

LaMoS: Enabling Efficient Large Number Modular Multiplication through SRAM-based CiM Acceleration

Haomin Li¹, Fangxin Liu^{1,2,*}, Chenyang Guan¹, Zongwu Wang^{1,2}, Li Jiang^{1,2,*}, Haibing Guan¹

1. School of Computer Science, Shanghai Jiao Tong University, 2. Shanghai Qi Zhi Institute

*Corresponding Author {haominli,liufangxin, ljiang_cs}@sjtu.edu.cn

Abstract—Barrett’s algorithm is one of the most widely used methods for performing modular multiplication, a critical nonlinear operation in modern privacy computing techniques such as homomorphic encryption (HE) and zero-knowledge proofs (ZKP). Since modular multiplication dominates the processing time in these applications, computational complexity and memory limitations significantly impact performance. Computing-in-Memory (CiM) is a promising approach to tackle this problem. However, existing schemes currently suffer from two main problems: 1) Most works focus on low bit-width modular multiplication, which is inadequate for mainstream cryptographic algorithms such as elliptic curve cryptography (ECC) and the RSA algorithm, both of which require high bit-width operations; 2) Recent efforts targeting large number modular multiplication rely on inefficient in-memory logic operations, resulting in high scaling costs for larger bit-widths and increased latency. To address these issues, we propose LaMoS, an efficient SRAM-based CiM design for large-number modular multiplication, offering high scalability and area efficiency. First, we analyze the Barrett’s modular multiplication method and map the workload onto SRAM CiM macros for high bit-width cases. Additionally, we develop an efficient CiM architecture and dataflow to optimize large-number modular multiplication. Finally, we refine the mapping scheme for better scalability in high bit-width scenarios using workload grouping. Experimental results show that LaMoS achieves a 7.02× speedup and reduces high bit-width scaling costs compared to existing SRAM-based CiM designs.

Index Terms—Modular Multiplication, Computing-in-Memory, Privacy Computing

I. INTRODUCTION

Privacy computing, a key approach to ensuring data security, has gained significant attention in recent years due to rising concerns over privacy and personal data protection on the Internet [1]–[6]. It encompasses a variety of applications, such as homomorphic encryption (HE) [7], which safeguards the privacy of user data and models, and zero-knowledge proofs (ZKP) [8], [9], an emerging cryptographic protocol that allows a verifier to confirm the truth of a statement without revealing any additional information. These applications rely on public key cryptography algorithms like elliptic curve cryptography (ECC) [10] and RSA [11], where modular multiplication plays a critical role in preventing overflow.

This work is supported by the National Key Research and Development Program of China (2024YFE0204300), the National Natural Science Foundation of China (Grant No.62402311), and Natural Science Foundation of Shanghai (Grant No.24ZR1433700). Fangxin Liu and Li Jiang are the corresponding authors.

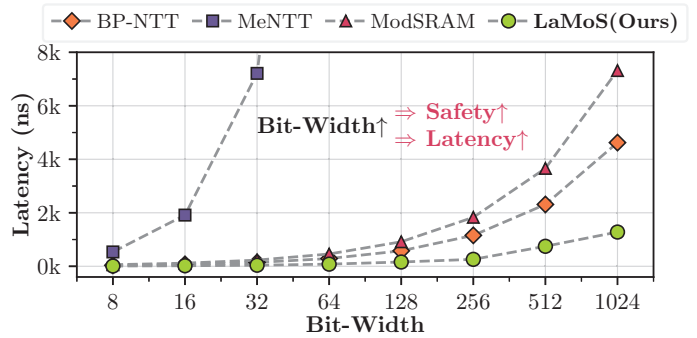


Fig. 1. Performance Comparison with previous works over various bit-widths.

Notably, modular multiplication in these algorithms requires extremely high bit-width implementations—ECC demands at least 224 bits [12], while RSA typically requires 1024 bits or more [13]. In most cryptographic schemes, higher bit-width operations correspond to greater security [14]. However, this high bit-width processing poses significant challenges for efficient hardware implementation: 1) Efficient algorithms like Barrett [15] and Montgomery [16] rely on fast multipliers, and implementing them for high bit-widths increases latency and area complexity. 2) Modular multiplication generates numerous intermediate results, placing heavy demands on storage and memory bandwidth.

Previous works on FPGA-based modular multiplication have focused on optimizing the use of DSPs and LUTs to enhance computational efficiency [17]–[23]. However, these approaches overlook the significant data handling overhead characteristic of real-world privacy computing applications. For instance, PipeZK [18] requires nearly 3 TB/s bandwidth at 100 MHz to compute a ZKP scheme, which is impractical for current systems. To address this issue, computing-in-memory (CiM) [24]–[28] has emerged as a promising platform to reduce data handling overhead. It has been widely explored for accelerating privacy computing applications, including advanced encryption standard (AES) and homomorphic encryption [29]–[35].

However, these acceleration designs primarily target cryptographic algorithms with low bit-width requirements, like post-quantum cryptography (PQC) [36], which typically only demand bit-widths of 64 or less, allowing modular multiplication to be implemented with minimal overhead. In contrast, many privacy computing applications rely on cryptographic

algorithms with much higher bit-width requirements, such as ZKP based on ECC, which requires several hundred bits for modular multiplication.

As illustrated in Figure 1, we compare the performance of different works [30]–[32] across various bit-widths¹. It is clear that extending existing designs naively to higher bit-widths introduces significant latency overhead, as these schemes are not optimized for large-number modular multiplication. The latest work, ModSRAM [30], focuses on accelerating large number modular multiplication for ECC. However, the performance of the proposed accelerator when integrated into SRAM in-memory logic is inadequate in two key areas:

- **High latency and resource inefficiency:** For 256-bit modular multiplication, ModSRAM requires 767 cycles at 400 MHz, which is insufficient for privacy computing applications. This high latency also negatively impacts the CiM design, as it locks the SRAM array for too long. For a 256-bit implementation, the 6KB (64×256) SRAM-based CiM array, where operands are stored, remains locked for 767 cycles, restricting its use solely to modular multiplication. In other words, high latency amplifies resource waste in CiM arrays.
- **High scaling cost:** Since the implementation relies on parallel bit-wise logic operations in SRAM, scaling ModSRAM to higher bit-widths is extremely costly. For example, increasing from 256-bit to 512-bit would require more than 3,000 cycles to complete a modular multiplication, unless additional SRAM arrays are introduced to maintain higher parallelism in the bit-wise logic.

To address the high latency and scaling challenges in large-number modular multiplication, we propose **LaMoS**, an efficient SRAM-based CiM design. First, we analyze the key operations in Barrett’s modular multiplication algorithm and develop a mapping scheme that divides the computation of large number multiplications into smaller workloads, which are mapped onto the SRAM-based CiM macro. Then, we design an architecture and dataflow that utilizes parallelism across CiM macros to achieve low-latency large-number modular multiplication. Next, we introduce a workload grouping-based mapping optimization to improve performance when scaling to higher bit-widths by reducing redundant CiM macro usage.

Overall, our work has the following contributions:

- We design LaMoS, a high-efficiency SRAM-based CiM architecture for large-number modular multiplication. By exploiting workload parallelism across multiple CiM macros, LaMoS significantly reduces latency for large-number modular multiplication.
- We propose a novel mapping strategy for large-number multiplication within SRAM-based CiM macros. By partitioning the large-number multiplication into smaller, manageable workloads and mapping them efficiently onto the macros, we enable high-performance parallel computation, which forms the basis for an optimized Barrett’s modular multiplication.

¹These works do not report performance across all bit-widths, so we scaled the results according to the complexity of the designs.

Algorithm 1: Barrett’s Modular Multiplication [15].

Data: three n -bit unsigned numbers A , B , and M

Result: $A \times B \bmod M$

1	Offline:	
2	$M' \leftarrow \lfloor \frac{2^{2n}}{M} \rfloor$	Precompute
3	Online:	
4	$C \leftarrow A \times B$	Multiplier
5	$u \leftarrow \lfloor \frac{C}{2^{n-1}} \rfloor \times M'$	Multiplier
6	$E \leftarrow \lfloor \frac{u}{2^{n+1}} \rfloor$	Shift
7	$P \leftarrow E \times M$	Multiplier
8	$T \leftarrow C - P$	Subtractor
9	$R \leftarrow T \text{ or } T - M \text{ or } T - 2M$	Subtractors

- We optimize the workload mapping process to minimize latency as bit-widths increase. Through grouping workloads and eliminating redundant computations, LaMoS can maximize CiM macro utilization. Evaluations demonstrate that LaMoS delivers high area efficiency, achieving a $6\times$ speedup over ModSRAM and significantly reducing scaling costs for 1024-bit and 2048-bit operations.

II. BACKGROUNDS

A. Modular Multiplication

Let A , B and M be n -bit unsigned inputs, where $C = AB$ is a $2n$ -bit unsigned value. Modular multiplication computes the remainder R of C divided by M :

$$R = AB \bmod M = C \bmod M \quad (1)$$

Barrett’s algorithm is one of the most commonly used methods for modular multiplication, as it replaces the slow and costly division operation with several multiplications. The detailed steps of Barrett’s algorithm are outlined in Algorithm 1. First, $C = A \times B$ is computed by an n -bit multiplier. Then, the quotient E is estimated by multiplying the lower $n + 1$ bits of C by M' , where M' is the precomputed reciprocal of modulo M , and then truncating the result to $n + 1$ bits. Since the modulo M is often static in a privacy computing application, M' can be computed offline as $\lfloor \frac{2^{2n}}{M} \rfloor$. After that, EM is subtracted from C to obtain an intermediate value T . Finally, T is subtracted by M until the final result lies within the range $[0, M)$. This final reduction step typically requires at most two subtractions, which can be efficiently implemented using cascaded subtractors and a multiplexer.

B. Cryptographic CiM

Several cryptographic CiM accelerators have been developed based on both SRAM [31], [32] and ReRAM array [33]–[35], with a focus on reducing the heavy data movement overhead commonly found in privacy computing applications. These accelerators typically target cryptographic schemes with lower bit-width requirements, such as AES and PQC. They often optimize at the algorithmic level, focusing on techniques

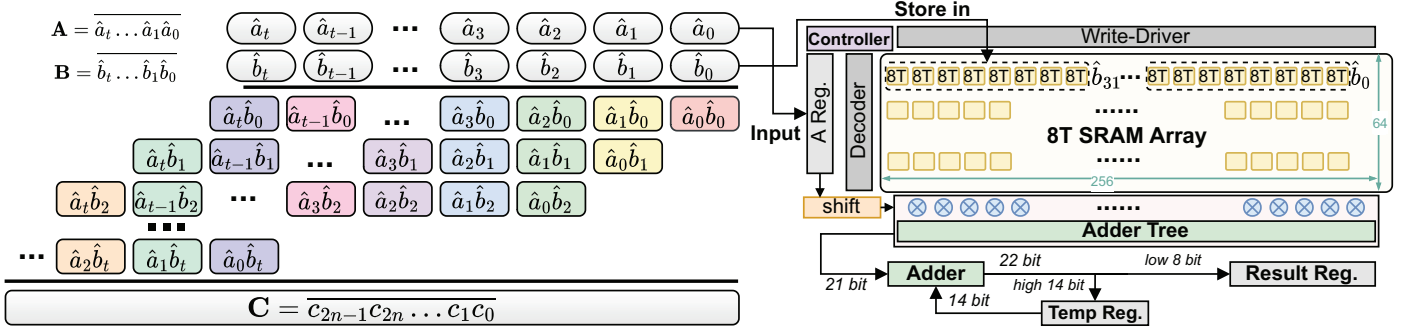


Fig. 2. Large Number Multiplication Workload Mapping onto the SRAM Macro.

like the number theoretic transform (NTT), or at the application level, for example, in homomorphic encryption. ModSRAM [30] is a more recent design that addresses large-number modular multiplication in ECC, which demands higher bit-width representations ranging from 224 to 512 bits. ModSRAM utilizes SRAM in-memory logic to implement modular multiplication based on a modified interleaved algorithm. However, its performance remains suboptimal. For an n -bit modular multiplication, ModSRAM requires $3n + 1$ cycles to complete the calculation, which falls short of the requirements for privacy computing applications.

C. Motivation

Recent CiM architectures primarily focus on cryptographic schemes such as AES and PQC, which represent only a small subset of cryptographic applications. Many widely used schemes like ECC and RSA require extremely high bit-width implementations, demanding support for large-number arithmetic operations. Additionally, in most cryptographic schemes, increasing the bit-width enlarges the key space, making the system harder to crack and thereby enhancing security.

Thus, accelerating high bit-width arithmetic operations, particularly large-number modular multiplication, is crucial for cryptography. However, existing CiM designs generally target low bit-width multiplication, or they are difficult to optimize for larger bit-widths, leading to significant area or latency overhead as the bit-width scales. While SRAM-based CiM designs implement bit-wise in-memory logic for large-number modular multiplication, like ModSRAM, they suffer from inefficiency and high scaling costs.

In summary, a more efficient solution with lower scaling costs for large-number modular multiplication is still missing in above works, which we believe will be the trend for future privacy computing accelerators. SRAM-based CiM macros with multiply-and-accumulate (MAC) capabilities [37] offer greater computational power, potentially leading to reduced latency and scaling costs compared to SRAM-based CiM designs using logic operations. Furthermore, Barrett's algorithm, which relies on multiplication, has a lower latency overhead compared to interleaved algorithms. Therefore, we propose an efficient and scalable modular multiplication architecture based on Barrett's algorithm, harnessing the MAC capabilities of SRAM CiM macros to achieve superior performance.

III. PROPOSED DESIGN

A. Workload Segmentation and Mapping

To address the computational demands of Barrett's algorithm, which involves three large-number multiplications (orange part in Algorithm 1), we propose a systematic approach for mapping these operations onto SRAM-based CiM MAC macros. Our design leverages the 64×256 8T SRAM macro, known for its capability to execute 32 8-bit MAC operations per cycle. We decompose the large-number multiplication into several smaller workloads, each comprising 32 8-bit multiplication operations. These workloads are efficiently mapped onto the SRAM macro, allowing the large-number multiplication to be executed with minimal latency and using a few low bit-width accumulators.

Specifically, we perform the following transformation on the large numbers A and B , involved in the multiplication, by expressing them as a weighted sum of 8-bit numbers:

$$\begin{aligned}
 A &= \sum_{i=0}^{n-1} (2^i a_i) = \sum_{i=0}^{\frac{n}{8}-1} \sum_{j=0}^7 (2^{8i+j} a_{8i+j}) = \sum_{i=0}^{\frac{n}{8}-1} \hat{a}_i \times 2^{8i} \\
 B &= \sum_{i=0}^{n-1} (2^i b_i) = \sum_{i=0}^{\frac{n}{8}-1} \sum_{j=0}^7 (2^{8i+j} b_{8i+j}) = \sum_{i=0}^{\frac{n}{8}-1} \hat{b}_i \times 2^{8i}
 \end{aligned} \tag{2}$$

where $\hat{a}_i = \sum_{j=0}^7 (2^j a_{8i+j})$ and $\hat{b}_i = \sum_{j=0}^7 (2^j b_{8i+j})$, both of which are 8-bit numbers. Next, the multiplication can be represented as a weighted sum of the products of these 8-bit multiplications:

$$\begin{aligned}
 C &= A \times B = \sum_{i=0}^{\frac{n}{8}-1} \sum_{j=0}^{\frac{n}{8}-1} \hat{a}_i \times 2^{8i} \times \hat{b}_j \times 2^{8j} \\
 &= \sum_{i=0}^{\frac{n}{8}-1} \sum_{j=0}^{\frac{n}{8}-1} \hat{a}_i \times \hat{b}_j \times 2^{8(i+j)}
 \end{aligned} \tag{3}$$

Figure 2 illustrates the vertical structure of the multiplication between A and B , where each entry in the vertical layout corresponds to the product of two 8-bit numbers \hat{a}_i and \hat{b}_j . Each multiplication $\hat{a}_i \hat{b}_j$ can be regarded as a separate workload. Therefore, a multiplication of two n -bit numbers can be viewed as the accumulation of t^2 workloads, where $t = \lceil \frac{n}{8} \rceil$. Additionally, within the vertical equation, workloads in the same column share the same weighting. For example, the results of $\hat{a}_2 \hat{b}_0$, $\hat{a}_1 \hat{b}_1$, $\hat{a}_0 \hat{b}_2$ all have the same weighting factor of 2^{16} . This allows us to map the workloads in each




Fig. 3. Input Timing Flow for 256-bit multiplication with single/multiple SRAM macro(s). Number of SRAM macros is set to 3 in the example.

column to a single MAC operation within the CiM macro, thereby optimizing the overall computation.

Figure 2 also illustrates the mapping process for a 256-bit multiplication. We begin by mapping the operand B onto a row of SRAM as 32 8-bit numbers. During each cycle, to perform a single MAC operation of large-number multiplication, we provide the necessary inputs from A for one column of workloads.

The input stream configuration for a single MAC macro is detailed in Figure 3. In the first cycle, the inputs consist of 31 $8'b0$ s and \hat{a}_0 . In the second cycle, the inputs are 30 $8'b0$, \hat{a}_1 , and \hat{a}_0 . Such a workload mapping ensures that the SRAM-based CiM MAC macro is utilized efficiently for workload computation. For a 256-bit multiplication, this approach allows all workloads to be completed within 63 cycles on a single MAC macro.

The peripheral accumulation circuit for the MAC results² is detailed in Figure 2. Each cycle generates a 21-bit MAC result, where the low-order 8 bits of this 21-bit value directly contribute to the final result. This is because the weighting difference between MAC results of adjacent cycles is 2^8 , making the low-order 8 bits of each cycle's result the significant part of the final result.

During each cycle, the 21-bit MAC result is split into two parts: the low-order 8 bits and the high-order 13 bits. The low-order 8 bits are added to the final result register. To handle the high-order bits, the high-order 13 bits of the current cycle's result are added to the high-order 14 bits of the 22-bit addition result (with 1 carry bit) from the previous cycle using a 21-bit adder. The result of this addition provides the high-order part for the current cycle, with its low-order 8 bits stored in the final result register and the high-order 14 bits kept in a temporary register to be carried over and added in the next cycle. This process ensures the correct accumulation of results across all cycles.

B. Architecture Design

Based on the proposed mapping scheme, we design LaMoS, an efficient and scalable architecture for CiM-based large-number modular multiplication, as shown in Figure 4.

For single large-number multiplications, LaMoS uses a scalable parallel design with multiple CiM macros. To perform modular multiplication, we incorporate peripheral circuits and dataflows that optimize computation efficiency. We use a parallel acceleration approach with multiple macros. Multiple

²The MAC result at each cycle is a $8 + 8 + \log_2(32) = 21$ -bit value.

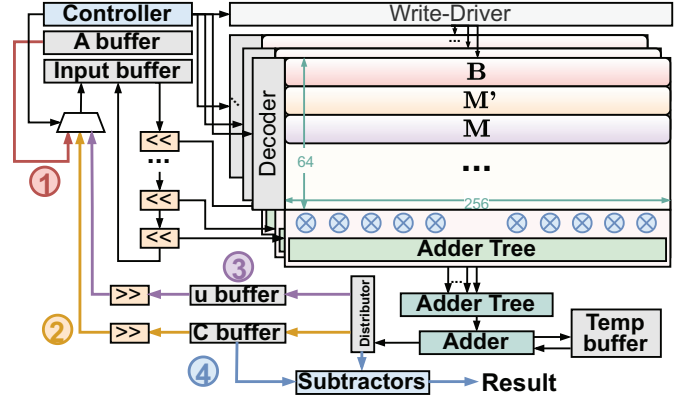


Fig. 4. LaMoS Architecture and Dataflow for efficient large number modular multiplication.

columns of the workload, as illustrated in Figure 3, are distributed across these macros. The value of B is replicated and stored in each macro. In each cycle, each macro processes one column of the workload independently. For example, in a 256-bit multiplication with three macros, during the first cycle, macro #1 processes 31 $8'b0$ and \hat{a}_0 , macro #2 processes 30 $8'b0$ along with \hat{a}_1 and \hat{a}_0 , and macro #3 processes 29 $8'b0$ along with \hat{a}_2 , \hat{a}_1 , and \hat{a}_0 .

The peripheral circuits designed to support large-number modular multiplication are illustrated in Figure 4. LaMoS includes a low-bit-width adder and temporary registers to accumulate computation results from each cycle. An adder tree aggregates results from parallel macros, with minimal overhead given that a small number of macros (between 2 and 8) is sufficient for optimal latency. For input data handling, we use a shift array to provide data with different shifts to the appropriate macro each cycle. We incorporate buffers, a multiplexer (MUX), and a distributor to manage the three different multiplications.

- First multiplication ($A \times B$): The MUX directs A to the input buffer for computation. The resulting product C is stored in the C buffer via the distributor.
- Second multiplication ($\lfloor \frac{C}{2^{n-1}} \rfloor \times M'$): The value of C from the C buffer is shifted and sent to the input buffer. The result u is stored in the u buffer.
- Third multiplication ($E \times M$): The value u from the u buffer is shifted to obtain E , which is then sent to the input buffer. Finally, the distributor routes the result P to a subtractor, where P is subtracted from C to obtain T . The subtractor then performs additional subtractions of M from T to complete the refinement process.

C. Mapping Optimization for Scaling to Higher Bit-Width

As LaMoS is scaled to handle higher bit-widths (e.g., beyond 256 bits), the latency of computations increases significantly. Figure 5(a) demonstrates the input stream when mapping high bit-width multiplications onto the architecture. For instance, in a 1024-bit multiplication with LaMoS configured with 4 macros, the 1024-bit number is expressed as a weighted sum of 128 8-bit numbers. Since each 64×256 macro supports MAC operations on only 32 8-bit inputs, the input stream

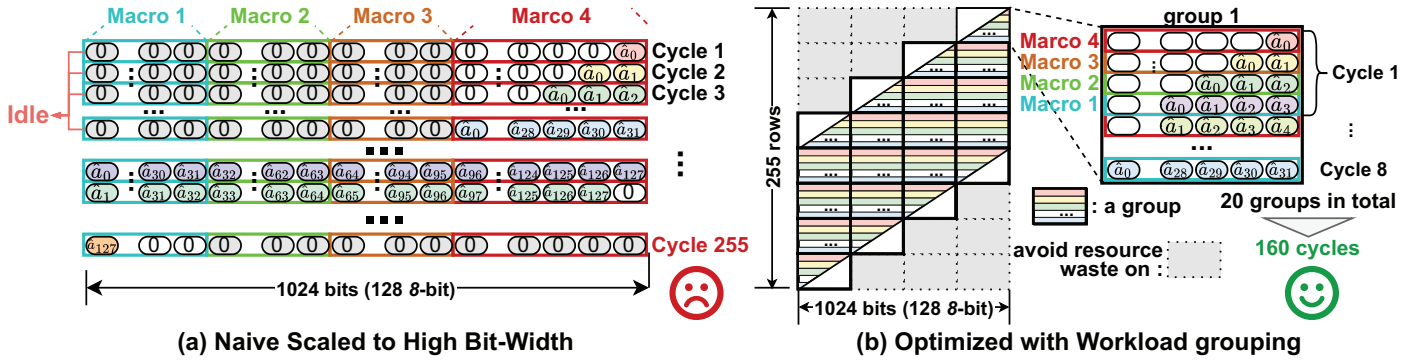


Fig. 5. LaMoS Execution for high bit-width modular multiplication. (a) Naively Scaling. (b) Optimization with workload grouping.

must be divided into 4 slices for parallel processing by the 4 macros. Consequently, 255 cycles are required to complete the multiplication.

To enhance performance, we analyze the workload distribution and observe that some macros are often idle. For example, in the first 92 rows of the input stream, macro 1 consistently processes 0s. This idleness results from varying workloads per column in Figure 2, leading to excessive padding with 0s. Figure 5(b) shows a 255-row by 1024-bit input stream for a 1024-bit multiplication, highlighting large amounts of padded 0s (gray and white regions). The gray regions correspond to cycles where all inputs to a macro are 0s, causing the macros to be idle.

To this end, we propose a grouping-based mapping optimization. We divide the input streams into groups of 32-row by 128-bit each, ensuring that a single macro processes exactly one row per group. Redundant groups (gray region) with only 0s are discarded. Non-redundant groups are processed by multiple macros in parallel, requiring only 8 cycles per group, given 4 macros. This optimization reduces the workload by 37.5% for 1024-bit multiplications and eliminates macro idleness, achieving a final latency of 160 cycles.

IV. EVALUATIONS

A. Evaluation Methodology

Implementation and Simulation: We implemented LaMoS using Verilog and synthesized the RTL with Synopsys Design Compiler [38] at 400MHz, targeting the TSMC 28nm standard library to evaluate the area. For the SRAM-based CiM macro configuration, we referred to the data reported in [37]. To assess latency, we developed a cycle-accurate simulator that models the micro-architecture behavior across various bit-widths. Balancing area and latency, we evaluated LaMoS configured with $2 \times 64 \times 256$ CiM macros.

Baselines: We compare the LaMoS with the following designs: (1) SRAM-based CiM designs: MeNTT [31] and BP-

NTT [32]; (2) ReRAM-based CiM designs: RM-NTT [33], CryptoPIM [34], and X-Poly [35]; (3) ModSRAM [30], our primary baseline, as it is the only SRAM-based CiM design focused on large-number modular multiplication and only report performance on low bit-width operations, we scale their results to high bit-widths based on the theoretical complexity of each design³.

B. Performance

Table I compares LaMoS with the baseline architectures. The cycle counts of the baselines are scaled to a 256-bit configuration for fair comparison. To evaluate overall performance, we use the $Latency \times Area$ metric, where lower values indicate better performance. Unlike existing SRAM-based CiM architectures, LaMoS is based on Barrett's algorithm, which results in extremely low latency, requiring only 104 cycles for a 256-bit modular multiplication. With the proposed mapping optimization, LaMoS can efficiently scale to support arbitrary bit-width computations. In terms of overall performance, LaMoS demonstrates a $3\times$ improvement over the latest architecture, ModSRAM. Additionally, compared to BP-NTT, which uses the efficient Montgomery algorithm for modular multiplication, LaMoS achieves $2.5\times$ higher area efficiency by employing a more specialized architecture and dataflow that fully leverages the computational power of CiM macros.

For ReRAM-based architectures, directly scaling modular multiplication to higher bit-widths on crossbars leads to significant area overheads, both in terms of the number of crossbars and the required ADCs. It is important to note that our proposed workload mapping and optimization techniques are also well-suited for ReRAM-based CiM macros. As such, designing efficient ReRAM-based architectures for modular multiplication will be a key focus of our future research.

Figure 6 provides the area breakdown of LaMoS, where the SRAM array occupies the largest portion. The near-memory circuits account for less than 35% of the total area, demonstrating that our design incurs minimal overhead when extending from the CiM macros. Furthermore, the simplicity of the additional circuits highlights that LaMoS can be practically implemented with ease.

³BP-NTT utilizes Montgomery's modular multiplication, so we account for the additional overhead of the Montgomery transformation.




Fig. 6. LaMoS Area Breakdown.

TABLE I
COMPARISON ON MODULAR MULTIPLICATION IN PIM DESIGNS.

Reference	LaMoS (ours)	ModSRAM [30]	MeNTT [31]	BP-NTT [32]	RM-NTT [33]	CryptoPIM [34]	X-Poly [35]
Application Type	ECC & RSA	ECC	PQC NTT	PQC NTT	HE NTT	PQC NTT	PQC NTT
Computation Method	Barrett	Direct	Direct	Montgomery	Montgomery	Montgomery/Barrett	Barrett
technology	28nm	65nm	45nm	28nm	28nm	45nm	45nm
Cell Type	8T SRAM	8T SRAM	6T SRAM	6T SRAM	ReRAM	ReRAM	ReRAM
Array size	2×64×256	64×256	4×162×256	4×256×256	64×4×128×128	512×512	16×128×128
Frequency (MHz)	400	420	151	3.8k	400	909	400
Bitwidth	Any Bit-width	256	14/16/32	2/4/8/16/32/64	14/16	16/32	16
Cycles	104	767	66049	4395	-	-	-
Area (mm ²)	0.11	0.053	0.36	0.063	-	0.152	0.27
Latency×Area (μs × mm ²)	0.029	0.09	157.47	0.073	-	-	-




Fig. 7. Detailed analysis of contributions.




Fig. 9. Design Space Exploration.

C. Ablation Study

Figure 7 presents the results of the ablation study for LaMoS. The “Serial” configuration uses a single CiM macro, while the “Parallel” configuration employs two CiM macros. Mapping optimization is applied when the bit-width exceeds 256 bits. For all bit-widths, the parallel LaMoS design achieves at least a 1.9× speedup compared to the serial LaMoS. This demonstrates the effectiveness of parallel acceleration in our architecture, which is further explored in Section IV-E. For bit-widths greater than 256 bits, LaMoS with mapping optimization achieves a speedup of 1.32× and 1.59× over the non-optimized version. This improvement is attributed to the workload grouping strategy, which minimizes idle time in the CiM macros and enhances the overall execution efficiency of LaMoS.

D. Scalability Analysis

To evaluate the scalability of LaMoS, we compare its latency and area efficiency with that of ModSRAM across different bit-widths, as shown in Figure 8. We use the same macro configuration as detailed in Table I to ensure a consistent basis for performance comparison at higher bit-widths. For a 2048-bit width, ModSRAM experiences a latency exceeding 110,000 ns and an area efficiency of less than 200,000 OPS/mm². This performance degradation is primarily due to: (1) the triple-linear theoretical complexity of ModSRAM’s bit-logic-based modular multiplication algorithm with increasing bit-width, and

(2) the need for multiple cycles to perform bit logic operations at high bit-widths, which can be completed in a single cycle at lower bit-widths, further amplifying latency issues. In contrast, LaMoS maintains a latency of less than 9,000 ns and an area efficiency greater than 2,000,000 OPS/mm² at the 2,048 bits. These improvements are attributed to LaMoS’s efficient workload mapping and optimization techniques, as well as its multi-macro parallel design.

E. Design Space Exploration

We conduct a design space exploration for LaMoS to assess its scalability and performance across various configurations. Figure 9 illustrates the area results for LaMoS with different numbers of CiM macros and the corresponding performance at varying bit-widths. The area of LaMoS increases linearly with the number of CiM macros, demonstrating minimal overhead for scaling up the architecture to support higher parallelism. With 4 CiM macros, LaMoS achieves fewer than 2,000 cycles for modular multiplications across various bit-widths. When configured with 8 CiM macros, LaMoS can complete a 256-bit modular multiplication in just 32 cycles. This scaling capability highlights LaMoS’s effectiveness in maintaining high performance and efficiency as the number of macros increases.

V. CONCLUSION

In this paper, we propose LaMoS, a scalable SRAM-based CiM architecture for efficient large number modular multiplication. We design a workload segmentation and mapping scheme to map multiplication onto CiM MAC macros and develop LaMoS for parallel acceleration. Additionally, we optimize the mapping to avoid unnecessary macro idles and improve LaMoS’s performance at high bit-widths. Extensive experiments demonstrate that LaMoS achieves significant performance advantages compared to existing CiM designs.




Fig. 8. Comparison with ModSRAM [30] on various bit-width.

REFERENCES

- [1] Z. Xiao and Y. Xiao, "Security and privacy in cloud computing," *IEEE communications surveys & tutorials*, vol. 15, no. 2, pp. 843–859, 2012.
- [2] L. Li, Y. Fan, M. Tse, and K.-Y. Lin, "A review of applications in federated learning," *Computers & Industrial Engineering*, vol. 149, p. 106854, 2020.
- [3] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowledge-Based Systems*, vol. 216, p. 106775, 2021.
- [4] F. Liu, H. Li, Y. Chen, T. Yang, and L. Jiang, "Hyperattack: An efficient attack framework for hyperdimensional computing," in *2023 60th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2023, pp. 1–6.
- [5] H. Li, F. Liu, Y. Chen, and L. Jiang, "Hyperfeel: An efficient federated learning framework using hyperdimensional computing," in *2024 29th Asia and south Pacific design automation conference (ASP-DAC)*. IEEE, 2024, pp. 716–721.
- [6] H. Li, F. Liu, Z. Wang, N. Yang, S. Huang, X. Liang, H. Guan, and L. Jiang, "Attack and defense: Enhancing robustness of binary hyperdimensional computing," *ACM Transactions on Architecture and Code Optimization*, 2025.
- [7] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Computing Surveys (Csur)*, vol. 51, no. 4, pp. 1–35, 2018.
- [8] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof-systems," in *Providing sound foundations for cryptography: On the work of shafi goldwasser and silvio micali*, 2019, pp. 203–225.
- [9] U. Feige, A. Fiat, and A. Shamir, "Zero knowledge proofs of identity," in *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, 1987, pp. 210–217.
- [10] N. Kobitz, "Elliptic curve cryptosystems," *Mathematics of computation*, vol. 48, no. 177, pp. 203–209, 1987.
- [11] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.
- [12] L. Chen, D. Moody, A. Regenscheid, and A. Robinson, "Digital signature standard (dss)," 2023.
- [13] J. Bos, M. Kaihara, T. Kleinjung, A. K. Lenstra, and P. L. Montgomery, "On the security of 1024-bit rsa and 160-bit elliptic curve cryptography," 2009.
- [14] B. Schneier, *Applied cryptography: protocols, algorithms, and source code in C*. john wiley & sons, 2007.
- [15] P. Barrett, "Implementing the rivest shamir and adleman public key encryption algorithm on a standard digital signal processor," in *Conference on the Theory and Application of Cryptographic Techniques*. Springer, 1986, pp. 311–323.
- [16] P. L. Montgomery, "Modular multiplication without trial division," *Mathematics of computation*, vol. 44, no. 170, pp. 519–521, 1985.
- [17] M. Langhammer and B. Pasca, "Efficient fpga modular multiplication implementation," in *The 2021 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 2021, pp. 217–223.
- [18] 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 *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2021, pp. 416–428.
- [19] A. Opasatian and M. Ikeda, "Lookup table modular reduction: A low-latency modular reduction for fast ecc processor," in *2023 IEEE Symposium in Low-Power and High-Speed Chips (COOL CHIPS)*. IEEE, 2023, pp. 1–6.
- [20] B. Devlin, "Blockchain acceleration using fpgas—elliptic curves, zk-snarks, and vdfs," *ZCASH Foundation*, 2019.
- [21] E. Öztürk, "Design and implementation of a low-latency modular multiplication algorithm," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 67, no. 6, pp. 1902–1911, 2020.
- [22] S. Gribok, M. Langhammer, and B. Pasca, "Fpga modular multipliers using hybrid reduction techniques," in *34th International Conference on Field-Programmable Logic and Applications-FPL 2024*, 2024.
- [23] F. Liu, H. Li, Z. Wang, B. Zhang, M. Zhang, S. Yan, L. Jiang, and H. Guan, "Allmod: Exploring area-efficiency of lut-based large number modular reduction via hybrid workloads," in *2025 62nd ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2025, pp. 1–7.
- [24] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "A modern primer on processing in memory," in *Emerging Computing: From Devices to Systems: Looking Beyond Moore and Von Neumann*. Springer, 2022, pp. 171–243.
- [25] A. Sebastian, M. Le Gallo, R. Khaddam-Aljameh, and E. Eleftheriou, "Memory devices and applications for in-memory computing," *Nature nanotechnology*, vol. 15, no. 7, pp. 529–544, 2020.
- [26] F. Liu, H. Li, B. Zhu, Z. Wang, Z. Song, H. Guan, and L. Jiang, "Asdr: Exploiting adaptive sampling and data reuse for cim-based instant neural rendering," in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, 2025, pp. 18–33.
- [27] Z. Wang, F. Liu, N. Yang, S. Huang, H. Li, and L. Jiang, "Compass: Sram-based computing-in-memory snn accelerator with adaptive spike speculation," in *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2024, pp. 1090–1106.
- [28] F. Liu, H. Li, N. Yang, Y. Chen, Z. Wang, T. Yang, and L. Jiang, "Paap-hd: Pim-assisted approximation for efficient hyper-dimensional computing," in *2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2024, pp. 46–51.
- [29] D. Reis, H. Geng, M. Niemier, and X. S. Hu, "Imcrypto: an in-memory computing fabric for aes encryption and decryption," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 30, no. 5, pp. 553–565, 2022.
- [30] J. Ku, J. Zhang, H. Shan, S. Samudrala, J. Wu, Q. Zheng, Z. Li, J. Rajendran, and Y. Chen, "Modsrn: Algorithm-hardware co-design for large number modular multiplication in sram," in *2024 61th ACM/IEEE Design Automation Conference (DAC)*, 2024.
- [31] D. Li, A. Pakala, and K. Yang, "Mennt: A compact and efficient processing-in-memory number theoretic transform (ntt) accelerator," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 30, no. 5, pp. 579–588, 2022.
- [32] J. Zhang, M. Imani, and E. Sadredini, "Bp-ntt: Fast and compact in-sram number theoretic transform with bit-parallel modular multiplication," in *2023 60th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2023, pp. 1–6.
- [33] Y. Park, Z. Wang, S. Yoo, and W. D. Lu, "Rm-ntt: An rram-based compute-in-memory number theoretic transform accelerator," *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits*, vol. 8, no. 2, pp. 93–101, 2022.
- [34] H. Nejatollahi, S. Gupta, M. Imani, T. S. Rosing, R. Cammarota, and N. Dutt, "Cryptopim: In-memory acceleration for lattice-based cryptographic hardware," in *2020 57th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2020, pp. 1–6.
- [35] M. Li, H. Geng, M. Niemier, and X. S. Hu, "Accelerating polynomial modular multiplication with crossbar-based compute-in-memory," in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. IEEE, 2023, pp. 1–9.
- [36] D. J. Bernstein and T. Lange, "Post-quantum cryptography," *Nature*, vol. 549, no. 7671, pp. 188–194, 2017.
- [37] H. Zhu, B. Jiao, J. Zhang, X. Jia, Y. Wang, T. Guan, S. Wang, D. Niu, H. Zheng, C. Chen *et al.*, "Comb-mcm: Computing-on-memory-boundary nn processor with bipolar bitwise sparsity optimization for scalable multi-chiplet-module edge machine learning," in *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 65. IEEE, 2022, pp. 1–3.
- [38] Synopsys, <https://www.synopsys.com/community/university-program/teaching-resources.html>, [Online].