

# APACHE: A Processing-Near-Memory Architecture for Multi-Scheme Fully Homomorphic Encryption

Lin Ding<sup>1</sup>, Song Bian<sup>2</sup>, Penggao He<sup>1</sup>, Yan Xu<sup>1</sup>, Gang Qu<sup>3</sup>, and Jiliang Zhang<sup>1</sup>

<sup>1</sup>Hunan University

<sup>2</sup>Beihang University

<sup>3</sup>University of Maryland, College Park

**Abstract**—Fully Homomorphic Encryption (FHE) allows one to outsource computation over encrypted data to untrusted servers without worrying about data breaching. Since FHE is known to be extremely computationally-intensive, application-specific accelerators emerged as a powerful solution to narrow the performance gap. Nonetheless, due to the increasing complexities in FHE schemes per se and multi-scheme FHE algorithm designs in end-to-end privacy-preserving tasks, existing FHE accelerators often face the challenges of low hardware utilization rates and insufficient memory bandwidth. In this work, we present APACHE, a layered near-memory computing hierarchy tailored for multi-scheme FHE acceleration. By closely inspecting the data flow across different FHE schemes, we propose a layered near-memory computing architecture with fine-grained functional unit design to significantly enhance the utilization rates of both computational resources and memory bandwidth. In addition, we propose a multi-scheme operator compiler to efficiently schedule high-level FHE computations across lower-level functional units. In the experiment, we evaluate APACHE on various FHE applications, such as Lola MNIST, HELR, fully-packed bootstrapping, and fully homomorphic processors. The results illustrate that APACHE outperforms the state-of-the-art ASIC FHE accelerators by 2.4× to 19.8× over a variety of operator and application benchmarks.

## I. INTRODUCTION

With the increasing popularity of cloud-based storage and computation outsourcing, fully homomorphic encryption (FHE) [47] rises as one of the most promising solutions to secure data confidentiality while maintaining its usability. Equipped with recent advances in fundamental cryptographic designs, modern FHE allows an extremely broad spectrum of complicated algorithms to be directly applied to ciphertexts, without the service provider ever knowing the decryption key.

To further enhance its usability, increasing attention focuses on designing protocols that adopt multi-scheme FHE constructions. In such case, multiple FHE schemes, e.g., TFHE [16] and CKKS [14], are used to (jointly or separately) establish privacy-preserving computation protocols, such as federated learning [18], [27], [43], [78], fully homomorphic processors [10], [11], [48], private medicine [5], [24], [40], and private data inquiry [15], [29], [44]. Here, schemes such as CKKS and BFV [21] are often used to evaluate polynomial functions (e.g., additions, multiplications), whereas TFHE-like schemes [16], [20] are employed to evaluate non-polynomial functions over ciphertexts.

Corresponding author: Jiliang Zhang

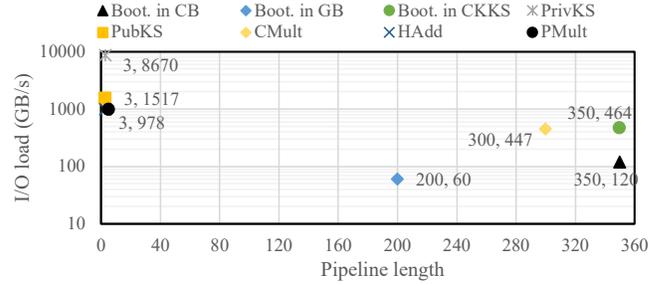


Fig. 1. Evaluation of I/O load in the pipelined accelerator, where the TFHE parameters are set according to [7], [16]. The bandwidth demand of CKKS operators is cited from [77].

While protocols based on multi-scheme FHE can be much more expressive and usable, we face new computational and communicational obstacles against the associated architectural designs. Here, we first categorize basic FHE algorithms into data-heavy operators and computation-heavy operators, as illustrated in Figure 1. Here, we see that multi-scheme-FHE-based protocols need to include a much wider range of fundamental FHE operators to support complex computations across different privacy-preserving tasks. However, as also sketched in Figure 1, due to varying ciphertexts parameters and evaluation key sizes (e.g., bootstrapping keys, key-switching keys, rotation keys), FHE operators of different schemes exhibit distinct data-heavy and computation-heavy characteristics, necessitating new explorations in the architectural design space of FHE accelerators. In addition, we point out that separating accelerator designs for BFV/CKKS-like and TFHE-like schemes can suffer from low hardware utilization rates. As Figure 2 illustrates, in privacy-preserving database applications [7], the runtime of TFHE-like scheme dominates the overall protocol latency. On the other hand, protocols like privacy-preserving neural network inference can be implemented solely based on BFV/CKKS-like schemes, without the need of instantiating any TFHE operator. Consequently, the effective acceleration of multi-scheme FHE protocols demands further studies on a unified architecture for the acceleration of both computation-heavy and data-heavy operators across different FHE schemes.

Based on the above analyses, we derive three essential design principles to enable the multi-scheme applicability of FHE accelerators: 1) layered memory access hierarchy for the

TABLE I  
QUALITATIVE COMPARISONS BETWEEN HOMOMORPHIC ACCELERATORS.

	[61]	[55]	[62]	[37]	[38]	[77]	[44]	[32]	[6]	[9]	[76]	[71]	[4]	[33]	[70]	Ours
TFHE-like <sup>a</sup>	△	✓	✗	✗	✗	△	✗	✓	✓	✗	△	✗	✗	✗	✗	✓
I/O bus load	High	Low	High	High	High	High	Low	High	Low	High	Low	High	High	High	High	Low
Configurability	✗	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✓	✓	✓	✓
Parallelism	▲	▲	▲	▲	▲	✗	▲	▲	✗	✗	✗	✗	▲	▲	▲	✓

▲ denotes that this accelerator has paralleled multiple cores but does not support the parallelism at the level of accelerators; <sup>a</sup> [6], [32], [55] only support the gate bootstrapping function of TFHE, while the proposed APACHE is capable of implementing both of gate bootstrapping and circuit bootstrapping functions. △ marks such works that claim to support at least one logic FHE scheme but without any evidence and description for the implementation.

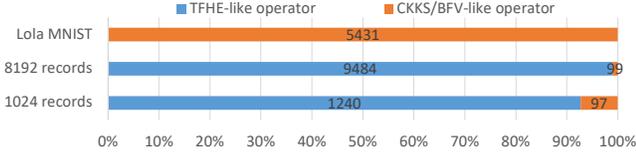


Fig. 2. Breakdown of “TPC-H Query 6” in the homomorphic database [7] and encrypted Lola-MNIST [8] based on CKKS. We set the record number in a database as 8192 and 1024. The metrics are “execution time in seconds per query” and “execution time in microseconds” for “TPC-H Query 6” and Lola MNIST, respectively.

minimization of I/O load, 2) reconfigurable functional unit (FU) design for the support of flexible operator granularity, and 3) effective operator scheduling for the maximization of data parallelism. Using such critiques, we evaluated the qualitative characteristics of existing FHE accelerators in Table I. and draw two main insights. First, prior accelerators [2], [33], [37], [38], [43], [44], [61], [62], [76], [77] mostly depend on a two-level memory hierarchy, namely, the off-chip and of-chip memory levels with low and high communication bandwidths, respectively. For instance, existing BFV/CKKS accelerators, such as [37], [38], [61], [62], utilize High-Bandwidth Memory (HBM) to buffer key and ciphertexts data, where HBM can provide a maximum of roughly 2 TB/s of I/O throughput. However, as Figure 1 highlights, a throughput of 8 TB/s or above is required to fully exhaust the computation resources of a fully pipelined circuit bootstrapping unit for TFHE. It is obvious that, while simply mounting additional memory bandwidth can solve the I/O bottleneck, due to the constraint on the area and the computational resources, the expensive bandwidth added are mostly wasted when executing other FHE operators. Second, the interconnect topology of existing architectures are generally fixed for a given FHE scheme without operator-level task scheduling, resulting in less optimal resource utilization rates. For example, in [36], [38], [61], [62], a BConv unit is explicitly instantiated to accelerate the base conversion algorithm in the BFV/CKKS bootstrapping procedure. Yet, since TFHE-like FHE schemes generally do not use the residual number system (RNS), the BConv unit is completely unused when accelerating TFHE, leaving as much as 30% of chip area unoccupied during task execution. Hence, a natural question to ask is that, how can we co-design the memory and computation architectures to support multi-scheme FHE operators, where both the memory bandwidth and the computing resources can be fully exploited?

**Our Contribution:** In this work, we propose APACHE, a processing-near-memory architecture designed for general-

purpose FHE acceleration. First, we leverage a three-level memory hierarchy to fully exploit the internal bandwidth of the DIMM units [35], [72], [75]. By lifting the large-volume communication from external I/O loads, we are able to avoid mounting an excessive amount of expensive HBM units. Moreover, we propose a new circuit topology to interconnect functional units, such that a high hardware utilization rate can be retained across the diverse set of FHE operators. The key contributions of our work can be summarized as follows.

- **Multi-Level PNM Architecture:** Our key observation is that, as illustrated in Figure 1, data-heavy operators generally have low computational circuit depth. Based on such insight, we allocate different functional units into a total of three memory levels: the external I/O level, the near-memory level, and the in-memory level. By placing computation-light but memory-heavy units closer (or even into) the memory die, we are able to reduce the data external I/O bandwidth by up to  $3.15 \times 10^5$  times.
- **Configurable Internconnect Topology:** We propose a configurable interconnect topology with fine-grained functional unit designs to simultaneously ensure multi-scheme operator support and high hardware utilization rates. we show that the utilization rates of the (I)NTT FUs can be kept at least 90% across multi-scheme FHE tasks, while that of existing accelerators only range from 50% to 85%.
- **Operator- and Task-level Scheduling:** To fully exploit the proposed hardware topology, we design a scheduler that extracts control and data flow of multi-scheme FHE operators to determine the proper resource allocations and the associated datapath configurations. Moreover, we also exploit task-level operator parallelism to further take advantage of the massive parallel processing capability of multi-channel dual in-line memory modules (DIMMs).
- **Implementation and Evaluation:** We thoroughly evaluate the performance of APACHE on a set of multi-scheme FHE operator and application benchmarks. The results show that APACHE achieves  $2.4\times$  to  $19.8\times$  speedup, compared to the state-of-the-art FHE accelerators.

The remaining part of this paper is organized as follows. In Section II, we first provide basics on multi-scheme FHE ciphertext types and operator. Then, we overview our APACHE architecture in Section III. Next, we show how to design and configure the underlying pipelined circuits and memory layouts in Section IV. Based on the lower-level functional units, we detail the scheduling strategy in Section V. Subsequently, we provide the detailed experiment setup and results

in Section VI, and present the related works in Section VII. Finally, we conclude our work in Section VIII.

## II. PRELIMINARIES

Similar to existing works [7], we divide existing FHE schemes into two primary categories: BFV/CKKS-like and TFHE-like. In what follows, we provide a brief summary on the fundamental ciphertext types and key homomorphic operators for multi-scheme FHE.

### A. Notations

For basic notations, we use  $\mathbb{B}$ ,  $\mathbb{R}$ , and  $\mathbb{Z}$  to denote the set of  $\{0, 1\}$ , real numbers, and integers, respectively. Let  $q/Q/Q'$  be different sizes of ciphertext modulus,  $p$  is for the plaintext modulus, and  $\mathbb{Z}_q/\mathbb{Z}_Q/\mathbb{Z}_{Q'}$  refer to the set of integers modulo  $q/Q/Q'$ . We use  $R$  and  $R_Q$  to represent  $\mathbb{Z}[X]/(X^N + 1)$  and  $\mathbb{Z}[X]/(X^N + 1) \bmod Q$ , respectively, for some ciphertext modulus  $Q$  and polynomial degree  $N$ . Throughout this paper, we use bold lowercase letters (e.g.,  $\mathbf{a}$ ) for vectors, tilde lowercase letters (e.g.,  $\tilde{a}$ ) for polynomials, and bold uppercase letters (e.g.,  $\mathbf{A}$ ) for matrices.

### B. Multi-Scheme FHE Ciphertext

Most existing lattice-based FHE construct ciphertexts based on either (or both) the learning with error (LWE) and ring learning with error (RLWE) hardness problems [46]. Based on their algebraic properties, we summarize three fundamental types of ciphertext formats that cover most FHE applications, namely, the LWE ciphertext, the RLWE ciphertext, and the RGSW ciphertext.

- **LWE Ciphertext:** Here, we use a secret key  $\mathbf{s} \in \mathbb{B}^n$  to encrypt a single integer message  $m \in \mathbb{Z}_p$ , and get a LWE ciphertext as follows,

$$\text{LWE}_{\mathbf{s}}^{n,q}(m) = (b, \mathbf{a}) = (\langle -\mathbf{a}, \mathbf{s} \rangle + \Delta \cdot m + e, \mathbf{a}), \quad (1)$$

where  $\mathbf{a}$  is chosen uniformly at random,  $e$  is a noise sampled from discretized Gaussian, and  $\Delta$  is a scale factor to protect the least significant bits of the message  $m$  from the noises.

- **RLWE Ciphertext:** A RLWE ciphertext can be defined by the following equation,

$$\text{RLWE}_{\tilde{\mathbf{s}}}^{N,Q}(\tilde{m}) = (\tilde{b}, \tilde{\mathbf{a}}) = (\langle -\tilde{\mathbf{a}}, \tilde{\mathbf{s}} \rangle + \Delta \cdot \tilde{m} + \tilde{e}, \tilde{\mathbf{a}}). \quad (2)$$

Here, a vector of messages  $\tilde{m} \in R_p$  is encrypted under the secret key  $\tilde{\mathbf{s}} \in \mathbb{B}[X]/(X^N + 1)$ . Similar to the LWE ciphertext,  $\tilde{\mathbf{a}} \in R_Q$  is also uniformly and randomly sampled from discretized Gaussian.  $\Delta$  is a scaling factor.

- **RGSW Ciphertext:** A RGSW ciphertext is a matrix consisting of  $2d$  RLWE ciphertexts and is noted as  $\text{RGSW}_{\tilde{\mathbf{s}}}^{N',Q'}(m)$ .

### C. Computational Optimizations for FHE

To boost the computation of FHE schemes, the current FHE implementations have adopted two main technique: number theoretic transform (NTT) and residue number system (RNS).

**NTT and Polynomial Multiplication:** In FHE, the most expensive routine is the polynomial multiplication. In most

cases, we adopt NTT to reduce the complexity of polynomial multiplication to  $\mathcal{O}(N \log N)$ , by computing  $\tilde{a} \cdot \tilde{b} = \text{INTT}(\text{NTT}(\tilde{a}) \circ \text{NTT}(\tilde{b}))$ , where  $\circ$  is coefficient-wise multiplication.

**RNS and Base Conversion:** In CKKS/BFV-like schemes, RNS is used to decompose polynomials with the large moduli  $Q$  which can be hundreds to thousands of bits to a set of sub-polynomials that are of smaller moduli  $\{q_i\}$ . CKKS/BFV-like FHE schemes often need to raise or reduce the moduli, known as the Modup and Moddown operators. To carry out Modup and Moddown, a large number of RNS base conversion (BConv) operations have to be invoked. Equation (3) shows the process of generating a new subpolynomial  $[\tilde{a}]_{p_j}$  from  $\tilde{a}$  of moduli  $Q$ .

$$\text{BConv}([\tilde{a}]_Q, p_j) : [\tilde{a}]_{p_j} = \left( \sum_{i=0}^{L-1} [[\tilde{a}]_{q_i} \cdot \hat{q}_i^{-1}]_{q_i} \cdot \hat{q}_i \right) \bmod p_j, \quad (3)$$

where  $L$  is the number of modulus  $q_i$  of  $Q$ . Based on BConv, Modup and Moddown can be defined by Equation (4) and Equation (5), respectively.

$$\text{Modup}([\tilde{a}]_Q, P \cdot Q) : [\tilde{a}]_{P \cdot Q} = \text{BConv}([\tilde{a}]_Q, p_j), \quad (4)$$

$$\text{Moddown}([\tilde{a}]_{P \cdot Q}, Q) :$$

$$[\tilde{a}]_{q_j} = ([\tilde{a}]_{q_j} - \text{BConv}([\tilde{a}]_{P \cdot Q}, q_j)) \cdot P^{-1} \bmod q_j, \quad (5)$$

where  $j \in [0, M)$  and  $M$  is the number of sub-modulus of  $P$ .

### D. Multi-Scheme FHE Operators

In this subsection, we provide a detailed explanation of the BFV/CKKS-like and TFHE-like homomorphic operators. Meanwhile, we decompose such homomorphic operators into lower-level operators, including (inverse) number theoretic transform ((I)NTT), modular addition (MAdd), modular multiplication (MMult), and coefficient automorphism. Considering the computation overhead and data size, we further classify such homomorphic operators into data-heavy and computation-heavy operators, as Table II illustrates.

1) *Homomorphic BFV/CKKS-like Operators:* The homomorphic arithmetic evaluation is conducted on the RLWE( $\tilde{m}$ ) ciphertext over  $\mathbb{R}_Q$ . The basic homomorphic arithmetic operators are homomorphic addition (HAdd) and homomorphic multiplication (HMult).

- **Homomorphic Addition (HAdd):** HAdd of two ciphertexts only requires MAdd between the ciphertext polynomials:  $\text{RLWE}(\tilde{m}_0 + \tilde{m}_1) = (\tilde{b}_0 + \tilde{b}_1, \tilde{a}_0 + \tilde{a}_1) \bmod Q$ .

- **Key switching (KeySwitch):** We use KeySwitch to transform a cipher polynomial  $\tilde{a}$  encrypted by  $\tilde{\mathbf{s}}$  to that of  $\tilde{\mathbf{s}}'$ . KeySwitch includes Modup and Moddown, requiring a large number of MMult, MAdd, NTT, and INTT operations.

- **Homomorphic Multiplication (HMult):** HMult supports the plaintext-ciphertext multiplication (PMult) and ciphertext-ciphertext multiplication (CMult). For PMult calculations, we have  $\text{RLWE}(\tilde{m}_0 \cdot \tilde{m}_1) = (\tilde{b}_0 \cdot \tilde{b}_1, \tilde{a}_0 \cdot \tilde{a}_1) \bmod Q$ . For CMult, we compute  $\text{RLWE}(\tilde{m}_0 \cdot \tilde{m}_1) = \text{KeySwitch}(\tilde{a}_0 \cdot \tilde{a}_1, \mathbf{evk}) + (\tilde{b}_0 \cdot \tilde{b}_1, \tilde{a}_0 \cdot \tilde{b}_1 + \tilde{a}_1 \cdot \tilde{b}_0) \bmod Q$ .

TABLE II  
DECOMPOSITION AND CLASSIFICATION OF HOMOMORPHIC OPERATORS.

Homomorphic Operators		Basic Functional Units				Pipeline Depth <sup>1</sup>	Cached Key Size	Operand Bitwidth <sup>2</sup>	Input Sym. <sup>3</sup>	Operator Type
		NTT	MA	MM	Auto.					
TFHE-like [41], [48]	CMUX	✓	✓	✓	✓	≤ 350	None	32, 64	✓	Computation
	PrivKS		✓	✓		≤ 3	1.8 GB	64	✗	Data
	PubKS		✓	✓		≤ 3	79 MB	32	✗	Data
	Gate Boot. <sup>4</sup>	✓	✓	✓	✓	≤ 350	37 MB	32	✓	Computation
	Circuit Boot. <sup>4</sup>	✓	✓	✓	✓	≤ 350	196 MB	32, 64	✓	Computation
BFV- and CKKS-like [37]	HAdd		✓			≤ 3	None	≤ 32	✓	Data
	HMult	✓	✓	✓		≤ 300	120 MB	≤ 32	✓	Computation
	CKKS Boot.	✓	✓	✓	✓	≤ 350	≈ 1 GB	≤ 32	✓	Both

<sup>1</sup>Commonly, a full-pipelined NTT unit is comprised of 150 to 250 stages of circuits. The pipeline lengths of MA and MM are generally less than three and five, respectively. As designed in [61], a 128-lane automorphism module needs at least 63 stages of circuits. The overall circuit depth is estimated based on the above parameters. <sup>2</sup>Operands of 128 bits or above are typically split into integers of 32 bits or less for efficient calculation. <sup>3</sup>Input Symmetry means that all the input operands have the same bitwidth. <sup>4</sup>Excluding the built-in key switching procedures.

- **Homomorphic Rotate (HRot)**: HRot performs a circular left shift by  $r$  slots on a ciphertext RLWE =  $(\tilde{b}, \tilde{a})$ , by executing  $\text{HRot}(\text{RLWE}, r, \mathbf{evk}^{(r)}) = \text{KeySwit}(\psi_r(\tilde{a}), \mathbf{evk}_{rot}^{(r)}) + (\psi_r(\tilde{b}), \tilde{0})$ , where  $\mathbf{evk}^{(r)}$  is a KeySwit key for HRot and  $\psi_r(\cdot)$  represents the automorphism, i.e., the  $i^{\text{th}}$  coefficient of a polynomial moves to the slot of  $i \cdot (5^r) \bmod N$ . Low-level operators of HRot includes  $\psi_r(\cdot)$  besides these of KeySwit.

- **BFV/CKKS Bootstrapping**: The above HMult operations only support the limited number of homomorphic multiplications. For the deeper multiplicative depth, the bootstrapping is proposed to recover the level budgets of homomorphic multiplications. bootstrapping is the most complex and expensive operation in both BFV and CKKS. Due to the space limit, we refer readers to [37], [61] for further details on the exact arithmetic characteristics of BFV and CKKS bootstrapping.

2) *Homomorphic TFHE-like Operators*: In TFHE, the most fundamental homomorphic operators are the controlled multiplexer (CMUX) and key switching. Moreover, based on both operators, gate bootstrapping (GB) and circuit bootstrapping (CB) have been built to enable various homomorphic functionalities in constructing diverse applications.

- **Controlled Multiplexer (CMUX)**: This operator is the homomorphic counterpart of a multiplexer. It outputs a new RLWE ciphertext which shares the same plaintext as either input RLWE ciphertexts for  $\text{ct}_0$  or  $\text{ct}_1$ , depending on the given RGSW ciphertext of  $\mathbf{C}$  encrypting a selection bit. The CMUX function can be represented as  $\text{CMUX}(\text{ct}_0, \text{ct}_1, \mathbf{C}) = \mathbf{C} \boxtimes (\text{ct}_0 - \text{ct}_1) + \text{ct}_1$ , where  $\boxtimes$  presents an external product.

- **Key Switching (KS)**: To switch keys and ciphertext formats, TFHE has defined two key-switching flavours, namely the public functional key switching (PubKS) and private functional key switching (PrivKS). And the main computation producers can be characterized by the two following equations,

$$\text{PubKS}(f, \mathbf{KS}, \mathbf{c}) = (0, f(b^1, \dots, b^p)) - \sum_{i=1}^n \sum_{j=1}^t \hat{a}_{i,j} \cdot \mathbf{KS}_{i,j}, \quad (6)$$

$$\text{PrivKS}(\mathbf{KS}^{(f)}, \mathbf{c}) = - \sum_{z=1}^p \sum_{i=1}^{n+1} \sum_{j=1}^t \hat{c}_{i,j}^{(z)} \cdot \mathbf{KS}_{z,i,j}^{(f)}, \quad (7)$$

where  $\mathbf{c}$  is the set of  $p$  LWE ciphertexts, and  $f$  is some Lipschitz continuous function.  $\hat{a}_{i,j}$  is the  $j^{\text{th}}$  bit of  $a_i$ , in

which  $a_i = f(a_i^{(1)}, \dots, a_i^{(p)})$ , while  $\hat{c}_{i,j}^{(z)}$  is the  $j^{\text{th}}$  coefficient in the  $z^{\text{th}}$  LWE ciphertext.  $\mathbf{KS}$  and  $\mathbf{KS}^{(f)}$  are the key-switching keys, each of which consists of either LWE ciphertexts or RLWE ciphertexts.

- **Bootstrapping of TFHE**: The operations inherently introduce a certain degree of noise into the resulting ciphertext, thereby increasing the risk of decryption failure. To guarantee the correct decryption, bootstrapping is designed to remove the noise at the end of each operation. We can combine bootstrapping and PubKS to construct various homomorphic logic gates (HomGate), such as HomAND, HomOR, and HomXOR. Similarly, by jointly using bootstrapping and PrivKS, a circuit bootstrapping (CB) functionality can be achieved, which provides a higher precision and enhanced flexibility for constructing complex and advanced applications, as exemplified by such [7], [48].

### III. MOTIVATION AND ARCHITECTURAL OVERVIEW

#### A. Challenges and Motivations

Before delving into the architectural details, we first provide a thorough analyses on the computation and memory characteristics for existing multi-scheme FHE operators, and discuss the primary challenges as well as key insights for APACHE.

As summarized in Table II, we focus on the most frequently used operators for both BFV/CKKS-like and TFHE-like schemes, and classify the operators into two types, namely, computation-heavy and data-heavy. Here, we point out two main challenges for accelerating multi-scheme FHE. First, the amount of cache data vary by as much as  $48\times$  between the operators. As a result, existing accelerators with the typical on-chip and off-chip memory architecture has to either cache all the data on chip (which can be impractical for large evaluation keys with up to 1.8 GB of data) or incur excessive off-chip I/O bandwidth that are left unused for computation-heavy operators. Second, while most high-level operators can be broken down to the a set of basic functional units, (i.e., NTT, modular addition and multiplication, and automorphism), the exact operator composition and control path do vary from one to the other. However, since existing FHE accelerators generally target on only one FHE scheme, their circuit topologies are fixed at design time (e.g., for [6], [32], [36], [37], [61], FUs are

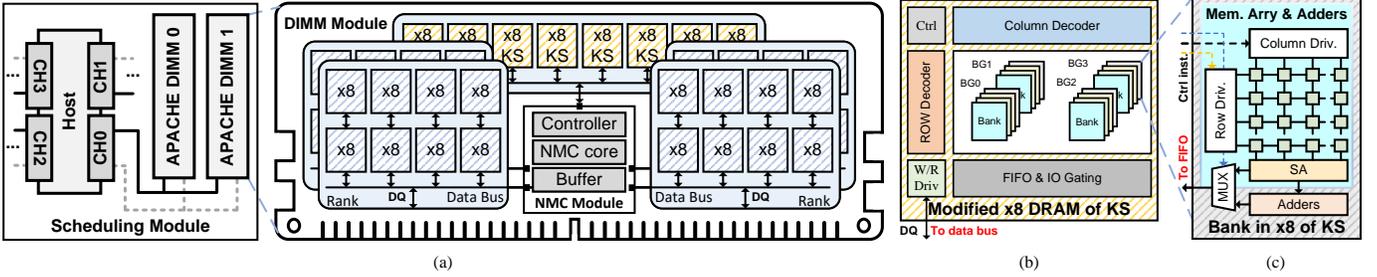


Fig. 3. (a) Overview of APACHE structure; (b) Modified  $\times 8$  DRAM chip with a hierarchy of memory array-to-bank-to-bank group; and (c) Modified memory array for computing KS. Here, the dashed and solid lines represent the control flow and data flow, respectively.

in the pre-determined order of NTT, modular multiplication, modular addition, and automorphism). Consequently, some of the FUs are under-utilized, or completely unused, when executing certain operators (e.g., PrivKS and PubKS does not use the NTT units).

Motivated by the above challenges, our main goal in this work is thus to decouple and decompose data-heavy and computation-heavy FHE operators based using a multi-level memory hierarchy with configurable FU interconnects, so that we can achieve high hardware utilization rates while significantly reducing I/O bandwidth.

### B. Architectural Overview

As illustrated in Figure 3(a), APACHE is comprised of two main modules: the scheduling module and the DIMM module. When a particular FHE program is received by a host CPU, the control and data flow graphs of the FHE operators are first extracted by the scheduler module. The scheduler then breaks down the data dependency between the operators and construct the task queue, where operator-level parallelism is extracted to exploit the computational capabilities of the multi-channel DRAM module. Next, the operators are fed into the DRAM module, where the majority of the data buffering and task execution happen. Lastly, the execution results for each of the operators in the task are naturally stored inside the DRAM module for the execution of the subsequent operators. Hence, the overall goal of APACHE is to significantly reduce the off-chip I/O loads for multi-scheme FHE acceleration by properly allocating functional units into the DRAM module.

To carry out highly complex cryptographic operations inside the DRAM module, computing and storage elements are carefully organized to form three fundamental memory levels: the external I/O level, the near-memory computing level, and the in-memory computing level. In what follows, we briefly summarize the memory and computing topology in each of the layers.

① **I/O Level:** First of all, on the I/O level, we expose the set of operators that constitute the application programming interface (API) of the overall accelerator, including all the FHE operators described in Section II-D. When executing an FHE operator, we assume that all ciphertexts data and evaluation keys are pre-loaded into the DIMMs, which we believe is a reasonable assumption since the same thing holds true for the regular computation model.

② **Near-Memory Computing Level:** The NMC level is where most of the low-level FUs locate. We have three main components here: i) the standard memory chips, ii) the NMC module, and iii) the modified in-memory computing module. First, for the native memory ranks, we employ the standard  $\times 8$  DRAM chips, where the internal architecture of these DRAMs remain unchanged. The only difference here is that the data bus of each rank is connected separately to one of the data buffers inside the NMC module.

By parallelizing the data bus of multiple DRAM ranks, APACHE provides large internal bandwidth to the NMC module, such that an excessive amount of I/O communication can be avoided. Inside the NMC module, we incorporate configurable FUs implementing the respective underlying operators, and the FUs are structured with a flexible interconnect topology to accommodate various data flows of FHE operators efficiently.

③ **In-Memory Computing Level:** As sketched in Figure 3(b) and 3(c), we place application-specific circuits on the in-memory computing level to mainly reduce the bandwidth bottleneck introduced by the TFHE key switching algorithms, namely, the PrivKS and PubKS operators. As illustrated in Table II, the sizes of keys needed to be loaded on-chip is extremely large for PrivKS and PubKS, while the computation circuit is extremely shallow (only a couple of adders). Hence, we insert the accumulation adders at the bank level of the  $\times 8$  DRAM chip, where the evaluation keys are pre-loaded.

## IV. DESIGN EXPLORATION OF FUNCTIONAL UNITS

To design a near-computing module for multi-scheme FHE, we see two significant design challenges. First, as illustrated in Figure 4(a) and (b), TFHE-like and BFV/CKKS-like schemes exhibit distinct data and control flows. Consequently, regardless of which scheme we use as the reference for the NMC design, there will always be inefficiencies when the other scheme is used. Second, within the BFV/CKKS-like schemes, many tasks only adopt the HAdd and PMult operators without invoking NTT. As a result, if a the NTT, MMult, and MAdd circuits pipelined together, the NTT modules become idle when only HAdd and PMult are called. To tackle the above challenges, we propose a flexible NMC module design with configurabilities on both the interconnect level and the internal FU level.

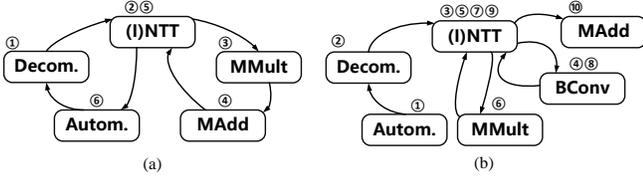


Fig. 4. Dataflow of (a) CMUX of TFHE and (b) HRot of CKKS. Here, BConv consists of MMult and MAdd. It is stressed that steps 2 to 9 can also represent the dataflow of CMult with KeySwitch.

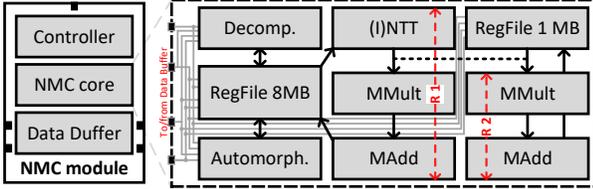


Fig. 5. The topology of NMC module. Dashed line stands for wires with transistors to control whether to link (I)NTT FU with MMult FU.

### A. The NMC Module and its Interconnection

Here, we first describe the instantiated FUs in the NMC module, and then discuss how the FUs need to be interconnected to form the overall NMC architecture. As outlined in Figure 5, the NMC module is composed of three subcomponents, namely, the interconnect controller, the NMC core, and the data buffer. During the execution of an FHE operator, the associated data are first loaded into the data buffer. Then, the data path in the NMC core is properly configured by the controller to execute the exact computations. Within the NMC core, we have six FUs: the register file, the decomposition unit Decomp, the (I)NTT unit (I)NTT, the modular multiplication unit MMult, the modular addition MAdd, and the automorphism unit Automorph. The (I)NTT module can implement the NTT and INTT operators by feeding different twiddle factors, respectively.

While the designs of the FUs themselves are close to those of the existing works [37], [55], [61], [62], the interconnects need to be completely re-designed due to the complex arithmetic characteristics of the multi-scheme FHE operators. First, we observe that the (I)NTT-MMult-MAdd routine is one of the most frequently used control flows for both BFV/CKKS and TFHE, as depicted in Figure 4 (a) and (b). Furthermore, it is observed that automorphism and decomposition operators are comparatively less utilized in both schemes, thus requiring lower bandwidth. Based on this observation, we propose a configurable interconnect topology for our NMC modules, as depicted in Figure 5. This topology incorporates two separate pipelines: the (I)NTT-MMult-MAdd routine and the MMult-MAdd routine. The former routine, coupled with a central register file of 8MB, is capable of executing operators such as CMUX of HomGate for TFHE-like schemes and CMult of BFV/CKKS-like schemes. Conversely, the latter routine can efficiently compute HAdd and PMult without interrupting the (I)NTT pipeline. Notably, the central register file ensures adequate internal bandwidth for the (I)NTT-MMult-MAdd routine, while the register file of 1MB supplies operands

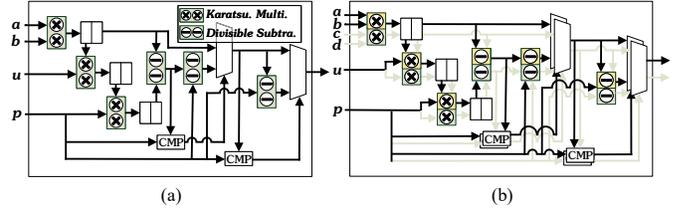


Fig. 6. The proposed configurable modular multiplier, (a) working as a 64-bit modular multiplier, and (b) working as two parallel 32-bit modular multiplier. Here, we employ the Karatsuba multiplier and divisible subtractor to provide the configurable ability.

to the MMult-MAdd routine. Concretely, we can formulate the utilization rate of the (I)NTT FU for a single fixed pipeline routine and our configurable interconnect topology by Equation (8) and Equation (9), respectively.

$$Utl_{NTT} = \frac{\mathbf{T}_{ALL} - \mathbf{T}_{nonNTT}}{\mathbf{T}_{ALL}}, \quad (8)$$

$$Utl'_{NTT} = \frac{R1.\mathbf{T}_{ALL} - R1.\mathbf{T}_{nonNTT}}{R1.\mathbf{T}_{ALL} \cup R2.\mathbf{T}_{ALL}}, \quad (9)$$

where  $\mathbf{T}_{ALL}$  stands for the overall latency of a routine, and  $R1.\mathbf{T}_{ALL} \cup R2.\mathbf{T}_{ALL}$  is the union of timing segments of the two pipeline routines R1 and R2. The union runtime is no bigger than  $\mathbf{T}_{ALL}$ . Meanwhile, it holds  $R1.\mathbf{T}_{nonNTT} \geq 0.5 \times \mathbf{T}_{nonNTT}$  (resp.  $R1.\mathbf{T}_{nonNTT} = \mathbf{T}_{nonNTT}$ ) for HAdd/PMult (resp. CMUX). Thus, we can observe a significant improvement in the utilization rate of the (I)NTT FU for both CKKS/BFV-like and TFHE-like schemes for the proposed interconnect topology.

### B. FU Configuration in the NMC Module

As illustrated in Table II, a general-purpose FHE accelerator must accommodate operands of 32 and 64 bits. However, existing FHE accelerators predominantly confine modular multiplication, addition, and (I)NTT circuits to fixed bitwidths. For example, BTS [38], ARK [37], and Strix [55] uses a fixed bit width of 64 for the (I)NTT/FFT designs. Such limitations inevitably result in notable resource under-utilization when handling 32-bit operands. To accommodate the pivot bit-widths, we have designed and integrated the FUs with configurable bitwidth within our APACHE.

**(1) Configurable MMult and MAdd circuits (Figure 6):** MMult [64], [67] and MAdd [31], [49] play the critical roles in the construction of various low-level operators and homomorphic operators, such as (I)NTT, HMult, CMUX, and high-level operators such as bootstrapping. Therefore, all modular multiplication and addition units in the proposed architecture have configurable operand bit width that can be flexibly switched between the 64-bit operand mode and 32-bit operand mode, where one execution of the multiplier (or adder) produces one 64-bit output or two 32-bit outputs.

The structure of the proposed configurable MMult is depicted in Figure 6. First, we point out that a 64-bit adder can be easily divided into two 32-bit adders by cutting the pipeline at the carry bit. Second, our key observation is that the Karatsuba multiplier processes a  $k$ -bit multiplication using three  $\frac{k}{2}$ -bit

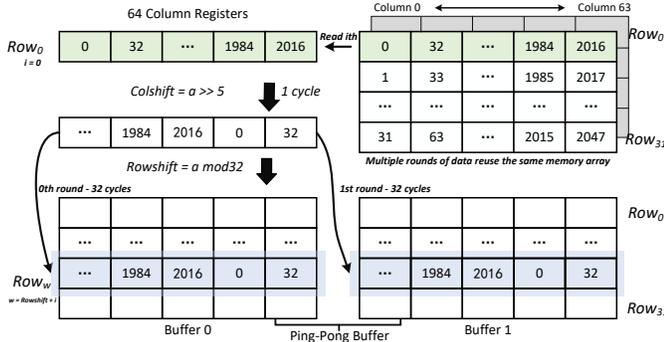


Fig. 7. The structure of automorphism module.

multipliers. Consequently, by modifying the structure of a 64-bit Karatsuba multiplier, it becomes feasible to compute either a 64-bit multiplication or at least two parallel 32-bit multiplications. More concretely, we can use three  $k$ -bit multipliers and three  $k$ -bit subtractors to construct a  $k$ -bit MMult [77]. Hence, we arrange our configurable 64-bit adders and 64-bit multiplier to construct the 64-bit configurable MMult as shown in Figure 6. Lastly, we can simply use the 64-bit configurable adder and subtractor to implement a configurable MAdd FU.

**(2) Configurable NTT FU:** The computation of (I)NTT [26], [53], [73], [79] is akin to the Discrete Fourier Transform (DFT) [69], [74], but operates over a integer polynomial ring. In existing FPGA and ASIC accelerators of FHE, the operand bitwidths of the customized (I)NTT circuits are specific to 28 bits, 32 bits, 64 bits, or others. Nevertheless, due to the substantial hardware overhead associated with implementing (I)NTT, it is impractical to directly incorporate multiple (I)NTT FUs for all possible operand bitwidths. Hence, the most viable solution is to devise the configurable (I)NTT module. In this work, we utilize our configurable MMult and MAdd circuits to construct the (I)NTT FUs. By such an approach, we can ensure that one configurable (I)NTT circuit can be flexibly configured into either one high bit-width (up to 64 bits) (I)NTT FU or two parallel low bit-width (each supporting up to 32 bits) (I)NTT FUs.

**(3) Automorphism (Figure 7):** In addition to supporting multiple operand bitwidths, the automorphism unit needs to address a new design challenge posed by the distinct coefficient rotation rules in the CKKS and TFHE encryption schemes. Specifically, in CKKS, the coefficient shift is determined by  $\tau = i \cdot k \bmod N$ . [61] has illustrated that implementing automorphism for BGV and CKKS can be achieved effectively by utilizing four permute/transpose operations within SRAM buffers. In contrast, in TFHE, we have a fixed automorphism of  $\tau = i + k \bmod 2N$  for computing the operation of  $X^{-a_i} \cdot \text{ACC}$ . It is observed that automorphism for TFHE can be implemented using only a few shift registers and a single SRAM buffer. Thus, we can see that there is an inevitable disparity between the automorphism operators employed in CKKS and TFHE.

## V. DATA AND TASK SCHEDULING

As mentioned, for the general acceleration of multi-scheme FHE, the design of an efficient and effective scheduling

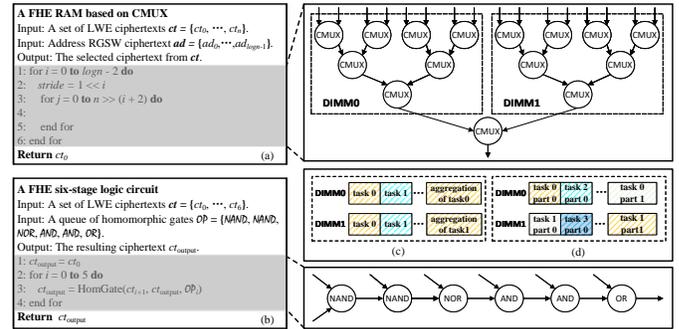


Fig. 8. Two demos for basic cases of homomorphic evaluation tasks: (a) without specific data dependency and (b) with specific data dependency. Task-level schedule schemes: (c) for the case without specific data dependency and (d) for the case with specific data dependency. Here, arrows stand for the data flow, and circles represent the homomorphic operators.

strategy is essential to facilitate data reuse and minimize unnecessary data movement [6], [61]. Moreover, in order to exploit the massive parallelism of multiple DIMM cards, the mapping and aggregation intermediate results has to be carefully treated to avoid excessive inter-DIMM communication.

To address the operator and data scheduling issue, we define scheme-specific scheduling strategies at both the FHE application level (i.e., task-level) and multi-scheme operator level. Based on the scheduled data and control sequences, we generate the micro-instructions to configure the respective FUs and the overall execution process.

### A. Task-Level Scheduling

For task-level scheduling, consider two basic cases of homomorphic evaluation tasks: those without a specific ciphertext sequence and those with. In both cases, the plaintext content remains undisclosed to the service provider. It is important to emphasize that more complex tasks can be constructed by the combinations of the above two cases.

When there is no explicit data dependency, the ciphertext can be processed concurrently in multiple APACHE DIMMs. Figure 8 (a) shows a demo involving two APACHE DIMMs in the absence of specific ciphertext dependency. In this case, the same homomorphic operators (i.e., tree-shape CMUX operators in the example of Figure 8 (a)) would be executed within each DIMM, independently. Then, only the local result is transmitted to another DIMM to obtain the aggregation outcome. In another case, the homomorphic task exhibits a specific order, as illustrated in Figure 8 (b). If the amount of ciphertext operands does not exceed the storage capacity of a single DIMM, we would process the task in a single APACHE DIMM. Conversely, if the storage capacity is exceeded, the local result of DIMM 0 is transmitted to DIMM 1 as the input to launch the subsequent evaluation.

By implementing the above scheduling strategy, primary data movements occur within each DIMM. Only small local results are communicated across APACHE DIMMs via the host bus. Although our scheme significantly reduces the IO load compared to AISC, FPGA, and GPU, the data transmission still hinders the underlying circuits to maintain the full pipeline work for a single task. This is because branch

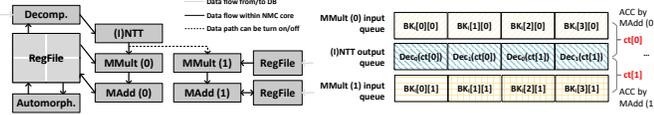


Fig. 9. Data flow of TFHE Bootstrap. in NMC module.

prediction [12], [30] and out-of-order execution mechanisms [45], [63] of the host CPU render that the local results are delivered at a random time point and in order. An important insight is that it is unnecessary to complete one end-to-end task before commencing the next one. Instead, servers commonly receive massive tasks, and thus we can perform other tasks while awaiting the delivery of local results, as Figure 8 (c) and (d) shown. Namely, using the data scheduling strategy and the task breakdown, we can ensure that our APACHE DIMMs provide the capability of parallelism and keep computing fully pipelined.

### B. Operator-level Scheduling

Once the task sequence is determined, the next step is to organize the homomorphic operators in a systematic manner to achieve high FU utilization rates and data reuse. As mentioned in Section V-A, at this stage, if the homomorphic operators that share the same evaluation key (e.g.,  $\mathbf{evk}$  of CKKS or PrivKS key of TFHE), such operators are clustered to be executed together. Based on the operator clustering results, we determine the operator batching size and the exact order of operator execution. Finally, the scheduler generates the sequences of the micro-instructions along with the associated data to the DIMM module. In the follows, we detail the operator-level scheduling for BFV/CKKS-like and TFHE-like schemes in a fine-grained manner.

**Scheduling CKKS/BFV-like Operators:** For BFV/CKKS-like schemes, we observe that directly executing operators according to the task-level schedule can lead to low resource utilization rates. For example, as Figure 4(b) shows, CMult and HRot can be decomposed into a sequence of Modup ((1)NTT③-MAdd④-(1)NTT⑤-MMult⑥) and Moddown ((1)NTT⑦-BConv⑧-(1)NTT⑨). If we map the (1)NTT③ and MAdd④ operators to the (1)NTT-MMult-MAdd routine for the Modup operation, the MMult and MAdd FUs experience pipeline bubbles after the (1)NTT execution in ⑤. Similarly, pipeline bubbles occur for (1)NTT and MMult/MAdd during the execution of KeySwitsh and Moddown. To mitigate this issue, the scheduler rearranges and divides the underlying operators into three groups: ((1)NTT③-MAdd④), ((1)NTT⑤-MMult⑥), and ((1)NTT⑦-BConv⑧). This approach can significantly reduce the occurrence of pipeline bubbles, generating only one instance of MAdd bubble when executing (1)NTT⑤-MMult⑥. Moreover, to increase the reuse of operand data (such as the Modup base for (1)NTT③-MAdd④ and the evaluation key for MMult⑥), we batch operators on the group level, instead of the operator level.

**Scheduling for TFHE-like Operators:** In the TFHE scheme, the bootstrapping key needs to be reused for gate and circuit bootstrapping. To reduce the memory bandwidth

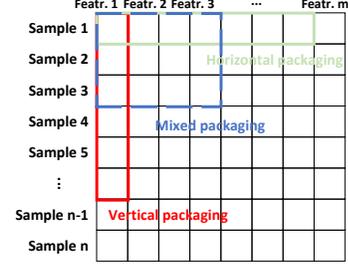


Fig. 10. Three packing methods of RLWE ciphertext.

demand, we adopt the design in [6], where a batch of ciphertexts are processed by the  $i^{th}$  CMUX in parallel before iterating to the  $i+1^{th}$  CMUX, to ensure the fully reuse of  $\mathbf{BK}_i$ . For a single CMUX, the data path is configured as illustrated in Figure 9 to keep the utilization rates of the (1)NTT FU as high as possible. Concretely, a CMUX needs to decompose the input polynomial into  $d$  sub-polynomials ( $d = 2$  in Figure 9). The decomposed polynomials fed into the (1)NTT FU followed by a multiplication of  $\mathbf{BK}_i$  in the subsequent MMult FU, where  $\mathbf{BK}_i$  is the  $i$ -th decomposed TFHE bootstrapping key. Since one decomposed polynomial needs is multiplied by two different shares of  $\mathbf{BK}_i$ , the two separate MMult-MAdd pipeline routines shown in Figure 9 are simultaneously activated, where the MAdd completes the computation of a CMUX by accumulating the  $d$  partial products. When the last accumulation of the batch of  $i^{th}$  CMUX is completed, we then call the (1)NTT FU again to perform the INTT and generate the final result. Finally, we note that our scheduling scheme is compatible with a wide range values of  $d$  with a high resource utilization. Whereas, in many existing TFHE accelerators, such as MATCHA [32], FPT [6], and Strix [55], the NTT and INTT FUs are configured specifically to  $d : 1$ , which leads to a under-utilization when  $d$  varies.

### C. Data Packaging and Parallelism Extraction

As discussed in Section II, all FHE schemes operate over three different types of ciphertexts: LWE, RLWE, and RGSW. When Equation (10) holds, we pack  $t$  LWE ciphertexts into an RLWE for an efficient transfer across DIMMs,

$$\mathbf{T}_{\text{Pack}(t\text{-LWE})} + \text{RLWE} \cdot \mathbf{T}_{\text{Transfer}} \leq t \cdot \text{LWE} \cdot \mathbf{T}_{\text{Transfer}}. \quad (10)$$

Here, we take the RLWE ciphertext as an example to illustrate how we manage data across DIMMs. As Figure 10 shows, there are three data packing methods of RLWE ciphertext: vertical packing, horizontal packing, and mixed packing.

**Vertical Packing:** Vertical packing is commonly used in applications that analyze and compare input samples within a given dimension (i.e., number of samples). For this class of applications, we ensure that the RLWE ciphertexts of the same dimension are stored within the same DIMM as much as possible. In this way, the computation of different dimensions of input samples can be parallelized in multiple DIMMs. Subsequently, the computed results of all dimensions are passed onto a DIMM for the final data aggregation.

**Horizontal Packing:** This method packs all the features of the same sample in an RLWE ciphertext. However, in

TABLE III  
APACHE DIMM CONFIGURATION.

Memory capacity	64 GB	DRAM chip	8Gb x8 3200 MT/s
Rank per DIMM	8	Clock frequency	1600 MHz
NMC per DIMM	1	tRCD-tCAS-tRP	22-22-22

many applications, the number of plaintext samples rarely reaches the lattice dimension, i.e.,  $2^{12}$  to  $2^{16}$ . Hence, an RLWE ciphertext can be made to contain multiple plaintext samples. The data delivery overhead of this type of packing depends on requirements of the application. For example, when computing the iteration of the K-means clustering [60], we need to transfer the RLWE ciphertexts of  $K$  centers and  $K$  distance sums, where each sum is the distance between a sample and a center. However, if the Euclidean distance is required for all sample pairs, the communication overhead approaches  $\mathcal{O}(\#sample^2)$ .

**Mixed Packing:** Mixed packing divides a multi-dimensional data matrix into multiple sub-matrices and encrypts one or more sub-matrices into an RLWE ciphertext. Similar to vertical packing, our APACHE system stores the RLWE ciphertexts of the same features in a DIMM to achieve high parallelism efficiency.

## VI. IMPLEMENTATION AND EVALUATION

### A. Implementation and Setup

In this section, we present the configurations used in the APACHE DIMM. The simulation framework is designed with the following components: 1) **Simulation of FHE functionality:** A behavioral level simulator is developed based on the TFHE and CKKS schemes. This simulator not only performs FHE operations but also generates the physical memory access trace required for further analysis. 2) **Trace-based Memory Simulator:** To simulate the cache and DRAM behaviors in each DIMM, we employ NVsim [19], CACTI [50], and trace-based simulators Ramulator [39], all of which are the widely-used tools for memory analysis. In this work, we set the cache and memory fabricated under the 22 nm technology node. 3) **NMC Module Analysis:** For the NMC module, we estimate its latency, area, and power characteristics by employing Synopsys Design Compiler with 22 nm technology library at a frequency of 1 GHz. By combining the above four components along with a workload distribution by our compiler, we are able to estimate the end-to-end performance of APACHE. The APACHE DIMM configuration and simulation of hardware overhead are illustrated in Table IV.

### B. Baselines and Benchmarks

For comparison purposes, we evaluate APACHE with the state-of-the-art FPGA [77], as well as six ASIC FHE accelerators [32], [37], [38], [55], [61], [62]. Meanwhile, we have done a evaluation for Circuit bootstrapping on the Strix [55] and Morphling [54] with each max bandwidth available for the PrivKS key. To facilitate comprehensive evaluation and comparison, we incorporate various CKKS and TFHE

TABLE IV  
AREA AND THERMAL DESIGN POWER (TDP) OF NMC MODULER OF A APACHE DIMM, WITH BREAKDOWN BY COMPONENT.

Component	Area [mm <sup>2</sup> ]	Power [W]
64-point (I)NTT $\times$ 4	13.04	6.28
Automorphism $\times$ 2	2.4	0.6
Decomposition $\times$ 2	0.03	0.02
Modular Multiplier $\times$ 256 $\times$ 2	5.0	3.01
Modular Adder $\times$ 256 $\times$ 2	0.36	0.39
Adders in each $\times$ 8 DRAM	0.12	0.02
Regfile (8 + 1 MB)	14.4	1.01
Data Buffer (16 MB)	25.6	1.8
Total NMC module	60.95	13.14

operators, along with five distinct FHE applications. All assessments adhere to a 128-bit security assumption if no additional declaration is provided. Specifically, in the evaluations of homomorphic operators, TFHE parameters conform to [7], [16], while those of CKKS remain consistent with [36].

1) *Mini-benchmarks:* To assess the performance of APACHE on individual FHE operators across different schemes, we adopt CMult, HAdd, PMult, KeySwit, and HRot, HomGate, and circuit bootstrapping as the set of mini-benchmarks.

2) *CKKS benchmarks:* In this experiment, we employ three CKKS-based schemes for evaluation: Logistic Regression (LR), Lola MNIST [8], and fully-packed bootstrapping. First, for LR, we use HELR [27] where a 196-element weight vectors is trained by 32 iterations. The reported metric is the average execution time per iteration. Second, we execute the Lola MNIST [8], a homomorphic Neural Network based on CKKS, with the same parameters as in [62]. Lastly, we benchmark the fully-packed CKKS bootstrapping [1], [13].

3) *TFHE and mixed benchmarks:* For TFHE benchmarks, we use the homomorphic circuits in [48] to benchmark the performance of end-to-end TFHE applications. In particular, [48] incorporates a five-stage pipeline processor with a various of logic gates such as HomNAND and HomXOR, along with 512 bytes of ROM and RAM that are built out of CMUX’s. HE<sup>3</sup>DB [7] is the state-of-the-art homomorphic database based on both CKKS and TFHE schemes, supporting the logic and arithmetical data querying and analysis. We performed an evaluation of our APACHE on “TPC-H Query 6” (i.e., a classical database workload benchmark) of HE<sup>3</sup>DB.

### C. Performance and Comparison

As shown in Table V and Figure 11, we present the performance of APACHE architecture for in terms of FHE operators and applications. Results are showcased for parallel acceleration using 2/4/8 DIMMs and are compared against multiple FHE accelerators, including Poseidon [77], F1 [61], CraterLake [62], BTS [38], ARK [37], SHARP [36], Strix [55], Morphling [54], and MATCHA [32].

**Performance of FHE Operators:** First of all, Table V shows the remarkable fact that only our APACHE can support CKKS and TFHE, while all the comparison work can only perform one of both. Moreover, our APACHE architecture

TABLE V

THROUGHPUT COMPARISON OF CKKS OPERATORS IS WITH  $N = 10^{16}$  AND  $L = 44$ . THE HOMGATE-I, HOMGATE-II, CIRCUIT BOOTSTRAPPING ARE SET TO 80-BIT, 110-BIT, AND 128-BIT SECURITY PARAMETERS, RESPECTIVELY. WE ADOPT “OPERATORS PER SECOND” (OP./S) AS THE METRIC.

	PMult	HAdd	CMult	Rotation	Keyswit.	HomGate-I	HomGate-II	CircuitBoot.
Poseidon [77]	14.6K	13.3K	273	302	312	-	-	-
F1 [61]	N/A	N/A	N/A	N/A	N/A	-	-	-
MATCHA [32]	-	-	-	-	-	10K	-	-
Strix [55]	-	-	-	-	-	74.7K	39.6K	$\leq 2.6K$
Morphling [54]	-	-	-	-	-	147K	78.7K	$\leq 7.4K$
APACHE $\times 2$	355K	355K	6.5K	6.8K	7.4K	500K	264K	49.6K
APACHE $\times 4$	708K	708K	13.1K	13.6K	14.8K	1,000K	528k	99.2K

The performance of CKKS operators with  $N = 2^{16}$  and  $L = 44$  is not provided by state-of-the-art FHE accelerators, such as F1 [61], CraterLake [62], BTS [38], and SHARP [36]. We use N/A to mark an accelerator that supports a specific FHE operator but without reported performance.

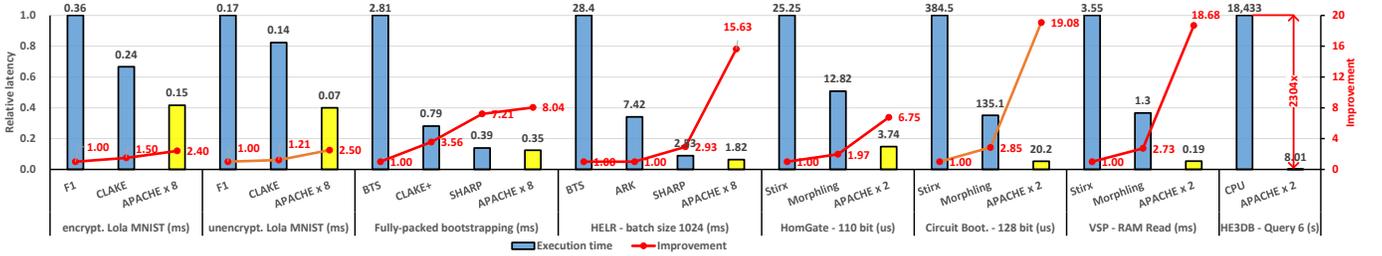


Fig. 11. Performance comparison of APACHE with state-of-the-art FHE accelerators. To ensure similar area overhead with the comparison accelerators, we report the performance of APACHE  $\times 2$  and  $\times 8$  for TFHE and CKKS benchmarks, respectively. The number of records is set to  $2^{14}$  for HE<sup>3</sup>DB [7].

achieves significant advantages over comparable solutions, when executing TFHE operators. Even with only two parallel APACHE DIMMs, we achieve remarkable performance, surpassing the recent work Strix (resp. Morphling) by  $6.69\times$  (resp.  $3.4\times$ ) and  $6.75\times$  (resp.  $3.38\times$ ) for 80-bit and 110-bit security, respectively. When implementing the circuit bootstrapping of 128-bit security, our APACHE attains the  $19.08\times$  and  $6.7\times$  speedup than Strix and Morphling, respectively.

The significant edge emanates from various design considerations. Primarily, our operator-level scheduling ensures the high utilization of (I)NTT FUs, which is significantly higher than existing works, where the NTT/FFT units often remain idle. Additionally, while the parts of accelerators employ fixed 64-bit FFT/NTT (e.g., 64-bit FFT adopted in Strix), our APACHE allows each 64-bit NTT FU to be configured into dual 32-bit NTT FUs, thereby enhancing resource utilization for 32-bit HomGate. Moreover, we decouple PrivKS and PubKS from computation clusters and process both directly within specific  $\times 8$  ranks, reducing bandwidth load by  $3.15 \times 10^5$  times and  $3.05 \times 10^4$  times, respectively. In contrast, Strix (resp. Morphling) takes around 24 ms (resp. 7.7 ms) for loading PrivKS key of 1.8 GB when batching 64 LWE ciphertexts for circuit bootstrapping.

**Full-system Performance:** As Figure 11 shows, our APACHE with 8 DIMMs achieves a speedup of  $2.4\times$  and  $2.5\times$  for Lola MNIST with encrypted and unencrypted weights, respectively. Moreover, our APACHE DIMM  $\times 8$  is  $8.04\times$  and  $15.63\times$  faster than the baseline of BTS for the CKKS schemes of fully-packed bootstrapping and HELR, respectively. APACHE has achieved such improvement mainly due to the flexible interconnect topology and fine-grained operator scheduling. Concretely, such real applications of CKKS commonly include a large number of PMult and HAdd

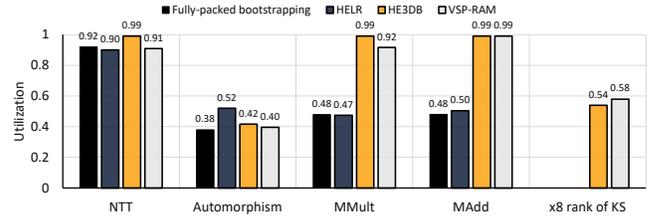


Fig. 12. Resource utilization of APACHE.

independent of NTT. Our APACHE allows both operators to be (partially or totally) executed in the pipeline routine 2. Thus, our APACHE significantly improves the utilization of NTT FUs and the overall performance. When APACHE carries out the TFHE-based VSP, we observe  $18.68\times$  and  $6.8\times$  speedup against Strix and Morphling, respectively. The VSP relies on expensive circuit bootstrapping to generate GSW-format addresses. Due to the optimization mismatch in circuit bootstrapping over Strix and Morphling, our APACHE exhibits a more competitive performance in VSP. Furthermore, Figure 11 demonstrates that our APACHE is the only work that supports HE<sup>3</sup>DB [7] based on both TFHE and CKKS schemes, suppressing CPU by  $2,304\times$ .

#### D. Architectural Analysis

**Utilization Rates of FUs:** As presented in Figure 12, our APACHE demonstrates high utilization for hardware resource. The utilization of the (I)NTT FU always keeps above 90%, while that of existing accelerators only range from 50% to 85%. The high utilization rate indicates that the proposed interconnect topology and operator scheduling strategy can effectively reduce the running time of NTT-independent operations. Moreover, the high utilization rates of (I)NTT also

means that our the NMC module is supplied with sufficient bandwidth, so that the proposed framework can almost fully exploit the hardware resources of the (I)NTT FU. In addition, we observe that the utilization of the in-memory computing KS module is around 50%. The effective utilization of the KS module indicates that our multi-level storage architecture effectively reduces the bandwidth caused by loading the key switching key.

**Remark on Data Communication across DIMMs:** In addition to the circuit utilization and local bandwidth load, another impact determinant of APACHE performance lies in the latency associated with data scheduling across DIMMs. It is pointed out that our APACHE searches the aggregation point to minimize the data volume that needs to be addressed. Consequently, the time required for data propagation across DIMMs is substantially smaller than that needed by local computations over each DIMM. For example, a APACHE DIMM need 0.38 ms to read a LWE ciphertext from the VSP RAM, in which 512 LWE ciphertexts are stored. By contrast, forwarding the readout LWE cipher to another DIMM via the host only requires 0.31 us with given 30GB/s of I/O bandwidth. That is, our APACHE can cover the communication latency within computation duration.

## VII. RELATED WORKS

**Near/in-memory Computing for FHE Acceleration:** Recent studies have proposed to use near/in-memory computing (N/IMC) to optimize FHE by reducing data transfers [25], [26], [44], [57], [65], [76]. For instance, ReRAM-based logic gates were used in [26] to create extensive pipeline architectures for polynomial operations and the full FHEW scheme [20], [25]. Alternatively, [57] and [65] deployed logic circuits adjacent to SRAM sense amplifiers for the schooltext polynomial and NTT computations, respectively. However, these N/IMC approaches have no discussed dividing and conquering FHE operators as data-heavy and computation-heavy traits, leading to lower performance versus ASICs or restricted compatibility with multi-scheme FHE tasks.

**GPU-based FHE Acceleration:** Research involving GPUs has explored accelerating FHE [3], [4], [22], [23], [33], [34], with notable efforts like using NVIDIA Tesla V100 in [33] achieving a 257 $\times$  speedup over CPUs. Other initiatives include integrating cryptographic functions into GPU operations [66] and developing open-source libraries [17], [52], [68] for agile FHE development. However, as the general-purpose platform, GPUs lack the specialized hardware of FHE operators, thereby leading to inferior performance compared to FPGA and ASIC.

**ASIC-based FHE Acceleration:** F1 [61] is the first programmable FHE ASIC designed for CKKS/BFV-like schemes, surpassing CPU by over 17,000 $\times$ . Building upon F1, Crater-Lake [62] enables unbounded multiplicative depth by architecture optimization for bootstrapping. Further advancements include BTS [38] and ARK [37], which enhance performance through optimal parameter selection and key reuse, respectively. More recently, SHARP [36] built a 36-bit FHE accelerator, attaining high performance with a smaller area and power consumption. Meanwhile, with much effort paid

into TFHE acceleration, we also see that the performance of HomGate has improved from 10,000 OP/S [32] to 74,000 OP/S [55] to the current 147,615 OP/S [54]. However, all the above ASIC chips are efficient just on either CKKS/BFV-like schemes or TFHE-like ones.

**FPGA-based FHE Acceleration:** Previous studies have explored FPGA-based acceleration of FHE for specific HE operations [42], [43] or leveled HE schemes [56], [58], [59]. Recently, literature has introduced FPGA-based approaches for FHE [2], [6], [28], [36], [51], [77]. For instance, FAB [2] leverages a well-designed datapath to optimize key-switching and effectively utilize its 43 MB on-chip storage. Poseidon [77] focuses on fine-grained decomposition of FHE operators, achieving efficient resource utilization. Moreover, FPT [6] stands out as the fastest FPGA accelerator for HomGate, attaining a 946 $\times$  speedup over CPU. However, FPGAs generally under-perform when compared to the latest ASIC accelerators of FHE, due to limited on-chip resources.

## VIII. CONCLUSION

In this section, we present a near-memory computing architecture for the acceleration of both BFV/CKKS-like and TFHE-like FHE schemes. By classifying the key operators of both two FHE lanes into data-heavy and computation-heavy operators, we point out that, rather than simply piling up computing resources, allocating the FUs on the correct memory level is the key in effectively accelerating multi-scheme FHE. Based on such insight, we propose APACHE that allocates fully pipelined circuits on both near-memory and in-memory levels with configurable FU topology to improve the overall hardware utilization rates. Through rigorous experiments, we show that APACHE can reduce the external I/O bandwidth by as much as  $3.15 \times 10^5$ . Finally, we evaluate APACHE with multiple FHE applications and prove that it behaves better than state-of-the-art ASIC accelerators.

## REFERENCES

- [1] “Lattigo v5,” Online: <https://github.com/tuneinsight/lattigo>, Nov. 2023, ePFL-LDS, Tune Insight SA.
- [2] R. Agrawal, L. de Castro, G. Yang, C. Juvekar, R. T. Yazicigil, A. P. Chandrakasan, V. Vaikuntanathan, and A. Joshi, “FAB: an fpga-based accelerator for bootstrappable fully homomorphic encryption,” in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2023, pp. 882–895.
- [3] S. Akleylek, Ö. Dagdelen, and Z. Y. Tok, “On the efficiency of polynomial multiplication for lattice-based cryptography on gpus using CUDA,” in *Cryptography and Information Security in the Balkans - Second International Conference*, vol. 9540, 2015, pp. 155–168.
- [4] A. A. Badawi, B. Veeravalli, C. F. Mun, and K. M. M. Aung, “High-performance FV somewhat homomorphic encryption on gpus: An implementation using CUDA,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, no. 2, pp. 70–95, 2018.
- [5] R. Banno, K. Matsuoka, N. Matsumoto, S. Bian, M. Waga, and K. Sue-naga, “Oblivious online monitoring for safety LTL specification via fully homomorphic encryption,” in *34th International Conference on Computer Aided Verification (CAV)*, vol. 13371, 2022, pp. 447–468.
- [6] M. V. Beirendonck, J. D’Anvers, and I. Verbauwhede, “FPT: a fixed-point accelerator for torus fully homomorphic encryption,” *IACR Cryptol. ePrint Arch.*, p. 1635, 2022.
- [7] S. Bian, Z. Zhang, H. Pan, R. Mao, Z. Zhao, YierJin, and ZhenyuGuan., “He<sup>3</sup>db: An efficient and elastic encrypted database via arithmeticand-logic fully homomorphic encryption,” in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2023.

- [8] A. Brutzkus, O. Elisha, and R. Gilad-Bachrach, "Low latency privacy preserving inference," in *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- [9] X. Cao, C. Moore, M. O'Neill, E. O'Sullivan, and N. Hanley, "Optimised multiplication architectures for accelerating fully homomorphic encryption," *IEEE Transactions on Computers*, vol. 65, 2016.
- [10] A. Chatterjee and K. M. M. Aung, *FURISC: FHE Encrypted URISC Design*, 2019, pp. 87–115.
- [11] A. Chatterjee and I. Sengupta, "Furisc: Fhe encrypted urisc design," Cryptology ePrint Archive, Paper 2015/699, 2015. [Online]. Available: <https://eprint.iacr.org/2015/699>
- [12] C. Chen, C. Shen, and J. Zhang, "Lightweight and secure branch predictors against spectre attacks," in *27th Asia and South Pacific Design Automation Conference (ASPDAC)*, 2022, pp. 25–30.
- [13] J. H. Cheon, K. Han, and M. Hhan, "Faster homomorphic discrete fourier transforms and improved FHE bootstrapping," *IACR Cryptol. ePrint Arch.*, 2018.
- [14] J. H. Cheon, A. Kim, M. Kim, and Y. S. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *23rd International Conference on the Theory and Applications of Cryptology and Information Security*, vol. 10624, 2017, pp. 409–437.
- [15] J. H. Cheon, M. Kim, and M. Kim, "Optimized search-and-compute circuits and their application to query evaluation on encrypted data," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 1, pp. 188–199, 2016.
- [16] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, "Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds," in *22nd International Conference on the Theory and Application of Cryptology and Information Security*, vol. 10031, 2016, pp. 3–33.
- [17] W. Dai and B. Sunar, "cube: A homomorphic encryption accelerator library," in *Cryptography and Information Security in the Balkans - Second International Conference*, vol. 9540, 2015, pp. 169–186.
- [18] R. Dathathri, O. Saarikivi, H. Chen, K. Laine, K. E. Lauter, S. Maleki, M. Musuvathi, and T. Mytkowicz, "CHET: an optimizing compiler for fully-homomorphic neural-network inferencing," in *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2019, pp. 142–156.
- [19] X. Dong, C. Xu, Y. Xie, and N. P. Jouppi, "Nvsim: A circuit-level performance, energy, and area model for emerging nonvolatile memory," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, no. 7, pp. 994–1007, 2012.
- [20] L. Ducas and D. Micciancio, "FHEW: bootstrapping homomorphic encryption in less than a second," in *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2015.
- [21] J. Fan and F. Vercauteren, "Somewhat practical fully homomorphic encryption," *IACR Cryptol. ePrint Arch.*, p. 144, 2012. [Online]. Available: <http://eprint.iacr.org/2012/144>
- [22] S. Fan, Z. Wang, W. Xu, R. Hou, D. Meng, and M. Zhang, "Tensorfhe: Achieving practical computation on encrypted data using gpgpu," in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2023, pp. 922–934.
- [23] S. Fan, Z. Wang, W. Xu, R. Hou, D. Meng, and M. Zhang, "Tensorfhe: Achieving practical computation on encrypted data using gpgpu," in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2023, pp. 922–934.
- [24] D. Froelicher, J. R. Troncoso-Pastoriza, J. L. Raisaro, M. A. Cuendet, and J. P. Hubaux, "Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption," *Nature Communications*, vol. 12, no. 1, 2021.
- [25] S. Gupta, R. Cammarota, and T. Rosing, "Memfhe: End-to-end computing with fully homomorphic encryption in memory," 2022.
- [26] S. Gupta and T. S. Rosing, "Invited: Accelerating fully homomorphic encryption with processing in memory," in *ACM/IEEE Design Automation Conference (DAC)*, 2021, pp. 1335–1338.
- [27] K. Han, S. Hong, J. H. Cheon, and D. Park, "Logistic regression on homomorphic encrypted data at scale," in *The 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019, pp. 9466–9471.
- [28] M. Han, Y. Zhu, Q. Lou, Z. Zhou, S. Guo, and L. Ju, "Coxhe: A software-hardware co-design framework for fpga acceleration of homomorphic computation," in *Proceedings of the 2022 Conference & Exhibition on Design, Automation & Test in Europe (DATE)*, 2022.
- [29] HELib, "Cuda-accelerated fully homomorphic encryption library," 2019. [Online]. Available: [https://github.com/homenc/HELlib/tree/master/examples/BGV\\_country\\_db\\_lookuparchivedathttps://perma.cc/U2MW-QLRJ](https://github.com/homenc/HELlib/tree/master/examples/BGV_country_db_lookuparchivedathttps://perma.cc/U2MW-QLRJ).
- [30] Intel. (2018) Intel® 64 and IA-32 Architectures Optimization Reference Manual. [Online]. Available: <https://www.intel.com/content/dam/doc/manual/64-ia-32-architectures-optimization-manual.pdf>
- [31] F. Jafarzadehpour, A. S. Molahosseini, A. A. E. Zarandi, and L. Sousa, "Efficient modular adder designs based on thermometer and one-hot coding," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 9, pp. 2142–2155, 2019.
- [32] L. Jiang, Q. Lou, and N. Joshi, "MATCHA: a fast and energy-efficient accelerator for fully homomorphic encryption over the torus," in *ACM/IEEE Design Automation Conference (DAC)*, 2022, pp. 235–240.
- [33] W. Jung, S. Kim, J. H. Ahn, J. H. Cheon, and Y. Lee, "Over 100x faster bootstrapping in fully homomorphic encryption through memory-centric optimization with gpus," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, no. 4, p. 114–148, 2021.
- [34] A. Khedr and G. Gulak, "Homomorphic processing unit (hpu) for accelerating secure computations under homomorphic encryption," 2019.
- [35] H. Kim, J. Mu, C. Yu, T. T.-H. Kim, and B. Kim, "A 1-16b reconfigurable 80kb 7t sram-based digital near-memory computing macro for processing neural networks," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 70, no. 4, pp. 1580–1590, 2023.
- [36] J. Kim, S. Kim, J. Choi, J. Park, D. Kim, and J. H. Ahn, "Sharp: A short-word hierarchical accelerator for robust and practical fully homomorphic encryption," in *Proceedings of the 50th Annual International Symposium on Computer Architecture (ISCA)*, 2023.
- [37] J. Kim, G. Lee, S. Kim, G. Sohn, M. Rhu, J. Kim, and J. H. Ahn, "ARK: fully homomorphic encryption accelerator with runtime data generation and inter-operation key reuse," in *55th IEEE/ACM International Symposium on Microarchitecture (Micro)*, 2022, pp. 1237–1254.
- [38] S. Kim, J. Kim, M. J. Kim, W. Jung, J. Kim, M. Rhu, and J. H. Ahn, "BTS: an accelerator for bootstrappable fully homomorphic encryption," in *The 49th Annual International Symposium on Computer Architecture (ISCA)*, 2022, pp. 711–725.
- [39] Y. Kim, W. Yang, and O. Mutlu, "Ramulator: A fast and extensible dram simulator," *IEEE Computer Architecture Letters*, vol. 15, no. 1, pp. 45–49, 2016.
- [40] H. Ku, W. Susilo, Y. Zhang, W. Liu, and M. Zhang, "Privacy-preserving federated learning in medical diagnosis with homomorphic re-encryption," *Computer Standards & Interfaces*, vol. 80, 2022.
- [41] C. Lee, S. Min, J. Seo, and Y. Song, "Faster TFHE bootstrapping with block binary keys," in *Proceedings of the 2023 ACM Asia Conference on Computer and Communications Security (ASIACCS)*, 2023, pp. 2–13.
- [42] E. Lee, J. Lee, J. Lee, Y. Kim, Y. Kim, J. No, and W. Choi, "Low-complexity deep convolutional neural networks on fully homomorphic encryption using multiplexed parallel convolutions," in *International Conference on Machine Learning (ICML)*, 2022, pp. 12403–12422.
- [43] J.-W. Lee, H. Kang, Y. Lee, W. Choi, J. Eom, M. Deryabin, E. Lee, J. Lee, D. Yoo, Y.-S. Kim, and J.-S. No, "Privacy-preserving machine learning with fully homomorphic encryption for deep neural network," *IEEE Access*, vol. 10, pp. 30039–30054, 2022.
- [44] J. Lin, L. Liang, Z. Qu, I. Ahmad, L. Liu, F. Tu, T. Gupta, Y. Ding, and Y. Xie, "Inspire: In-storage private information retrieval via protocol and architecture co-design," in *Proceedings of the 49th Annual International Symposium on Computer Architecture (ISCA)*, 2022, p. 102–115.
- [45] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J. Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, "Meltdown: Reading Kernel Memory from User Space," in *USENIX Security Symposium*, 2018, pp. 973–990.
- [46] V. Lyubashevsky, C. Peikert, and O. Regev, "On ideal lattices and learning with errors over rings," in *29th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2010.
- [47] C. Marcolla, V. Sucasas, M. Manzano, R. Bassoli, F. H. P. Fitzek, and N. Aaraj, "Survey on fully homomorphic encryption, theory, and applications," *Proceedings of the IEEE*, vol. 110, pp. 1572–1609, 2022.
- [48] K. Matsuoka, R. Banno, N. Matsumoto, T. Sato, and S. Bian, "Virtual secure platform: A five-stage pipeline processor over TFHE," in *USENIX Security Symposium*, 2021, pp. 4007–4024.
- [49] A. S. Molahosseini, A. Asadpoor, A. A. E. Zarandi, and L. Sousa, "Towards efficient modular adders based on reversible circuits," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018.
- [50] N. Muralimanohar, R. Balasubramonian, and N. Jouppi, "Optimizing nuca organizations and wiring alternatives for large caches with cacti 6.0," in *IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2007, pp. 3–14.
- [51] K. Nam, H. Oh, H. Moon, and Y. Paek, "Accelerating n-bit operations over tfhe on commodity cpu-fpga," in *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2022.

- [52] nucypher, “Nufhe, a gpu-powered torus fhe implementation,” <https://github.com/nucypher/nufhe>, 2019.
- [53] R. Paludo and L. Sousa, “Ntt architecture for a linux-ready risc-v fully-homomorphic encryption accelerator,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 69, no. 7, pp. 2669–2682, 2022.
- [54] Prasetyo, A. Putra, and J.-Y. Kim, “Morphling: A throughput-maximized tthe-based accelerator using transform-domain reuse,” in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2024.
- [55] A. Putra, Prasetyo, Y. Chen, J. Kim, and J.-Y. Kim, “Strix: An end-to-end streaming architecture with two-level ciphertext batching for fully homomorphic encryption with programmable bootstrapping,” in *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2023.
- [56] B. Reagen, W. Choi, Y. Ko, V. T. Lee, H. S. Lee, G. Wei, and D. Brooks, “Cheetah: Optimizing and accelerating homomorphic encryption for private inference,” in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2021, pp. 26–39.
- [57] D. Reis, J. Takeshita, T. Jung, M. T. Niemier, and X. S. Hu, “Computing-in-memory for performance and energy-efficient homomorphic encryption,” *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 28, no. 11, pp. 2300–2313, 2020.
- [58] M. S. Riaz, K. Laine, B. Pelton, and W. Dai, “HEAX: an architecture for computing on encrypted data,” in *Architectural Support for Programming Languages and Operating Systems, Lausanne, (ASPLOS)*, 2020, pp. 1295–1309.
- [59] S. Rixner, W. J. Dally, B. Khailany, P. R. Mattson, U. J. Kapasi, and J. D. Owens, “Register organization for media processing,” in *Proceedings of the Sixth International Symposium on High-Performance Computer Architecture (HPCA)*, 2000, pp. 375–386.
- [60] L. Rovida, “Fast but approximate homomorphic k-means based on masking technique,” *Int. J. Inf. Sec.*, 2023.
- [61] N. Samardzic, A. Feldmann, A. Krastev, S. Devadas, R. G. Dreslinski, C. Peikert, and D. Sánchez, “F1: A fast and programmable accelerator for fully homomorphic encryption,” in *54th Annual IEEE/ACM International Symposium on Microarchitecture (Micro)*, 2021, pp. 238–252.
- [62] N. Samardzic, A. Feldmann, A. Krastev, N. Manohar, N. Genise, S. Devadas, K. Eldefrawy, C. Peikert, and D. Sánchez, “Craterlake: a hardware accelerator for efficient unbounded computation on encrypted data,” in *The 49th Annual International Symposium on Computer Architecture (ISCA)*, 2022, pp. 173–187.
- [63] C. Shen, C. Chen, and J. Zhang, “Micro-architectural cache side-channel attacks and countermeasures,” in *26th Asia and South Pacific Design Automation Conference, (ASPAC)*, 2021, pp. 441–448.
- [64] D. Soni, M. Nabeel, H. Gamil, O. Mazonka, B. Reagen, R. Karri, and M. Maniatakos, “Design space exploration of modular multipliers for asic fhe accelerators,” in *24th International Symposium on Quality Electronic Design (ISQED)*, 2023, pp. 1–8.
- [65] J. Takeshita, D. Reis, T. Gong, M. Niemier, X. S. Hu, and T. Jung, “Accelerating finite-field and torus fhe via compute-enabled (s)ram,” *IEEE Transactions on Computers*, pp. 1–14, 2023.
- [66] S. Tan, B. Knott, Y. Tian, and D. J. Wu, “Cryptgpu: Fast privacy-preserving machine learning on the GPU,” in *IEEE Symposium on Security and Privacy (SP)*, 2021.
- [67] W. Tan, A. Wang, X. Zhang, Y. Lao, and K. K. Parhi, “High-speed VLSI architectures for modular polynomial multiplication via fast filtering and applications to lattice-based cryptography,” *IEEE Transactions on Computers*, vol. 72, no. 9, pp. 2454–2466, 2023.
- [68] vernamlab, “Cuda-accelerated fully homomorphic encryption library,” <https://github.com/vernamlab/cuFHE>.
- [69] J. Wang, C. Xiong, K. Zhang, and J. Wei, “Fixed-point analysis and parameter optimization of the radix-2<sup>k</sup> pipelined FFT processor,” *IEEE Transactions on Signal Processing*, vol. 63, no. 18, pp. 4879–4893, 2015.
- [70] W. Wang, Z. Chen, and X. Huang, “Accelerating leveled fully homomorphic encryption using GPU,” in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2014, pp. 2800–2803.
- [71] W. Wang, X. Huang, N. Emmart, and C. Weems, “Vlsi design of a large-number multiplier for fully homomorphic encryption,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 22, no. 9, pp. 1879–1887, 2014.
- [72] X. Wang, J. Yang, Y. Zhao, X. Jia, R. Yin, X. Chen, G. Qu, and W. Zhao, “Triangle counting accelerations: From algorithm to in-memory computing architecture,” *IEEE Transactions on Computers*, vol. 71, no. 10, pp. 2462–2472, 2022.
- [73] Z. Wang, P. Li, R. Hou, Z. Li, J. Cao, X. Wang, and D. Meng, “He-booster: An efficient polynomial arithmetic acceleration on gpus for fully homomorphic encryption,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 34, no. 4, pp. 1067–1081, 2023.
- [74] S.-Y. Wu, K.-Y. Chen, and M.-D. Shieh, “Efficient vlsi architecture of bluestein’s fft for fully homomorphic encryption,” in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2022, pp. 2242–2245.
- [75] X. Xie, P. Gu, Y. Ding, D. Niu, H. Zheng, and Y. Xie, “Mpu: Memory-centric simt processor via in-dram near-bank computing,” *ACM Transactions on Architecture and Code Optimization*, vol. 20, 2023.
- [76] Y. Yang, H. Lu, and X. Li, “Poseidon-ndp: Practical fully homomorphic encryption accelerator based on near data processing architecture,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 12, pp. 4749–4762, 2023.
- [77] Y. Yang, H. Zhang, S. Fan, H. Lu, M. Zhang, and X. Li, “Poseidon: Practical homomorphic encryption accelerator,” in *IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2023.
- [78] P. Zhang, T. Huang, X. Sun, W. Zhao, H. Liu, S. Lai, and J. K. Liu, “Privacy-preserving and outsourced multi-party k-means clustering on dependable and secure computing,” *IEEE Transactions on Dependable and Secure Computing*, vol. 20, no. 3, 2023.
- [79] Y. Zhang, S. Wang, X. Zhang, J. Dong, X. Mao, F. Long, C. Wang, D. Zhou, M. Gao, and G. Sun, “Pipezk: Accelerating zero-knowledge proof with a pipelined architecture,” in *48th ACM/IEEE Annual International Symposium on Computer Architecture (ISCA)*, 2021, pp. 416–428.