MeNTT: A Compact and Efficient Processing-in-Memory Number Theoretic Transform (NTT) Accelerator

Dai Li*, *Student Member, IEEE,* Akhil Pakala*, *Student Member, IEEE,* and Kaiyuan Yang, *Member, IEEE*

Abstract—Lattice-based cryptography (LBC) exploiting Learning with Errors (LWE) problems is a promising candidate for post-quantum cryptography. Number theoretic transform (NTT) is the latency- and energy- dominant process in the computation of LWE problems. This paper presents a compact and efficient in-MEmory NTT accelerator, named MeNTT, which explores optimized computation in and near a 6T SRAM array. Specifically-designed peripherals enable fast and efficient modular operations. Moreover, a novel mapping strategy reduces the data flow between NTT stages into a unique pattern, which greatly simplifies the routing among processing units (i.e., SRAM column in this work), reducing energy and area overheads. The accelerator achieves significant latency and energy reductions over prior arts.

Index Terms—Post-Quantum Cryptography, PQC, Learning with Errors, LWE, Processing-in-Memory, PIM, SRAM, NTT.

I. INTRODUCTION

Most contemporary public-key cryptographic primitives like RSA and Elliptic Curve Cryptography (ECC) rely on the difficulty to solve integer factorization and discrete algorithms. However, with the advent of quantum computers, these problems are expected to be solvable, such as using Shor's algorithm [1] in polynomial time. Therefore cryptographic algorithms that are resistant to potential attacks using quantum computers are being studied by researchers around the globe. Lattice-based Cryptography (LBC) has emerged as abprime candidate among the Post-quantum protocols.

Lattice-based protocols have come to light because of the hardness of inherent Learning with Errors (LWE) problems [2], [3]. Module LWE and Ring LWE are two primary variants of the LWE problem. Ring LWE secures the message using polynomial operations between secret key and public key along with error addition (Fig. 1). Polynomial multiplication is performed using Number theoretic transform (NTT) [4], an FFT like structure except that the operations performed are modular arithmetic. In a typical Homomorphic/PQC system, sampling and NTT are the two main operations in Ring LWE schemes and they take a similar amount of time [5]. Recent studies have significantly improved the energy and area efficiency of random samplers, e.g., our recent work MePLER [6] achieved 20.6-pJ per sampler energy efficiency and 85.9-MSample/s constant throughput in custom 65-nm

*These authors contributed equally.
The authors are with the Department of Electrical and Computer Engineering, Rice University, Houston, TX, 77005, USA. email: {dl37, arp5, kyang}@rice.edu. (Corresponding author: Kaiyuan Yang)
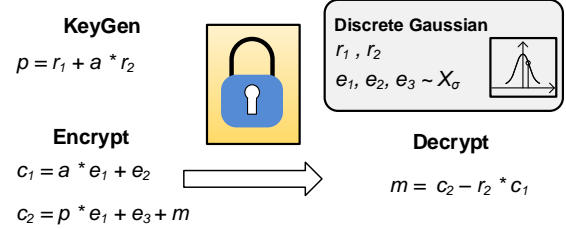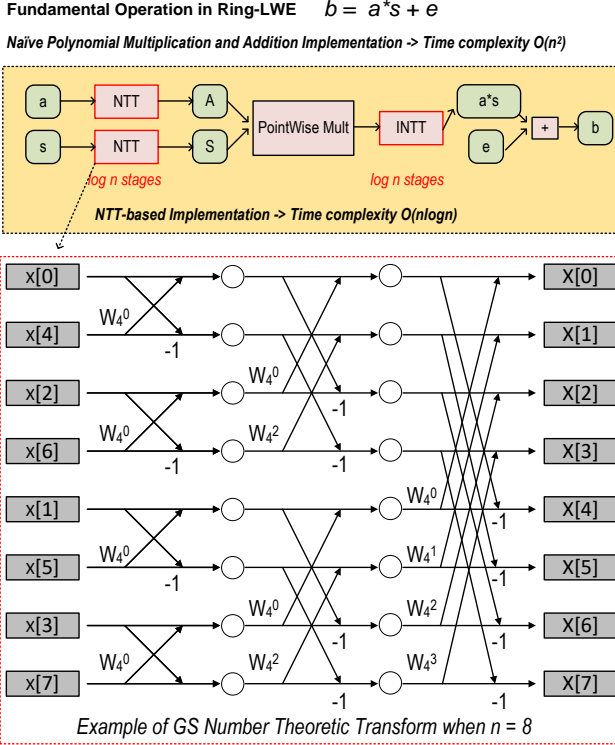

Fig. 1. The basic scheme of lattice-based cryptography.

hardware. This paper thus focuses on NTT acceleration in Ring LWE With n being power of 2, time complexity for computing NTT is $O(n \log n)$. Theoretically, n/2 modulo multiplications can be performed in parallel for each NTT. However, it requires accessing elements in parallel, which proves to be the bottleneck. Most of the existing NTT designs in FPGA [7] or digital ASIC have been optimizing the dataflow to enhance the overall performance and efficiency for a certain configuration of butterfly units and memory banks. The hardware configuration underlines a fundamental trade-off between area and efficiency. For example, [8], [9] minimize the area overhead by sequentially accessing and sharing a butterfly unit, while [10] attempt to layout and operate several butterfly units in parallel. Both approaches could not completely address the issue of designing a compact and parallel accelerator.

Processing-in-Memory (PIM) [11] is an emerging technology for memory-constrained computation, owing to its capabilities of highly parallel computing with amortized energy for memory accesses and logic operations. PIM also enables enhanced data locality, avoiding frequent data transfers to and from the computing units. These additional capabilities, while retaining the compact nature of the memory, make it a promising technology for accelerating NTT [12].

While PIM accelerators based on beyond-CMOS memory devices, such as ReRAM or MRAM, hold great promise for future memory-centric computing with superb density and non-volatility. SRAM-based PIM accelerators solely based on mainstream CMOS technologies undoubtedly represent a clear path towards reliable mass productions and robust operations. SRAM PIM allows low-voltage read, write, and logic operations for energy savings, and could always take advantage of the latest CMOS processes. Further, the larger footprint of SRAMs is amortized by the relatively large PIM peripherals in practical implementations.

In this paper, we present MeNTT, an in-6T-SRAM

Fig. 2. The overall flow of Ring-LWE and polynomial multiplication (NTT).

Algorithm 1 NTT Multiplication with Cooley-Tukey Method

Require: Given polynomial $a \in R_q$, n-th root of unity w_n in Z_q

Ensure: Polynomial $\hat{a} = NTT(a)$ such that $\hat{a} \in R_q$

1:   $\hat{a} \leftarrow PolyBitRev(a)$
2:   **for** (stage =1; stage$\leq \log_2 n$; stage = stage+1) **do**
3:      $m \leftarrow 2^{stage}$
4:      $w_m \leftarrow w_n^{n/m}$
5:      **for** (k =0; $k < n$; k=k+m) **do**
6:        $w \leftarrow 1$
7:        **for** (j =0; $j < m/2$;j=j+1) **do**
8:          $t \leftarrow w.\hat{a}[k+j+m/2] \bmod q$
9:          $u \leftarrow \hat{a}[k+j]$
10:        $\hat{a}[k+j] \leftarrow u + t \bmod q$
11:       $\hat{a}[k+j+m/2] \leftarrow u - t \bmod q$
12:       $w \leftarrow w.w_m \bmod q$
13:      **end for**
14:     **end for**
15: **end for**
16: **return** $\hat{a}$

is a power of 2 [13], q is a prime number. Ring LWE [14] states that it is difficult to find secret vector $\mathbf{s} \in R_q$ given the pair (**a,b**). Here, a*s is a polynomial multiplication, and is performed by transforming both a and b in NTT domain (Fig. 2).

B. Number Theoretic Transform (NTT)

Let a,s are two polynomials sampled from R_q, whose coefficients are in range [0,q) where q is prime number. We denote NTT coefficients of polynomial ($a = a_{n-1}x^{n-1} + a_{n-2}x^{n-2} + ... + a_0$) as $\hat{a}_{n-1}, \hat{a}_{n-2}, .., \hat{a}_0$ respectively.

$$b = INTT((NTT(a) * NTT(s)) \tag{1}$$

Polynomial multiplication of b=a*s is done using eq 1. After transforming both vectors a and s in NTT domain, NTT coefficients of b are calculated by coefficient wise multiplication of a and s. Original coefficients are calculated back by transforming $\hat{b}$ using Inverse NTT.

$$\hat{b} = \sum_{i=0}^{N-1} (\hat{a}_i * \hat{s}_i)x^i \tag{2}$$

There are two popular variants of butterfly optimizations for the acceleration of polynomial multiplication: Cooley-Tukey [15] and Gentleman-Sande. The former is adopted in MeNTT and described in Algorithm 1. INTT computation is similar to NTT except that the twiddle factors are w^{-n} instead of w^n.

C. Processing in Memory for Cryptographic Acceleration

Like many other memory-centric computation problems, such as deep learning, NTT computing faces the "memory walls" because of the energy and throughput bottleneck between logic and memory [16]. To alleviate this problem, a new computing paradigm with in- and near-memory computing

NTT/INTT accelerator with boosted area and energy efficiency. Our key contributions include:

- We develop a novel protocol to perform bit-serial modular addition/subtraction/multiplication arithmetic in 6T SRAM array with massive parallelism and less steps.
- We propose a mapping strategy to optimize the dataflow between NTT and INTT stages and dramatically reduce the routing overhead for large NTT/INTT.
- The MeNTT accelerator is rigorously evaluated in TSMC 65nm LP process, through a combination of transistor-level post layout simulation for the memory and post-layout for digital logic by Design compiler

The rest of the paper is organized as follows. Section II provides the necessary background for Ring LWE and NTT. Section III discusses the implementation of modular arithmetic in 6T SRAM, the data flow between the NTT/INTT stages and routing technique. Section IV shows the evaluation results and comparison with prior arts. Section V concludes the paper.

II. BACKGROUND AND RELATED WORK

This section covers mathematical background of the Ring LWE problem and the NTT/INTT algorithm based polynomial multiplication. It also provides a brief overview on recent PIM research, especially those for cryptographic acceleration.

A. Ring Learning with Errors (LWE)

Let the pair of vectors (**a,b**) be related by the equation $b = a * s + e$, where **a** is a randomly sampled vector from R_q, **e** is a error vector sampled from a Gaussian distribution. Here $R_q = Z_q[x]/(x^n + 1)$ is a ring of polynomial where n

A_i	B_i	$BL_i = A_i \& B_i$	$BLB_i = \sim(A_i \mid B_i)$
0	0	0	1
0	1	0	0
1	0	0	0
1	1	1	0

Truth Table of Processing-in-SRAM

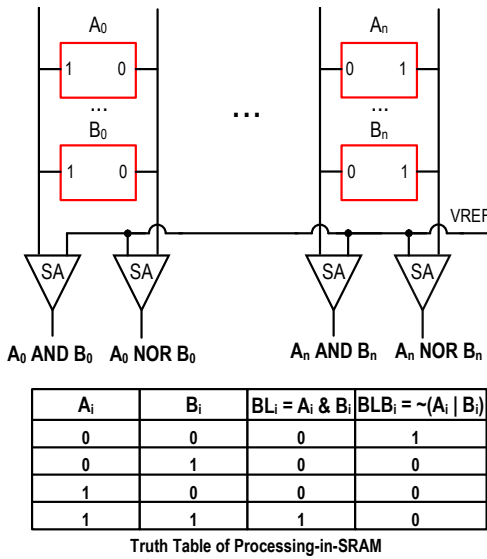Fig. 3. Basic bit-wise logic operations in 6T SRAM by accessing two rows simultaneously [25].

emerged to reduce data movement, amortize memory access energy, and energy-efficient mixed-signal logic operation within a memory array [11].

In cryptographic computing, the requirement on PIM is different from that of machine learning applications because of its zero tolerance to compute errors. Thus, bit-parallel and bit-serial operations are more suitable than the lossy computing mechanisms in current, charge, or voltage domains [17], [18], [19], [20]. Recent Digital in-SRAM architectures [21], [22] have been designed to compute with full precision and high parallelism. However, these architectures are designed to compute MAC operations and are not suitable to output multiple modular arithmetic in parallel. Bit-serial in-SRAM logic performed by accessing two words simultaneously (Fig. 3) is first utilized to modern cryptography accelerators by [23], which performs very wide word bitwise logic and finite field arithmetic in memory. However, it does not provide the high parallelism expected for NTT acceleration. Performing arithmetic in bit-serial fashion was introduced in [24], which is promising for NTT and modular arithmetic because it achieves full accuracy computing with massive parallelism, by spending more clock cycles in each arithmetic operation. [12] leveraged a similar technique with projected high-performance and high-density ReRAM devices. It employs an unfolded architecture with straightforward control and routing schemes, which is possible with high-density ReRAM devices and slim peripherals, but will lead to an unacceptably large area if implemented on SRAM. On the other hand, recent ASIC NTT accelerator [10] has tried computing multiple butterfly operations in parallel and moving buffers closer to the computing units. However,the extensive use of registers for local buffering comes with a high area overhead. Thus, achieving the desired in-memory computing performance and efficiency gains with compact physical footprint and cost is the primary goal of this work.

III. Data Storage and Arithmetic Flow in MeNTT

MeNTT performs complete polynomial multiplication com-

putation in NTT and INTT with all the modular arithmetic steps in and near a 6T SRAM array, as shown in Fig. 4. The SRAM array functions both as data storage and a computing unit. Inter-column router route data at the end of one stage into corresponding columns for the computation of the next stage. SRAM peripheral and controller are specifically designed for in- and near-memory modular arithmetic. During each stage in the radix-2 NTT, each column works like a butterfly unit, thus enabling massively paralleled computing to support a large number of points. The operands and intermediate results are all stored in the same array with an allocation strategy shown in Fig. 5. As a result, the width n of the array represents the maximum number of points of NTT, while the number of rows is related to the supported bit width, as shown in Fig. 5. The modular addition, subtraction, and multiplication are performed in a bit-serial manner, similar to the generic bit-serial logic achieved in [24], but with significantly reduced steps and energy-optimized for modular arithmetic. The high parallelism enabled by the bit-serial approach improves the overall performance and energy efficiency of NTT operations with a large number of points. In our implementation, a single SRAM bank with 162 by 1024 cells is designed to support 1024-point and 32-bit NTT operations. The physical layout of the wide SRAM array can be folded as shown in Fig. 4, in order to reduce word line and inter-column routing length, and maintain a proper aspect ratio. 6T SRAM is desired because of its maturity and high density.

Modular arithmetic differs from regular arithmetic as it pays more attention to overflow and underflow issues. Most previous works take the strategy to calculate the result in an integer field, then reduce it to the desired finite field. Barrett reduction and Montgomery reduction (Algorithm 2, Algorithm 3) are the two most common approaches to reduce numbers into a finite field. MeNTT proposes a reduce-on-the-fly technique for modular addition, subtraction, and multiplication.

A. Near-Memory Bit-serial Comparator and Reduction

We execute addition or accumulation from LSB to MSB in our bit-serial arithmetic operation. In order to keep the final result within field range, our reduction scheme utilizes the bit-serial comparator (Fig. 6) to keep track of the overflow condition of the result from the last operation. The comparison between temporary result and field limit 'q' starts from LSB to MSB cycle by cycle as bit-serial addition, subtraction or multiplication takes place. Reduction is applied based on the comparison result (Fig. 4) by applying a subtraction logic together with the adder. Only when reduction is required, the subtraction of q is enabled. For addition and subtraction, a 'raw' result is calculated initially to evaluate an overflow. Then a 'real' step is done to calculate the actual result with reduction. For multiplication, the partial sum is compared with q, while reduction is applied in the next round of accumulation with proper scaling.

B. In-Memory Modular Addition/Subtraction

The proposed modular addition and subtraction consist of a regular round of addition/subtraction with overflow/underflow
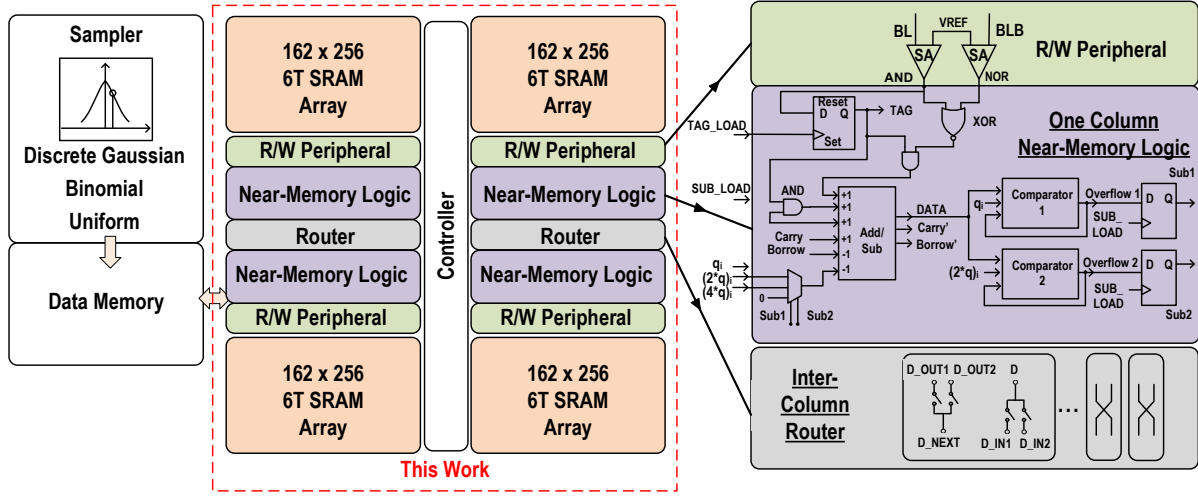
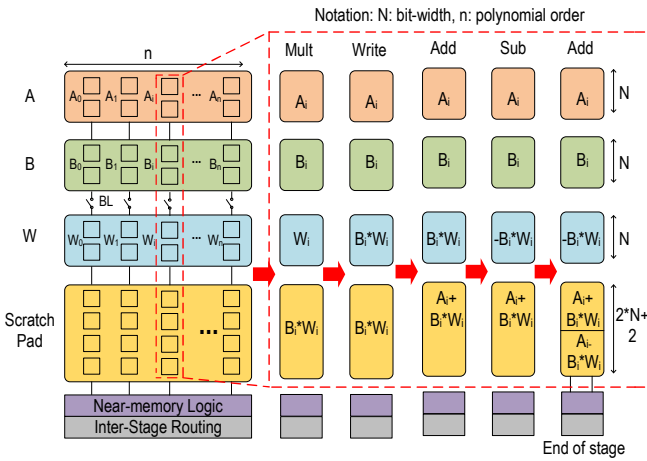Fig. 4. Diagrams of MeNTT and custom circuitry for the key peripherals.

Fig. 5. Data arrangement and operation sequence in an NTT stage.
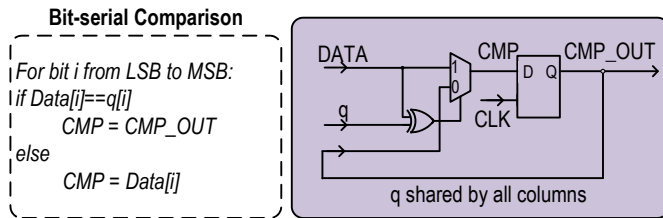

Fig. 6. Bit-serial comparator operation in column peripheral.

detection called "Trial Add and Compare Phase", and a round of modular reduction called "Modular Addition Phase", as depicted in (Fig. 7). The WL driver activates corresponding bits of two operands A and B sequentially. The value on BL will be A AND B, while the value on BLB will be A NOR B. These values are read using traditional sense-amplifier and given as inputs to near-memory. The peripheral maintains one bit Carry for the addition, and acts like a full adder which adds A and B up sequentially and writes back to the array. In "Trial Add and Compare Phase", in the meantime of addition, the peripheral reads one bit of q, and makes a comparison sequentially. An overflow bit is set if the sum of A and B is

Algorithm 2 Barrett Modular Multiplication

Require: : $x, y \in Z_q$, q, m and k such that $m = \lfloor 2^k/q \rfloor$
Ensure: z = (x.y) mod q
1: $z \leftarrow x \cdot y$
2: $t \leftarrow (z \cdot m) \gg k$
3: $z \leftarrow z - (t \cdot q)$
4: **if** $(z \geq q)$ **then**
5: $\quad z \leftarrow z - q$
6: **end if**
7: **return** z

larger than q. In that case, column will calculate A+B-q during "Modular Addition Phase". Otherwise, column only calculates A+B to avoid timing side-channel leakage

Modular subtraction is done similar to modular addition by transforming one of the input B into it's 2's complement first. This can be done by setting the initial carry to 1 and using the value from BLB as input. During "trial add and compare" phase after 2's complement, an underflow bit is recorded. Based on the underflow bit value, addition of q is done during second-round subtraction to keep the result value correct. Addition requires 2*(N+1) cycles while the subtraction requires an extra N+1 cycles for 2's complement calculation.

C. In-Memory Bit-Serial Modular Multiplication

Modular multiplication is the most time-consuming part of the NTT. Traditionally, modular multiplication is performed by computing the raw product first, then going through a Barrett or Montgomery reduction(Fig. 9). In previous PIM works [24], this involves three normal multiplications and also takes extra redundant cycles as well as space. Although in some special cases [12] this reduction can be simplified, there is no general optimization for the modular multiplication using traditional reduction approaches.

This work proposes a fast bit-serial multiplication scheme to complete a modular multiplication in $(N+1)^2$ cycles (Fig. 8). The overall algorithm is described in Algorithm 5. The
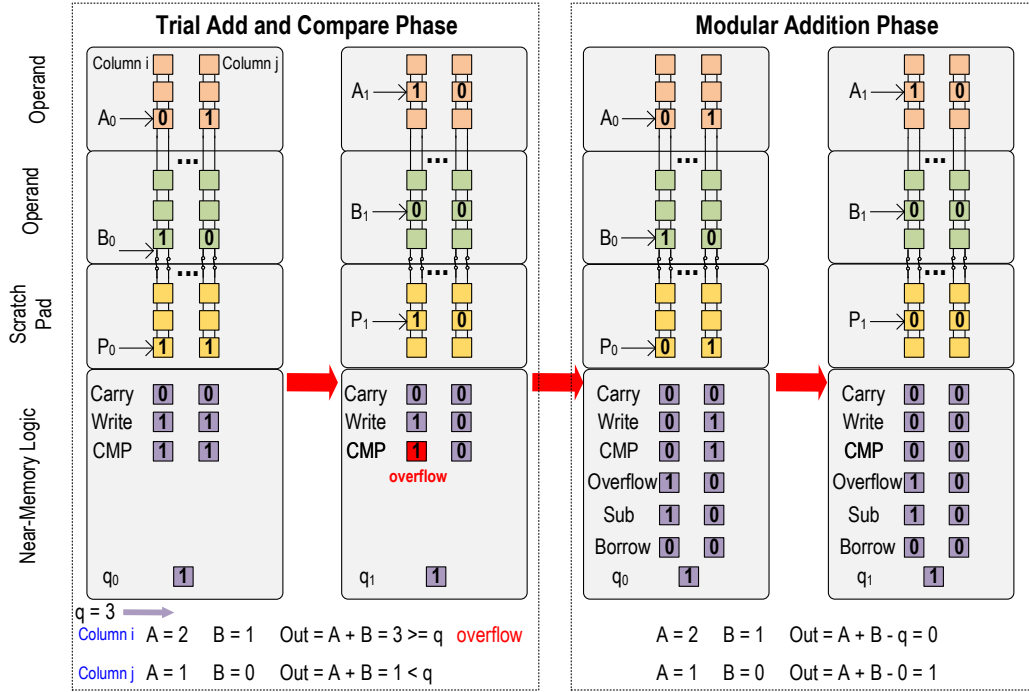
Fig. 7. Illustration of in-memory two-bit modular addition.

Algorithm 3 Montgomery Modular Multiplication

Require: : $x, y \in Z_q$, q, r and k such that $r > q$, $gcd(r, q) = 1$, $k = \frac{r(r^{-1} \bmod n) - 1}{n}$

Ensure: : $c = x \cdot y \bmod q$

1: $\bar{a} = a \cdot r \bmod q$
2: $\bar{b} = b \cdot r \bmod q$
3: $x = \bar{a} \cdot \bar{b}$
4: $s = x \cdot k \bmod r$
5: $t = x + s \cdot q$
6: $u = \frac{t}{r}$
7: **if** $(u \geq q)$ **then**
8:     $\bar{c} = u - q$
9: **else**
10:    $\bar{c} = u$
11: **end if**
12: $c = (\bar{c} \cdot r^{-1} \bmod n)$
13: **return** c

Algorithm 4 Modular Addition

Require: $x, y \in Z_q$

Ensure: $z = x + y \bmod q$

1: Trial Add and Compare Phase:
2: $s \leftarrow x + y$
3: $cmp \leftarrow s \geq q$
4: Modular Addition Phase:
5: **if** $(cmp = 1)$ **then**
6:    $z \leftarrow x + y - q$
7: **else**
8:    $z \leftarrow x + y$
9: **end if**
10: **return** z

---

multiplication is decomposed into shift and add operations. There is a Tag bit which is similar to the approach of [24]. The shift is done by choosing a different address for rows in each round. The addition is controlled by the Tag bit from operand B and computed in the peripheral with carry and borrow. We take advantage of the bit-serial computation flow and compare every partial sum with q and 2*q. A reduction will be performed in the next cycle to make sure the partial sum is contained in the correct range. The reduction is enabled by the subtraction and borrow circuits, and controlled by the two overflow bits. In this way, the modular multiplication can be completed in parallel with the shift-and-add itself with very little space and time overhead. In order for the algorithm to be

valid, operand A has to be smaller than q/2. Therefore we use 2*q instead of q as the field constant. Then MeNTT reduce the partial sum below q with an extra cycle. A switch is controlled by the Tag to separate the operand A from influencing the BL readout result. It decides whether the partial sum is added by an operand or kept the same. The reduction with q or 2q is done in peripheral with simple digital logic shown in Fig. 6.

D. NTT and INTT Dataflow

Fig. 2 shows the CT radix-2 style NTT for the acceleration of polynomial multiplication in a ring. Each stage of NTT contains groups of butterfly operation which can be decomposed to a multiplication, an addition and a subtraction. Each stage has own grouping of points depending on the index.

1) Intra-stage flow: In a single stage, the data arrangement is shown in Fig. 5. The N-bit operands A and B are stored sequentially in same column followed by the twiddle factor

Algorithm 5 Modular Multiplication

Require: $x, y \in Z_q$, where y_k represents k-th bit of N-bit y
Ensure: z = (x.y) mod q

1: $psum = 0, overflow\_4q = 0, overflow\_2q = 0$
2: **for** $(k = N - 1; k \geq 0; k = k - 1)$ **do**
3:     **for** $(j = 0; j < N; j = j + 1)$ **do**
4:         **if** $(overflow\_4q = 1)$ **then**
5:             $psum_j \leftarrow (psum_j \ll 1) + (x_j * y_k) - q_{j-2} +$
            $carry - borrow$
6:         **else if** $(overflow\_2q = 1)$ **then**
7:             $psum_j \leftarrow (psum_j \ll 1) + (x_j * y_k) - q_{j-1} +$
            $carry - borrow$
8:         **else**
9:             $psum_j \leftarrow (psum_j \ll 1) + (x_j * y_k) + carry -$
            $borrow$
10:         **end if**
11:         $write_j = LSB(psum_j)$
12:         $carry = (psum_j > 1)?1 : 0$
13:         $borrow = (psum_j < 0)?1 : 0$
14:         $cmp_{2q} = (psum_j == q_{j-i})?cmp_{2q} : psum_j$
15:         $cmp_{4q} = (psum_j == q_{j-2})?cmp_{2q} : psum_j$
16:     **end for**
17:     **if** $(cmp_{4q} == 1)$ **then**
18:         $overflow\_4q \leftarrow 1$
19:     **else**
20:         $overflow\_4q \leftarrow 0$
21:     **end if**
22:     **if** $(cmp_{2q} == 1)$ **then**
23:         $overflow\_2q \leftarrow 1$
24:     **else**
25:         $overflow\_2q \leftarrow 0$
26:     **end if**
27: **end for**
28: **if** $(psum \geq q)$ **then**
29:     $psum \leftarrow psum - q$
30: **end if**
31: $z \leftarrow psum$
32: **return** z
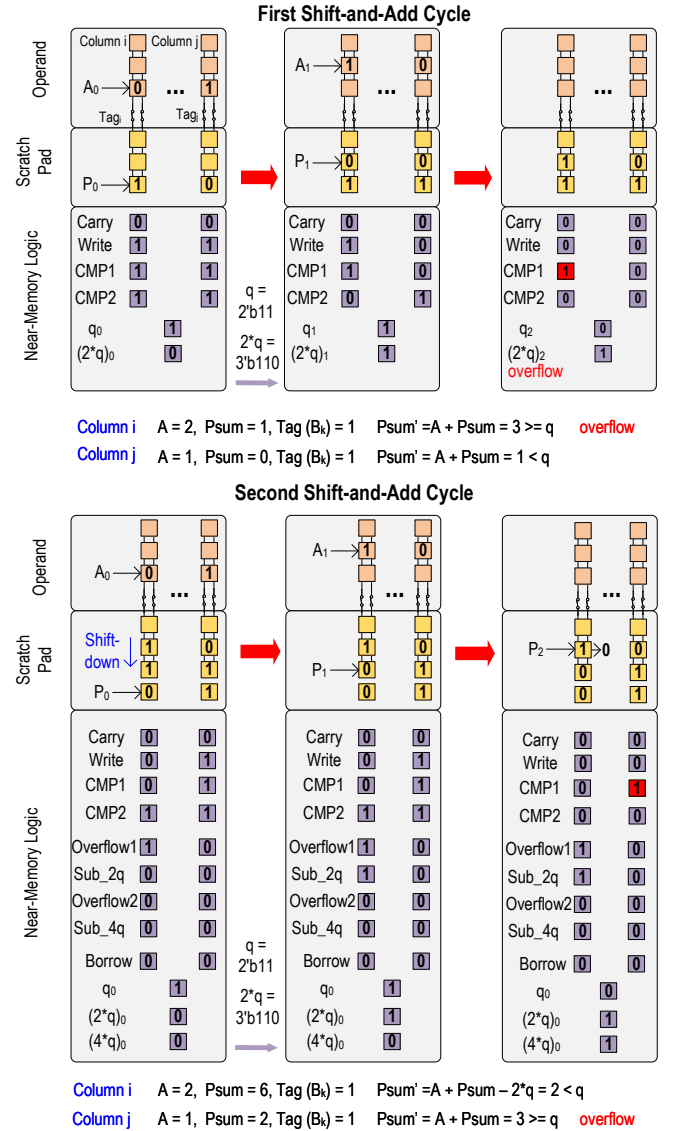
Fig. 8. Illustration of in-memory 2-bit modular multiplication (q is 2'b10, Tag is 1 for two cycles).

Fig. 9. Comparison of conventional modular multipliers and the proposed one in MeNTT.

W. The results of addition, subtraction and multiplication are computed and stored in the scratchpad area. Therefore the temporary results for an NTT stage can be computed in sequential order and updated into the operand area for follow-up computations. All the operations addition, subtraction, copying, inverting and multiplication required for a stage can be completed inside a single column in a bit-serial manner. As the algorithm for modular multiplication requires extra bits for each column, the size of scratchpad is set to 2*N+2. This space can be fully utilized to store the final outputs of a stage i..e, A+W*B and A-W*B.

2) Inter-stage data movement: Once the computation of a single stage is done, the data needs to be read out and written into the operands area for next stage. Traditional in-place NTT has complicated data routing, varying from stage to stage. This makes the data movement different in different stages of NTT and INTT. A configurable crossbar routing is required to enable various address matching, adding to the area and energy overheads. In MeNTT we take advantage of unconventional data arrangement in SRAM for routing. Instead of reading column-wise, we read in traditional row-wise fashion i.e, one bit of each output. We switch these output bits and write back to different addresses. We propose a mapping strategy that makes this data movement between stages constant. Fig. 10 shows a complete data movement flow in a polynomial multiplication including a stage to stage routing in NTT, point-wise multiplication and INTT. The data is stored in a specific column with address. The physical
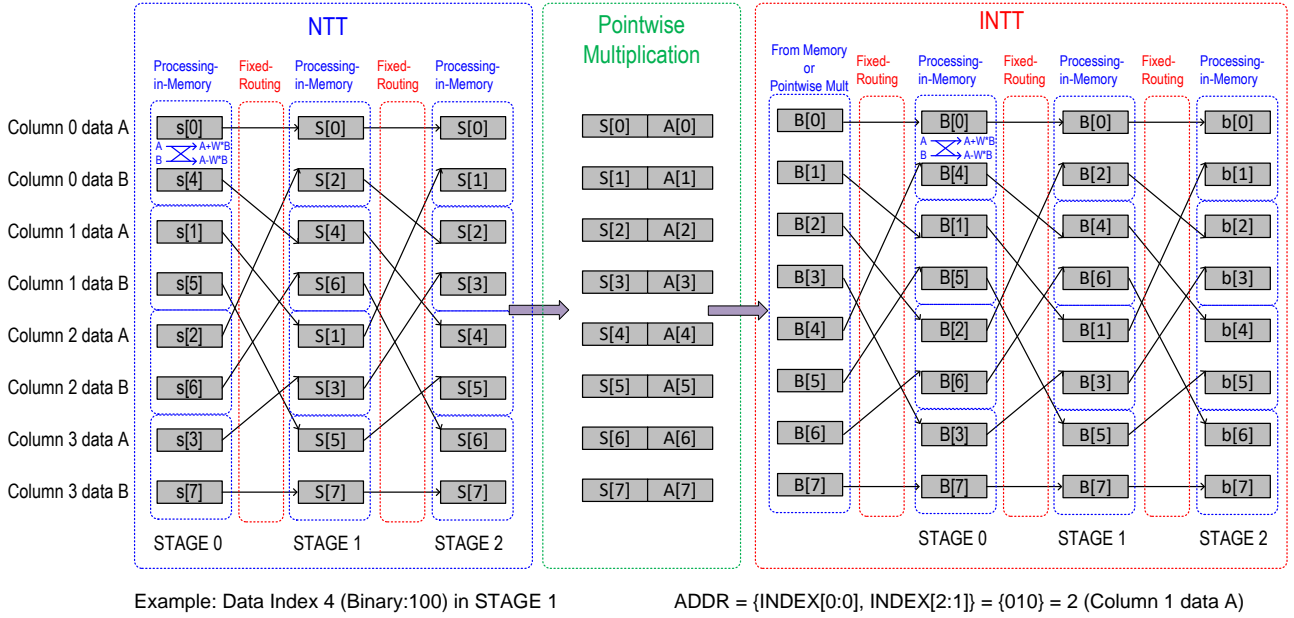
Fig. 10. MeNTT dataflow (for computing b = INTT(NTT(A)·NTT(S)) = a*b) with a novel mapping strategy.

address is a function of the current stage (s) and its original index, as described below:

$$addr = \{index[log_2 n - s - 1 : 0], index[msb : log_2 n - s]\} \quad (3)$$

For example, data S[4] in Fig. 10 has index '100'. In stage 1 of NTT, the actual address is {'0','10'} according to the mapping, which leads to address of 2, translating to Column-1, data A. The key observation for this approach is that the actual physical output-input address mappings are the same for each stage. In the example of 8-point NTT, the output of address 4 (Column 2 data A) always goes to the input address 1 (Column 0 data B) in every stage. Therefore the crossbar connections can be reduced to simple switches. This schedule make it possible to construct a single-bank processing-in-memory block to compute the NTT process iteratively. MeNTT takes 4*N cycles for routing between stages, where N is bit width of the operand. In conventional digital architectures which use traditional SRAM and read row-wise, these cycles depend linearly on higher valued polynomial order.

INTT follows a similar mapping strategy and shares the same physical routing. The point-wise multiplication between NTT and INTT can be completed in place with precomputed NTT result of another polynomial. In such a scheme, the whole polynomial multiplication operation is performed in a single SRAM bank with bit-wise modular arithmetic and constant stage to stage routing. The MeNTT therefore utilizes memory reuse for optimal area efficiency and memory bandwidth, making it suitable for resource-constraint applications.

IV. EVALUATION AND DISCUSSION

We evaluated MeNTT in TSMC 65nm LP CMOS. We implemented the digital circuits design using verilog followed by synthesis, auto-place and route (APR) by Synopsys Design Compiler. We evaluated the design using post-layout SPICE simulation of SRAM and digital circuits to estimate the overall
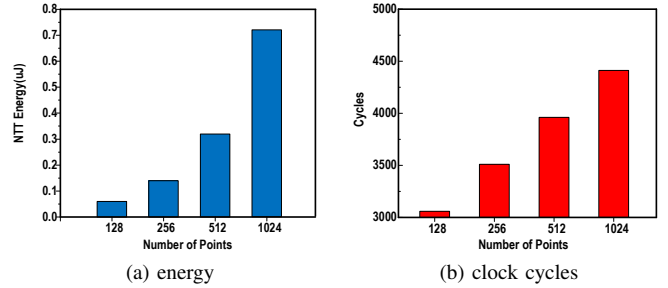

Fig. 11. MeNTT energy and clock cycles for different number of points at N=14b.

system accurately. Different configurations of bit width and polynomial order are evaluated to compare our work with prior arts for different protocols. Evaluation results are shown in Fig. 11, Fig. 12 and Fig. 13. MeNTT is configurable for different bit width and polynomial order. Column-wise power-gating makes it possible to accommodate NTT of smaller polynomial order with energy saving. The read/write address and peripheral instructions are sent from the control block. Thanks to the routing schedule, the 6-T SRAM for PIM can process arbitrary polynomial order within the size constraint. Good scalability is maintained as SRAM can be extended to more columns in MeNTT. Table I compares MeNTT with software, FPGA, ASIC, and previous PIM designs for NTT computation. For FPGA works, the energy and area are normalized to 65nm technology based on the gate count reported in the original paper combined with our SRAM area and energy model in 65nm node. For ASIC works, the area and energy are normalized to 65nm technology by scaling from the original technology node. For software works, the energy consumption is evaluated by the corresponding processor's reported latency, frequency, and power.
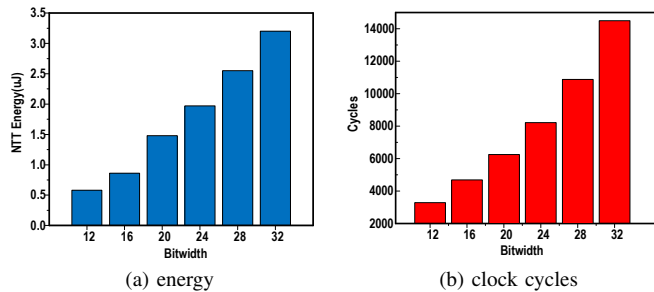
Fig. 12. MeNTT energy and clock cycles for different bit widths at a polynomial order of 1024.

Fig. 13. MeNTT area and energy values breakdown for different modules at n=1024,N=32b.

A. Comparisons with software solutions

As discussed in previous sections, traditional software implementation has an obvious bottleneck for Ring-LWE computations with the increase of polynomial order and coefficient's bit width. The energy cost and latency are higher than hardware approaches by several orders. Under current Von-Neumann architecture, the highly parallel butterfly operations in NTT can only be executed in a serially. While modular arithmetic usually takes only one or few cycles each, the data transfer between memory and processor involves, cache and main memory read/write, which can cost 10s of cycles, making the whole process extremely expensive. From Table I, one can observe the significant energy and latency overhead in X86 based software approach. Despite the fact that a CPU can run at a higher speed than custom hardware, it is not the perfect solution for low-power and high-performance security applications. Also, software solutions are not well guarded against side-channel attacks. Since all operations are executed in ALU in serial, it is relatively easy for an attacker to retrieve power and timing information leaks for a side-channel attack. Due to the low efficiency of memory access in software approaches, hardware acceleration is widely desired and adopted for FFT and NTT applications.

B. Comparison with FPGA solutions

FPGA-based hardware approaches exhibit improved performance, thanks to the customized architecture and datapath to elevate parallelism and efficiency for NTT computation. To support the modular arithmetic in NTT flow, DSP blocks [7] or custom multipliers are required to provide sufficient computing capability. Optimized datapaths are designed to accommodate the stage by stage NTT and INTT in a pipelined [26] approach which achieved higher throughput with a considerable area overhead. As shown in Table I, FPGA solutions outperform the software approach in throughput and latency, but their energy and area costs are significantly higher than custom hardware. Therefore FPGA is more generally more suitable for prototyping and deployment in the cloud.

C. Comparison with ASIC solutions

The disadvantages of software and FPGA approaches inspired exploration in custom hardware solutions. ASIC implementations 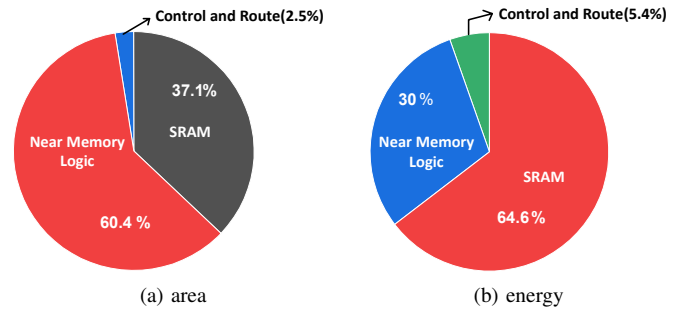make use of standard SRAM or registers and compute modular arithmetic in digital circuits [8]. Although ASICs usually perform better than FPGAs in terms of speed, energy, and area, they come with much higher costs. Another major drawback for this approach is the amortized BL energy spent on reading data from SRAM. The processing speed is also limited by traditional memory bandwidth. While embedding registers for local storage in each computing unit increases throughput by increasing parallelism as in [10], the area overhead and limited scaling potential are the main drawbacks.

In terms of scalability to larger bit width and q, MeNTT executes modular arithmetic in bit-serial and word-parallel order while traditional software, FPGA, and ASIC approaches carry out butterfly unit operations word by word. Thus, MeNTT will achieve even higher throughput and energy efficiency than the other solutions, when the polynomial order gets higher.

ASIC solutions require higher design and fabrication cost, and need extra overhead to handle different schemes with different polynomial order, q and bit width. MeNTT provides higher degree of configurability by providing general modular arithmetic in a highly parallel computing approach. The computation and data movement can be reprogrammed with minimal effort by changing WL accessing sequence and peripheral configurations. The SRAM rows and columns can be gated in different cryptography schemes for higher throughput and energy savings. MeNTT also provides higher normalized energy and area efficiency, as shown in Table I, mainly benefiting from in- and near-memory computation.

D. Comparisons with existing PIM solutions

Compared to previous PIM studies that focus on leveraging parallel word-serial or bit-serial operation for general-purpose arithmetic in SRAM [23] and ReRAM [12], [24], optimized modular arithmetic and dataflow help MeNTT outperforms prior PIM works in energy and area efficiency as well as latency. MeNTT masks the inherent modular reduction cycles in the modular multiplication operation itself, whereas other designs use Barret or Montgomery reduction techniques which cause an extra area and latency overheads.

While [12] achieved higher throughput by introducing multiple pipelines to break the data path into shorter pieces, our approach focuses on designing ultra-compact single SRAM bank implementation, which is more applicable to resource-constrained applications. Note that the reported latency for

TABLE I
COMPARISON TABLE WITH STATE-OF-THE-ART NTT ACCELERATORS.

Method	This Work				ISSCC 2019 [8]			CICC 2018 [10]		DAC 2020 [12]			ICASSP 2020 [26]			DAC 2020 [12]		
Method	Bit-serial 6T SRAM				Serial Butterfly			Parallel Butterfly		Bit-serial RRAM			FPGA			Software		
Technology	65nm				40nm			40nm		RRAM (45nm)			Virtex-6 (40nm)			X86 (gem5)		
Frequency (MHz)	151				72			300		909			184			2000		
Bitwidth	14	14	14	16	13	13	14	13	14	16	16	16	16	16	16	16	16	16
Polynomial Order	256	512	1024	1024	256	512	1024	256	512	256	512	1024	256	512	1024	256	512	1024
Supported Bitwidth and Polynomial orders	n $\leq$32K, N $\leq$32				n $\leq$1K, N $\leq$24			n $\leq$2K, N$\leq$32		n $\leq$32K, N $\leq$32			n $\leq$1K, N $\leq$16			-		
Latency (us)	23	26	29	34.3	17.9	39.2	85.5	0.533	1.6	68.7	75.9	83	21.6	47.6	102	85	169	349
Throughput (NTT/sec)	42k	38k	35k	29k	56k	25k	12k	1.8M	609k	553k	553k	553k	46k	21k	10k	12k	5.9k	2.9k
Energy (nJ/NTT)	144	324	720	868	166	411	894	31	96	2580	5020	10040	2150	5280	12520	570k	1179k	2483k
Normalized Energy (nJ/NTT)[a]	144	324	720	868	780	1932	4202	146	451	2580	5020	10040	2150	5280	12520	-	-	-
Area (mm²)	0.36				0.28			1.4		-			-			-		
Normalized Area (mm²)[a]	0.36				0.74			3.7		-			-			-		

a. Normalized to 65nm LP Process (nominal VDD at 1.2V), assuming ideal scaling of transistor capacitance and area
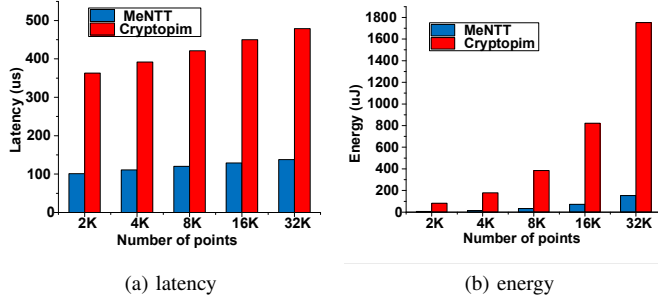

(a) latency (b) energy

Fig. 14. Comparison of latency and energy values for MeNTT vs Cryptopim for different polynomial orders at 32 bit width.

[12] is for specially selected q with very small hamming weight. The latency varies a lot with different choices of schemes and q. Other generic choices of q will incur high latency overhead. Last but not least, while the MeNTT design is implemented and evaluated with CMOS devices and SRAMs, the proposed modular arithmetic protocol and dataflow techniques are generic to any bit-serial PIM fabrics and can be easily adopted in PIM with emerging memories.

V. CONCLUSION

In conclusion, this paper presents MeNTT, a novel PIM architecture for NTT acceleration. With the proposed bit-serial modular arithmetic protocol and mapping strategy, it achieves superior efficiency and throughput with a compact footprint. A fully functional mixed-signal implementation of the system verifies its feasibility in physical design, and provides realistic estimation of its performance for comparison.

REFERENCES

[1] P. W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer," *arXiv:quant-ph/9508027*, Jan. 1996, arXiv: quant-ph/9508027.

[2] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *J. ACM*, vol. 56, no. 6, Sep. 2009.

[3] V. Lyubashevsky, C. Peikert, and O. Regev, "On Ideal Lattices and Learning with Errors over Rings," in *Advances in Cryptology – EUROCRYPT 2010*, H. Gilbert, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 1–23.

[4] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–297, May 1965.

[5] A. C. Mert, E. Öztürk, and E. Savaş, "Design and Implementation of Encryption/Decryption Architectures for BFV Homomorphic Encryption Scheme," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 28, no. 2, pp. 353–362, Feb. 2020.

[6] D. Li, Y. He, A. Pakala, and K. Yang, "MePLER: A 20.6-pJ side-channel-aware in-memory CDT sampler," in *2021 Symposium on VLSI Circuits*, 2021.

[7] T. Pöppelmann and T. Güneysu, "Area optimization of lightweight lattice-based encryption on reconfigurable hardware," in *2014 IEEE International Symposium on Circuits and Systems (ISCAS)*, Jun. 2014, pp. 2796–2799, iSSN: 2158-1525.

[8] U. Banerjee, A. Pathak, and A. P. Chandrakasan, "2.3 An Energy-Efficient Configurable Lattice Cryptography Processor for the Quantum-Secure Internet of Things," in *2019 IEEE International Solid- State Circuits Conference - (ISSCC)*, Feb. 2019, pp. 46–48, iSSN: 2376-8606.

[9] S. S. Roy, F. Vercauteren, N. Mentens, D. D. Chen, and I. Verbauwhede, "Compact Ring-LWE based Cryptoprocessor," Tech. Rep. 866, 2013.

[10] S. Song, W. Tang, T. Chen, and Z. Zhang, "LEIA: A 2.05mm² 140mW lattice encryption instruction accelerator in 40nm CMOS," in *2018 IEEE Custom Integrated Circuits Conference (CICC)*. San Diego, CA: IEEE, Apr. 2018, pp. 1–4.

[11] N. Verma, H. Jia, H. Valavi, Y. Tang, M. Ozatay, L. Chen, B. Zhang, and P. Deaville, "In-Memory Computing: Advances and Prospects," *IEEE Solid-State Circuits Magazine*, vol. 11, no. 3, pp. 43–55, 2019, conference Name: IEEE Solid-State Circuits Magazine.

[12] 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.

[13] J. Buchmann, F. Göpfert, T. Güneysu, T. Oder, and T. Pöppelmann, "High-Performance and Lightweight Lattice-Based Public-Key Encryption," in *Proceedings of the 2nd ACM International Workshop on IoT Privacy, Trust, and Security - IoTPTS '16*. Xi'an, China: ACM Press, 2016, pp. 2–9.

[14] V. Lyubashevsky, "Lattice signatures without trapdoors," *EUROCRYPT 2012*, p. 738.

[15] S. S. Roy, F. Vercauteren, N. Mentens, D. D. Chen, and I. Verbauwhede, "Compact Ring-LWE Cryptoprocessor," in *Cryptographic Hardware and Embedded Systems – CHES 2014*, L. Batina and M. Robshaw, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 371–391.

[16] M. Horowitz, "1.1 Computing's energy problem (and what we can do about it)," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb. 2014, pp. 10–14.

[17] Z. Chen, Z. Yu, Q. Jin, Y. He, J. Wang, S. Lin, D. Li, Y. Wang, and K. Yang, "CAP-RAM: A Charge-Domain In-Memory Computing 6T-SRAM for Accurate and Precision-Programmable CNN Inference," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 6, pp. 1924–1935, Jun. 2021.

[18] M. Kang, M.-S. Keel, N. R. Shanbhag, S. Eilert, and K. Curewitz, "An energy-efficient VLSI architecture for pattern recognition via deep embedding of computation in SRAM," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 8326–8330.

[19] H. Valavi, P. J. Ramadge, E. Nestler, and N. Verma, "A 64-Tile 2.4-Mb In-Memory-Computing CNN Accelerator Employing Charge-Domain Compute," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 6, pp. 1789–1799, Jun. 2019.

[20] A. Shafiee, A. Nag, N. Muralimanohar, R. Balasubramonian, J. P. Strachan, M. Hu, R. S. Williams, and V. Srikumar, "Isaac: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars," in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, 2016, pp. 14–26.

[21] H. Kim, T. Yoo, T. T.-H. Kim, and B. Kim, "Colonnade: A Reconfigurable SRAM-Based Digital Bit-Serial Compute-In-Memory Macro for Processing Neural Networks," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 7, pp. 2221–2233, Jul. 2021.

[22] Y.-D. Chih, P.-H. Lee, H. Fujiwara, Y.-C. Shih, C.-F. Lee, R. Naous, Y.-L. Chen, C.-P. Lo, C.-H. Lu, H. Mori, W.-C. Zhao, D. Sun, M. E. Sinangil, Y.-H. Chen, T.-L. Chou, K. Akarvardar, H.-J. Liao, Y. Wang, M.-F. Chang, and T.-Y. J. Chang, "16.4 An 89TOPS/W and 16.3TOPS/mm2 All-Digital SRAM-Based Full-Precision Compute-In Memory Macro in 22nm for Machine-Learning Edge Applications," in *2021 IEEE International Solid- State Circuits Conference (ISSCC)*, vol. 64, Feb. 2021, pp. 252–254.

[23] Y. Zhang, L. Xu, K. Yang, Q. Dong, S. Jeloka, D. Blaauw, and D. Sylvester, "Recryptor: A reconfigurable in-memory cryptographic cortex-m0 processor for iot," in *2017 Symposium on VLSI Circuits*, 2017, pp. C264–C265.

[24] C. Eckert, X. Wang, J. Wang, A. Subramaniyan, R. Iyer, D. Sylvester, D. Blaaauw, and R. Das, "Neural Cache: Bit-Serial In-Cache Acceleration of Deep Neural Networks," in *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, Jun. 2018, pp. 383–396, iSSN: 2575-713X.

[25] S. Jeloka, N. B. Akesh, D. Sylvester, and D. Blaauw, "A 28 nm Configurable Memory (TCAM/BCAM/SRAM) Using Push-Rule 6T Bit Cell Enabling Logic-in-Memory," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 4, pp. 1009–1021, Apr. 2016.

[26] H. Nejatollahi, S. Shahhosseini, R. Cammarota, and N. Dutt, "Exploring Energy Efficient Quantum-resistant Signal Processing Using Array Processors," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020, pp. 1539–1543, iSSN: 2379-190X.
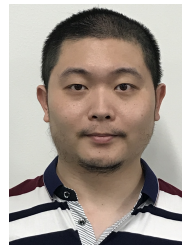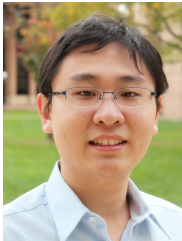
Dai Li received the B.S. and M.S. degrees in electronic engineering from Tsinghua University, Beijing, China, in 2010 and 2013, respectively, and the PhD degree in electrical and computer engineering from Rice University, Houston, TX, USA, in 2021. He is now at Google.

His research interests include very large-scale integration (VLSI) circuits, hardware security, mixed-signal integrated circuits, and low-power circuits.

Akhil Pakala received his Bachelor's and Master's degree from the Indian Institute of Technology Madras, Chennai, India, in 2019. From 2019-2020, he has worked with Samsung Semiconductor Research and Development, India on SerDes PHY IP. He is currently pursuing his doctoral degree at Rice University, Houston, Texas.

His current research interests include designing digital and mixed-signal circuits for security and machine learning applications.

Kaiyuan Yang (S'13-M'17) received the B.S. degree in Electronic Engineering from Tsinghua University, Beijing, China, in 2012, and the Ph.D. degree in Electrical Engineering from the University of Michigan, Ann Arbor, MI, in 2017. His Ph.D. research was recognized with the 2016-2017 IEEE Solid-State Circuits Society (SSCS) Predoctoral Achievement Award.

He is an Assistant Professor of Electrical and Computer Engineering at Rice University, Houston, TX. His research interests include low-power digital/analog/mixed-signal integrated circuit and system design for secure and intelligent micro-systems, hardware security, and bioelectronic applications. Dr. Yang received a number of best paper awards from major conferences in various fields, including the Best Paper Award at the 2021 IEEE Custom Integrated Circuit Conference (CICC), Distinguished Paper Award at the 2016 IEEE International Symposium on Security and Privacy (Oakland), Best Student Paper Award (1st place) at the 2015 IEEE International Symposium on Circuits and Systems (ISCAS), the Best Student Paper Award finalist at the 2019 IEEE Custom Integrated Circuit Conference (CICC), and the 2016 Pwnie Most Innovative Research Award Finalist. His research was selected as "Research Highlight" of Communication of ACM in 2017.