

MultPIM: Fast Stateful Multiplication for Processing-in-Memory

Orian Leitersdorf, Ronny Ronen, *Fellow, IEEE*, and Shahar Kvatinsky, *Senior Member, IEEE*

Abstract—Processing-in-memory (PIM) seeks to eliminate computation/memory data transfer using devices that support both storage and logic. Stateful logic techniques such as IMPLY, MAGIC and FELIX can perform logic gates within memristive crossbar arrays with massive parallelism. Multiplication via stateful logic is an active field of research due to the wide implications. Recently, RIME has become the state-of-the-art algorithm for stateful single-row multiplication by using memristive partitions, reducing the latency of the previous state-of-the-art by $5.1\times$. In this paper, we begin by proposing novel partition-based computation techniques for broadcasting and shifting data. Then, we design an in-memory multiplication algorithm based on the carry-save add-shift (CSAS) technique. Finally, we develop a novel stateful full-adder that significantly improves the state-of-the-art (FELIX) design. These contributions constitute MultPIM, a multiplier that reduces state-of-the-art time complexity from quadratic to linear-log. For 32-bit numbers, MultPIM improves latency by an *additional* $4.2\times$ over RIME, while even slightly reducing area overhead. Furthermore, we optimize MultPIM for full-precision matrix-vector multiplication and improve latency by $25.5\times$ over FloatPIM matrix-vector multiplication.

I. INTRODUCTION

THE von Neumann architecture separates computation and memory in computing systems. Each has significantly improved in recent years, leading to the data-transfer between them becoming a bottleneck (*memory wall* [1]). *Processing-in-Memory* (PIM) aims to nearly eliminate this data-transfer by using devices that support both storage and logic.

Processing-in-memory can be implemented using *memristors* [2], two-terminal devices with variable resistance. Their resistance may represent binary data by being set to either low-resistive state (LRS) or high-resistive state (HRS). A high-density memory can be built using a memristor crossbar array structure [3]. Uniquely, the resistance of a memristor can be controlled via an applied voltage, enabling stateful logic to be performed within the crossbar array. While there remain various challenges with memristive memory and stateful logic, promising ongoing research has experimentally demonstrated stateful logic [4], [5] and proposed solutions for reliable operation [6]–[8]. Therefore, we assume the widely-accepted stateful-logic model [9] and focus on algorithmic aspects. Examples of stateful logic techniques include IMPLY [10], MAGIC [11] and FELIX [12], which can also be performed in parallel along rows/columns. Hence, single-row computation algorithms are advantageous as they can be repeated along

all rows with the exact same latency. Additional parallelism can arise from memristive partitions [12] which dynamically divide the crossbar array using transistors. In this paper, we propose novel partition-based techniques for efficiently broadcasting/shifting data amongst partitions.

Multiplication is fundamental for many applications, *e.g.*, convolution and matrix-multiplication. Initially, only non-crossbar-compatible and non-single-row algorithms [13]–[18] for in-memory multiplication were considered. Yet, these algorithms only support multiplying two numbers per *crossbar*, rather than per *crossbar row* – which would enable paralleled element-wise vector multiplication. The first in-row multiplication algorithm was proposed by Haj-Ali *et al.* [19], and was later utilized in IMAGING [20] for image processing and in FloatPIM [21] for deep neural networks. This algorithm requires $O(N^2)$ latency and $O(N)$ memristors, where N is the width of each number. Recently, RIME [22] improved the latency by $5.1\times$ for $N = 32$ via memristive partitions [12], while slightly reducing area (*i.e.* memristor count) as well. The asymptotic latency/area remains at $O(N^2)$ and $O(N)$ (respectively). RIME is based on Wallace tree computation using $N - 1$ partitions in a single row, each partition representing a full-adder unit. The bottleneck of RIME is the partial product computation and data-transfer between partitions (as they occur serially), accounting for 81% of the latency.

In this paper, we speedup multiplication using three methods. First, we propose novel partition-based computation techniques for broadcasting/shifting data amongst partitions. Second, we replace the Wallace tree with a carry-save-add-shift (CSAS) multiplier [23]–[25]. Lastly, we propose a novel full-adder design that significantly improves the previous state-of-the-art (FELIX [12]). The final algorithm, coined MultPIM, achieves an asymptotic latency of $O(N \log N)$ with $O(N)$ area. For $N = 32$, MultPIM achieves a $4.2\times$ improvement in latency over RIME (that is, $21.1\times$ over Haj-Ali *et al.*) while maintaining constant partition count and even slightly reducing area. This paper contributes the following:

- *Partition Techniques*: Introduces novel techniques for broadcasting/shifting data amongst partitions.
- *Full Adder*: Proposes a full-adder design that improves the previous state-of-the-art (FELIX [12]) by up to 33%.
- *MultPIM*: Proposal of an efficient parallel multiplier that replaces quadratic time complexity with linear-log. We show latency improvement of $4.2\times$ and slight area reduction over the previous state-of-the-art (RIME [22]).
- *Matrix-vector multiplication*: We present an *optimized* implementation of MultPIM in matrix-vector multiplication that improves latency by $25.5\times$ over FloatPIM [21].

Manuscript received June 23, 2021; revised September 6, 2021; accepted September 20, 2021. This work was supported in part by the European Research Council through the European Union’s Horizon 2020 Research and Innovation Programme under Grant 757259, and in part by the Israel Science Foundation under Grant 1514/17. (*Corresponding author: Orian Leitersdorf.*)

Orian Leitersdorf, Ronny Ronen, and Shahar Kvatinsky are with the Technion - Israel Institute of Technology, Haifa 3200003, Israel (e-mail: orianl@campus.technion.ac.il; shahar@ee.technion.ac.il).

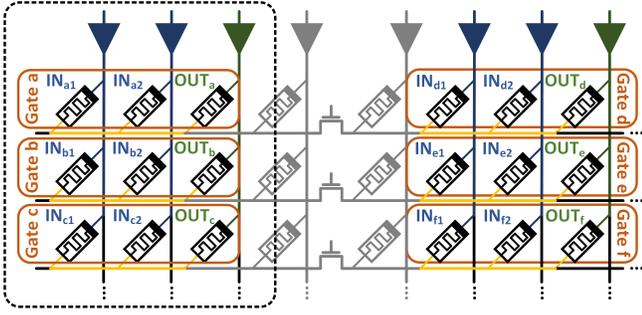


Fig. 1. Dashed: Memristive crossbar array with simultaneous in-row MAGIC NOR operations. Column partitions [12] increase parallelism, e.g., performing all of the highlighted MAGIC gates in a single clock cycle.

II. BACKGROUND

A. Stateful Logic

Memristive crossbar arrays have horizontal wordlines, vertical bitlines, and memristors at crosspoints. Stateful logic employs the same memristors for logic. This is possible by exploiting the unique property of memristors (voltage-controlled variable resistance). IMPLY [10], MAGIC [11] and FELIX [12] are such techniques, computing logic gates by applying voltages along either bitlines or wordlines. Together, they support logic gates such as NOT, NOR, OR, NAND, and Minority3¹. Further, an AND with the previous output cell value can be performed by skipping initialization [12], [26].

Stateful logic support massive-parallelism. The same in-row logic gate can be repeated along rows while still being performed in a single clock cycle, as seen in the dashed portion of Figure 1. Essentially, in a single cycle we can perform a single element-wise logic operation on columns of a crossbar. Hence, memristive computation algorithms are typically limited to a single row of memristors as this allows repetition of the algorithm along many rows (e.g., for vector operation) in the same latency [27]. This parallelism can be further increased through memristive partitions [12]. These transistors divide the crossbar into *partitions* and can be dynamically set to either non-conducting (for parallel operation amongst the partitions, see Figure 1) or conducting (for logic between partitions).

B. Carry-Save Add-Shift (CSAS)

The carry-save add-shift (CSAS) technique [23]–[25] utilizes a carry-save adder [28] for multiplication. The technique stores two numbers, the current sum and the current carry, and adds the partial products to these numbers using the carry-save adder. This can be more efficient than a traditional *shift-and-add* multiplier as carry propagation is avoided in intermediate steps. Rather than shift the partial products, the CSAS technique shifts the sum – effectively emulating moving full adders (FAs). Figure 2 details the overall circuit. In N stages, this circuit produces the lower N bits of the product. The top N bits can be computed as the sum of the final sum and carry numbers. These two numbers can be added by feeding zero partial products to the FAs for N stages, or with a regular adder (e.g., ripple carry) [23]–[25].

¹Haj-Ali *et al.* [19] assumes NOT/NOR, RIME [22] assumes NOT/NOR/NAND/Min3 and MultPIM assumes NOT/Min3 (fair comparison to RIME). MultPIM with other gates is included on the repository.

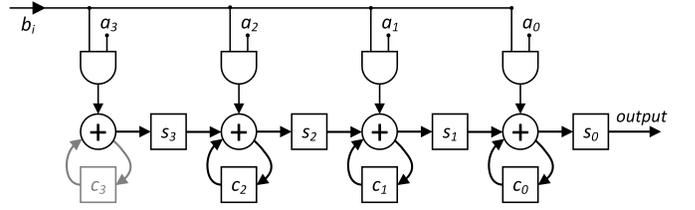


Fig. 2. Four-bit CSAS multiplier [25]. In each cycle, a single bit from input b is fed to compute that corresponding partial product; a carry-save adder adds this partial-product to the current sum/carry. Notice that c_3 is always zero. Latches are squares and full-adders are circles.

III. PARTITION TECHNIQUES

In this section, we introduce two novel partition techniques. The first technique involves *broadcasting* a single bit from one partition to k partitions in $\log_2(k)$ cycles, and the second technique involves *shifting* bits across k partitions in two cycles. Throughout this section, we assume k consecutive partitions and p_i refers to the i^{th} partition.

For simplicity, we do not discuss initialization cycles in this section and we also assume the existence of a *copy* gate: similar to MAGIC NOT, but without negation. Note that the final MultPIM implementation accounts for initialization cycles and does not require a *copy* gate.

A. Broadcasting Technique

Assume that partition p_1 contains a bit that we want to transfer to all of the other partitions. The naive approach illustrated in Figure 3(a) will perform the operation serially: copying the bit from the first partition to each of the others, one at a time, for a total of $k - 1$ clock cycles. In terms of area, this naive approach requires one memristor from each partition (no extra intermediate memristors are necessary).

We propose a novel recursive technique for solving this task, dynamically selecting the partition transistors. We begin by copying the bit from p_1 to $p_{k/2+1}$. Then, we set the transistor between $p_{k/2}$ and $p_{k/2+1}$ to non-conducting and proceed recursively *in parallel* with $p_1, \dots, p_{k/2}$ and $p_{k/2+1}, \dots, p_k$. This technique requires a total of $\log_2 k$ cycles and only one memristor per partition (no extra intermediate memristors necessary), as shown in Figure 3(b).

B. Shift Technique

Assume that each partition begins with its own bit, and that we want to shift these bits between the partitions: the bit from p_1 moves to p_2 , the bit from p_2 moves to p_3 , ..., the bit from p_{k-1} moves to p_k . RIME performs this transfer in $k - 1$ cycles as shown in Figure 3(c). In terms of area, this technique requires no additional intermediate memristors.

We propose a novel technique involving only two steps: copying from all odd partitions to even partitions, and then copying from all even partitions to odd partitions. This technique is demonstrated in Figure 3(d), utilizing exactly 2 clock cycles in total. Note that we can replace the copy gate with any other logic gate (i.e., storing multiple input bits in each partition, and storing the output of the logic gate on the inputs of the i^{th} partition in the $i + 1^{\text{th}}$ partition). This concept is utilized in Section IV-B to optimize full-adder logic.

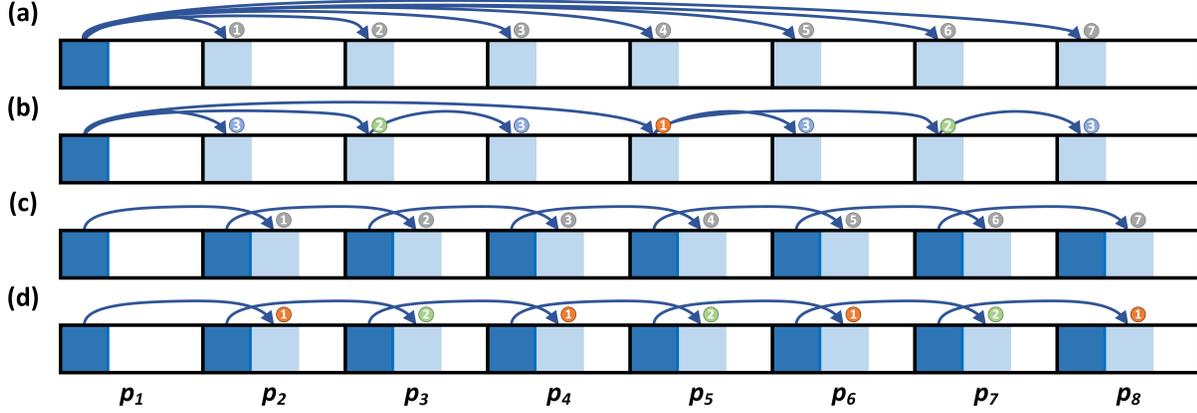


Fig. 3. (a) The naive solution to the broadcasting task, requiring $k - 1$ cycles, and (b) the proposed solution requiring $\log_2 k$ cycles. (c) The naive solution to the shift task, requiring $k - 1$ cycles, and (d) the proposed solution requiring 2 cycles. Circled numbers represent the cycle number.

IV. MULTPIM: FAST STATEFUL MULTIPLIER

In this section, we combine the CSAS multiplier with the two novel techniques from Section III to introduce MultPIM. We begin by describing the general algorithm concept, and then continue by providing various optimizations for latency and area. Throughout this section, let $a = (a_{N-1} \dots a_0)_2$ and $b = (b_{N-1} \dots b_0)_2$; we are interested in computing $a \cdot b$. Recall that p_i is the i^{th} partition, and let $p_i.x$ represent the variable x stored in p_i (single bit).

A. General Algorithmic Concept

The general concept involves using N full-adders in parallel (each in a partition), similar to the CSAS technique (see Figure 2). Each partition stores one bit of a throughout the entire computation (*i.e.*, partition p_i stores a_{N-i}). In addition, each partition stores carry/sum bits (similar to CSAS latches).

Following the CSAS technique, the computation begins with N stages in which b is fed into the system. For the i^{th} stage in the first N stages, we perform the following:

- Copy b_i to all of the partitions using the technique from Section III-A in $\log_2 N$ cycles.
- Compute the partial product in all of the partitions in parallel (similar to AND gates in CSAS).
- Compute full-adder in each of the partitions in parallel, using the stored carry/sum bits and the partial product bit. The new sum/carry replace the old sum/carry bits.
- Shift the sum bits amongst the partitions using the technique from Section III-B. Lowest bit is stored as output.

We choose³ to proceed by feeding another N zeros for b to propagate the stored carries. That is, the algorithm continues with another N stages as follows:

- Compute half-adder in each of the partitions in parallel, using the stored carry-bit and the stored sum-bit. The new sum/carry bits replace the old ones.
- Shift the sum bits amongst the partitions using the technique from Section III-B. Lowest bit is stored as output.

²MultPIM also supports different widths for a and b .

³A regular adder can be implemented instead in p_{N+1} . During that time, partitions p_0, p_1, \dots, p_N could compute the product of a different independent pair of numbers as part of a multiplication pipeline.

Overall, N stages with latency $O(\log_2 N)$ and another N stages with latency $O(1)$. Hence, total latency is $O(N \log_2 N)$. Each partition requires $O(1)$ memristors, so we require $O(N)$ memristors in total. The above stages are shown in Figure 4.

B. Implementation and Optimizations

Algorithm 1 details the steps of the computation. Note that the usage of $\forall i$ in the “*In parallel*” lines indicates that the computation is performed in parallel on all partitions. The for loops in the algorithm are evaluated serially. We detail here various specific optimizations to the algorithm.

1) *Full Adder*: The state-of-the-art⁴ (FELIX [12]) requires 6 cycles (without init.), assumes NOT/OR/NAND/Min3, and requires 2 intermediates. Our novel full-adder is based on:

$$C_{out} = \text{Min}_3'(A, B, C_{in}), \quad (1)$$

$$S_{out} = \text{Min}_3(C_{out}, C'_{in}, \text{Min}_3(A, B, C'_{in})). \quad (2)$$

The improvement over FELIX [12] originates from using C_{out} for computing S . These expressions enable 5 cycles, assuming only NOT/Min3, and requiring 3 intermediate memristors⁵. Further, if the not of an input is also given, only 4 cycles are required (*i.e.*, no need to compute C'_{in})⁶. The latter is utilized for Lines 6-7 and Lines 10-11 by storing both C, C' and performing the sum computation as part of shift.

2) *Lines 4-5*: Performing the Section III-A algorithm with NOT (rather than the theoretical *copy*) results in some partitions receiving b_k and others receiving b'_k . The partitions that receive b_k perform Line 5 using no-initialization NOT (see Section II-A) of the stored a'_i into b_k , resulting in $(a'_i)' \cdot b_k = a_i \cdot b_k$. Those that receive b'_k perform Line 5 using $\text{Min}_3(a'_i, b'_k, 1) = a_i \cdot b_k$. Thus, Line 5 requires 1 cycle.

3) *Partitions*: Note that p_0 and p_{N+1} can be merged with p_1 and p_N (respectively) to reach a total of N partitions. Furthermore, since the top carry bit is always zero (see Figure 2), then we can use $N - 1$ partitions rather than N .

⁴The full-adder proposed by RIME [22] requires 7 cycles. Note that our novel full-adder is *inspired* by the expressions from RIME.

⁵6 cycles, assuming NOT/Min3, and only 2 intermediate memristors is possible with re-use. Therefore, FELIX [12] is replaced completely.

⁶This enables N -bit addition with $5N$ cycles and $3N + 5$ memristors using only NOT/Min3, compared to $7N$ and $3N + 2$ from FELIX (including init.).

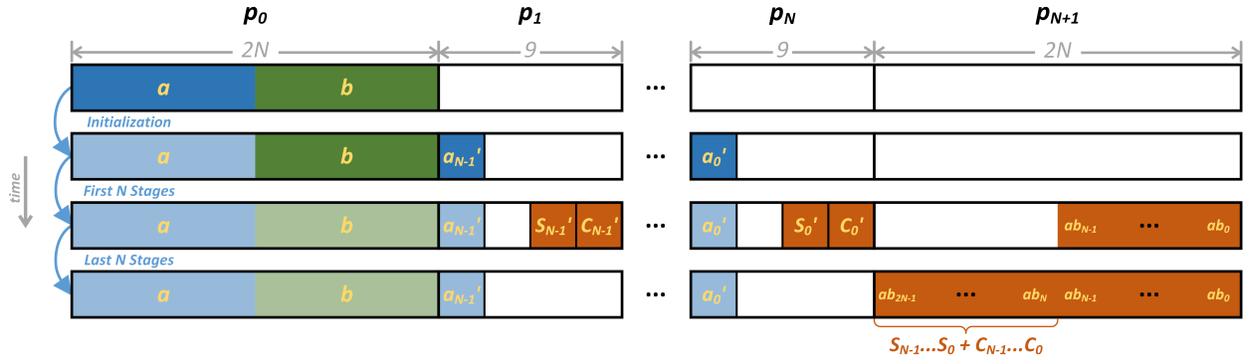


Fig. 4. The main steps of the MultPIM algorithm. Note that the last N bits of the product are the sum of $S_{N-1} \dots S_0$ and $C_{N-1} \dots C_0$; this sum can either be computed via the *Last N Stages*, or by using a conventional adder. Faded-out cells indicate values no longer used.

Algorithm 1 MultPIM

Input: a, b stored in p_0 (start of the row)

Output: $a \cdot b$ stored in p_{N+1} (end of the row)

Initialization:

- 1: $\forall i : p_i.c, p_i.s \leftarrow 0$ {In parallel, init. carry/sum}
- 2: **for** $i = 1$ to N **do** $p_i.a \leftarrow p_0.a_{N-i}$ {Store a_{N-i} in p_i }

First N Stages:

- 3: **for** $k = 1$ to N **do**
- 4: $\forall i : p_i.b = b_k$ {Using Section III-A}
- 5: $\forall i : p_i.ab = p_i.a \cdot p_i.b$ {In parallel}
- 6: $\forall i : p_i.s, p_i.c = FA(p_i.s, p_i.c, p_i.ab)$ {In parallel}
- 7: $\forall i : p_{i+1}.s = p_i.s$ {Using Section III-B}
- 8: **end for**

Last N Stages:

- 9: **for** $k = 1$ to N **do**
- 10: $\forall i : p_i.s, p_i.c = HA(p_i.s, p_i.c)$ {In parallel}
- 11: $\forall i : p_{i+1}.s = p_i.s$ {Using Section III-B}
- 12: **end for**

V. RESULTS

We evaluate MultPIM for single-row N -bit multiplication. We compare MultPIM to Haj-Ali *et al.* [19] and RIME [22] in terms of latency, area (memristor count), and partition count. We also present MultPIM-Area that prioritizes area over latency via additional re-use [27]. The results are verified by a custom cycle-accurate simulator.

A. Latency

We evaluate the latency of the MultPIM algorithm in clock cycles. The algorithm begins with N cycles at the start to copy a . Then N stages which feed b through the full-adders, with each stage requiring $\log_2 N + 8$ cycles ($\log_2 N + 1$ for Lines 4-5, 5 cycles for Lines 6-7, and 1 initialization cycle). Finally, N stages at the end, each requiring 6 cycles (5 for Lines 10-11 and 1 initialization cycle). Overall, $N \log_2 N + 14 \cdot N$ cycles. In Table I, we compare this latency with the previous works, demonstrating $4.2\times$ improvement over the previous state-of-the-art (RIME) for the common case of $N = 32$.

B. Area

The exact number of memristors required for MultPIM is evaluated here. The computation row contains $2N$ memristors

TABLE I
LATENCY (CLOCK CYCLES)

Algorithm	Expression	$N = 16$	$N = 32$
Haj-Ali <i>et al.</i> [19]	$13 \cdot N^2 - 14 \cdot N + 6$	3110	12870
RIME [22]	$2 \cdot N^2 + 16 \cdot N - 19$	749	2541
MultPIM	$N \cdot \log_2 N + 14 \cdot N + 3$	291	611
MultPIM-Area	$N \cdot \log_2 N + 23 \cdot N + 3$	435	899

TABLE II
AREA (# MEMRISTORS)

Algorithm	Expression	$N = 16$	$N = 32$
Haj-Ali <i>et al.</i> [19]	$20 \cdot N - 5$	315	635
RIME [22]	$15 \cdot N - 12$	228	468
MultPIM	$14 \cdot N - 7$	217	441
MultPIM-Area	$10 \cdot N$	160	320

for storing the inputs, $2N$ memristors for storing the outputs, and N full-adder units each requiring 10 memristors total. Hence, MultPIM requires $2 \cdot N + 2 \cdot N + 10 \cdot N = 14 \cdot N$ memristors. Table II compares this with the previous works, showing a slight improvement over the state-of-the-art (RIME). Note that MultPIM and RIME both require $N - 1$ partitions⁷.

C. Logic Simulation

We verify the results of the algorithm with a custom cycle-accurate simulator⁸. The simulator models the crossbar array, and has an interface for performing *operations* in-memory. The MultPIM algorithm is tested by first writing the inputs to the crossbar, then allowing MultPIM to perform in-memory *operations*, and finally verifying the output. The simulator counts the exact number of *operations* that MultPIM uses (including initializations), verifying the theoretical analysis.

VI. MATRIX-VECTOR MULTIPLICATION

Here, we optimize MultPIM for matrix-vector multiplication. Formally, let \mathbf{A} be an $m \times n$ matrix and let \mathbf{x} be a vector of dimension n , we are interested in computing $\mathbf{A}\mathbf{x}$. Each element in the matrix/vector is a fixed-point number with N bits, and the data elements are stored horizontally.

The multiplication is performed by duplicating \mathbf{x} along rows (see Figure 5), multiplying each column of the matrix

⁷The evaluation of exact partition overhead is left for future work. Regardless, MultPIM and RIME require the same number of partitions.

⁸Available at <https://github.com/oleitersdorf/MultPIM>.

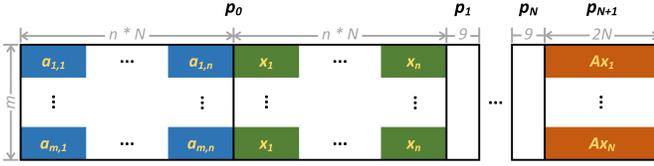


Fig. 5. Matrix-vector multiplication with an optimized MultPIM multiplier. Matrix A is shown in blue, vector x in green, and Ax in orange. The partitions are only used along columns, with the same overhead as MultPIM.

TABLE III
MATRIX MULTIPLICATION ($n = 8, N = 32$)

Algorithm	Latency (Clock Cycles)	Area (Min. Crossbar Dim.)
FloatPIM	109616	$m \times 1723$
MultPIM	4292	$m \times 965$
MultPIM-Area	6204	$m \times 778$

with each column of the duplicated vector matrix, and then adding the results horizontally. Essentially, each row performs an inner product between the stored row of A and x (e.g., $Ax_1 = a_{1,1} \cdot x_1 + \dots + a_{1,n} \cdot x_n$ in the first row). A similar concept is used in FloatPIM [21] for fixed-point matrix-multiplication. The naive solution replaces only the fixed-point multiplication algorithm in FloatPIM with MultPIM (i.e., compute $a_{1,1} \cdot x_1, \dots, a_{1,n} \cdot x_n$ by using MultPIM n times, and sum using an adder). That provides only $9.5\times$ latency improvement to FloatPIM as addition becomes non-negligible.

Instead, we optimize MultPIM to compute the sum while computing the products and further reduce product latency. The optimized algorithm receives numbers a, b (N -bit) and s_i, c_i ($2N$ -bit), and computes s_o, c_o ($2N$ -bit) such that $s_o + c_o = a \cdot b + s_i + c_i$. This algorithm performs only *Initialization* and *First N Stages*, thus *reducing* latency compared to regular MultPIM. This is achieved by initializing the sum fields of the full-adders to the lower N bits of s_i (rather than zero) and feeding p_1 the upper bits of s_i and c_i . The value of $s_o + c_o$ at each run of MultPIM is the sum of the products until that point. At the end, the sum $s_o + c_o$ is computed once.

The results of the optimized matrix-vector multiplication are summarized in Table III for $n = 8, N = 32$, verified by the logic simulator. We achieve $25.5\times$ latency and $1.8\times$ area improvement over FloatPIM matrix-vector multiplication, utilizing 33 partitions. In the general case, latency is improved from $n \cdot (13N^2 + 12N + 6)$ to $n \cdot (N \log_2 N + 11N + 9) + 4N - 4$ cycles, and area is improved from $m \times (4nN + 22N - 5)$ to $m \times (2nN + 14N + 5)$ memristors, with $N + 1$ partitions.

VII. CONCLUSION

We present MultPIM: a novel partition-based in-memory multiplication algorithm that improves the state-of-the-art latency complexity from quadratic to linear-log, specifically by $4.2\times$ for 32-bit. The improvement is based on the carry-save add-shift technique, two novel memristive-partition computation techniques, and an improvement to the state-of-the-art full-adder. Furthermore, we optimize MultPIM for matrix-vector multiplication and achieve $25.5\times$ latency and $1.8\times$ area improvements over FloatPIM matrix-vector multiplication by computing addition while performing multiplication. Correctness is verified via a cycle-accurate simulator.

REFERENCES

- [1] A. Pedram, S. Richardson, M. Horowitz, S. Galal, and S. Kvatinsky, "Dark memory and accelerator-rich system optimization in the dark silicon era," *IEEE Design & Test*, 2017.
- [2] L. Chua, "Memristor—the missing circuit element," *IEEE Transactions on Circuit Theory*, vol. 18, no. 5, pp. 507–519, 1971.
- [3] S. Kvatinsky, E. G. Friedman, A. Kolodny, and U. C. Weiser, "The desired memristor for circuit designers," *IEEE CAS Magazine*, 2013.
- [4] B. Hoffer, V. Rana, S. Menzel, R. Waser, and S. Kvatinsky, "Experimental demonstration of memristor-aided logic (MAGIC) using valence change memory (VCM)," *IEEE Transactions on Electron Devices*, 2020.
- [5] Z. Sun, E. Ambrosi, A. Bricalli, and D. Ielmini, "Logic computing with stateful neural networks of resistive switches," *Advanced Materials*, 2018.
- [6] J. Xu, Y. Zhan, Y. Li, J. Wu, X. Ji, G. Yu, W. Jiang, R. Zhao, and C. Wang, "In situ aging-aware error monitoring scheme for imply-based memristive computing-in-memory systems," *IEEE TCAS-I*, 2021.
- [7] P. Liu, Z. You, J. Wu, B. Liu, Y. Han, and K. Chakrabarty, "Fault modeling and efficient testing of memristor-based memory," *IEEE Transactions on Circuits and Systems I: Regular Papers*, pp. 1–12, 2021.
- [8] S. Swami and K. Mohanram, "Reliable nonvolatile memories: Techniques and measures," *IEEE Design & Test*, 2017.
- [9] J. Reuben, R. Ben-Hur, N. Wald, N. Talati, A. H. Ali, P.-E. Gaillardon, and S. Kvatinsky, "Memristive logic: A framework for evaluation and comparison," in *PATMOS*, 2017.
- [10] J. Borghetti, G. S. Snider, P. J. Kuekes, J. J. Yang, D. R. Stewart, and R. S. Williams, "Memristive switches enable 'stateful' logic operations via material implication," *Nature*, vol. 464, no. 7290, pp. 873–876, 2010.
- [11] S. Kvatinsky, D. Belousov, S. Liman, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser, "MAGIC—memristor-aided logic," *IEEE Transactions on Circuits and Systems II: Express Briefs*, 2014.
- [12] S. Gupta, M. Imani, and T. Rosing, "FELIX: Fast and energy-efficient logic in memory," in *ICCAD*, 2018, pp. 1–7.
- [13] J. Yu, R. Nane, I. Ashraf, M. Taouil, S. Hamdioui, H. Corporaal, and K. Bertels, "Skeleton-based synthesis flow for computation-in-memory architectures," *Transactions on Emerging Topics in Computing*, 2020.
- [14] M. Imani, S. Gupta, and T. Rosing, "Ultra-efficient processing in-memory for data intensive applications," in *2017 54th ACM/EDAC/IEEE Design Automation Conference (DAC)*, 2017, pp. 1–6.
- [15] L. Guckert and E. E. Swartzlander, "Dadda multiplier designs using memristors," in *ICICDT*, 2017, pp. 1–4.
- [16] D. Radakovits, N. TaheriNejad, M. Cai, T. Delaroche, and S. Mirabbasi, "A memristive multiplier using semi-serial imply-based adder," *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2020.
- [17] L. Guckert and E. E. Swartzlander, "Optimized memristor-based multipliers," *IEEE TCAS-I*, vol. 64, no. 2, pp. 373–385, 2017.
- [18] S. Shin, K. Kim, and S.-M. Kang, "Resistive computing: Memristors-enabled signal multiplication," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 60, no. 5, pp. 1241–1249, 2013.
- [19] A. Haj-Ali, R. Ben-Hur, N. Wald, and S. Kvatinsky, "Efficient algorithms for in-memory fixed point multiplication using MAGIC," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018.
- [20] A. Haj-Ali, R. Ben-Hur, N. Wald, R. Ronen, and S. Kvatinsky, "IMAGING: in-memory algorithms for image processing," *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2018.
- [21] M. Imani, S. Gupta, Y. Kim, and T. Rosing, "FloatPIM: In-memory acceleration of deep neural network training with high precision," in *Annual International Symposium on Computer Architecture*, 2019.
- [22] Z. Lu, M. T. Arafat, and G. Qu, "RIME: A scalable and energy-efficient processing-in-memory architecture for floating-point operations," in *Asia and South Pacific Design Automation Conference*, 2021.
- [23] S. Sunder, F. El-Guibaly, and A. Antoniou, "Two's-complement fast serial-parallel multiplier," *IEE Proceedings-Circuits, Devices and Systems*, vol. 142, no. 1, pp. 41–44, 1995.
- [24] R. Richards, *Arithmetic Operations in Digital Computers*, ser. University series in higher mathematics. New York, 1955.
- [25] Gnanasekaran, "A fast serial-parallel binary multiplier," *IEEE Transactions on Computers*, vol. C-34, no. 8, pp. 741–744, 1985.
- [26] N. Peled, R. Ben-Hur, R. Ronen, and S. Kvatinsky, "X-MAGIC: Enhancing PIM using input overwriting capabilities," in *VLSI-SoC*, 2020.
- [27] R. Ben-Hur, R. Ronen, A. Haj-Ali, D. Bhattacharjee, A. Eliahu, N. Peled, and S. Kvatinsky, "SIMPLER MAGIC: Synthesis and mapping of in-memory logic executed in a single row to improve throughput," *IEEE TCAD*, 2020.
- [28] M. Vlăduțiu, *Computer arithmetic: algorithms and hardware implementations*. Springer Science & Business Media, 2012.