\mathfrak{s}),Q}^N(\tau_j(\mathfrak{m}))$; and applies the RLWE key switching from the secret key $\mathfrak{s}(X^j)$ to $\mathfrak{s}(X)$.

Homomorphic Trace (HomoTrace). For the polynomial ring R_N and its residue ring $R_{N,Q}$, the tower of finite fields is expressed as $E = K_N \geq K_{N/2} \geq \dots \geq K_1 = K$, where K_n denotes the $(2n)$ -th cyclotomic field for a power-of-two n . Then, the trace Tr_{K_N/K_n} can be defined as $\text{Tr}_{K_N/K_n}(\mathfrak{m}(X)) = \sum_{i=0}^{n-1} \mathfrak{m}_{i,N/n} \cdot X^{iN/n}$, which can be evaluated through $\log_2 N/n$ automorphism operations. The homomorphic trace algorithm Tr_{K_N/K_n} is described in Appendix B.4.

Repacking (Repack). Given n RLWE ciphertexts $\text{RLWE}_{\mathfrak{s},Q}^N(\mathfrak{m}_i) = (\mathfrak{a}_i, \mathfrak{b}_i)$ for $i \in [0, n-1]$, the repacking algorithm can pack the first items of ciphertexts $(\mathfrak{m}_{0,0}, \dots, \mathfrak{m}_{N-1,0})$ into a new RLWE ciphertext $\tilde{\mathfrak{m}} = \sum_{i=0}^{n-1} \mathfrak{m}_{i,N/n} \cdot X^{iN/n}$, where $n|N$ ($n < N$) are referred to as sparse packing (SRepack). This paper uses the methods [14, 38] that need to perform the HomaAuto and Homotrace operations. Detailed algorithms and noise growth are presented in Appendix B.4.

2.3 Functional Bootstrapping

Functional bootstrapping (FBS) [9] can evaluate the arbitrary function in bootstrapping. Specifically, given a ciphertext $c \in \text{LWE}_{\mathfrak{s},p/q}^n(m)$, FBS can output a new ciphertext that encrypts $f(m)$ as shown in Alg. 1. For the arbitrary function $f : \mathbb{Z}_p \rightarrow \mathbb{Z}_{p_{\text{out}}}$, we can embed it into a polynomial testP as shown in line 1. Note that we set $p_{\text{out}} = p$ by default, unless otherwise specified. Then, the modulus switching is used to convert the ciphertext from q to $2N$. After performing the lines 3-6 through n CMux operations (referred to as blind rotation in [20]), we obtain the RLWE ciphertext $\text{RLWE}(\text{testP} \cdot X^{\tilde{m}})$, where $\tilde{m} = \Delta \cdot m + e$ and $\Delta = \frac{2N}{p}$. Then, the desired result is preserved in the constant term of the polynomial. Finally,

Algorithm 1 Functional Bootstrapping (FBS).

Input:

LWE ciphertext $c = (\mathfrak{a}, \mathfrak{b}) \in \text{LWE}_{\mathfrak{s},p/q}^n(m)$.

Bootstrapping key $\text{Bsk} : \{\text{RGSW}_{\mathfrak{s},Q}(s_i)\}$ for $i \in [0, n-1]$.

Output:

LWE ciphertext of $f(m)$.

- 1: Set $\text{testP} = -\sum_{i=0}^{N-1} \left\lfloor \frac{Q}{p_{\text{out}}} \right\rfloor \cdot f\left(\left\lfloor \frac{ip}{2N} \right\rfloor\right) \cdot X^{N-i-1}$
 - 2: $(\mathfrak{a}, \mathfrak{b}) = \text{ModSwitch}_{q \rightarrow 2N}(c)$
 - 3: $\text{acc} = \text{testP}(X) \cdot X^{\mathfrak{b}}$,
 - 4: for $i = 0$ to $n-1$ do
 - 5: $\text{acc} = \text{CMux}(\text{acc}, X^{-a_i} \cdot \text{acc}, \text{Bsk}_i)$,
 - 6: end for
 - 7: **return** $\text{SampleExtraction}(\text{acc})$
-

we can use the sample extraction to obtain the LWE ciphertext of $f(m)$. Essentially, the bootstrapping performs LWE decryption on the exponent of X .

3 Revisiting Previous Batched Bootstrapping

Guimarões et al. [33] present AmortBoot that aims to bootstrap an RLWE ciphertext instead of the LWE ciphertext, where the RLWE encrypts a polynomial $\mathfrak{m}(X) = \sum_{i=0}^{n-1} m_i \cdot X^i$. In detail, given the ciphertext $\text{RLWE}_{\mathfrak{s},p/q}^n(\mathfrak{m})$, this procedure can generate n ciphertexts, that is $\text{RLWE}(\text{testP}_i \cdot X^{\tilde{m}_i})$ for $i \in [0, n-1]$. This multiplication $\text{testP}_i \cdot X^{\tilde{m}_i}$ subsequently results in the encryption of $f(m_0), \dots, f(m_{n-1})$.

To obtain n ciphertext $\text{RLWE}(\text{testP}_i \cdot X^{\tilde{m}_i})$ simultaneously, one needs to compute $\mathfrak{m}(X) = \mathfrak{b}(X) - \mathfrak{a}(X) \cdot \mathfrak{s}(X)$ on the exponent of X in a single bootstrapping. AmortBoot exploits the sparsity of the RLWE secret key, which accelerates the polynomial multiplication process. When the secret key is chosen from a sparse binary key distribution with a Hamming weight of h ($h < n$)¹, RLWE decryption formula can be expressed as

$$\mathfrak{b} - \mathfrak{a} \cdot \mathfrak{s} = \mathfrak{b} - \mathfrak{a} \sum_{j \in \text{id}x(\mathfrak{s})} X^j = \mathfrak{b} - \sum_{j \in \text{id}x(\mathfrak{s})} \left(\sum_{i=0}^{n-1} \mathfrak{a}(X)^i \right) \tilde{\odot} X^j, \quad (1)$$

where $j \in \text{id}x(\mathfrak{s})$ represents the set containing the degrees of nonzero coefficients of \mathfrak{s} , n is the degree of RLWE, and $\tilde{\odot}$ represents a monomial-by-polynomial multiplication. For instance, if $\mathfrak{s}(X) = 1 + 1 \cdot X^5 + 1 \cdot X^7$, we have $j = (0, 5, 7)$. For the cleartext, the Eq. 1 can be computed in $O(hn)$ operations. However, since the index j is related to the secret key, it is encrypted into the bootstrapping key. After that, AmortBoot [33] rewrites the Eq. 1 as

$$\left((\mathfrak{b} \tilde{\odot} X^{-j_0} - \mathfrak{a}) \tilde{\odot} X^{j_0 - j_1} - \mathfrak{a} \right) \tilde{\odot} X^{j_1 - j_2} \dots - \mathfrak{a} \tilde{\odot} X^{j_h}, \quad (2)$$

which can be efficiently computed in the exponent of X by using external products. Note that since \mathfrak{a} is a public poly-

¹As many notable FHE schemes [1, 19, 33, 37] have done.

nomial, subtracting it in the exponent can be achieved by multiplying all X^{-a_i} . Moreover, AmortBoot [33] constructs a sub-procedure MPmul to evaluate $\tilde{\odot}$. This process can be generalized to compute the multiplication $u \cdot X^v \pmod{X^n + 1}$ in the exponent, and v can be instantiated for all possible values of $j_i - j_{i+1}$. After that, Eq. 2 can be reduced to

$$\text{MPmul} \left(\dots \text{MPmul} \left(\text{MPmul} \left(b, X^{-\tilde{j}_0} \right) - a, X^{\tilde{j}_0 - \tilde{j}_1} \right) \dots - a, X^{\tilde{j}_h} \right). \quad (3)$$

Note that, during the homomorphic computation process, all coefficients of u are encrypted into RLWE(X^{u_i}) as inputs, while $v \in [-n, n]$ is encrypted and encapsulated within the bootstrapping key, denoted as Bsk. Essentially, multiplying by X^v is to rotate the coefficients of u by v positions. For this positive v rotations², the multiplication $u \cdot X^v \pmod{X^n + 1}$ has the following map

$$(u_0, u_1, \dots, u_{n-1}) \mapsto \underbrace{(-u_{n-v}, -u_{n-v+1}, \dots, -u_{n-1}, u_0, u_1, \dots, u_{n-v-1})}_{\text{negate } v \text{ positions}}.$$

Because the RLWE is instantiated in the cyclotomic polynomial $\mathbb{Z}[X]/\langle X^n + 1 \rangle$, we also need to multiply v of them by -1 in some positions.

In this process, one can process each coefficient in sequence in this map. More precisely, MPmul further decomposes v as $v = \sum_{i=0}^{\ell-1} 2^i \cdot v_i$, and sequentially multiplies $X^{2^0 v_0}, \dots, X^{2^{\ell-1} v_{\ell-1}}$ for each coefficient (position). It is clear that multiplying by $X^{2^i v_i}$ is equal to evaluating the following equation:

$$f(u_0, \dots, u_{n-1}) \mapsto \begin{cases} (u_0, u_1, \dots, u_{n-1}) & \text{if } v_i = 0, \\ (-u_{n-2^i}, \dots, -u_{n-1}, u_0, \dots, u_{n-2^i-1}) & \text{if } v_i = 1. \end{cases}$$

Then, according to the input $v_i \in \{0, 1\}$, one can use the CMux to pick out the corresponding value. Hence, for the output coefficient range $2^i \leq j \leq n-1$, given the $\text{Bsk}_i = \text{RGSW}(v_i)$, the homomorphic computation for this step is shown as follows

$$\hat{c}_j = (c_{j-2^i} - c_j) \square \text{Bsk}_i + c_j = \begin{cases} c_{j-2^i}, & \text{if } v_i = 0; \\ c_j, & \text{else } v_i = 1. \end{cases}$$

For the first 2^i coefficients, we first apply the $\text{HomoAuto}_{\tau_{-1}}$: $X \mapsto X^{-1}$ to get the encryption of $X^{-u_{n-2^i+j}}$, and perform the CMux as

$$\hat{c}_j = (\text{HomoAuto}_{-1}(c_{n-2^i+j}) - c_j) \square \text{Bsk}_i + c_j. \quad (4)$$

This process involves $n\ell$ CMux and $\ell(2^\ell - 1)$ automorphism operations as shown in [33], where $2^\ell < n$ in practical parameters.

²Followed by a similar method for handling negative values.

4 Design Overview of BatchBoot

In this section, we propose some novel batched bootstrapping algorithms that support functional evaluations for multiple messages. Afterwards, we conduct a comprehensive evaluation of the correctness and noise growth of these algorithms.

4.1 Efficient MPmul Procedure

This section proposes a more efficient MPmul (EMPMul) procedure. At a fine-grained level, external product \square and HomoAuto involve gadget decomposition and FFT (used in accelerating polynomial multiplication). For example, the evaluation of $\hat{c}_j = (c_{j-2^i} - c_j) \square \text{Bsk}_i + c_j$ can be broken down into five steps. 1) Apply the gadget decomposition on $c_{j-2^i} - c_j$ with the coefficient representation; 2) Perform FFTs on the result of the decomposition; 3) Evaluate Hadamard multiplication with Bsk_i , where Bsk_i is preprocessed into FFT form; 4) Perform the inverse FFT (iFFT) to recover the coefficient representation; 5) Compute homomorphic addition with c_j . In addition, in Eq. 4, i.e., $\hat{c}_j = (\text{HomoAuto}_{-1}(c_{n-2^i+j}) - c_j) \square C_i + c_j$, HomoAuto_{-1} also needs to perform additional FFTs (see Sec. 2.2), and then this implementation can be summarized as

$$\text{HomoAuto}_{-1}(c) \square \mapsto \underbrace{\text{Auto} + \text{Decom} + \text{FFT} + \text{Had} + \text{iFFT}}_{\text{HomoAuto}} + \underbrace{\text{Decom} + \text{FFT} + \text{Had} + \text{iFFT}}_{\text{External product}}, \quad (5)$$

Since FFTs exhibit a computational complexity of $O(n \log n)$, as opposed to the $O(n)$ complexity of Hadamard multiplication, FFT operations represent the main computational bottleneck in Eq. 5. Overall, the method of Sec. 3 requires approximately $2(2d+2) \cdot n\ell$ FFTs, where d is the length of the gadget decomposition.

Our method. We design the EMPmul to evaluate two-input bits for v instead of a single-input bit, which can reduce the number of time-consuming FFTs from $2(2d+2) \cdot n\ell$ to $(d+1) \cdot n\ell$. For simplicity, we focus on the case $v > 0$ ($v < 0$ is shown in Appendix C.1), and then the sequential multiplication $X^{2^i v_i}$ and $X^{2^{i+1} v_{i+1}}$ has the map as:

$$f(u_0, \dots, u_{n-1}) \mapsto \begin{cases} (u_0, \dots, u_{n-1}) & \text{if } v_{i+1} = (0, 0) \\ (-u_{n-2^i}, \dots, -u_{n-1}, u_0, \dots, u_{n-2^i-1}) & \text{if } v_{i+1} = (1, 0) \\ (-u_{n-2^{i+1}}, \dots, -u_{n-1}, u_0, \dots, u_{n-2^{i+1}-1}) & \text{if } v_{i+1} = (0, 1) \\ (-u_{n-2^i-2^{i+1}}, \dots, -u_{n-1}, u_0, \dots, u_{n-2^i-2^{i+1}-1}) & \text{if } v_{i+1} = (1, 1) \end{cases} \quad (6)$$

Let us first consider these non-negative positions, i.e., $2^i + 2^{i+1} \leq j \leq n-1$. The multi-bit CMux operation can be employed directly to evaluate the function as

Table 1: Setup of bootstrapping keys for multi-bit CMux.

Bsk	Bsk _{i,0}	Bsk _{i,1}	Bsk _{i,2}	Bsk _{i,3}	Value
Bsk ⁺	RGSW(0)	RGSW(0)	RGSW(0)	RGSW(1)	$v_i = 0; v_{i+1} = 0$
	RGSW(1)	RGSW(0)	RGSW(0)	RGSW(0)	$v_i = 1; v_{i+1} = 0$
	RGSW(0)	RGSW(1)	RGSW(0)	RGSW(0)	$v_i = 0; v_{i+1} = 1$
	RGSW(0)	RGSW(0)	RGSW(1)	RGSW(0)	$v_i = 1; v_{i+1} = 1$
Bsk ⁻	RGSW ^{τ₋₁} (0)	RGSW ^{τ₋₁} (0)	RGSW ^{τ₋₁} (0)	RGSW(1)	$v_i = 0; v_{i+1} = 0$
	RGSW ^{τ₋₁} (1)	RGSW ^{τ₋₁} (0)	RGSW ^{τ₋₁} (0)	RGSW(0)	$v_i = 1; v_{i+1} = 0$
	RGSW ^{τ₋₁} (0)	RGSW ^{τ₋₁} (1)	RGSW ^{τ₋₁} (0)	RGSW(0)	$v_i = 0; v_{i+1} = 1$
	RGSW ^{τ₋₁} (0)	RGSW ^{τ₋₁} (0)	RGSW ^{τ₋₁} (1)	RGSW(0)	$v_i = 1; v_{i+1} = 1$

$$\hat{c}_j = c_{j-2^i} \boxtimes \text{Bsk}_{i,0}^+ + c_{j-2^{i+1}} \boxtimes \text{Bsk}_{i,1}^+ + c_{j-2^{i-2^{i+1}}} \boxtimes \text{Bsk}_{i,2}^+ + c_j \boxtimes \text{Bsk}_{i,3}^+ \quad (7)$$

$$= \begin{cases} c_j, & \text{if } v_i = 0; v_{i+1} = 0; \\ c_{j-2^i}, & \text{if } v_i = 1; v_{i+1} = 0; \\ c_{j-2^{i+1}}, & \text{if } v_i = 1; v_{i+1} = 0; \\ c_{j-2^i-2^{i+1}}, & \text{if } v_i = 1; v_{i+1} = 1 \end{cases}$$

to obtain the desired result, where bootstrapping keys are set up in Tab. 1. However, the situation is more complex for the first $2^i + 2^{i+1}$ indexes, since the additional automorphism $X \mapsto X^{-1}$ needs to be performed. To eliminate the time-consuming FFT induced by automorphism, we introduce the following method.

Parametrized External Product. We first review the parametrized external product proposed by Bernard et al. [7]. Specifically, the variant RGSW ciphertext associated with the automorphism τ_{-1} is defined as

$$\text{RGSW}_{\mathfrak{s}}^{\tau_{-1}}(\mathbf{m}_1) = (\text{RLWE}'_{\mathfrak{s}}(-\tau_{-1}(\mathfrak{s}) \cdot \mathbf{m}_1), \text{RLWE}'_{\mathfrak{s}}(\mathbf{m}_1)).$$

Then, given ciphertext $c = (\mathbf{a}, \mathbf{b}) \in \text{RLWE}_{\mathfrak{s}}(\mathbf{m}_2)$, the variant external product associated with τ_{-1} , denote as $\boxtimes_{\tau_{-1}}$, is evaluated as:

$$\begin{aligned} & \text{RLWE}(\mathbf{m}_2) \boxtimes_{\tau_{-1}} \text{RGSW}_{\mathfrak{s}}^{\tau_{-1}}(\mathbf{m}_1) \\ &= \tau_{-1}(\mathbf{a}) \odot \text{RLWE}'(-\tau_{-1}(\mathfrak{s}) \cdot \mathbf{m}_2) + \tau_{-1}(\mathbf{b}) \odot \text{RLWE}'(\mathbf{m}_2) \\ &= \text{RLWE}(-\tau_{-1}(\mathbf{a}) \cdot \tau_{-1}(\mathfrak{s}) \cdot \mathbf{m}_2 + \tau_{-1}(\mathbf{b}) \cdot \mathbf{m}_2) \\ &= \text{RLWE}(\tau_{-1}(\mathbf{m}_1) \cdot \mathbf{m}_2 + \tau_{-1}(\mathfrak{e}_1) \cdot \mathbf{m}_2). \end{aligned}$$

In fact, this step integrates the automorphism τ_{-1} and multiplication into a single external product operation, thereby avoiding the additional RLWE key switching (including FFTs) induced by the automorphism. Nevertheless, even so, the problem remains fundamentally unsolved. This is because both operations $c \boxtimes$ and $\text{HomoAuto}_{-1}(c) \boxtimes$ may be applied simultaneously to the input ciphertext c , such as $c_{n-2^i}, c_{n-2^{i+1}}, c_{n-2^i-2^{i+1}}$ in Eq. 6. We observe that the Auto_{-1} can penetrate both gadget decomposition and FFT operations, enabling us to deal with Auto_{-1} in the FFT domain.

Then, we can perform Eq. 8 instead of Eq. 5, where the external product and automorphism share the instance of the Decom and FFT operations.

$$\text{Auto}_{-1}(c) \boxtimes_{\tau_{-1}} \mapsto \text{Decom} + \text{FFT} + \begin{cases} \text{Auto} + \text{HadM}(\text{Bsk}^-) \\ \text{HadM}(\text{Bsk}^+) \end{cases} + \text{iFFT} \quad (8)$$

Note that the penetration technique (termed as Hoisting [8]) has been applied in the BFV and CKKS schemes, but we integrate this technique into TFHE for the first time to minimize FFT computations. Moreover, for the remaining intervals of the map of Eq. 6, we can further break it down into three intervals $0 \sim 2^i - 1, 2^i \sim 2^{i+1} - 1, 2^{i+1} \sim 2^i + 2^{i+1} - 1$. For instance, for the first 2^i indexes, by integrating the multi-bit CMux and the parametrized external product technique, we can perform

$$\hat{c}_j = c_{n-2^i+j} \boxtimes_{\tau_{-1}} \text{Bsk}_{i,0}^- + c_{n-2^{i+1}+j} \boxtimes_{\tau_{-1}} \text{Bsk}_{i,1}^- + c_{n-2^i-2^{i+1}+j} \boxtimes_{\tau_{-1}} \text{Bsk}_{i,2}^- + c_j \boxtimes \text{Bsk}_{i,3}^+ \quad (9)$$

$$= \begin{cases} c_j, & \text{if } v_i = 0; v_{i+1} = 0; \\ -c_{n-2^i+j}, & \text{if } v_i = 1; v_{i+1} = 0; \\ -c_{n-2^i+1+j}, & \text{if } v_i = 1; v_{i+1} = 0; \\ -c_{n-2^i-2^{i+1}+j}, & \text{if } v_i = 1; v_{i+1} = 1. \end{cases}$$

to deal with this case. The remaining indices and complete procedural steps are outlined in Alg. 2, where the symbol \boxtimes represents the external product computed within the FFT domain. Note that our approach executes the (i)FFT only once for all input and output ciphertexts c_0, \dots, c_{n-1} as shown in lines 3 and 18, thus minimizing the FFTs in the iterative layers and eliminating the extra FFTs induced by automorphisms. Finally, we show the noise growth in the lemma 4.1.

Lemma 4.1. Input a list $c_i \in \text{RLWE}(X^{u_i})$ for $0 \leq i \leq n$ with error variance $\text{Var}(c_{in})$, bootstrapping $\text{Bsk}_{i,j}^{\pm}$ for $0 \leq i \leq \ell$, Alg. 2 outputs a list for the encryption for $\mathfrak{t}(X) = u(X) \cdot X^v \bmod X^n + 1$, and its variance satisfies

$$\mathbb{V}_{\text{EMP}} < \text{Var}(c_{in}) + 2\ell \cdot \mathbb{V}_{\text{EP}},$$

where \mathbb{V}_{CMux} is associated with Bsk. The complete proof is present in Appendix C.2.

Table 2: Computation cost for different input bits of EMPmul.

Input-bits	$\delta = 1$	$\delta = 2$	$\delta = 3$	$\delta = 4$
FFT	$2(d+1)n\ell$	$(d+1)n\ell$	$2(d+1)n\ell/3$	$(d+1)n\ell/2$
Hadamard-Mult	$4dn\ell$	$8dn\ell$	$12dn\ell$	$16dn\ell$

Remark 1. Our Alg. 2 (EMPmul) can be extended to larger δ -input bits for v inputs by using the δ -bit CMux operation. Note that larger inputs reduce FFT calls but increase Hadamard

Algorithm 2 Efficient Monomial-polynomial multiplication (EMPMul).

Input:

A list $c_i \in \text{RLWE}(X^{u_i})$ for $0 \leq i \leq n$ representing a polynomial $u(X)$ encrypted in the exponent.

The bootstrapping key $\text{Bsk}_{i,j}$ for $0 \leq i < \ell/2 - 1$ and $0 \leq j \leq 3$ as shown in Tab 1, which is associated with an encrypted value v .

Output:

A list $\hat{c}_i \in \text{RLWE}(X^{w_i})$ for polynomial $t(X) = u(X) \cdot X^v \pmod{X^n + 1}$, where $0 \leq i \leq n$.

```

1: for  $i = 0$  to  $\ell/2 - 1$  do
2:   for  $j = 0$  to  $n - 1$  do
3:      $\hat{c}_j = \text{FFT} \circ \text{Decom}(c_j)$ 
4:   end for
5:   for  $j = 0$  to  $2^{2i} - 1$  do
6:      $\hat{c}_j = c_{n-2^{2i}+j} \boxtimes \tau_{-1} \text{Bsk}_{2i,0}^-$ 
7:        $+ c_{n-2^{2i}+1+j} \boxtimes \tau_{-1} \text{Bsk}_{2i,1}^-$ 
8:        $+ c_{n-2^{2i}-2^{2i}+1+j} \boxtimes \tau_{-1} \text{Bsk}_{2i,2}^- + c_j \boxtimes \text{Bsk}_{2i,3}^+$ 
9:   end for
10:  for  $j = 2^{2i}$  to  $2^{2i+1} - 1$  do
11:     $\hat{c}_j = c_{j-2^{2i}} \boxtimes \text{Bsk}_{2i,0}^+$ 
12:       $+ c_{n-2^{2i}+1+j} \boxtimes \tau_{-1} \text{Bsk}_{2i,1}^-$ 
13:       $+ c_{n-2^{2i}-2^{2i}+1+j} \boxtimes \tau_{-1} \text{Bsk}_{2i,2}^- + c_j \boxtimes \text{Bsk}_{2i,3}^+$ 
14:  end for
15:  for  $j = 2^{2i} + 2^{2i+1}$  to  $n - 1$  do
16:     $\hat{c}_j = c_{j-2^{2i}} \boxtimes \text{Bsk}_{2i,0}^+ + c_{j-2^{2i}+1} \boxtimes \text{Bsk}_{2i,1}^+$ 
17:       $+ c_{j-2^{2i}-2^{2i}+1} \boxtimes \text{Bsk}_{2i,2}^+ + c_j \boxtimes \text{Bsk}_{2i,3}^+$ 
18:  end for
19:  for  $j = 0$  to  $n - 1$  do
20:     $\hat{c}_j = \text{iFFT}(c_j)$ 
21:  end for
22: return  $(c_0, \dots, c_{n-1}) \leftarrow (\hat{c}_0, \dots, \hat{c}_{n-1})$ 

```

multiplications and key sizes, and the computation cost analysis is provided in Tab. 2. In consideration of the trade-off between performance and key sizes, we use $\delta = 2$ as the default parameter in all experimental evaluations (corresponding to Alg. 2).

4.2 Sparse Batched Bootstrapping

Building on the EMPmul procedure, it is easy to construct the BATCHBOOT by the Eq. 3 as shown in Alg. 3, which bootstraps n messages under the full packing case. However, scenarios may arise where bootstrapping a subset of \underline{n} messages ($\underline{n} \mid n$) is required, analogous to the thin bootstrapping mechanism used in BFV/CKKS [17, 35]. We note that the computational cost of thin bootstrapping will decrease if \underline{n} is reduced. However, in the bootstrapping philosophy of Sec. 3, the computational overhead does not change under the sparse packing case.

To address this issue, we propose an efficient batched bootstrapping for the sparse-aware packing. In detail, given the ciphertext $\text{RLWE}_{\mathfrak{s},q}^n(\mathbf{m})$ for $\mathbf{m} = \sum_{i=0}^{\underline{n}-1} \mathbf{m}_{i \cdot n/\underline{n}} \cdot X^{i \cdot n/\underline{n}}$, we can use the RLWE-to-MLWE extraction to generate the MLWE ciphertext as $(\mathbf{a}_0, \dots, \mathbf{a}_{k-1}, \mathbf{b}) \in \text{MLWE}_{\mathfrak{s},q}^{k,\underline{n}}(\mathbf{m})$, where $k = n/\underline{n}$, and $\mathbf{S} = (\mathfrak{s}_0, \dots, \mathfrak{s}_{k-1}) = \text{RingSwitch}(\mathfrak{s})$. According to the MLWE decryption, we have $\mathbf{m} = \mathbf{b} - \sum_{i=0}^{k-1} \mathbf{a}_i \cdot \mathfrak{s}_i$. Note that this extraction process preserves the sparsity for the secret key, i.e., $h = \sum_{i=0}^{k-1} h_w(\mathfrak{s}_i)$. Then, we can rewrite the decryption formula as

$$\mathbf{b} - \sum_{i=0}^{k-1} \mathbf{a}_i \cdot \mathfrak{s}_i = \mathbf{b} - \sum_{i=0}^{k-1} \sum_{j \in \text{idx}(\mathfrak{s}_i)} \left(\sum_{t=0}^{\underline{n}-1} \mathbf{a}(X) \right) \tilde{\odot} X^j.$$

This process enables sequentially performing $h + k$ EMPmul through the MLWE decryption, but with a lower degree of the polynomial $\underline{n} = n/k$. To elaborate further, let the output of $i - 1$ -th iteration be \mathbf{acc} , the i -th iteration can be evaluated as follows

$$\left((\mathbf{acc} \tilde{\odot} X^{-j_{i,0}} - \mathbf{a}) \tilde{\odot} X^{j_{i,0} - j_{i,1}} - \mathbf{a} - \tilde{\odot} X^{j_{i,1} - j_{i,2}} \dots - \mathbf{a} \right) \tilde{\odot} X^{j_{i,h_w(\mathfrak{s}_i)}}. \quad (10)$$

After that, we can set $\mathbf{v}_i = (v_{i,0}, \dots, v_{i,h_w(\mathfrak{s}_i)}) = (-j_{i,0}, j_{i,0} - j_{i,1}, \dots, j_{i,h_w(\mathfrak{s}_i)})$ for $0 \leq i \leq k - 1$, and encapsulate all \mathbf{v}_i as the bootstrapping BSK_i as shown in Tab. 1, where $\text{BSK}_{i,j} = \{ \text{Bsk}_{i,j,0}^\pm, \dots, \text{Bsk}_{i,j,3}^\pm \}$. We show the detailed steps in Alg. 3, and analyze the computational complexity for the full packing and sparse packing cases as follows.

- **Full packing** case necessitates $h + 1$ EMPmul, and each EMPmul involves $O(v \log n)$ external products. Since $v < n$, the overall computation overhead result in $O((h + 1)n \log n)$ external product operations.
- **Sparse packing** case needs to perform the $h + k$ EMPmul, but with a lower degree of \underline{n} , resulting in a total of $O((h + k) \cdot \underline{n} \log \underline{n})$ external product operations. Note that in general settings, we have $k < h < n$.

Remark 2. After performing the accumulation process, it outputs \underline{n} RLWE ciphertext as $\text{RLWE}_{\mathfrak{s},Q}^N(\mathbf{m}_j) = (\mathbf{a}_j, \mathbf{b}_j)$ as shown in line 7 and line 17. If one wants to proceed with the next batched bootstrapping, these $m_{0,0}, \dots, m_{\underline{n}-1,0}$ need to be

Algorithm 3 Batched Bootstrapping

Input:

The RLWE ciphertext $c \in \text{RLWE}_{\mathfrak{s},q}^n(m)$ under a sparse secret key \mathfrak{s} , where $m(X) = \sum_{i=0}^{n-1} m_i X^{i \cdot \frac{n}{d}}$. The bootstrapping key BSK encrypting the binary decomposition for $\mathbf{v}_0, \dots, \mathbf{v}_h$ as shown in Eq. 10 and Tab 1 under a secret key $\tilde{\mathfrak{s}}$. The automorphism keys $\text{Auk}_i : \text{RLWE}'_{\tilde{\mathfrak{s}},Q}(\tau_i(\tilde{\mathfrak{s}}))$ used in RLWE packing.

The RLWE key switching key $\text{Rlk} : \text{RLWE}'_{\tilde{\mathfrak{s}}}(\tilde{\mathfrak{s}})$. n test polynomials testP_i that encode the functions f_0, \dots, f_{n-1} as shown in Alg. 1.

Output:

The RLWE ciphertext that encrypts new message $m' = \sum_{i=0}^{n-1} f_i(m_i) X^{i \cdot \frac{n}{d}}$.

/ Full Packing Case $n = n^*$ */*

- 1: $(\mathbf{a}, \mathbf{b}) \leftarrow \text{ModSwitch}_{q \rightarrow 2N}(c)$
 - 2: Let $\mathbf{acc}_0 = (c_0, \dots, c_{n-1})$, where $c_i = \text{testP}_i \cdot X^{b_i}$
 - 3: for $i = 0$ to h do
 - 4: $\mathbf{acc}_i \leftarrow \text{EMPMul}(\mathbf{acc}_i, \text{BSK}_i)$
 - 5: $\mathbf{acc}_i = (\text{acc}_{i,0} \cdot X^{-a_0}, \dots, \text{acc}_{i,n-1} \cdot X^{-a_{n-1}})$
 - 6: end for
 - 7: $\mathbf{acc}_h \leftarrow \text{EMPMul}(\mathbf{acc}_h, \text{BSK}_h)$
 - 8: $c' = \text{Repack}(\mathbf{acc}_h, \text{Auk})$
 - 9: $c' = \text{RLWE.KeySwitch}_{\tilde{\mathfrak{s}} \rightarrow \mathfrak{s}}(c', \text{Rlk})$
 - 10: **return** c'
- /* Sparse Packing Case $n|n^*$ */*
- 11: $(\mathbf{a}_0, \dots, \mathbf{a}_{k-1}, \mathbf{b}) \leftarrow \text{RtM.Ext}_{n \rightarrow k, n}(\text{ModSwitch}_{q \rightarrow 2N}(c))$
 - 12: Let $\mathbf{acc}_{0,0} = (c_0, \dots, c_{n-1})$, where $c_i = \text{testP}_i \cdot X^{b_i}$
 - 13: for $i = 0$ to $k-1$ do
 - 14: for $j = 0$ to $hw(\mathfrak{s}_i)$ do
 - 15: $\mathbf{acc}_{i,j} \leftarrow \text{EMPMul}(\mathbf{acc}_{i,j}, \text{BSK}_{i,j})$
 - 16: $\mathbf{acc}_{i,j} = (\text{acc}_{i,j,0} \cdot X^{-a_{i,0}}, \dots, \text{acc}_{i,j,n-1} \cdot X^{-a_{i,n-1}})$
 - 17: end for
 - 18: $\mathbf{acc}_{i,hw(\mathfrak{s}_i)} \leftarrow \text{EMPMul}(\mathbf{acc}_{i,hw(\mathfrak{s}_i)}, \text{Bsk}_{i,hw(\mathfrak{s}_i)})$
 - 19: $c' = \text{SRepack}(\mathbf{acc}_{k-1,hw(\mathfrak{s}_{k-1})}, \text{Auk})$
 - 20: $c' = \text{RLWE.KeySwitch}_{\tilde{\mathfrak{s}} \rightarrow \mathfrak{s}}(c', \text{Rlk})$
 - 21: **return** c'
-

repacked into a new RLWE ciphertext under the sparse secret key. We adopted the sparse-secret encapsulation technique of AmortBoot [33]. In detail, the RLWE repacking procedure first packs these ciphertexts into an RLWE encryption under secret key $\tilde{\mathfrak{s}}$, and subsequently performs key-switching to the secret key \mathfrak{s} associated with the input RLWE, which possesses a lower Hamming weight. Since the computational overhead of this step has a negligible effect on the overall performance, we offer the algorithm description in Appendix B.4. Finally, we show the noise growth in Lemma 4.2.

Lemma 4.2. Input a ciphertext $\text{RLWE}_{\mathfrak{s},q}^n(m)$, some function f_i , and the associated key BSK, Auk, Rlk, the batched bootstrapping outputs a new RLWE, which can be further categorized into the following two cases:

1. Full packing outputs encryption of $m' = \sum_{i=0}^{n-1} f_i(m_i) X^i$, and the error variance satisfies $V_{\text{Boot}} < (h+1) \cdot V_{\text{EMP}} + V_{\text{Pack}} + V_{\text{RKS}}$.
2. Sparse packing outputs encryption of $m' = \sum_{i=0}^{n-1} f_i(m_i) X^{i \cdot \frac{n}{d}}$, where the error variance of satisfies $V_{\text{SBoot}} < (h+k) \cdot V_{\text{EMP}} + V_{\text{SPack}} + V_{\text{RKS}}$.

Here $V_{\text{EMP}}, V_{\text{Pack}}, V_{\text{SPack}}, V_{\text{RKS}}$ are associated with BSK, Auk, and Rlk. Note that in the sparse packing, the sparsity parameter k increases noise growth. However, in practical parameter settings where $k \ll h$, the resulting additional noise growth is negligible compared to that under the full packing case. The complete proof is present in Appendix C.2.

4.3 Batched Circuit Bootstrapping

Existing batched bootstrapping, including Alg. 3, focuses on the functional bootstrapping that can only evaluate the single-input function, i.e., $f(x) : \mathbb{Z}_p \rightarrow \mathbb{Z}_{p_{\text{out}}}$, and has low message precision ($p < 2^8$). When higher precision is required, its efficiency decreases significantly due to the necessity of enlarging the system parameters. We note that the circuit bootstrapping [21] and emerging variants [38, 40, 51] effectively support a wider variety of functions without parameter inflation, such as the higher precision via multi-input function, e.g. $f(x_1, x_2) : \mathbb{Z}_p \times \mathbb{Z}_p \rightarrow \mathbb{Z}_{p_{\text{out}}}$. We propose the BATCHBOOT, a batched circuit bootstrapping that integrates the method of [40] (called EP-CBS). In particular, the workflow of EP-CBS can be divided into two subroutines, ciphertext conversion from $\text{LWE}(m)$ to $\text{RGSW}(X^{\tilde{m}})$ and LUT evaluation with the external product tree. Our main insight is that BATCHBOOT can be used in the first stage to obtain better amortized results. Our approach is presented in the following two steps.

Step 1: Half.BatchCBS. The ciphertext conversion from $\text{LWE}(m)$ to $\text{RGSW}(X^{\tilde{m}})$ is constructed in EP-CBS through an expensive blind rotation as shown in Alg. 1. Based on the BATCHBOOT, we construct a batched ciphertext conversion from $\text{RLWE}_{\mathfrak{s},q}^n(m)$ to $\text{RGSW}(X^{\tilde{m}_i})$, for $i \in [0, n-1]$, denoted as Half.BatchCBS. Specifically, we begin by converting the ciphertext modulus from q to $\frac{2N}{d}$ and scaling the ciphertext by a factor of \bar{d} , adhering to the encoding technique presented in EP-CBS [40]. This operation clears the $\log_2 \bar{d}$ least significant bits (LSBs) of the encoding space. These zeroed bits subsequently allow for the gadget vector to be embedded within the initial accumulator. To illustrate, for $\bar{d} = 2$, the n initial accumulators are configured as $\text{acc}_i = \bar{B}_0 \cdot X^{b_i} + \bar{B}_1 \cdot X^{b_i+1}$.

By performing BATCHBOOT without repacking, it outputs n RLWE ciphertexts, that are $c_i = \text{RLWE}_{\tilde{\mathfrak{s}},Q}^N(\bar{B}_0 X^{\tilde{m}_i} + \bar{B}_1 X^{\tilde{m}_i+1})$, for $0 \leq i \leq n-1$. After that, we can perform the ciphertext conversion of EP-CBS [40]

$$\text{RLWE}(\bar{B}_0 X^{\tilde{m}_i} + \bar{B}_1 X^{\tilde{m}_i+1}) \xrightarrow{\text{HomoTrace}} \text{RLWE}'(X^{\tilde{m}_i}) \xrightarrow{\text{SchemeSwitch}} \text{RGSW}(X^{\tilde{m}_i}).$$

Algorithm 4 Batched Circuit Bootstrapping

Input:

The RLWE ciphertext $c \in \text{RLWE}_{\mathfrak{s}, p/q}^n(\mathfrak{m})$ under a sparse secret key \mathfrak{s} , where $\mathfrak{m}(X) = \sum_{i=0}^{n-1} m_i X^i$.

The bootstrapping key BSK used in Alg. 3.

The automorphic key Auk as shown in Alg. 5.

The scheme switching key Ssk : $\text{RLWE}'_{\mathfrak{s}}(\tilde{\mathfrak{s}}^2)$.

Gadget vector $\mathfrak{g} = (\overline{B}_0, \dots, \overline{B}_{\overline{d}-1})$ corresponding to the output RGSW.

A multi-input function $f : [0, p-1]^\beta \rightarrow [0, p-1]^\gamma$.

Output:

Some ciphertexts of $f(m_0, \dots, m_{\beta-1})$.

- 1: Let $(\mathfrak{a}, \mathfrak{b}) = \text{ModSwitch}_{q \rightarrow 2N/\overline{d}}(c) \cdot \overline{d} \pmod{2N}$
 - 2: for $i = 0$ to $n-1$ do
 - 3: $\text{acc}_i = \overline{B}_0 \cdot X^{\mathfrak{b}_i} + \dots + \overline{B}_{\overline{d}-1} \cdot X^{\mathfrak{b}_i + \overline{d} - 1}$
 - 4: end for
 - 5: $(c_0, \dots, c_{n-1}) \leftarrow \text{BatchBoot}(\text{acc}_0, \dots, \text{acc}_{n-1}, \text{BSK})$ /*
 Perform lines 2-6 of Alg. 3 */
 - 6: for $i = 0$ to $n-1$ do
 - 7: for $j = 0$ to $\overline{d}-1$ do
 - 8: $c_{i,j} = c_i \cdot X^{-j}$
 - 9: $\overline{C}_{i,j+\overline{d}} = \text{HomoTrace}.\text{Tr}_{R_N/R_{N/\overline{d}}}(c_{i,j}, \text{Auk})$
 - 10: $\overline{C}_{i,j} = \text{Scheme.Switch}(\overline{C}_{i,j+\overline{d}}, \text{Ssk})$
 - 11: end for
 - 12: end for
 - 13: for $i = 0$ to $\gamma-1$ do
 - 14: $c'_i = \text{EP.Tree}(\overline{C}_0, \dots, \overline{C}_{\beta-1})$ /* Perform external product tree in Alg. 8 *, where $\beta < n$ */
 - 15: end for
 - 16: **return** $c'_0, \dots, c'_{\gamma-1}$
-

as shown in lines 6-12 of Alg. 4. Here, since \overline{B}_0 and \overline{B}_1 fall into distinct equivalence classes, the HomoTrace (HT) $\text{Tr}_{R_N/R_{N/2}}$ ³ is employed to generate the $\text{RLWE}(\overline{B}_0 \cdot X^{\overline{m}_i})$. Subsequently, the term \overline{B}_1 is rotated to align with \overline{B}_0 (via multiplication by X^{-i}), and the HomoTrace is repeated to produce a second RLWE ciphertext for \overline{B}_1 . This procedure generalizes to any length \overline{d} , ultimately yielding a gadget RLWE ciphertext $\text{RLWE}'(X^{\overline{m}_i})$. Finally, the secret key $\tilde{\mathfrak{s}}$ is embedded into the ciphertext $\text{RLWE}(-B_i \tilde{\mathfrak{s}} \cdot X^{\overline{m}_i})$ through \overline{d} times scheme switching as shown in Appendix B.3, and the $2\overline{d}$ ciphertexts make up the ciphertext $\text{RGSW}(X^{\overline{m}_i})$.

Step 2: External product tree. Upon receiving RGSW ciphertexts $\text{RGSW}(X^{\overline{m}_i})$ generated by Half.BatchCBS, we use the external product tree proposed in EP-CBS to compute the multi-input function, such as $f(m_0, m_1)$. Given the ciphertexts $\text{RGSW}(X^{\overline{m}_0})$ and $\text{RGSW}(X^{\overline{m}_1})$ through the batched circuit bootstrapping, one can first perform p external products at the first level. The results are then repacked into a new RLWE

³We show the algorithm for general case Tr_{R_N/R_n} and noise growth in Appendix B.4, and more details can be referred to [38, 40].

ciphertext, corresponding to the second level of the test polynomial form. The final result is obtained after completing the last level.

Alg. 8 formalizes this for a generic function $f : [0, p-1]^\beta \rightarrow [0, p-1]^\gamma$. For each sub-function $f : [0, p-1]^\beta \rightarrow [0, p-1]$, the multi-valued bootstrapping method [13] (as shown in Appendix A) can be used at the first level,⁴ which reduces $p^{\beta-1}$ external products. Hence, this process involves $\gamma \cdot \sum_{i=0}^{\beta-2} p^i = \gamma \cdot \frac{p^{\beta-1}-1}{p-1}$ external products, and $\gamma \cdot \sum_{i=0}^{p-2} p^i = \gamma \cdot \frac{p^{\beta-1}-1}{p-1}$ repacking operations. We omit the associated correctness analysis, which can be obtained directly from the external product and repacking steps. For a more detailed analysis, please refer to scheme [40], and the noise growth is presented in Appendix C.2. Finally, we adopt the notation $\text{BatchCBS}_{p,\beta,\gamma}$ for configurations parameterized by message precision p , input size β in EP.Tree, and the number invocations γ for EP.Tree step.

5 Application for Batched Bootstrapping

In this section, we show two practical applications for the proposed batch bootstrapping: private set intersection and 8-bit FHE instruction sets.

5.1 Application I: Private Set Intersection

5.1.1 Universal Techniques for PSI

In PSI settings, sender holds a set X of size N_x , and the receiver holds a set of size N_y . Both sets consist of σ -bit strings. Like most FHE-based PSI protocols, our protocol is particularly powerful against unbalanced PSI, where the sender's set is much larger than the receiver's set, i.e. $N_y \ll N_x$. The ideal PSI functionality outputs $X \cap Y$ to the receiver and nothing to the sender. Hashing-to-bins can reduce the total number of comparisons and the cost of each comparison in a PSI protocol [25, 43]. Elements are distributed into a predetermined number of bins, and comparisons are only performed between client and server elements in the same bin. Since the bin in which an element is placed is determined using hashes of that element, it is necessary to pad each bin with dummy elements up to a predetermined size ω to prevent the load of the bins from leaking information about the dataset. The hash functions required in PSI are as follows:

- *Cuckoo hashing* [46]. Cuckoo hashing can place N_y elements in $b = \zeta N_y$ bins, for some constant ζ , and each bin holds at most one element. Typically, the placement of each element is determined using $k = 3$ hash functions, and $b \approx 1.27N_y$.

⁴Details of this part can be found in Sec. 4 of [40].

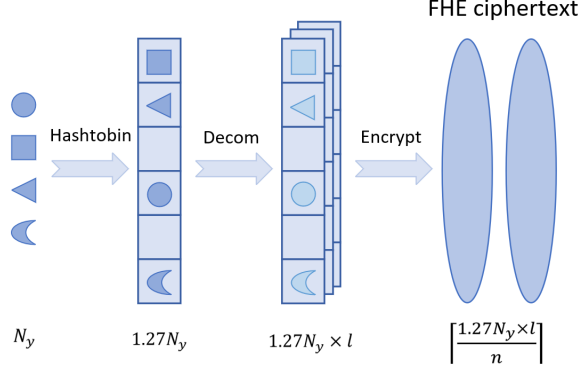


Figure 1: Dataset Preprocessing for client.

- *Permutation-based Hashing* [25, 43]. Permutation hashing can reduce the length of the elements being compared from σ to $\sigma' = \sigma - c$, where $c = \lceil \log_2 b \rceil$.

5.1.2 PSI protocol with Batched Bootstrapping

Our PSI protocol involves three stages: dataset preprocessing, server computation, and client decryption.

Step 1: Dataset Preprocessing. The parties first agree on FHE and PSI parameters such as $n, p, p_{out}, q = 2N$, and $b, \omega, \sigma, \sigma'$. Then, the client and server can asynchronously preprocess their datasets to prepare them for the next stage. Specifically, the preprocessing of the server is as follows: use permutation hash and insert these items into a two-dimensional hash table \mathcal{B}^x using h hash functions H_1, \dots, H_h , where \mathcal{B}^x has b bins and each bin has ω elements, labeled as $B^x[i][j] \in \{0, 1\}^{\sigma'}$ for $i \in [0, b-1], j \in [0, \omega-1]$; and then decompose all elements of \mathcal{B}^x with the base p , i.e., $x = \sum_{i=0}^{\ell-1} p^i \cdot x_i$ ($\ell = \log_p 2^{\sigma'}$) to obtain ℓ sub-tables $\mathcal{B}_0^x, \dots, \mathcal{B}_{\ell-1}^x$, where p is the plaintext modulus of FHE; finally, encode each $B^x[i][j]$ as a polynomial $\text{TB}_{i,j}^k = X^{-\Delta} \cdot \mathcal{B}_k^x[i][j] \sum_{i=-\Delta/2}^{\Delta/2} X^{-i}$, where $\Delta = \frac{N}{p}$.

On the other hand, the client first performs the permutation with Cuckoo hash that outputs a one-dimensional hash table \mathcal{B}^y with $1.27N_y$ elements. After that, it decomposes all elements to obtain ℓ sub-tables $\mathcal{B}_0^y, \dots, \mathcal{B}_{\ell-1}^y$. Finally, the client concatenates each sub-table and fills it with $\lceil 1.27 \cdot N_y / n \rceil$ plaintext polynomials, then encrypts them into RLWE ciphertext with modulus q as shown in Fig. 1. Note that through the decomposition, we can deal with large data sizes $\sigma' > \log_2 p$, which also means that we need to compare x and y with a match being successful when all corresponding x_i and y_i are equal for $i \in [0, \ell-1]$.

Step 2: PSI Computation. In this phase, the client generates the evaluation keys as shown in Alg. 3, and then transmits the evaluation keys and ciphertexts to the server. The server receives the ciphertexts and performs batched bootstrapping for all ciphertexts. This process produces the intermediate ciphertexts corresponding to the elements of sub-tables $\mathcal{B}_0^y, \dots, \mathcal{B}_{\ell-1}^y$

as $\text{RLWE}(X^{\mathcal{B}_k^y[i]})$ as shown in lines 1-7 of Alg. 3, where $\mathcal{B}_k^y[i] = \Delta \cdot \mathcal{B}_k^y[i] + e$, and e is noise of RLWE that satisfies $|e| < \frac{\Delta}{2}$. Then, the server evaluates the equivalence operation

$$\begin{aligned} \text{RLWE}(X^{\mathcal{B}_k^y[i]}) \cdot \text{TB}_{i,j}^x &= \text{RLWE}(X^{\mathcal{B}_k^y[i]}) \cdot X^{-\Delta \cdot \mathcal{B}_k^x[i][j]} \sum_{i=-\Delta/2}^{\Delta/2} X^{-i} \\ &= \text{RLWE}\left(X^{\Delta \cdot \mathcal{B}_k^y[i] - \Delta \cdot \mathcal{B}_k^x[i][j]} \sum_{i=-\Delta/2}^{\Delta/2} X^{e-i}\right) \\ &= \text{RLWE}(m_b + *). \end{aligned}$$

for all $j \in [0, \omega-1], i \in [0, b-1], k \in [0, \ell-1]$, and uses the sample extraction to obtain LWE ciphertexts of m_b . Note that we can conclude that if $\mathcal{B}_k^x[i]$ equals $\mathcal{B}_k^y[i][j]$, $m_b = 1$; otherwise, it equals 0.

Subsequently, the server carries out the aggregation operations by LWE additions. The first step is to sum all elements of each bin as $c_{k,i} = \sum_{j=0}^{\omega-1} c_{k,i,j}$ with the result being an encrypted value of either 0 or 1. The second aggregation step sums all sub-tables as $c_i = \sum_{k=0}^{\ell-1} c_{k,i}$ to get b LWE ciphertexts. At this moment, the encrypted value m_i satisfies $m_i \in \{0, \ell-1\}$, where $m_i = \ell$ corresponds to a successful intersection result. In addition, to prevent information leakage, the server selects a random $r_i \leftarrow \mathbb{Z}_{p_{out}}^*$ and computes $c^i = (c_i - \ell) \cdot r_i$, for $i \in [0, n-1]$. This step serves to mask the other information, and a successful intersection implies that the encryption yields 0. To reduce the communication, the server uses Repack operation to integrate these LWE ciphertexts into an RLWE ciphertext. Finally, we apply bootstrapping once again to refresh the noise, ensuring that the final output noise is independent of the noise generated during the PSI computation step, as shown in [48]. Notably, we avoid the use of noise flooding techniques [15, 16, 43], as it results in an undesirable expansion of the system parameters.

Step 3: Decryption. In the decryption stage, the client decrypts the ciphertext received from the server to obtain the indicator vector. The indicator vector identifies whether a match has been found in the server set for each element in the client set. The client then extracts the indices encrypted as 0 and generates the intersection result.

5.1.3 Analysis of Our PSI Protocol

Security Analysis. Our PSI protocol operates in the semi-honest model, i.e., the client and server follow the protocol but may try to infer extra information. Client input privacy is guaranteed due to the use of homomorphic encryption. On the one hand, the random number r is selected to prevent information leakage of each bin. On the other hand, considering that the noise level in homomorphic encryption texts may reveal additional information about the server dataset, we use the bootstrapping method (also known as sanitisation, see [48]) to refresh the noise level, thereby protecting circuit privacy.

⁵ $r_i \in [0, p_{out}-1]$, and p_{out} is set to a prime number greater than ℓ .

Communication Complexity. The total communication complexity consists of the request and response messages. In the request process, we have $\ell = \lceil \frac{2^{\sigma'}}{p} \rceil$ tables, each containing $1.27 \cdot N_y$ elements. These elements are then encrypted, resulting in a total of $n_{\text{req}} = \lceil \ell \cdot 1.27 \cdot N_y / n \rceil$ ciphertexts. In addition, since we use RLWE symmetric encryption, the part of a can be replaced by random seeds, thus reducing the ciphertext size by half. In the response phase, the server only returns $n_{\text{res}} = \lceil 1.27 \cdot N_y / n \rceil$ ciphertexts, after applying modulus switching to scale the ciphertext modulus down to an input modulus q . A key advantage of our protocol is the compact FHE parameter set, notably a 12-bit ciphertext modulus q . Consequently, the overall communication amount is minimized to about 15 KB. This optimization is a unique feature of our TFHE bootstrapping design and is infeasible for any past or future PSI implementation relying on BFV.

Computation Complexity. The computational cost of our protocol comprises $n_{\text{req}} + n_{\text{res}}$ BATCHBOOT processes, $\ell \omega \cdot b$ scalar multiplications for RLWE, and $\ell \omega \cdot b + \ell b$ LWE ciphertext additions. In addition, we set the temporary ciphertext modulus to the machine word size (i.e., $Q = 2^{64}$), leading to a substantial improvement in computational efficiency. We present a detailed performance and comparison analysis in Sec. 6.2.

5.2 Application II: FHE Instruction Sets

The FHE instruction set can handle addition, multiplication arithmetic operations, and logical operations. Zama [12] achieves these universal instruction sets based on the non-batch TFHE scheme, and further designs a privacy-preserving fhEVM coprocessor [26] for blockchain. Currently, several TFHE-based techniques can be utilized to construct these processors, including gate bootstrapping [12], tree-based bootstrapping [49], and circuit bootstrapping [50]. Nevertheless, these methods are limited to processing only one encrypted and one instruction set at a time. We employ BATCHCBOOT for batch computation of the instruction set, specifically targeting 8-bit messages as in state-of-the-art methods [49, 50]. We customize each instruction according to its operational characteristics, thereby ensuring efficient execution following the approach adopted in [49]. The proposed methodology comprises the following steps.

Message encoding and encryption. For the 8-bit instruction set, we aim to compute $f(u, v)$, such as the $\text{ADD}(u, v) = u + v \pmod{2^8}$, where $u, v \in \mathbb{Z}_{2^8}$. Given $2n$ messages $\mathbf{u} = (u_0, \dots, u_{n-1})$ and $\mathbf{v} = (v_0, \dots, v_{n-1})$, we first decompose each message into two 4-bit, i.e., $u_i = u_{i,L} + 16 \cdot u_{i,M}$, which means that $p = 2^4$ in FHE system parameters. After that, we can encrypt these decomposed messages into four RLWE ciphertexts $(c_{u,L}, c_{u,M}, c_{v,L}, c_{v,M})$.

Instruction Sets Evaluation. In the evaluation of the

general 16-bit-to-8-bit LUT, i.e. BatchCBS_{4,4,2}, we perform Half.BatchCBS on these ciphertexts to obtain n ciphertext tuples, and each tuple containing ciphertexts $\text{RGSW}(X^{\tilde{u}_{i,L}}), \text{RGSW}(X^{\tilde{u}_{i,M}}), \text{RGSW}(X^{\tilde{v}_{i,L}}), \text{RGSW}(X^{\tilde{v}_{i,M}})$. Then, we can construct two external product trees, i.e., $\beta = 4$ and $\gamma = 2$ in Alg. 8, to evaluate the general LUT. This process demands four Half.BatchCBS invocations, along with $2 \times 16^2 \times n = 512n$ instances of repacking and an equal number ($512n$) of external product operations. The

Trama et al. [49] point out that the number of homomorphic operations can be reduced based on the characteristics of the instruction set. For example, the AND instruction can be defined as $\text{AND}(u, v) = (u_M \& v_M, u_L \& v_L)$. Then, we can use two sub-procedures

$$\text{BatchCBS}_{4,2,1}(c_{u,M}, c_{v,M}, \text{AND})$$

$$\text{BatchCBS}_{4,2,1}(c_{u,L}, c_{v,L}, \text{AND})$$

to evaluate $u_M \& v_M$ and $u_L \& v_L$, respectively. Since the level of the external product tree is two (i.e., $\beta = 2$), we can reduce the number of external products operations and repacking steps from $512n$ to $2n$. We show the computational costs for different instruction sets in Tab 3, with detailed computation steps for each instruction outlined in [49].

Table 3: Computational costs for various 8-bit instruction sets.

Instructions	Half.BatchCBS	External product	Repacking
AND/XOR/OR	4	$2n$	$2n$
ADD/SUB	6	$4n$	$4n$
MUL	6	$4n$	$4n$
EQ	6	$3n$	$3n$
LT(E)/GT(E)	8	$5n$	$5n$
MIN/MAX	12	$10n$	$10n$

6 Evaluation

6.1 Methodology

Our algorithms mainly involve two ciphertext types: RLWE and RGSW, which serve as input ciphertext and bootstrapping keys. By default, the ciphertext modulus is set to 2^{64} , unless stated otherwise. The bootstrapping keys include bootstrapping key Bsk, automorphism key Auk, scheme switching key Ssk, and RLWE sparse switching key Rlk, each associated with the length d and base B used in the approximate gadget decomposition. We show the standard deviation of noise for different evaluation keys in Tab. 9 of Appendix C.3. We use the lattice estimator⁶ to evaluate the security, where the security levels of these parameters exceed 128 bits. For the sparse secret key, we adopt the strategy [33] where the key generation

⁶<https://github.com/malb/lattice-estimator>.

process has a high probability of satisfying $j_i - j_{i+1} < n/h$. The security level is determined by using the Lattice Estimator [2] using the default reduction model and dual hybrid attack with meet-in-the-middle optimizations from [18]. Note that Rlk involves a larger standard deviation because it uses sparse keys for encryption. Then, in Appendix D.1, we show the parameters of BATCHBOOT in Tab. 10 and BATCHCBOOT in Tab. 11, respectively. Finally, under these parameters, we can calculate a concrete decryption failure rate. The detailed calculation process can be found in Appendix C.4.

6.2 Performance for Batched Bootstrapping

We show some experimental results for BATCHBOOT, and all of our algorithms are executed in the MOSFHE [32] library. The benchmarking experiments were conducted on a cloud server equipped with an Intel(R) Xeon(R) Gold 6258R CPU @ 2.70GHz and 32 GB of RAM. The operating system was Ubuntu 22.04.4 LTS, and the code was compiled with g++ 14.2.0, which supports AVX-512 instruction set optimization. The single-threaded performance results of our Alg. 3 are presented in Tab. 4. We compare it with the state-of-the-art batched [33] and non-batched [12] implementations across various levels of message precision.

Table 4: Performance, key sizes, and decryption failure rate. Note that set $S_{\text{Boot}_{i,k}}$ is associated with the same set of parameters Boot_i and $\underline{n} = n/k$, which is used in sparse packing case.

Scheme	Sets	$\log_2 p$	Packed Mess.	Total Time	Amort. Time	Failure Prob.	Key Sizes	Speedup
TFHE-rs [12]	Def	2	1	9.8 ms	9.8ms	2^{-53}	76.3 MB	2.74 \times
AmortBoot [33]	Boot ₂	2	2048	16.57s	8.09 ms	2^{-120}	17.1 MB	2.2 \times
BatchBoot	Boot ₂	2	2048	7.32s	3.57 ms	2^{-120}	59.6 MB	–
BatchBoot	SBoot _{2,2}	2	1024	3.83 s	3.74 ms	2^{-120}	70.5 MB	–
TFHE-rs [12]	Def	4	1	13.6 ms	13.6 ms	2^{-53}	52.1 MB	3.77 \times
AmortBoot [33]	Boot ₄	4	2048	17.72s	8.65 ms	2^{-94}	18.4 MB	2.4 \times
BatchBoot	Boot ₄	4	2048	7.39s	3.6 ms	2^{-94}	64.1 MB	–
BatchBoot	SBoot _{4,2}	4	1024	3.86 s	3.71 ms	2^{-94}	75.8 MB	–
TFHE-rs [12]	Def	6	1	131.1 ms	131.1 ms	2^{-53}	488.5 MB	15.55 \times
AmortBoot [33]	Boot ₆	6	4096	36.53s	8.92 ms	2^{-64}	18.6 MB	1.05 \times
BatchBoot	Boot ₆	6	4096	34.57s	8.43 ms	2^{-64}	120.1 MB	–
BatchBoot	SBoot _{6,2}	6	2048	18.31 s	8.94 ms	2^{-64}	142 MB	–
TFHE-rs [12]	Def	8	1	1.1 s	1.1 s	2^{-53}	3.25 GB	43.8 \times
AmortBoot [33]	Boot ₈	8	4096	234.88 s	57.34 ms	2^{-62}	76.6 MB	2.27 \times
BatchBoot	Boot ₈	8	4096	102.86s	25.11 ms	2^{-62}	205.3 MB	–
BatchBoot	SBoot _{8,2}	8	2048	54.54 s	26.63 ms	2^{-62}	242.6 MB	–

Compared to TFHE-rs (using their default parameters⁷), we achieve gains of up to 43.8 times under 8-bit precision. After that, we implement the open source codes⁸ corresponding to scheme [33] in the same test environment. Our method achieves a speedup of up to 2.4 \times compared to [33] due to a drastic reduction in FFT operations. Interestingly, we can reduce the latency from 7.39 s to 3.83 s for 4-bit LUT

⁷<https://github.com/zama-ai/tfhe-rs>

⁸<https://github.com/antonioceg/Fast-Amortized-Bootstrapping>

while maintaining approximately the same amortization results through the sparse packing case, which is not feasible with [33]. Tab. 4 also shows the comparison of the bootstrapping key sizes for different schemes. Although the key sizes of BATCHBOOT are larger than [33], the performance gains of BATCHBOOT are substantial. Additionally, the bootstrapping key is generated only once and therefore does not become a bottleneck in practical applications. On the other hand, we note that BatchBoot exhibits excellent scalability in multi-core CPU environments, since FFT and external products in the EMPmul procedure are inherently independent operations. Multi-threading experiments show that latency decreased from 102.8 s to 19 s under the 8-bit message precision by using 8 threads as shown in Tab. 12, and more detailed performance results and comparisons are present in Appendix D.2.

Table 5: Performance comparison with the CKKS functional bootstrapping.

Scheme	Packed Mess.	4-4 LUT		6-6 LUT	
		Total Time	Amort. Time	Total Time	Amort. Time
CKKS [3]	65536	9.93 s	0.15 ms	12.7 s	0.19 ms
	2048	5.58 s	2.72 ms	6.97 s	3.4 ms
	512	5.35 s	10.4 ms	6.57 s	12.8 ms
BatchBoot	2048	7.39 s	3.6 ms	18.31 s	8.94 ms
	512	2.11 s	4.13 ms	5.4 s	10.7 ms

In addition, we compare BATCHBOOT with the state-of-the-art CKKS functional bootstrapping [3] by executing their open-source code⁹, where the performance results are summarized in Tab. 5. From the result, we can derive the fact that CKKS bootstrapping performs better in the fully packing mode. However, under some sparse packing (e.g., 512 slots), BATCHBOOT achieves 2.5 \times improvement under 4-bit LUT evaluation by leveraging better sparse packing techniques. Moreover, our scheme still offers two additional advantages over CKKS-based bootstrapping. Firstly, CKKS relies on a super-polynomial modulus, while BATCHBOOT is built on a polynomial-modulus RLWE assumption with stronger security as shown in Appendix A. Secondly, CKKS bootstrapping entails huge key sizes, typically on the order of gigabytes in [3, 5, 6, 29, 36]. Even though HERMES [4] incorporates some optimization techniques to reduce bootstrapping key sizes, it still requires a key size of 673 MB under functional bootstrapping parameters. In contrast, BATCHBOOT achieves a 3.2–11.3 \times reduction in key sizes. Finally, we also present the peak memory usage in single-threaded setting in Tab. 13 of Appendix D. For 2-bit and 4-bit LUT parameters, BATCHBOOT utilizes 360 MB of memory, while CKKS-based scheme [29] demands a minimum of 13.8 GB. This significant disparity underscores the superior memory efficiency of

⁹<https://github.com/openfheorg/openfhe-development/tree/main/src/pke>

Table 6: Performance comparison for PSI protocols.

N_y	N_x	PEPSI [43]			BatchBoot		
		ω	Time (s)	Comm.	ω	Time (s)	Comm.
2^{24}	1024	6710	9.6	4.21 MB	25616	13.82/0.7×	15 KB/287×
	2048	6710	10.5	4.21 MB	20288	13.85/0.75×	30 KB/143×
	4096	6710	19	8.32 MB	10367	13.92/1.36×	51 KB/170×
2^{28}	1024	100565	152.9	4.32 MB	397376	66.89/2.28×	15 KB/294×
	2048	100565	152.9	4.32 MB	313363	67.26/2.27×	30 KB/147×
	4096	100565	279.9	8.32 MB	157572	67.34/4.1×	51 KB/170×
2^{32}	1024	N/A	N/A	N/A	6308095	870.93	15 KB
	2048	N/A	N/A	N/A	4968820	873.67	30 KB
	4096	N/A	N/A	N/A	2487973	874.98	51 KB

BATCHBOOT.

6.3 Performance for PSI Protocol

For the PSI protocol, we run experiments with $\sigma = 32$ bit elements. By using the permutation-based hashing, we can obtain $\kappa = \lceil \log_2 b \rceil$. Thus, the hash size of elements is $\sigma' = \sigma - \kappa$ bits, where $b \geq 2^{11}$ in our setting. We use the parameter set PSI.Boot_6 , and split each element of N_y into four digits as detailed in Sec. 5.1. Tab. 6 shows a comparison of runtime and communication costs with the scheme [43]. Here the size N_y is set to 1024, 2048, and 4096, and the size of N_x is chosen as 2^{24} , 2^{28} and 2^{32} . The server bin load ω is determined by the server set size and the number of bins, where we adopt the method [43, 46] to obtain the upper bounds for ω [43, 46] and validate them experimentally via binning simulations. The results indicate that our protocol achieves a dramatic reduction of two orders of magnitude (143–294 \times). In addition, we also show the runtimes for PSI under 32 threads. From the experimental results, we observe that for some smaller parameter sets (e.g., $N_x = 2^{24}$ and $N_y = 2048$), our scheme incurs higher latency. In contrast, a performance gain of 4.1 times is achieved over PEPSI [43] when N_x is scaled to 2^{28} . In addition, our solution can also support larger datasets, such as $N_x = 2^{32}$, whereas existing solutions do not show execution results in this case. Therefore, our proposed method is particularly advantageous for deployment in bandwidth-sensitive environments or for processing substantial data volumes. Finally, we present further comparative analyses (such as comparisons with [25]) in Appendix A.

6.4 Performance for General LUT Evaluation

For the batched circuit bootstrapping, since Alg. 4 introduces larger noise growth than Alg. 3 (as shown in Lemma C.2), we need to reset the parameters. We add the parameters \bar{B} , \bar{d} , and the scheme switching key Ssk in Tab. 11. These elements are used in Alg. 4 and can support a message precision of $p = 2^4$ for each sub-block. We first show the performance with general LUTs using our batched circuit bootstrapping

Table 7: Amortization performance of general LUT.

Method	8-to-8 LUT	12-to-12 LUT	Speedup
TBC+ [49]	160.15 ms	4.6 s	12.06/134.3 \times
WHS+ [50]	112.74 ms	170.62 ms	8.5/4.9 \times
LSL+ [40]	630 ms	1.12 s	47/35 \times
BatchBoot	13.27ms	34.24ms	–

Table 8: Performance of 8-bit instruction sets.

Instructions	TBC+ [49]	WHS+ [50]	BatchCBoot	Speedup
AND/OR/XOR	184 ms	227 ms	28.14 ms	6.5 \times
ADD/SUB	493 ms	227 ms	42.68 ms	5.3 \times
MUL	504 ms	227 ms	42.68 ms	5.3 \times
EQ	276 ms	227 ms	42.21 ms	5.4 \times
LT(E)/GT(E)	437 ms	227 ms	56.75 ms	4 \times
MIN/MAX	825 ms	227 ms	68.3 ms	3.3 \times

as shown in Tab. 7. Under the 8-bit case, i.e., $\beta = \gamma = 2$ in Alg. 4, our method is **8.5–47** \times faster than the previous circuit bootstrapping methods (TBC+ [49], WHS+ [50], LSL+ [40]) in terms of amortized results based on the benchmarks they provided. Finally, we evaluate 8-bit FHE instruction sets using the methodology from Sec. 5.2. We implement 6 practical instruction sets, including AND, ADD, MUL, etc., where the computational overhead of instruction evaluation can be optimized according to its circuit complexity. The amortization analysis reveals a **3.3–6.5** \times performance improvement across these instruction sets.

7 Conclusion

This paper presents BATCHBOOT, a batched bootstrapping method for TFHE that processes multiple ciphertexts simultaneously. We design efficient homomorphic polynomial multiplication to reduce the number of FFT operations, and introduce an MLWE-based decryption paradigm for sparse packing, ensuring linear complexity in packing density. We extend the functionality through circuit bootstrapping, enabling the evaluation of more general functions. Overall, BATCHBOOT yields a 2.4 \times speedup over state-of-the-art batched schemes and a 43.8 \times improvement over non-batched implementations. Practical applications include a PSI protocol and an 8-bit FHE instruction set; our approach achieves a 294 \times reduction in communication overhead and a 4.1 \times performance gain over previous BFV-based PSI protocols, and up to 6.5 \times computational speedup for the instruction set.

Ethical Considerations

We have assessed the value of publishing this work against potential implications. 1. Identification of Stakeholders. We performed a comprehensive analysis of all parties who could be affected by our proposed BatchBoot algorithm, PSI protocols, and FHE instruction set design. The following stakeholders were identified: Researchers in privacy-preserving and cryptography: These individuals may use and extend our efficient BatchBoot algorithm. Users or devices requiring private Set Intersection can utilize our proposed PSI to reduce latency and communication overhead. Compiler and systems developers targeting encrypted computing: FHE instruction set helps separate programming from cryptographic details, reduces latency for encrypted operations, and can lower barriers to practical privacy-preserving analytics in regulated domains. 2. Potential Impacts on Stakeholders. Impacts during the research process: The research process poses minimal ethical risk because: All experiments use public benchmark datasets or random number which contain no personally identifiable information. No sensitive organizational or user data were accessed. These properties ensure that the research caused no direct harm to any stakeholder. 3. Mitigations. We release only the artifacts essential for benchmark reproduction, documenting all underlying assumptions and limitations. Production deployments should enforce rigorous key management, fine-grained access control, rate limiting, comprehensive audit logging, and periodic compliance reviews to mitigate residual risks. 4. Rationale for Conducting and Publishing the Research. The motivations for publishing our work are as follows: Advancing batched TFHE has tangible practical value, as it can reduce both computational and communication overhead in PSI protocols. Making the technology publicly available enables the academic community to validate, optimize, and securely build upon its foundations. Given these benefits—along with the minimal associated risks and the presence of robust mitigation measures—we believe that the overall ethical impact of conducting and publishing this research is substantially positive.

Open Science

Following the open science policy, we are dedicated to guaranteeing that our work is transparent and easily available. Our work facilitates the analysis of batch bootstrapping while ensuring data confidentiality and security, thus promoting open science and collaboration. The code for Batchboot is available at <https://doi.org/10.5281/zenodo.17936945>.

References

- [1] Albrecht, M., Chase, M., Chen, H., Ding, J., Goldwasser, S., Gorbunov, S., Halevi, S., Hoffstein, J., Laine, K., Lauter, K., et al.: Homomorphic encryption standard. In: Protecting privacy through homomorphic encryption, pp. 31–62. Springer (2022)
- [2] Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology* **9**(3), 169–203 (2015)
- [3] Alexandru, A., Kim, A., Polyakov, Y.: General functional bootstrapping using ckks. In: Annual International Cryptology Conference. pp. 304–337. Springer (2025)
- [4] Bae, Y., Cheon, J.H., Kim, J., Park, J.H., Stehlé, D.: HERMES: efficient ring packing using mlwe ciphertexts and application to transciphering. In: Annual International Cryptology Conference. pp. 37–69. Springer (2023)
- [5] Bae, Y., Cheon, J.H., Kim, J., Stehlé, D.: Bootstrapping bits with CKKS. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 94–123. Springer (2024)
- [6] Bae, Y., Kim, J., Stehlé, D., Suvanto, E.: Bootstrapping small integers with CKKS. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 330–360. Springer (2025)
- [7] Bernard, O., Joye, M.: Bootstrapping (t)FHE ciphertexts via automorphisms: Closing the gap between binary and gaussian keys. *Cryptology ePrint Archive*, Paper 2025/163 (2025), <https://eprint.iacr.org/2025/163>
- [8] Bossuat, J.P., Mouchet, C., Troncoso-Pastoriza, J., Hubaux, J.P.: Efficient bootstrapping for approximate homomorphic encryption with non-sparse keys. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 587–617. Springer (2021)
- [9] Boura, C., Gama, N., Georgieva, M., Jetchev, D.: Simulating homomorphic evaluation of deep learning predictions. In: International Symposium on Cyber Security Cryptography and Machine Learning. pp. 212–230. Springer (2019)
- [10] Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical GapSVP. In: Annual Cryptology Conference. pp. 868–886. Springer (2012)
- [11] Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Proceedings of the forty-fifth annual ACM symposium on Theory of computing. pp. 575–584 (2013)
- [12] Brenna, L., Singh, I.S., Johansen, H.D., Johansen, D.: TFHE-rs: A library for safe and secure remote computing using fully homomorphic encryption and trusted execution environments. *Array* **13**, 100118 (2022)

- [13] Carпов, S., Izabachène, M., Mollimard, V.: New techniques for multi-value input homomorphic evaluation and applications. In: *Topics in Cryptology—CT-RSA 2019: The Cryptographers’ Track at the RSA Conference 2019*, San Francisco, CA, USA, March 4–8, 2019, Proceedings. pp. 106–126. Springer (2019)
- [14] Chen, H., Dai, W., Kim, M., Song, Y.: Efficient homomorphic conversion between (ring) LWE ciphertexts. In: *International Conference on Applied Cryptography and Network Security*. pp. 460–479. Springer (2021)
- [15] Chen, H., Huang, Z., Laine, K., Rindal, P.: Labeled psi from fully homomorphic encryption with malicious security. In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. pp. 1223–1237 (2018)
- [16] Chen, H., Laine, K., Rindal, P.: Fast private set intersection from homomorphic encryption. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. pp. 1243–1255 (2017)
- [17] Cheon, J.H., Han, K., Kim, A., Kim, M., Song, Y.: Bootstrapping for approximate homomorphic encryption. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. pp. 360–384. Springer (2018)
- [18] Cheon, J.H., Hhan, M., Hong, S., Son, Y.: A hybrid of dual and meet-in-the-middle attack on sparse and ternary secret lwe. *IEEE Access* **7**, 89497–89506 (2019)
- [19] Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: *International Conference on the Theory and Application of Cryptology and Information Security*. pp. 409–437. Springer (2017)
- [20] Chillotti, I., Gama, N., Georgieva, M., Izabachene, M.: Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In: *Advances in Cryptology—ASIACRYPT 2016: 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I* 22. pp. 3–33. Springer (2016)
- [21] Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Faster packed homomorphic operations and efficient circuit bootstrapping for TFHE. In: *International Conference on the Theory and Application of Cryptology and Information Security*. pp. 377–408. Springer (2017)
- [22] Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: fast fully homomorphic encryption over the torus. *Journal of Cryptology* **33**(1), 34–91 (2020)
- [23] Chillotti, I., Ligier, D., Orfila, J.B., Tap, S.: Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for TFHE. In: *International Conference on the Theory and Application of Cryptology and Information Security*. pp. 670–699. Springer (2021)
- [24] Chung, H., Kim, H., Kim, Y.S., Lee, Y.: Amortized large look-up table evaluation with multivariate polynomials for homomorphic encryption. *Cryptology ePrint Archive* (2024)
- [25] Cong, K., Moreno, R.C., da Gama, M.B., Dai, W., Iliashenko, I., Laine, K., Rosenberg, M.: Labeled PSI from homomorphic encryption with reduced computation and communication. In: *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. pp. 1135–1150 (2021)
- [26] Dahl, M., Danjou, C., Demmler, D., Frederiksen, T., Ivanov, P., Joye, M., Rotaru, D.: fhevm: Confidential evm smart contracts using fully homomorphic encryption (2023)
- [27] De Micheli, G., Kim, D., Micciancio, D., Suhl, A.: Faster amortized fhew bootstrapping using ring automorphisms. p. 322–353. Springer-Verlag, Berlin, Heidelberg (2024)
- [28] 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)
- [29] Dumezy, J., Alexandru, A., Polyakov, Y., Clet, P.E., Chakraborty, O., Boudguiga, A.: Evaluating larger lookup tables using ckks. *IACR Cryptol. ePrint Arch.* **2025**, 1301 (2025), <https://api.semanticscholar.org/CorpusID:280038428>
- [30] Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. *Cryptology ePrint Archive* (2012)
- [31] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: *Proceedings of the forty-first annual ACM symposium on Theory of computing*. pp. 169–178 (2009)
- [32] Guimarães, A., Borin, E., Aranha, D.F.: Mosfhet: Optimized software for fhe over the torus. *Journal of Cryptographic Engineering* **14**(3), 577–593 (2024)
- [33] Guimarães, A., Pereira, H.V.: Fast amortized bootstrapping with small keys and polynomial noise overhead. In: *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*. pp. 2967–2981 (2025)

- [34] Guimarães, A., Pereira, H.V., Van Leeuwen, B.: Amortized bootstrapping revisited: Simpler, asymptotically-faster, implemented. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 3–35. Springer (2023)
- [35] Halevi, S., Shoup, V.: Bootstrapping for helib. *Journal of Cryptology* **34**(1), 7 (2021)
- [36] Kim, J., Noh, T.: Modular reduction in CKKS. *Cryptology ePrint Archive* (2024)
- [37] Kluczniak, K., Schild, L.: FDFB: Full domain functional bootstrapping towards practical fully homomorphic encryption. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2023**, Issue 1, 501–537 (2022)
- [38] Lee, K.H., Yoon, J.W.: Homomorphic field trace revisited : Breaking the cubic noise barrier. *Cryptology ePrint Archive*, Paper 2025/1088 (2025), <https://eprint.iacr.org/2025/1088>
- [39] 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. In: Advances in Cryptology–EUROCRYPT 2023: 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23–27, 2023, Proceedings, Part III. pp. 227–256. Springer (2023)
- [40] Li, Z., Shen, X., Lu, X., Wang, R., Zhao, Y., Wang, Z., Wei, B.: Leveled functional bootstrapping via external product tree. *Cryptology ePrint Archive*, Paper 2025/022 (2025), <https://eprint.iacr.org/2025/022>
- [41] Liu, Z., Wang, Y.: Amortized functional bootstrapping in less than 7 ms, with $\tilde{O}(1)$ polynomial multiplications. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 101–132. Springer (2023)
- [42] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. *Journal of the ACM (JACM)* **60**(6), 1–35 (2013)
- [43] Mahdavi, R.A., Lukas, N., Ebrahimiaghazani, F., Humphries, T., Kacsmar, B., Premkumar, J., Li, X., Oya, S., Amjadian, E., Kerschbaum, F.: Pepsi: Practically efficient private set intersection in the unbalanced setting. In: 33rd USENIX Security Symposium. pp. 6453–6470 (2024)
- [44] 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)
- [45] Micciancio, D., Sorrell, J.: Ring packing and amortized FHEW bootstrapping. *Cryptology ePrint Archive* (2018)
- [46] Pagh, R., Rodler, F.F.: Cuckoo hashing. *Journal of Algorithms* **51**(2), 122–144 (2004)
- [47] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)* **56**(6), 1–40 (2009)
- [48] Smart, N.P., Walter, M.: Error-simulatable sanitization for TFHE and applications. *IACR Communications in Cryptology* **2**(3) (2025)
- [49] Trama, D., Boudguiga, A., Clet, P.E., Sirdey, R., Ye, N.: Designing a general-purpose 8-bit (t) fhe processor abstraction. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2025**(2), 535–578 (2025)
- [50] Wang, R., Ha, J., Shen, X., Lu, X., Chen, C., Wang, K., Lee, J.: Refined leveled homomorphic evaluation and its application. In: Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security. pp. 2309–2323 (2025)
- [51] Wang, R., Wen, Y., Li, Z., Lu, X., Wei, B., Liu, K., Wang, K.: Circuit bootstrapping: faster and smaller. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 342–372. Springer (2024)

A Related Work

This section presents related work and provides some comparison results.

BFV/CKKS-based functional bootstrapping. A recent series of works [3, 5, 6, 24, 36, 41] develop BFV and discrete CKKS schemes to support amortized functional bootstrapping. These schemes use polynomial interpolation to evaluate arbitrary functions while reducing error based on the noise cleaning technique. Note that BFV/CKKS schemes are characterized by significant noise growth and large parameters, which often necessitate stronger security assumptions. More precisely, if one can efficiently solve (R)LWE instances in dimension N with modulus Q and discrete Gaussian error with parameter σ , then one can also solve hard problems such as the approximate shortest vector problem, GapSVP_γ , in dimension N with approximation factor $\gamma = O(N \cdot Q/\sigma)$ [11, 47]. To ensure correctness under higher noise in BFV/CKKS schemes, the ciphertext modulus Q must be chosen considerably larger than the error bound e , leading to increased γ for the underlying lattice problem that governs security. As γ grows, the associated lattice problems such as GapSVP_γ become easier to solve. Consequently, assuming the worst-case hardness of GapSVP_γ becomes a stronger security assumption as γ increases.

FHEW/TFHE-based batch functional bootstrapping. In Se. 6.2, we provide the comparison and analysis of the state-of-the-art scheme [33]. We note that there is another route for batch functional bootstrapping for FHEW/TFHE schemes [27, 34, 45]. These schemes use the NTT to evaluate polynomial multiplication in RLWE decryption without utilizing the sparsity of the secret key. This method can reduce the complexity to $O(n \log n)$. However, to keep the noise in approximate polynomial factors, the recursive depth is limited to a small constant. A concrete implementation is presented in [34], with the amortization result of 851 ms under 7 message precision bits. Our method significantly outperforms this, with a 34× speedup. Furthermore, their approach requires very large bootstrapping keys (65.3 GiB in total), our scheme reduces it to just 205.3 MB, a reduction by a factor of 584.

Other FHE-based PSI. The FHE-based PSI protocol [16] was proposed by Chen et al. in 2017, which focuses on unbalanced settings by using the BFV scheme. Their foundational idea was to represent the equality operation via an interpolation polynomial. This evaluation reveals an intersection if the client’s element exists in the server’s set, result equals to zero; otherwise, the result is a random number. To improve performance, Cong et al. [25] introduce some optimizations, such as incorporating extremal postage-stamp bases and Frobenius operation, to reduce multiplicative depth and support for elements of arbitrary bit length. Regarding parameter settings $n = 2^{28}$ and $m = 4096$, the overall latency of the protocol is 2133 s under 32 threads, which is 31 times slower than our PSI solution. However, their protocol achieves faster server-side online computation. Specifically, their method requires only 2.5 s for online time, which is 20.4× faster than our method (51 s in our method). Moreover, our approach yields a 357× reduction in communication overhead, substantially alleviating network bandwidth constraints and enhancing scalability for large-scale deployments.

B Fundamental Algorithms and Noise Growth

B.1 Notations

The lower-case regular bold letters indicate vectors in this paper. For a real number r , we write the floor, ceiling, and round functions as $\lfloor r \rfloor$, $\lceil r \rceil$, $\text{round}(r)$, respectively. We use \mathbb{Z}_Q to denote the ring of integers $\mathbb{Z}/Q\mathbb{Z}$ in the range $[-Q/2, Q/2) \cap \mathbb{Z}$. In addition, we use the cyclotomic ring $R_{N,Q} = \mathbb{Z}_Q[X]/\langle X^N + 1 \rangle$, where N is an integer to the power of 2, and its coefficients are from \mathbb{Z}_Q . The notation a_i or $a[i]$ is used to index the vector and polynomial coefficients. We use \mathfrak{s} or $\mathfrak{s}(X)$ to denote polynomials, and s_i represents the i -th coefficient of the polynomial, where its coefficient vector is $\text{Coef}(\mathfrak{s}) = (s_0, \dots, s_{N-1}) \in \mathbb{Z}_Q^N$. Furthermore, we use $\text{Var}(\mathfrak{a})$ to denote the variance of the Gaussian distribution for $\mathfrak{a} \in R$, and V_{step} to indicate the variance accumulated during step execution.

B.2 Noise Growth for GP, EP, and CMux

For the gadget product with the gadget vector $\mathfrak{g} = (B_0, \dots, B_{d-1})$, the error variance of the gadget product result is bounded by $V_{\text{GP}} = \frac{B^2}{12} dN \cdot \text{Var}(c_{\text{in}}) + \frac{N}{12} \cdot \left(\left\lceil \frac{Q}{B^d} \right\rceil \right)^2$ as shown in [22], where $\text{Var}(c_{\text{in}})$ is the error variance of the input gadget RLWE. Since the external product requires two gadget products, we can derive the noise variance as $V_{\text{EP}} = \frac{B^2}{6} dN \cdot \text{Var}(C_{\text{in}}) + \frac{N}{12} \cdot \left(\left\lceil \frac{Q}{B^d} \right\rceil \right)^2$, where $\text{Var}(C_{\text{in}})$ is the error variance of the input RGSW ciphertext. Furthermore, the CMux has an additional subtraction from the external product, so the noise variance is twice that of the first term of external product, i.e., $V_{\text{CMux}} = \frac{B^2}{3} dN \cdot \text{Var}(C_{\text{in}}) + \frac{N}{12} \cdot \left(\left\lceil \frac{Q}{B^d} \right\rceil \right)^2$.

B.3 Homomorphic Operations

Modulus Switching (MS). Given an RLWE ciphertext $c = (\mathfrak{a}, \mathfrak{b}) \in \text{RLWE}_{s,p/Q}^n(m)$ with error variance $\text{Var}(c_{\text{in}})$, the modulus switching algorithm computes $\text{ModSwitch}_{Q \rightarrow q}(c) = \left(\left\lfloor \frac{q}{Q} \cdot \mathfrak{a} \right\rfloor, \left\lfloor \frac{q}{Q} \cdot \mathfrak{b} \right\rfloor \right)$. It is easy to verify that the ciphertext is a correct ciphertext under modulo q , and its variance is $V_{\text{MS}} \leq \left(\frac{q}{Q} \right)^2 \cdot \text{Var}(c_{\text{in}}) + \frac{\|s\|_2^2 + 1}{12}$, where the factor $\frac{1}{12}$ is the standard deviation of a uniform distribution in $[-1/2, 1/2]$. For the binary secret key with h Hamming Weight, we have $V_{\text{MS}} \leq \left(\frac{q}{Q} \right)^2 \cdot \text{Var}(c_{\text{in}}) + \frac{h+1}{12}$.

RLWE key switching (RKS). The RLWE key switching can convert the ciphertext $\text{RLWE}_{s_1,Q}(m)$ to a new ciphertext $\text{RLWE}_{s_2,Q}(m)$. Given the $c = (\mathfrak{a}, \mathfrak{b}) \in \text{RLWE}_{s_1,Q}(m)$, and the key switching key $\text{Rk} = \text{RLWE}'_{s_2,Q}(s_1)$, the RLWE key switching operation performs $\text{RLWE.KeySwitch}(c, \text{Rk}) = (\mathfrak{0}, \mathfrak{b}) - \mathfrak{a} \odot \text{Rk}$, which outputs the ciphertext $\text{RLWE}_{s_2,Q}(m)$, where the correctness can be found in [22, 39]. The error variance V_{RKS} is the same as that of the gadget product V_{GP} .

Scheme switching (SS). We can use the scheme switching [27, 40, 51] to embed the secret key \mathfrak{s} into the message. Specifically, given the key switching key $\text{Ssk} = \text{RLWE}'_{s,Q}(s^2)$, and the ciphertext $c = (\mathfrak{a}, \mathfrak{b}) \in \text{RLWE}_{s,Q}(m)$, we can compute $\text{Scheme.Switch}(c, \text{Ssk}) = \mathfrak{a} \odot \text{Ssk} + (\mathfrak{b}, \mathfrak{0})$, which outputs ciphertext $\text{RLWE}_{s,Q}(-\mathfrak{s} \cdot m)$ as shown in [27, 51], where the error variance satisfies $V_{\text{SS}} \leq \frac{N}{2} \cdot \text{Var}(c_{\text{in}}) + V_{\text{RKS}}$.

RingSwitch (RS). Ring switching can switch the polynomial ring $R_{N,Q}$ to the subset and coset for polynomial $R_{N/k,Q}$, where both k and N are powers of 2, and $k|N$. Given a polynomial $m(X) = m_0 + m_1X + \dots + m_{N-1}X^{N-1}$, we define the ring switching as $(\mathfrak{m}_0, \dots, \mathfrak{m}_{k-1}) \leftarrow \text{RingSwitch}(m)$, where $\mathfrak{m}_i(X) = m_iX^i + m_{k+i}X^{k+i} + \dots + m_{(N/k-1)k+i}X^{N/k-1} \in R_{N/k,Q}$. This process does not increase noise, and please refer to work [4] for more details.

Algorithm 5 HomoTrace for Tr_{K_N/K_n}

Input:

The RLWE ciphertext $c = (a, b) \in \text{RLWE}_{s,Q}^N(\mathfrak{m}(X))$.
The automorphism keys $\text{Auk}_i = \text{RLWE}'_{s,Q}(\tau_i(\mathfrak{s}))$, where $i \in \{2^{\log N - k + 1} + 1\}$ and $k \in [1, \log N]$.

Output:

RLWE ciphertext $c' \in \text{RLWE}(\text{Tr}_{K_N/K_n}(\mathfrak{m}))$.
1: for $i = \log_2 n + 1$ to $\log_2 N$ do
2: $c' = \text{ModSwitch}_{Q \rightarrow Q/2}(c')$
3: $c' = \text{ModRaise}_{Q/2 \rightarrow Q}(c')$
4: $c' = c' + \text{HomoAuto}_{\tau_{2^{i+1}}}(c', \text{Auk}_{2^{i+1}})$,
5: end for
6: **return** c' .

B.4 HomoTrace and Repacking

Homomorphic Trace (HomoTrace). We show the HomoTrace operation for Tr_{K_N/K_n} ($n|N$) as shown in Alg. 5. Specifically, given the ciphertext $\text{RLWE}(\mathfrak{m}(X))$ for $\mathfrak{m}(X) = m_0 + \dots + m_{N-1}X^{N-1}$, this algorithm outputs a new ciphertext $\text{RLWE}(\text{Tr}_{K_N/K_n}(\mathfrak{m}))$, where $\text{Tr}_{K_N/K_n}(\mathfrak{m}(X)) = \sum_{i=0}^{n-1} m_{i \cdot N/n} \cdot X^{i \cdot N/n}$. Note that this algorithm also incurs an expansion by a factor of N/n . Here, we use the method of RevHomTrace [38] to remove the effect of this factor inflation by evaluating the change in ModSwitch and ModRaise between Q and $Q/2$. In addition, the error variance of HomoTrace is bound by $V_{\text{HT}} \leq \log_2(\frac{N}{n}) \cdot V_{\text{Auto}} + 4 \log_2(\frac{N}{n}) \cdot V_{\text{MS}} + \text{Var}(\text{Auk})$. For more details, please refer to [38].

Repacking. (Repack) Given n LWE ciphertexts or RLWE ciphertexts as $\text{RLWE}_{s,Q}^N(\mathfrak{m}_j) = (a_j, b_j)$, where $i \in [0, n-1]$ and $n|N$. We can pack these ciphertexts into an RLWE ciphertexts as shown in Alg. 6. Here, we consider two scenarios based on the number of ciphertexts n to be packed.

- Sparse packing case $n|N$. In this case, all ciphertexts are first integrated into an RLWE that encrypts $\sum_{j=0}^{n-1} m_{j,0} \cdot X^{j \cdot \frac{N}{n}} \cdot \sum_{k=0}^{\frac{N}{n}-1} *X^k$ through the repacking. Next, the Tr_{K_N/K_n} is used to eliminate redundancies Alg. 5. Therefore, the repacking process requires $\frac{n}{2} + \frac{n}{4} + \dots + 1 + \log(N/n) = n - 1 + \log(N/n)$ automorphisms, and the error variance can be summarized as $V_{\text{Pack}} = (\log_2 n + \log_2(N/n)) \cdot (V_{\text{Auto}} + 4V_{\text{MS}}) + \text{Var}(c_{\text{in}})$.

- Full packing case $n = N$. In this case, we apply the automorphism technique to combine them into an RLWE ciphertext, and this process requires $\frac{N}{2} + \frac{N}{4} + \dots + 1 = N - 1$ automorphisms. The error variance is bound by $V_{\text{Pack}} = \log_2 N \cdot (V_{\text{Auto}} + 4V_{\text{MS}}) + \text{Var}(c_{\text{in}})$ as shown in [38].

C Algorithm and Analysis for BatchBoot

C.1 EMPmul for negative number

Recall that the Alg. 2 presents the EMPmul procedure with the positive integer v . Here, the Alg. 7 shows the case for

Algorithm 6 RLWE ciphertext repacking (RLWE.Repack).

Input: n RLWE samples $c_j \in \text{RLWE}_{s,Q}^N(\mathfrak{m}_i)$, $j \in [0, n-1]$, where $n|N$.

The automorphism keys $\text{Auk}_i = \text{RLWE}'_{s,Q}(\tau_i(\mathfrak{s}))$, where $i \in \{2^{\log N - k + 1} + 1\}$ and $k \in [1, \log N]$.

Output: An RLWE ciphertext $ct' \in \text{RLWE}_s(\mathfrak{m})$.

1: for $i = 1$ to $\log n$ do
2: for $j = 0$ to $n/2^i - 1$ do
3: $\tilde{c}_j = c_j - X^{\frac{N}{2^i}} \cdot c_{j+n/2^i}$,
4: $\tilde{c}_j = \text{ModSwitch}_{Q \rightarrow Q/2}(c')$
5: $\tilde{c}_j = \text{ModRaise}_{Q/2 \rightarrow Q}(c')$
6: $c_j = \tilde{c}_j + \text{HomoAuto}_{\tau_i}(\tilde{c}_j, \text{Auk}) + X^{\frac{N}{2^i}} \cdot c_{j+n/2^i}$,
7: end for
8: end for
9: if $n < N$, **return** $ct' = \text{Eval.Tr}_{K_N/K_n}(c_0)$.
10: if $n = N$, **return** c_0 .

negative v , i.e. $v \in [-n, 0]$. For the negative, the encrypted message will be rotated to the left. This means the sub-loop order and output ciphertext indices must be reversed, while its noise growth remains the same as Alg. 2.

C.2 Correctness and noise growth

Proof of Lemma 4.1 Firstly, Alg. 2 performs $\ell/2$ iterations, each iteration involves four external products. Its correctness can be obtained directly from the multi-bit CMux gates and the parameterized external product. In addition, for the noise growth, Alg. 2 removes the noise from the automorphism, but introduces the noise growth of four external products. Since the noise variance of the external product is additive, after $\ell/2$ iterations, we can conclude that $V_{\text{EMP}} < \text{Var}(c_{\text{in}}) + 2\ell \cdot V_{\text{EP}}$.

Proof of Lemma 4.2. Alg. 3 first performs $h+k$ EMPmul operations (Alg. 2), with additive noise as stated in Lemma 4.2. The resulting ciphertexts are then repacked into a new RLWE ciphertext, with error variance analyzed in Sec. B.4. Finally, RLWE key switching adjusts the key to match the original secret key input, bounding the error variance by

$$\begin{aligned} V_{\text{Boot}} &< (h+k) \cdot V_{\text{EMP}} + V_{\text{Pack}} + V_{\text{RKS}} \\ &< 2(h+k)\ell \cdot V_{\text{CMux}} + \log_2 N \cdot (V_{\text{Auto}} + 4V_{\text{MS}}) + V_{\text{RKS}}. \end{aligned}$$

Lemma C.1. Input an RLWE ciphertext $ct \in \text{RLWE}_{s,q}^n(\mathfrak{m})$, the bootstrapping key Bsk , automorph key Auk , and RLWE key switching key Rvk . Alg. 4 outputs a n encryption for $\text{RGSW}_{s,Q}^N(X^{m_i})$ for $i \in [0, n-1]$, and its variance satisfies

$$V_{\text{CBoot}} \leq N(h+k)\ell \cdot V_{\text{CMux}} + \frac{N \cdot \log_2 \bar{d}}{2} \cdot V_{\text{Auto}} + V_{\text{RKS}}.$$

Proof of Lemma C.1. Alg. 4 first performs BatctBoot as shown in Alg. 3 without repacking, which outputs n RLWE

Algorithm 7 Efficient MPMul for negative number

Input:

A list $c_i \in \text{RLWE}(X^{u_i})$ for $0 \leq i \leq n$ representing a polynomial $u(X)$ encrypted in the exponent.

The bootstrapping key $\text{Bsk}_{i,j}$ for $0 \leq i < \ell/2 - 1$ and $0 \leq j \leq 3$ as shown in Tab 1, which is associated with an encrypted value $-v, (v > 0)$.

Output:

A list $\hat{c}_i \in \text{RLWE}(X^{w_i})$ for polynomial $w(X) = u(X) \cdot X^{-v} \bmod X^n + 1$, where $0 \leq i \leq n$.

```
1: for  $i = 0$  to  $\ell/2 - 1$  do
2:   for  $j = 0$  to  $n - 1$  do
3:      $\hat{c}_j = \text{FFT} \circ \text{Decom}(c_j)$ 
4:   end for
5:   for  $j = 0$  to  $n - 2^{2i} - 2^{2i+1} - 1$  do
6:      $\hat{c}_j = c_{j+2^{2i}} \boxtimes \text{Bsk}_{2i,0}^+ + c_{j+2^{2i+1}} \boxtimes \text{Bsk}_{2i,1}^+$ 
7:        $+ c_{j+2^{2i+2^{2i+1}}} \boxtimes \text{Bsk}_{2i,2}^+ + c_j \boxtimes \text{Bsk}_{2i,3}^+$ 
8:   end for
9:   for  $j = n - 2^{2i} - 2^{2i+1}$  to  $n - 2^{2i+1} - 1$  do
10:     $\hat{c}_j = c_{j+2^{2i}} \boxtimes \text{Bsk}_{2i,0}^+$ 
11:       $+ c_{j+2^{2i+1}} \boxtimes \text{Bsk}_{2i,1}^+$ 
12:       $+ c_{j+2^{2i+2^{2i+1}-n}} \boxtimes \tau_{-1} \text{Bsk}_{2i,2}^-$ 
13:       $+ c_j \boxtimes \text{Bsk}_{2i,3}^+$ 
14:   end for
15:   for  $j = n - 2^{2i}$  to  $n - 1$  do
16:     $\hat{c}_j = c_{j+2^{2i-n}} \boxtimes \tau_{-1} \text{Bsk}_{2i,0}^-$ 
17:       $+ c_{j+2^{2i+1-n}} \boxtimes \tau_{-1} \text{Bsk}_{2i,1}^-$ 
18:       $+ c_{j+2^{2i+2^{2i+1}-n}} \boxtimes \tau_{-1} \text{Bsk}_{2i,2}^-$ 
19:       $+ c_j \boxtimes \text{Bsk}_{2i,3}^+$ 
20:   end for
21:   for  $j = 0$  to  $n - 1$  do
22:      $\hat{c}_j = \text{iFFT}(c_j)$ 
23:   end for
24:    $(c_0, \dots, c_{n-1}) \leftarrow (\hat{c}_0, \dots, \hat{c}_{n-1})$ 
25: end for
26: return  $(c_0, \dots, c_{n-1})$ 
```

Algorithm 8 External Product Tree (EPT).

Input: β RGSW ciphertexts $C_k = \text{RGSW}(X^{\tilde{m}_k})$, for $k \in [0, \beta - 1]$.

$p^{\beta-1}$ polynomials testP_i that encode the LUT $f: [0, p_{\text{in}} - 1]^\beta \rightarrow [0, p_{\text{out}} - 1]$, where $t \in [0, p^\beta - 1]$, where $\text{testP}_i = -\sum_{j=0}^{N-1} \left\lfloor \frac{Q}{p_{\text{out}}} \right\rfloor \cdot f\left(t, \left\lfloor \frac{jp}{2N} \right\rfloor\right) \cdot X^{N-\frac{N}{p}-j-1}$

Output: An RLWE ciphertext, where its zeroth term corresponds to $f(m_0, \dots, m_{\beta-1})$.

```
1: for  $k = 0$  to  $\beta - 1$  do
2:   for  $i = 0$  to  $p^{\beta-k-1} - 1$  do
3:      $c_{k,i} = \text{testP}_i \boxtimes C_k$ 
4:   end for
5:   for  $i = 1$  to  $p^{\beta-k-2}$  do
6:      $\text{testP}_{i-1} = \text{RLWE.Repack}(c_{k,(i-1)p}, \dots, c_{k,ip-1}) \cdot \sum_{i=0}^{2N/p-1} -X^i$ 
7:   end for
8: end for
9: return  $\text{testP}_0$ .
```

ciphertexts. After that, it uses HomoTrace to extract the ciphertext associated with B_i . This step can generate the gadget RLWE ciphertext. Finally, the scheme switching is used to embed the secret key. The noise in these steps is additive, so we have

$$\begin{aligned} V_{\text{CBoot}} &< \frac{N}{2} ((h+k) \cdot V_{\text{EMp}} + V_{\text{HT}}) + V_{\text{RKS}} \\ &< N(h+k)\ell \cdot V_{\text{CMux}} + \frac{N \cdot \log_2 \bar{d}}{2} \cdot V_{\text{Auto}} + V_{\text{RKS}}. \end{aligned}$$

Lemma C.2. Input β RGSW ciphertext $C_k \in \text{RGSW}(X^{\tilde{m}_i})$ with the noise variance V_{in} , and the automorphism key A_{uk} , Alg. 8 outputs the ciphertext of $f(m_0, \dots, m_{\beta-1})$, where the variance satisfies

$$V_{\text{EPT}} \leq \beta \cdot V_{\text{EP}} \cdot V_{\text{in}} + (\beta - 1) \cdot V_{\text{Pack}}.$$

Proof of Lemma C.2. The correctness of Alg. 8 can be directly obtained from the correctness of the external product and functional bootstrapping. It is worth noting that this algorithm introduces a multiplicative noise growth. By carrying the variance of the CBS (lemma C.1) into the variance of the initial external product, we can obtain the final error variance batched circuit bootstrapping as follows $V_{\text{out}} < \beta \cdot V_{\text{EP}} \cdot (N(h+k)\ell \cdot V_{\text{CMux}} + \frac{N \cdot \log_2 \bar{d}}{2} \cdot V_{\text{Auto}} + V_{\text{RKS}}) + (\beta - 1) \cdot V_{\text{Pack}}$. For a more detailed noise analysis, please refer to [40].

C.3 Security and standard deviations

In this subsection, we show the standard deviations of the initial noise corresponding to various bootstrapping keys and ciphertexts in Tab. 9. All parameter settings satisfy the 128-bit security level through the lattice estimator.

Table 9: Security and standard deviations for FHE parameters using the lattice estimator.

Sets	h	N	Q	Std	λ
Boot ₂	39	2048	2^{64}	2^{-15}	> 128
Boot ₂	512	2048	2^{64}	2^{-53}	> 128
Boot ₄	42	2048	2^{64}	2^{-17}	> 128
Boot ₄	512	2048	2^{64}	2^{-53}	> 128
Boot ₆	33	4096	2^{64}	2^{-21}	> 128
Boot ₆	512	2048	2^{64}	2^{-53}	> 128
Boot ₈	34	4096	2^{64}	2^{-24}	> 128
Boot ₈	512	4096	2^{64}	2^{-56}	> 128

C.4 Failure Probability

Given a ciphertext with Gaussian-distributed noise e with standard deviation Std, we can compute the probability of decryption failure. Notice that we are leaving one bit of padding for the message to avoid the negacyclicity as done in [33], i.e., $m_i \in \mathbb{Z}_p$ and $\Delta = q/2^{p+1}$. Therefore, the message $m_i \in \mathbb{Z}_p$ is correctly bootstrapped with probability $\Pr[e < q/2^{p+2}] = \text{erf}\left(\frac{q/2^{p+2}}{\sqrt{2} \cdot \text{Std}_M}\right)$, where Std_M denotes the maximum standard deviation of noise throughout the evaluation process. For example, in Alg. 3, the maximum standard deviation satisfies $\text{Std}_M = (\text{V}_{\text{Boot}} + \text{V}_{\text{MS}})^{\frac{1}{2}}$. Finally, we present more detailed results for decryption failure rates in Tab. 4.

D Parameter Settings and Performance Results

D.1 FHE Parameter Settings

we present the parameters of BATCHBOOT in Tab. 10 and BATCHCBOOT in Tab. 11, respectively.

Table 10: Parameters of BATCHBOOT of Alg. 3 where the set PSI.Boot₆ is used in PSI.

Para. Sets	p	n	h	N	B_{Bsk}	d_{Bsk}	B_{Auk}	d_{Auk}	B_{Rlk}	d_{Rlk}
Boot ₂	2^2	2^{11}	39	2^{11}	2^{23}	1	2^{23}	1	2	12
Boot ₄	2^4	2^{11}	42	2^{11}	2^{23}	1	2^{23}	1	2	14
Boot ₆	2^6	2^{12}	33	2^{12}	2^{23}	1	2^{23}	1	2	17
Boot ₈	2^8	2^{12}	34	2^{13}	2^{23}	1	2^{23}	1	2	17
PSI.Boot ₆	2^6	2^{11}	39	2^{11}	2^{15}	2	2^{15}	2	2	12

Table 11: Parameters of BATCHCBOOT of Alg. 4.

Sets	p	n	h	N	B_{Bsk}	d_{Bsk}	B_{Auk}	d_{Auk}	B_{Rvk}	d_{Rvk}	\bar{B}	\bar{d}
CBoot ₄	2^4	2^{11}	39	2^{11}	2^{15}	2	2^{15}	2	2^{15}	2	2^4	4

D.2 Performance Results

This section shows more detailed implementation results. We first present the performance results for multithreaded execution in Tab. 12. Our experimental findings indicate that our algorithm directly benefits from an increased number of threads. Finally, we present the peak memory usage in Tab. 13.

Table 12: Multi-threaded execution result for Batchboot, where T represents the number of threads.

Params.	Packed Mess.	$T = 1$	$T = 2$	$T = 4$	$T = 8$	$T = 16$
Boot ₈	4096	102.8 s	55.3 s	29.2 s	19 s	14.7 s

Table 13: Peak memory usage for the evaluation of LUTs in the single-threaded setting, where the parameters use the fully packing parameters corresponding to Tab. 4.

	2-bit LUT	4-bit LUT	6-bit LUT	8-bit LUT
Peak RAM	360 MB	360 MB	1425 MB	3743 MB