

DS-CIM: Digital Stochastic Computing-In-Memory Featuring Accurate OR-Accumulation via Sample Region Remapping for Edge AI Models

Kunming Shao^{1,*,\dagger}, Liang Zhao^{2,3,*}, Jiangnan Yu^{1,*}, Zhipeng Liao⁴,
Xiaomeng Wang¹, Yi Zou^{2,3}, Tim Kwang-Ting Cheng¹, Chi-Ying Tsui¹

¹The Hong Kong University of Science and Technology, ²South China University of Technology,

³MoE Engineering Research Center of Design and Technology Co-Optimization of IC, ⁴Westlake University

*Equally Contributed Authors. †Email: kshaoaa@connect.ust.hk

Abstract—Stochastic computing (SC) offers hardware simplicity but suffers from low throughput, while high-throughput Digital Computing-in-Memory (DCIM) is bottlenecked by costly adder logic for matrix-vector multiplication (MVM). To address this trade-off, this paper introduces a digital stochastic CIM (DS-CIM) architecture that achieves both high accuracy and efficiency. We implement signed multiply-accumulation (MAC) in a compact, unsigned OR-based circuit by modifying the data representation. Throughput is enhanced by replicating this low-cost circuit 64 times with only a $1\times$ area increase. Our core strategy, a shared Pseudo Random Number Generator (PRNG) with 2D partitioning, enables single-cycle mutually exclusive activation to eliminate OR-gate collisions. We also resolve the 1s saturation issue via stochastic process analysis and data remapping, significantly improving accuracy and resilience to input sparsity. Our high-accuracy DS-CIM1 variant achieves 94.45% accuracy for INT8 ResNet18 on CIFAR-10 with a root-mean-squared error (RMSE) of just 0.74%. Meanwhile, our high-efficiency DS-CIM2 variant attains an energy efficiency of 3566.1 TOPS/W and an area efficiency of 363.7 TOPS/mm², while maintaining a low RMSE of 3.81%. The DS-CIM capability with larger models is further demonstrated through experiments with INT8 ResNet50 on ImageNet and the FP8 LLaMA-7B model.

I. INTRODUCTION

Deep neural networks (DNNs) and Transformer models are foundational to modern AI, but their computational demands present significant challenges for energy- and area-constrained edge devices [1]–[7]. Efficient matrix vector multiplication (MVM) is critical for their edge deployment. Approximate computing has been explored to improve efficiency, but the high accuracy requirements of Transformers limit the applicability of many methods.

As illustrated in Figure 1, OR-Gate-based stochastic computing (SC) is one of these approximate methods that offers low-cost and fault-tolerant computation. However, it is hampered by low throughput from long bitstream latencies and significant 1s saturation errors [8] during low-sparsity accumulation, restricting its use in modern AI applications [9]–[14]. Digital compute-in-memory (DCIM) is another method to enhance edge MVM by performing parallel computations directly in the SRAM array and reusing the weight, but its

This research was supported by ACCESS - AI Chip Center for Emerging Smart Systems, sponsored by InnoHK funding, Hong Kong SAR.

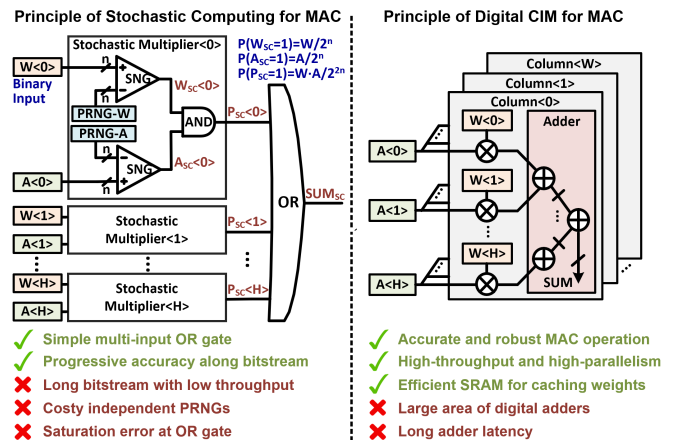


Fig. 1: Principles of Stochastic Computing and Digital CIM.

efficiency is often bottlenecked by the large and power-hungry digital adder tree used for accurate accumulation [15]–[17].

Analog CIM (ACIM) [18], [19] and approximate DCIM [20]–[22] sacrifice precision by utilizing low-precision ADCs or approximate adders for improved efficiency. ACIM employs efficient but PVT-sensitive ADCs for accumulation, while approximate DCIM uses simplified adders [21] that introduce logical errors. However, their accuracy cannot meet the requirement of transformer models.

SC can achieve much lower errors by increasing the bitstream length. Integrating SC with CIM is a promising alternative, yet existing works suffer from inefficiencies with large weight bit-widths, limited parallelism, and persistent saturation errors, particularly with signed operations [23]–[27].

To address these challenges, we propose our DS-CIM architecture. **First**, we introduce a novel method that shifts signed numbers to an unsigned representation, allowing for efficient, unipolar stochastic computing that supports both signed activations and weights. **Second**, we replicate a simple unipolar OR-MAC circuit to enhance parallelism by $64\times$, compensating for the throughput degradation from long bitstreams with only a $1\times$ area increase. **Third**, we use a shared PRNG and a data remapping technique to ensure each random sampling point activates at most one row, completely eliminating 1s saturation errors. **Finally**, we optimize the PRNG configuration to

minimize overall RMSE.

Post-layout simulations in a 40nm process validate our approach. Our precise DS-CIM1 achieves 669.7 TOPS/W and 117.1 TOPS/mm² with only 0.74% RMSE. The efficient DS-CIM2 reaches 3566.1 TOPS/W and 363.7 TOPS/mm² with 3.81% RMSE. DS-CIM demonstrates high efficiency, low error, and robust accuracy across all product sparsity levels (0-100%).

II. PRELIMINARIES

A. Stochastic Computing

Stochastic computing generates unary bitstreams by comparing binary numbers with random values, ensuring each position's probability of a '1' matches the binary value. This conversion simplifies high-bit-width binary values, reducing logical computation costs and enhancing hardware efficiency [9], [11]–[14].

Classic SC applications, such as multiplication and scaling, leverage the probabilistic nature of unary bitstreams for simpler, scalable hardware implementations. However, generating long unary bitstreams can increase latency, requiring careful management of the trade-off between efficiency and computation time based on application needs. There are also challenges that when applying SC on MAC operations, there will be distortion error due to the accumulation [8], [10].

B. Digital Computing-in-Memory

AI accelerators face the memory wall problem, where memory access speed limits overall performance. Digital Compute-In-Memory (CIM) addresses this by integrating processing within memory units, reducing data transfer needs and enhancing efficiency [15]–[17]. However, challenges remain, such as low computing density and throughput, and the significant overhead of adder trees. Overcoming these issues is crucial for maximizing the potential of Digital CIM in AI accelerators.

C. Approximate Computing-in-Memory

Some works have applied approximate computing techniques to CIMs, such as using analog signals for accumulation [18], designing approximate adders [21], and introducing stochastic computing [27]. However, these approaches face accuracy issues. Although high Root Mean Square Error can show good results in DNNs or event cameras, with the development of transformers, previous approximate computing methods are no longer applicable.

III. DS-CIM HARDWARE ARCHITECTURE

A. Overall Architecture

As shown in Fig. 2(a), DS-CIM integrates SC's PRNGs and Stochastic Number Generators (SNGs) into a DCIM framework, replacing costly multi-bit adders with efficient single-bit OR-based Multiply-Accumulate (OR-MAC) circuits. The macro has a 128x32 structure with 128 8-bit SRAMs, 128 SNGs, and 64 OR-MACs per column. An input SNG array converts binary activations to stochastic bitstreams. To support

data remapping (Sec. IV), the macro uses only two independent PRNGs: PRNGA for activation bitstreams and PRNGW for weight bitstreams. These PRNGs generate random numbers for the SNGs to convert binary data into stochastic bitstreams, A_{SC} and W_{SC} . A_{SC} is broadcast across columns, while W_{SC} is shared among the 64 OR-MACs within a column.

After multiplication, product bitstreams (P_{SC}) are accumulated via an OR gate, and the partial sum is added cycle-by-cycle in an accumulator. For balancing accuracy and efficiency, we designed two versions: the precise DS-CIM1 uses eight OR16-based MACs (OR-MAC16) per column, while the efficient DS-CIM2 uses two OR64-based MACs (OR-MAC64). Both perform accurate MAC operations but differ in their accumulation and data remapping granularities.

B. Unipolar OR-MAC for Signed MAC

SC employs a probabilistic process to perform unsigned operations. As depicted in Figure 3(a), for executing signed MAC in stochastic CIM (S-CIM), [27] proposes a sign-aware bipolar MAC circuit for signed weights but unsigned activations. This differential circuit separates the positive weight bitstream W_{SC+} from the negative weight bitstream W_{SC-} and performs MAC operations in two ways, with the final result being the difference between the two operations. However, this bipolar MAC scheme significantly increases overhead and limits parallelism and efficiency.

In contrast, DS-CIM introduces a unipolar MAC scheme for both signed activations and weights by converting signed numbers into unsigned ones and converting the output to signed numbers after the unsigned MAC. As shown in Eq. 1, the partial sum is logically derived from the MAC operation of two columns of 8-bit signed activations $x[i]$ and weights $w[i]$ represented in 2's complement. To convert a signed number to an unsigned one, we invert its sign bit, which is equivalent to adding 128. For example, we transform 1000_0011 (2's complement, -125) to 0000_0011 (unsigned, 3) by inverting the sign bit, and 0111_1111 (2's complement, 127) to 1111_1111 (unsigned, 255). Eq. 2 and 3 illustrate our method for decomposing signed multiplication into three terms, while Eq. 4 demonstrates how we partition the high-cost signed MAC into efficient unsigned MAC, addition, and subtraction.

$$psum = \sum_{i=0}^{H-1} x[i]w[i] \quad (1)$$

$$x'[i] = x[i] + 2^7 \quad \text{and} \quad w'[i] = w[i] + 2^7 \quad (2)$$

$$x[i]w[i] = x'[i]w'[i] - 2^7x[i] - 2^7w'[i] \quad (3)$$

$$psum^{\circledast} = \sum_{i=0}^{H-1} x'[i]w'[i]^{\circledast} - 2^7 \sum_{i=0}^{H-1} x[i]^{\circledast} - 2^7 \sum_{i=0}^{H-1} w'[i]^{\circledast} \quad (4)$$

In our DS-CIM hardware, term b can be efficiently calculated using our unsigned unipolar OR-MAC, term c is summed through SIMDs during runtime and term d is computed offline and stored as a look-up table. The overhead for term c is spatially amortized across multiple weight columns. Compared to

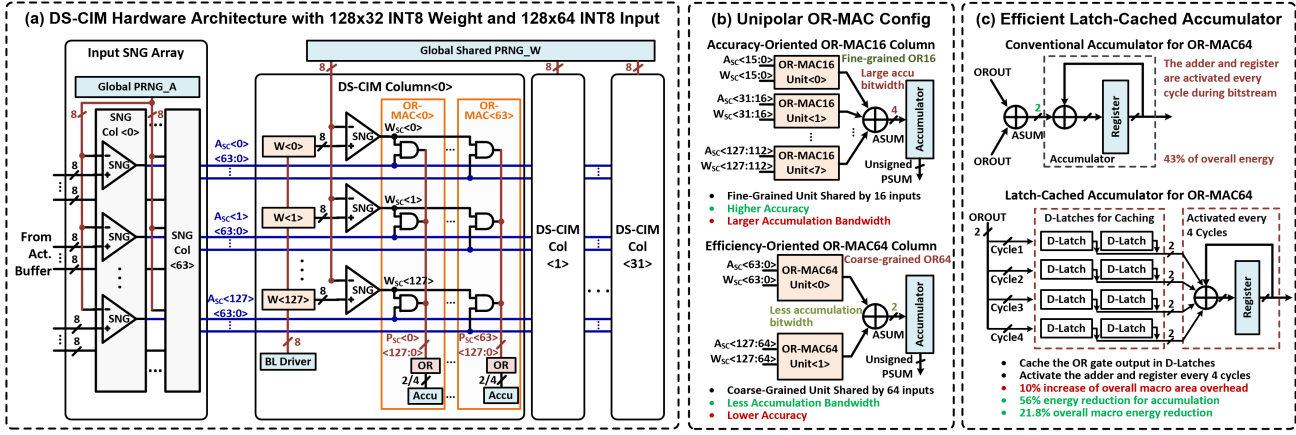


Fig. 2: (a) Overall architecture of proposed DS-CIM. (b) Demonstration of the proposed unipolar OR-MAC configurations with the accuracy-oriented OR-MAC16 configuration for DS-CIM1 and the efficiency-oriented OR-MAC64 configuration for DS-CIM2. (c) Efficient latch-cached accumulator for reducing the high energy consumption of the conventional accumulators.

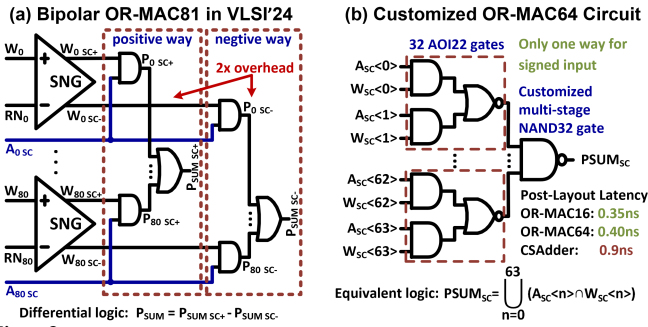


Fig. 3: (a) Bipolar OR-MAC81 for event camera in [27]. (b) Proposed unipolar OR-MAC64 circuit with only 0.4ns latency in 40nm process.

the substantial MAC overhead, the subtraction and amortized addition overheads are negligible.

Figure 2(b) illustrates our two unipolar OR-MAC circuit configurations. OR-MAC16 is accuracy-oriented, featuring finer OR granularity with a larger addition bitwidth, whereas OR-MAC64 is efficiency-oriented, characterized by coarser OR granularity and a smaller addition bitwidth. Figure 3(b) presents the schematic of our customized OR-MAC64 circuit, which is smaller and faster than carry-save adders [28].

C. Compute/Memory Ratio and the Area Efficiency

To achieve n -bit precision, SC requires a bitstream length of 2^n . However, for an 8-bit MAC operation, the delay of a 256-point SC is 32 times longer than that of an 8-bit serial DCIM, which only requires 8 cycles. Furthermore, in S-CIM that has one set of memory with only one set of MAC circuit, the overhead of the OR-MAC circuit is significantly lower than that of SRAMs and SNGs, leading to unbalanced computation and area efficiency.

To address this issue, we replicate the OR-MAC circuit within the same DS-CIM column [27]. By integrating 64 ultra-efficient OR-MAC circuits in a column, denoted as compute/memory ratio (CMR) = 64, the parallelism of DS-CIM is enhanced by 64 \times , boosting the throughput and area efficiency. Our post-layout data shows that DS-CIM with

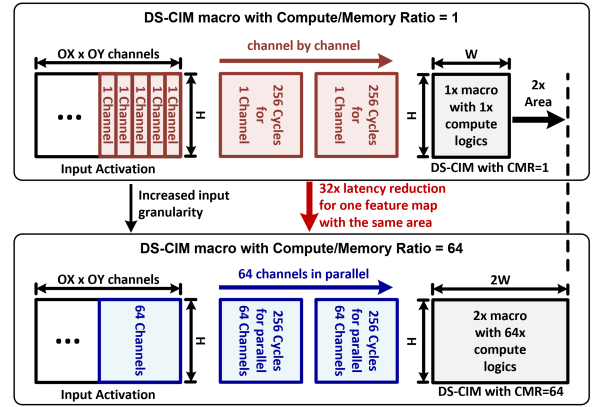


Fig. 4: An illustration of achieving 32x latency reduction and 32x compute density increment by increasing CMR from 1 to 64.




Fig. 5: Memory-friendly pipelined layer channel input timing diagram of the DS-CIM macro with CMR=64 and bitstream length=256.

CMR=64 incurs only a 1 \times area increase compared to CMR=1 designs. As illustrated in Figure 4, By 64 \times throughput with only 2 \times area, DS-CIM achieves a 32 \times latency reduction for a specific CONV layer. Moreover, DS-CIM features progressive precision, enabling a 4 \times extra improvement with a bitstream length=64. By replicating low-cost unipolar OR-MAC circuits, we have mitigated the throughput limitations inherent in SC, providing higher throughput.

As shown in Figure 5, since the activations and weights are temporally stationary during SC process, the SRAM bandwidth could be efficiently utilized for different input channels.

D. Latch-Cached Bitstream Accumulator

As shown in Figure 2(c), DS-CIM requires accumulating a long bitstream cycle by cycle. In DS-CIM, due to the

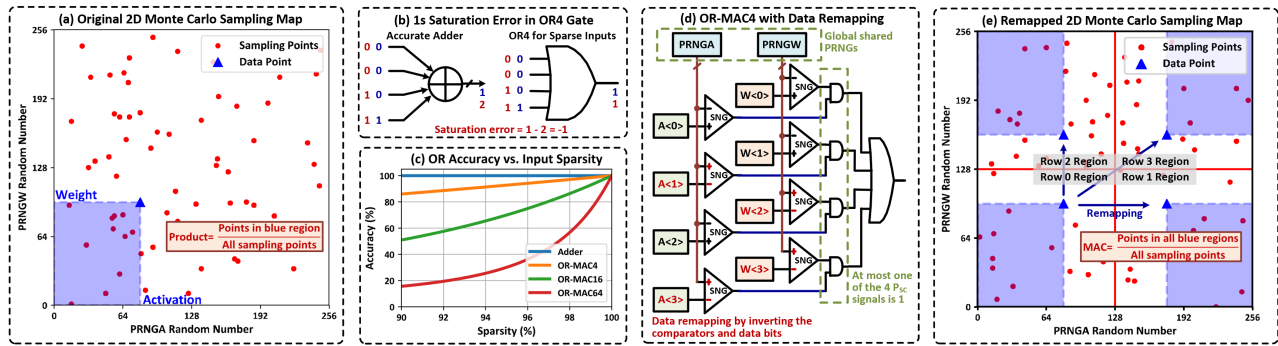


Fig. 6: (a) Representing stochastic multiplication as 2D Monte Carlo process. (b) 1s saturation error in the OR4 gate. (c) The relationship between conventional OR-MAC accuracy and the input product sparsity. (d) OR-MAC4 supporting data remapping with inverted data and comparators. (e) Data remapping in 2D Monte Carlo sampling map, where a sampling point appear at most in one of the four blue regions.

replication of low-cost OR-MAC units, the accumulator’s energy consumption accounts for 43% of the entire macro. For DS-CIM2, which has a smaller accumulation bitwidth, we propose a latch-cached accumulator. We put eight D-latches at the 2b output of the OR-MAC. The outputs are sequentially stored in latches, filling them every four cycles. On the fourth cycle, when the D-latches are full, an adder sums the four 2-bit cached numbers and accumulates the result with the previously stored value in registers. This approach activates the accumulator only once every four cycles, reducing accumulation energy by 56% and overall macro power consumption by 21.8%, with only 10% increase in area.

IV. ERROR ANALYSIS AND OPTIMIZATION

In Stochastic CIM (S-CIM) architecture, due to the nature of random sampling and accumulation, large errors in each row can cancel each other out, resulting in a relatively smaller MAC error. However, inherent errors, such as 1s saturation error and sampling point discrepancies, introduce systematic errors that affect S-CIM. In this section, we analyze the error sources in our DS-CIM and propose corresponding solutions.

A. Random Sampling and 1s Saturation Error

Activation and weight multiplication can be viewed as a 2D Monte Carlo sampling process that estimates the product by calculating the area formed by the two values. As depicted in Figure 6(a), a red sampling point, $(PRNGA, PRNGW)$, appears randomly in the 2D sampling map cycle by cycle. When a sampling point falls within the blue rectangle defined by weight and activation, the conditions of the activation and weight SNG comparators are met, causing both A_{SC} and W_{SC} to equal 1, resulting in an output product of 1. Consequently, the number of 1s in the product bitstream P_{SC} , which reflects the magnitude of the product of activation and weight, depends on how many random points fall within the rectangular area during the sampling process. Thus, the stochastic multiplication process is effectively a 2D Monte Carlo process [13].

Next, we consider S-CIM, where multiple independent Monte Carlo sampling processes occur across different rows. Previous work [27] relied on independent PRNGs for each row to generate independent P_{SC} s, which are then accumulated

together by an OR gate. However, the OR gate is particularly sensitive to the sparsity of P_{SC} . As shown in Figure 6(b), taking an OR4 gate as an example, if the input sparsity is high, most cycles will have all-0 inputs or just a single 1. Occasionally, there may be multiple inputs equal to 1. However, the OR gate can only output a single 1, leading to a **1s saturation error**. A detailed probabilistic analysis in Figure 6(c) reveals that OR accumulation is highly sensitive to product sparsity, with coarser-grained OR gates exhibiting even greater sensitivity to variations in sparsity, which restricts the applicability of OR gate-based S-CIM.

B. Data Remapping for 1s Saturation error

To address the 1s saturation error of OR accumulation, we re-examine the random independence between rows and discover that sharing PRNG between rows allows us to map all data points to the same 2D sampling map. By shifting and symmetry, we ensure that each sampling point can activate at most one row to output 1.

Specifically, as shown in Figure 6(a), we replace the strategy of deploying independent PRNGs for each row in [27] with a shared PRNGW for all weights and a shared PRNGA for all activations, enabling each row in the column to utilize the same sampling point. We illustrate this with the OR-MAC4 case. In Figure 6(e), we right-shift both activations and weights by 1 bit, causing the data distribution region enclosed by red lines to occupy only the *lower-left quarter* of the 2D sampling map.

We then invert all bits of the Row 1’s activation $A[1]$, making it spatially equivalent to a symmetry of 127.5, e.g., from 0000_0010 to 1111_1101. Simultaneously, we alter the decision direction of the second row’s activation SNG. This remaps the red data distribution region and blue Monte Carlo effective region of the second row from the *lower-left corner* to the *lower-right corner*. Similarly, we can position Row 2 in the *upper-left corner* by inverting $W[2]$ and its SNG, and Row 3 in the *upper-right corner* by inverting $A[3]$, $W[3]$, and the corresponding two SNGs. For a shared sampling point, **it will lie in at most one of the four blue rectangles**, ensuring that at most one of the four inputs to the OR4 gate will be 1. By sacrificing data bit-width through right-shifting, we can eliminate the 1s saturation error.

TABLE I: DS-CIM Performance on CNN Models

Bistream	DS-CIM1			DS-CIM2		
	64	128	256	64	128	256
RMSE	3.57%	2.03%	0.74%	3.81%	2.63%	0.84%
ResNet18@ CIFAR10(vs 94.54%)	90%	93.08%	94.45%	89.46%	92.46%	94.31%
ResNet50@ ImageNet(vs 80.82%)	79.45%	80.11%	80.67%	79.34%	79.9%	80.65%

Similarly, we can implement a 4×4 division of the 2D sampling map by right-shifting by 2 bits and inverting the corresponding data bits, with every 16 rows sharing an OR16 gate, forming an OR-MAC16 design, denoted as precise **DS-CIM1**. Furthermore, by right-shifting by 3 bits and remapping the data, we can divide the 2D sampling map into 8×8 segments, with every 64 rows sharing an OR64 gate, forming an OR-MAC64 design, denoted as efficient **DS-CIM2**.

C. Measurement of Optimized RMSEs

The distribution uniformity of sampling points, which is affected by PRNG types and initial seeds, is crucial to the SC accuracy [13]. To minimize the distribution discrepancy of sampling points, we collected mainstream 8-bit PRNGs and searched for optimal initial values for the two random number sequences of PRNGA and PRNGW. We identified optimal PRNG and initial value configurations for 64, 128, and 256 points that minimize the overall RMSE of OR-MAC16 and OR-MAC64 for different data distributions. These configurations ensure optimal RMSE for each application during runtime.

Table I illustrates the RMSE performance of our different OR-MAC circuits and bitstream lengths. Precise DS-CIM1 with bitstream length=256 achieves the lowest RMSE of 0.74% and Efficient DS-CIM2 with bitstream length=64 achieves the acceptable RMSE of 3.81%, which is better than the 4.03% and 6.76% reported in [21]. Due to greater precision loss from right-shifting, the RMSE of OR-MAC64 is slightly higher than that of OR-MAC16. Additionally, the RMSE significantly increases with shorter sequences due to greater discrepancies and insufficient sampling points.

The test results demonstrate that our unipolar OR-MAC circuits are resilient to variations in sparsity, exhibiting uniform errors for any magnitude distribution of the partial sum, making our DS-CIM compatible with MVM applications across flexible data distributions.

V. EXPERIMENTS AND VALIDATION

To demonstrate the advantages of DS-CIM in hardware efficiency, flexibility, and accuracy, we implemented the layout designs of DS-CIM1 and DS-CIM2 using a 40nm CMOS process. The customized SNGs and OR-MAC units were designed in Cadence Virtuoso and characterized into standard cells for synthesis in Cadence Liberate and Cadence Abstract, as the method introduced in [28]. After obtaining all the necessary digital standard cells, the digital gate-level netlist was synthesized using Synopsys Design Compiler. Placement and routing were executed in Cadence Innovus, and post-layout simulation was performed with Cadence Virtuoso.

For software evaluation, we conducted comprehensive experiments across quantized DNN and transformer models to assess the impact of stochastic error on model accuracy. The DS-CIM error pattern was added to the MVM results of the models. We tested ResNet18 model on CIFAR-10 and a larger ResNet50 on ImageNet for accuracy evaluation under INT8 quantization, respectively.

We also employed the open-source LLM-FP4 quantization framework [29] to apply the FP8 quantization to the LLaMA-7B model, assessing DS-CIM’s effectiveness on mainstream datasets. Additionally, following the method outlined in [30], FP8 activations and weights were aligned to INT8 with a granularity of 128 as inputs for DS-CIM.

A. Hardware Evaluations

We first evaluate the power performance of the two designs, as well as their sensitivity to signed operations. As shown in Figure 7, the data remapping ensures that there is at most a single 1 in the input of each OR gate, achieving ultra-low power consumption. Due to the coarser OR-MAC granularity, simpler addition and lower accumulation bandwidth compared to DS-CIM1, DS-CIM2 is well compatible with latch-cached accumulators and consumes less energy in the MAC circuit than DS-CIM1. In signed operations, we transform the signed data to unsigned [0, 255], which reduces the sparsity of the stochastic bitstream, and increases power consumption because of more 1s to be accumulated.

As depicted in Figure 7, our evaluation indicates that DS-CIM1 is sensitive to signed operations, whereas DS-CIM2, with OR-MAC64 and simple addition logic, demonstrates resilience to signed operations. Additionally, we minimize the power consumption of PRNGs by reusing them for all weights and activations. However, SNGs and accumulators still contribute significantly to dynamic power consumption. For the area, we observe that a high portion of adders in DS-CIM1, which also confirms that increasing the granularity of the OR-MAC reduces the overhead of adder and accumulator logic, resulting in higher efficiency. Additionally, by duplicating the OR-MAC column 64 times, we have amortized the overhead of the weight SNGs, keeping its area and power proportions within a reasonable range.

B. Accuracy Evaluations

Table I shows that the RMSE is influenced by OR-MAC granularity and bitstream length, revealing a trade-off between accuracy and efficiency. The conclusions align with our initial intuitions. **First**, increasing bitstream could enhance accuracy; however, this improvement comes at the cost of degraded efficiency due to additional computing cycles. **Second**, DS-CIM1 demonstrates an accuracy advantage, while DS-CIM2 exhibits superior efficiency. **Notably**, with a bitstream length=256, DS-CIM1 achieves an accuracy of 94.45%, with only a 0.09% drop compared to the accurate model, while DS-CIM2 attains an accuracy of 94.31%, reflecting a drop of 0.23%. We also showcase DS-CIM capability on larger ResNet50_Weights.IMAGENET1K_V2 on ImageNet,

TABLE II: DS-CIM Performance on FP8 Quantized LLaMA-7B Model

LLaMA-7B Model	BoolQ	PIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	Avg.	Accuracy Degradation
Full Precision	75.1%	78.7%	56.9%	69.9%	75.3%	41.9%	66.3%	0
FP8 Quantized	75.6%	78.2%	56.6%	70.2%	74.6%	40.7%	66.0%	0.3% from full precision
DS-CIM1	71.7%	74.7%	52.1%	67.4%	70.3%	39.4%	62.6%	3.4% from FP8
DS-CIM2	67.1%	70.7%	47.5%	65.4%	66.3%	35.6%	58.8%	7.2% from FP8

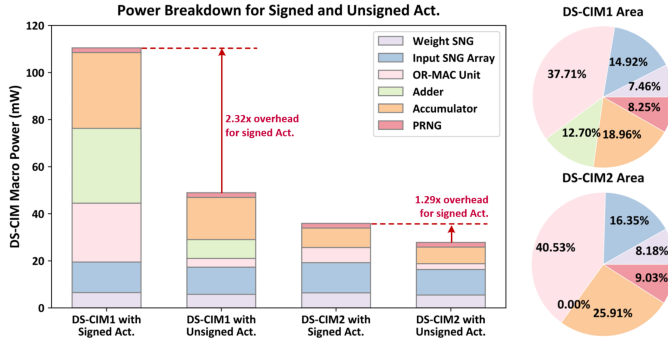


Fig. 7: Power and Area breakdown of DS-CIM1 and DS-CIM2 for signed and unsigned activations, where DS-CIM2 is equipped with the latch-cached accumulator.

TABLE III: Comparison with SOTA Approximate Works

	ISSCC'23 [22]	DIMCA1 [21]	ISSCC'24 [18]	VLSI'24 [27]	DS-CIM1	DS-CIM2
Technology	28nm	28nm	65nm	12nm	40nm	
Macro Type	Digital	Digital	Analog	Digital Stochastic	Digital Stochastic	
Precision	A: 2~8b W: 2~8b	A: 1~4b W: 1b	A: 4~8b W: 4~8b	A: 2/6b W: 6b	A: 8b W: 8b	
Array Size	16Kb	16Kb	80Kb	15Kb	32Kb	
Area (mm ²)	0.028	0.049	0.48	0.5	0.78	0.72
Voltage (V)	0.54~0.9	0.45~1.1	0.6~1.1	0.64	0.7~1.2	
Input	18% toggle	25% toggle	N/A	N/A	87.5% sparsity	
TOPS/mm ² (¹)	131.7	192.2	33.0	22.0	117.1 ⁽²⁾ ~468.4 ⁽³⁾	90.9 ⁽²⁾ ~363.7 ⁽³⁾
TOPS/W(¹)	4569.6	694.4	6652.8	928.8	669.7 ⁽²⁾ ~2677.2 ⁽³⁾	891.5 ⁽²⁾ ~3566.1 ⁽³⁾
Retraining	N/A	Yes	N/A	N/A	Not Needed	
ResNet18/20 @CIFAR-10 ⁽⁴⁾	91.55%	90.41%	91.70%	N/A	90% ⁽³⁾ ~94.45% ⁽²⁾	89.46% ⁽³⁾ ~94.31% ⁽²⁾
ResNet50 @ImageNet ⁽⁴⁾	74.5%	N/A	N/A	N/A	79.45% ⁽³⁾ ~80.67% ⁽²⁾	79.34% ⁽³⁾ ~80.65% ⁽²⁾
Transformer Models	N/A	N/A	ViT-Small @CIFAR10: 95.8%	N/A	LLaMA-7B @WG: 67.4% ⁽²⁾	LLaMA-7B @WG: 65.4% ⁽²⁾

⁽¹⁾ Scaled to 40nm, 1b input and 1b weight.

⁽²⁾ At bitstream length=256 for best accuracy.

⁽³⁾ At bitstream length=64 for best efficiency.

⁽⁴⁾ Baseline accuracies are incomplete, and the model versions may also be different.

which is an improved ResNet50 model by using a new training recipe, showing higher accuracy and robustness than previous versions. Particularly, with a bitstream length=64, DS-CIM1 experienced only a 1.37% drop in accuracy, while DS-CIM2 had just a 1.48% drop. This not only indicates that DS-CIM performs well on large DNNs but also highlights how continuously evolving algorithms are becoming increasingly favorable for approximate computing.

In Table II, we present the accuracy of the FP8 LLaMA-7B model across different datasets, along with using DS-CIM for inference. We aligns FP8 to INT8 [30], and applied the error pattern from DS-CIM with a bitstream length=256 to MAC outputs in the forward propagation of LLaMA-7B. Therefore, the error sources for DS-CIM here include two aspects: one from the process of aligning FP8 to INT8, and the other from the stochastic errors of DS-CIM. Table II shows that DS-

CIM maintains good accuracy in complex transformer models, where DS-CIM1 showing only a 3.4% average accuracy drop. This demonstrates that the application scenarios for DS-CIM have been further expanded.

The analysis shows that eliminating 1s saturation errors in OR accumulation, combined with mutual cancelation of errors across different rows, leads to a uniform and stable error distribution for DS-CIM. This enables a high accuracy comparable with digital adders, making DS-CIM suitable for various application scenarios.

C. Comparison with State-of-the-Art Works

Table III compares the proposed DS-CIM1 and DS-CIM2 with state-of-the-art approximate CIM designs. DS-CIM1 achieves a peak accuracy of 94.45% with a bitstream length of 256 and a comparable energy efficiency of 669.7TOPS/W. Meanwhile, DS-CIM2 demonstrates greater power efficiency, with an energy efficiency of 3566.1 TOPS/W and an area efficiency of 363.7 TOPS/mm² at a bitstream length of 64, while maintaining an accuracy of 89.46%. DS-CIM also supports larger models, such as ResNet50 on ImageNet and the transformer model LLaMA-7B. Its flexible bitstream scheduling allows DS-CIM to adjust precision and efficiency for different applications.

VI. CONCLUSION

In conclusion, the DS-CIM architecture effectively merges high accuracy and efficiency in stochastic DCIM, overcoming challenges related to bitstream lengths and traditional DCIM limitations. By executing signed MAC operations on a unipolar OR-MAC column and replicating low-cost OR-MAC circuits, we achieve a 32 \times increase in compute density with only a 1 \times extra overhead in area. Our stochastic process analysis and data remapping techniques effectively address the 1s saturation issue in OR accumulation, enhancing accuracy and resilience to variations in sparsity. DS-CIM1 achieves an impressive ResNet18 inference accuracy of 94.45% with a minimal RMSE of 0.74% on the CIFAR-10 dataset, while DS-CIM2 delivers exceptional energy efficiency of 3566.1 TOPS/W, maintaining a competitive RMSE of 3.81%. Experiments with INT8 ResNet50 on ImageNet and FP8 LLaMA-7B also comprehensively demonstrate DS-CIM's capability with large models. This work demonstrates the potential of DS-CIM in advancing energy-efficient, high-performance, and precision-adjustable approximate computing systems.

REFERENCES

- [1] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [2] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [3] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [4] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016.
- [5] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [6] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [7] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [8] Peter Schober, M Hassan Najafi, and Nima TaheriNejad. High-accuracy multiply-accumulate (mac) technique for unary stochastic computing. *IEEE Transactions on Computers*, 71(6):1425–1439, 2021.
- [9] Ethan G Rogers, Sohan Salahuddin Mugdho, Kshemal Kshemendra Gupte, and Cheng Wang. Stox-net: Stochastic processing of partial sums for efficient in-memory computing dnn accelerators. *arXiv preprint arXiv:2407.12378*, 2024.
- [10] Aokun Hu, Wenjie Li, Dongxu Lyu, and Guanghui He. Efficient parallel stochastic computing multiply-accumulate (mac) technique using pseudo-sobol bit-streams. *IEEE Transactions on Nanotechnology*, 2024.
- [11] Christiam F Frasser, Pablo Linares-Serrano, Iván Díez de los Ríos, Alejandro Morán, Erik S Skibinsky-Gitlin, Joan Font-Rosselló, Vincent Canals, Miquel Roca, Teresa Serrano-Gotarredona, and Josep L Rosselló. Fully parallel stochastic computing hardware implementation of convolutional neural networks for edge computing applications. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12):10408–10418, 2022.
- [12] Amogh Agrawal, Indranil Chakraborty, Deboleena Roy, Utkarsh Saxena, Saima Sharmin, Minsuk Koo, Yong Shim, Gopalakrishnan Srinivasan, Chamika Liyanagedera, Abhronil Sengupta, et al. Revisiting stochastic computing in the era of nanoscale nonvolatile technologies. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 28(12):2481–2494, 2020.
- [13] Armin Alaghi and John P Hayes. Fast and accurate computation using stochastic circuits. In *2014 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1–4. IEEE, 2014.
- [14] Wenlun Zhang, Shimpei Ando, Yung-Chin Chen, Satomi Miyagi, Shinya Takamaeda-Yamazaki, and Kentaro Yoshioka. Pacim: A sparsity-centric hybrid compute-in-memory architecture via probabilistic approximation. *arXiv preprint arXiv:2408.16246*, 2024.
- [15] Yu-Der Chih, Po-Hao Lee, Hidehiro Fujiwara, Yi-Chun Shih, Chia-Fu Lee, Rawan Naous, Yu-Lin Chen, Chieh-Pu Lo, Cheng-Han Lu, Haruki Mori, et al. 16.4 an 89tops/w and 16.3 tops/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)*, volume 64, pages 252–254. IEEE, 2021.
- [16] Hidehiro Fujiwara, Haruki Mori, Wei-Chang Zhao, Mei-Chen Chuang, Rawan Naous, Chao-Kai Chuang, Takeshi Hashizume, Dar Sun, Chia-Fu Lee, Kerem Akarvardar, et al. A 5-nm 254-tops/w 221-tops/mm² fully-digital computing-in-memory macro supporting wide-range dynamic-voltage-frequency scaling and simultaneous mac and write operations. In *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 65, pages 1–3. IEEE, 2022.
- [17] Hidehiro Fujiwara, Haruki Mori, Wei-Chang Zhao, Kinshuk Khare, Cheng-En Lee, Xiaochen Peng, Vineet Joshi, Chao-Kai Chuang, Shu-Huan Hsu, Takeshi Hashizume, et al. 34.4 a 3nm, 32.5 tops/w, 55.0 tops/mm² and 3.78 mb/mm² fully-digital compute-in-memory macro supporting int12× int12 with a parallel-mac architecture and foundry 6t-sram bit cell. In *2024 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 67, pages 572–574. IEEE, 2024.
- [18] Kentaro Yoshioka. 34.5 a 818-4094tops/w capacitor-reconfigured cim macro for unified acceleration of cnns and transformers. In *2024 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 67, pages 574–576. IEEE, 2024.
- [19] Kodai Ueyoshi, Ioannis A Papistas, Pouya Houshmand, Giuseppe M Sarda, Vikram Jain, Man Shi, Qilin Zheng, Sebastian Giraldo, Peter Vranx, Jonas Doevenspeck, et al. Diana: An end-to-end energy-efficient digital and analog hybrid neural network soc. In *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 65, pages 1–3. IEEE, 2022.
- [20] Mengjie Li, Hongyi Zhang, Siqi He, Haozhe Zhu, Hao Zhang, Jinglei Liu, Jiayuan Chen, Zhenping Hu, Xiaoyang Zeng, and Chixiao Chen. A 19.7 tflops/w multiply-less logarithmic floating-point cim architecture with error-reduced compensated approximate adder. In *2024 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2024.
- [21] Chuan-Tung Lin, Dewei Wang, Bo Zhang, Gregory K Chen, Phil C Knag, Ram Kumar Krishnamurthy, and Mingoo Seok. Dimca: An area-efficient digital in-memory computing macro featuring approximate arithmetic hardware in 28 nm. *IEEE Journal of Solid-State Circuits*, 2023.
- [22] Yifan He, Haikang Diao, Chen Tang, Wenbin Jia, Xiyuan Tang, Yuan Wang, Jinshan Yue, Xueqing Li, Huazhong Yang, Hongyang Jia, et al. 7.3 a 28nm 38-to-102-tops/w 8b multiply-less approximate digital sram compute-in-memory macro for neural-network inference. In *2023 IEEE International Solid-State Circuits Conference (ISSCC)*, pages 130–132. IEEE, 2023.
- [23] Mohsen Riahi Alam, M Hassan Najafi, Nima TaheriNejad, Mohsen Imani, and Lu Peng. Stochastic computing for reliable memristive in-memory computation. In *Proceedings of the Great Lakes Symposium on VLSI 2023*, pages 397–401, 2023.
- [24] Mohsen Riahi Alam, M Hassan Najafi, Nima Taherinejad, Mohsen Imani, and Raju Gottumukkala. Stochastic computing in beyond von-neumann era: Processing bit-streams in memristive memory. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 69(5):2423–2427, 2022.
- [25] Shady Agwa and Themis Prodromakis. Digital in-memory stochastic computing architecture for vector-matrix multiplication. *Frontiers in Nanotechnology*, 5:1147396, 2023.
- [26] Jiyue Yang, Tianmu Li, Wojciech Romaszkan, Puneet Gupta, and Sudhakar Pamarti. A 65nm 8-bit all-digital stochastic-compute-in-memory deep learning processor. In *2022 IEEE Asian Solid-State Circuits Conference (A-SSCC)*, pages 10–11. IEEE, 2022.
- [27] Jiyue Yang, Alexander Graening, Wojciech Romaszkan, Vinod K Jacob, Puneet Gupta, and Sudhakar Pamarti. A 278-514m event/s adc-less stochastic compute-in-memory convolution accelerator for event camera. In *2024 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*, pages 1–2. IEEE, 2024.
- [28] Kunming Shao, Fengshi Tian, Xiaomeng Wang, Jiakun Zheng, Jia Chen, Jingyu He, Hui Wu, Jinbo Chen, Xihao Guan, Yi Deng, et al. Synd-cim: A performance-aware digital computing-in-memory compiler with multi-spec-oriented subcircuit synthesis. In *2025 Design, Automation & Test in Europe Conference (DATE)*, pages 1–7. IEEE, 2025.
- [29] Shih-yang Liu, Zechun Liu, Xijie Huang, Pingcheng Dong, and Kwang-Ting Cheng. Llm-fp4: 4-bit floating-point quantized transformers. *arXiv preprint arXiv:2310.16836*, 2023.
- [30] Fengbin Tu, Yiqi Wang, Zihan Wu, Ling Liang, Yufei Ding, Bongjin Kim, Leibo Liu, Shaojun Wei, Yuan Xie, and Shouyi Yin. Redcim: Reconfigurable digital computing-in-memory processor with unified fp/int pipeline for cloud ai acceleration. *IEEE Journal of Solid-State Circuits*, 58(1):243–255, 2022.