

GAMA: High-Performance GEMM Acceleration on AMD Versal ML-Optimized AI Engines

Kaustubh Manohar Mhatre
Arizona State University
Tempe, USA
kmhatre@asu.edu

Endri Taka
The University of Texas at Austin
Austin, USA
endri.taka@utexas.edu

Aman Arora
Arizona State University
Tempe, USA
aman.kbm@asu.edu

Abstract—General matrix-matrix multiplication (GEMM) is a fundamental operation in machine learning (ML) applications. We present the first comprehensive performance acceleration of GEMM workloads on AMD’s second-generation AIE-ML architecture, which is specifically optimized for ML applications. Compared to AI-Engine (AIE), AIE-ML offers increased compute throughput and larger on-chip memory capacity. We propose a novel design that maximizes AIE-ML memory utilization, incorporates custom buffer placement within the AIE-ML and staggered kernel placement across the AIE-ML array, significantly reducing performance bottlenecks such as memory stalls and routing congestion, resulting in improved performance and efficiency compared to the default AMD’s compiler. We evaluate the performance benefits of our design at three levels: single AIE-ML, pack of AIE-ML’s and the complete AIE-ML array. GAMA achieves state-of-the-art performance, delivering up to 165 TOPS (85% of peak) for int8 precision and 83 TBFLOPS (86% of peak) for bfloat16 precision GEMM workloads. Our solution achieves 8.7%, 9%, 39% and 53.6% higher peak throughput efficiency compared to the state-of-the-art AIE frameworks AMA, MAXEVA, ARIES and CHARM, respectively.

Index Terms—AMD Versal, AI, ML, AI Engine, Matrix Multiplication, Hardware Acceleration, FPGA

I. INTRODUCTION

AMD Versal is a heterogeneous architecture that integrates Adaptive Intelligent Engines (AIE), programmable logic (PL) and processing system (PS) which consists of CPUs. Such architectures can enable high performance for computationally intensive applications such as machine learning (ML), which are mainly dominated by General Matrix Multiplications (GEMM). While prior research has extensively explored the performance of GEMM on the first-generation AI Engine (AIE) [3], no existing work has investigated its performance on the second-generation AIE-ML Engine [5].

Compared to AIE, the AIE-ML¹ architecture has larger AIE memory, 2x throughput for int8 precision, and supports bfloat16 precision. Design frameworks from prior work such as MaxEVA [14], AMA [8], CHARM [20] exploit the architecture of AIE achieving significant peak performance. These frameworks, however, have architectural dependencies that limit their performance gains to AIE and make them not directly adaptable to AIE-ML.

To efficiently exploit the available enhanced compute throughput of the AIE-ML, it is crucial to first maximize the single AIE performance and then ensure scaling to maximize

the AIE array utilization. At the single AIE level, maximizing performance requires running largest possible GEMM sizes, which in turn requires maximizing AIE memory utilization to fit those sizes. Naively maximizing AIE memory utilization often results in buffer allocation to the nearby AI engines, making it challenging to scale the design to the whole array. Restricting the buffers to a single AIE to facilitate scaling often results in stalls due to the compiler’s default buffer placement strategies. No prior work focuses on buffer placement since none of them tries to maximize AIE memory utilization.

Another critical factor to maximize performance is to maximize the utilization of the AIE array. Due to limited PL-AIE interface resources and numerous ways of connecting the AIEs, scaling to the AIE array is non-trivial. For example, scaling a single AIE design to the complete array might not be feasible due to limited PLIO resources. Generally, the single AIE design is scaled up to a group of AIEs which we call a “pack”, and then that pack is replicated across the AIE array. Extensive design space exploration is required to find the optimal connectivity within a pack. Coupled with the large design space, automatic AIE kernel placement and buffer allocation at the full array level can result in compilation failures or lead to inefficient solutions with lower AIE array utilization. In particular, when scaling designs to the complete array, we observe compilation failures due to placement errors and routing congestion. In this paper, we propose **GAMA**, a framework to address these challenges by systemically analyzing them and identifying solutions at three levels.

At the *single AIE level*, we propose a custom buffer placement algorithm that maximizes the AIE memory utilization (100% in some cases) and minimizes memory stalls while keeping the buffers in a single AIE. At the *pack level*, we determine the optimal connectivity and pack size that guarantees scalability and performance, while taking into account the constraints posed by the limited PL-AIE interface resources. Using custom buffer placement to keep buffers within the pack, we make our pack easily scalable. Finally, at the *array level* we replicate the pack by proposing a staggered placement pattern overcoming the compiler’s routing congestion issues resulting in 94% AIE array utilization. Furthermore, our custom buffer and kernel placement techniques, alleviate the efforts required by the compiler, resulting in a significant reduction in compilation time.

This work makes the following main contributions:

¹In this paper, for ease of reading, we use the word AIE to refer to AIE-ML.

- GAMA is the first comprehensive study and implementation of GEMM acceleration on the AMD Versal AIE-ML architecture. It adopts a scalability-first approach that emphasizes scalability over the performance of individual kernels, promoting regularity in kernel mapping to facilitate straightforward replication.
- GAMA maximizes the internal memory usage of the AIE, achieving up to **100%** utilization — the highest reported compared to all previous work. It effectively addresses memory stall issues through a novel custom buffer placement algorithm, reducing stalls by an average of **12%**, outperforming the standard compiler optimization strategies provided by AMD’s compiler.
- GAMA achieves very high AIE array utilization of **94%** while minimizing efficiency loss through a custom AIE kernel placement strategy that avoids PLIO routing congestion. Furthermore, custom buffer and kernel placement methods significantly accelerate compilation time by **6x**.
- This is the first work on GEMMs on Versal that supports multiple output precision (int8, int16, int32) for int8 inputs to accommodate various accuracy requirements and the first to show performance on bfloat16 on the AIE, achieving **86%** peak throughput utilization of the chip. For int8, we achieve up to **85%** of the chip’s peak throughput utilization, surpassing prior work’s performance by up to **39%**.
- We open-source GAMA for other researchers to leverage the framework at: <https://github.com/advent-lab/GAMA>.

II. RELATED WORK

Since the introduction of AI engines by AMD [2], a wide variety of workloads have been deployed on them [7], [9], [11], [13], [17]–[19], [23]. Additionally, several frameworks have mapped and analyzed the performance of GEMM on these AI engines. Table I summarizes the differences between these and GAMA. CHARM [20], [21] and AutoMM [24] accelerated GEMM on AIE using a resource conservative approach which enables efficient scaling but results in limited performance. CHARM and AutoMM show performance for the fp32 input and fp32 output (fp32-fp32) [20] and int8-int8 and int16-int16 [24] precisions, respectively. GAMA maximizes resource utilization (such as AIEs and AIE-PL interfaces) to achieve maximum performance. MaxEVA [14] (fp32-fp32, int8-int32), on the other hand, maximizes GEMM performance while using significant resources on the chip. It uses two different kernels for matrix multiplication and reduction (addition). This limits the overall efficiency of MaxEVA to 80% since only 80% of the AIEs in the design can perform the matrix multiplication operation (20% are used for addition). On the other hand, GAMA uses a single kernel that performs both matrix multiplication and reduction, similar to CHARM. GAMA utilizes 94% of the AIE array for matrix multiplication. It uses a cascade interface for partial sum transfer between AIEs, unlike MaxEVA which uses buffer sharing between nearby engines for reduction. AMA [8] uses a lower precision output (fp32-fp16, int8-int16) improving

performance at the cost of potential accuracy loss. MaxEVA and AMA provides AIE-only performance considering that input data is available in PL, unlike CHARM [20], which provides a full implementation including PL and PS. RSN-XNN [15] designs an overlay by proposing a reconfigurable stream network, leveraging the flexibility of FPGA.

Some prior work has focused on the compilation for AIEs. Vyasa [6] extends the Halide [12] DSL compiler to automatically generate code for AIEs, improving the programmability. ARIES [22] designs a code generator for AIE on Versal and AIE-ML on AMD Ryzen processors by leveraging the MLIR framework [10]. The ARIES framework supports AIE-ML on Ryzen CPUs, demonstrating performance results for ResNet instead of GEMM. Therefore, a direct comparison with ARIES is not feasible unless modifications are made to enable support for Versal devices. Wierse [16] perform performance analysis of the AIE’s communication interfaces.

Compared to prior work, GAMA is the first framework to show GEMM acceleration on AIE-ML. GAMA has the highest AIE memory utilization with custom buffer placement, a flexible pack size and a staggered kernel placement for efficient scaling.

III. BACKGROUND

The AMD Versal AI Engine is a heterogeneous architecture that combines AI Engines, Programmable Logic (PL), and Processing System (PS). The AI Engines are VLIW vector processors capable of executing multiple instructions per clock cycle. Connected to the AI Engines is an FPGA fabric (or PL) consisting of Lookup Tables (LUTs), flip-flops (FFs), Digital Signal Processing slices (DSPs), Block RAMs (BRAMs), and Ultra RAMs (URAMs). The AI Engines communicate with the PL and DRAM through two interfaces: the PL interface input/output (PLIO) for data transfer between the PL and AI Engines, and NoC interface tiles for data transfer between the AI Engines and DRAM. The AI Engine array architecture provides us with Buffer (an AIE can directly access a neighboring AIE’s memory), Cascade (direct AIE to AIE connectivity) and Via-switch connections to communicate from one AIE to another. The application processor can run a Linux operating system and can be used to control both PL and AI Engines. In this work, we use the 2nd generation of the AI Engine called AIE-ML. We further discuss the differences compared to AIE.

Figure 1 illustrates the Versal AIE-ML architecture. At the top are the AIE-ML cores that offer an int8 throughput of 256 MACs/cycle, which is double that of AIE. Additionally, AIE-ML features twice the memory of 64 KB compared to 32 KB in AIE. Although AIE-ML does not support fp32 computations, it incorporates support for bfloat16, contrasting AIE, which lacks bfloat16 but has fp32 support. AIE-ML also includes a wider cascade bus of 512 bits compared to 384 bits in AIE. Although cascade connectivity remains unidirectional, AIE-ML now supports vertical cascading as well. AIE-ML also provides us with addition memory storage in the form of memory tiles.

TABLE I
COMPARING PRIOR WORK FOCUSING ON GEMM ACCELERATION ON VERSAL ARCHITECTURE WITH THE PROPOSED GAMA FRAMEWORK

Work	AIE Ver- sion	Precision (ip-op)	Pack Size	Max used AIEs	Cascade reduction	for	AIE kernel placement	AIE buffer placement	AIE memory utilization
CHARM [20]	AIE	fp32-fp32	4 (fp32)	384/400 (96%)	Yes		Manual (Horizontal)	Custom	Low (75%)
AutoMM [24]	AIE	int16-int16, int8-int8	2 (int8)	288/400 (72%), 192/400 (48%)	Yes		Manual (Horizontal)	Custom	Low (75%)
MaxEVA [14]	AIE	fp32-fp32, int8-int32	3,4	400/400 (100%), 400/400 (100%)	No		Manual (Tetris)	Auto (Compiler default)	Low (75%)
AMA [8]	AIE	fp32-fp16, int8-int16	3,4	342/400 (85%), 342/342 (85%)	No		Manual (Horizontal)	Auto (Compiler default)	Low (75%)
RSN-XNN [15]	AIE	fp32-fp32	4	384/400 (96%)	Yes		Manual (Horizontal)	-	Low (75%)
ARIES [22]	AIE	fp32-fp32, int16-int16, int8-int8	-	352/400 (88%), 352/400 (88%), 320/400 (80%)	Yes		Manual (Horizontal)	Custom	Low (75%)
GAMA (Ours)	AIE-ML	int8-int32, int8-int16, int8-int8, bf16-bf16	2-38	288/304 (94%), 288/304 (94%), 288/304 (94%), 288/304 (94%)	Yes		Manual (Staggered horizontal)	Custom	High (97%)

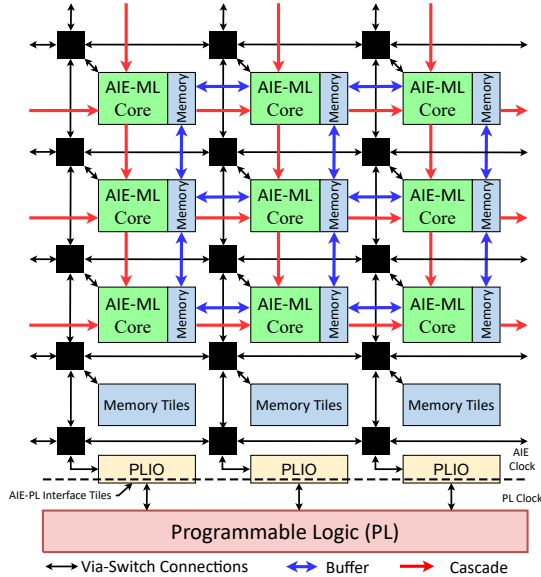


Fig. 1. AMD Versal AIE-ML Architecture.

IV. GAMA FRAMEWORK

In this section, we discuss the GAMA framework in detail. We assume that the input/output data is available in the PL and use the PLIO interface for communication. We use AMD’s Versal VE2802 device which has 304 engines in a grid of 8 rows x 38 columns. This device has 112 input PLIOs (PL to AIE) and 84 output PLIOs (AIE to PL). Each PLIO is configured to be 128 bit wide. Matrix multiplication workload is specified with the notation M, K, and N. Where, matrix A has dimensions $M \times K$, matrix B has dimensions $K \times N$, and the resulting matrix has dimensions $M \times N$.

A. Single AIE kernel design

Kernel size: Our kernels are designed using AMD APIs leveraging the generic (templated) matrix multiplication kernel

from AMD’s user API guide [1], thereby ensuring their portability across a variety of Versal devices. We support multiple sizes, i.e., the M, K, and N matrix dimensions, for a single kernel. Although our kernel code supports variety of M, K and N sizes, selecting the right size is crucial for high performance and AIE kernel compute utilization. The kernel size selection is governed by the input/output precision, the AIE memory limit and balance between compute and PL-AIE communication, which is defined by the compute-to-communication ratio γ . Equation 1 calculates the theoretical kernel compute cycles based on the workload dimension and the peak throughput of the chip for a specific precision. Equations 2, 3 and 4 calculate the cycles required to transfer the input matrix A and matrix B from PL to AIE and the output matrix C from AIE to PL, where $sizeof()$ depends on the precision used. Each AIE has two input PLIO and one output PLIO connections. This allows simultaneous pipelined read, compute, and write operations. Equation 5 shows the calculation for γ . A ratio < 1 means the workload is PLIO bandwidth bound, while > 1 means it is compute bound.

$$Compute_cycles = \frac{M \times K \times N}{Peak_MACs} \quad (1)$$

$$Comm_A = \frac{M \times K \times sizeof(input)}{\left(\frac{PLIO_width}{8}\right)} \quad (2)$$

$$Comm_B = \frac{K \times N \times sizeof(input)}{\left(\frac{PLIO_width}{8}\right)} \quad (3)$$

$$Comm_C = \frac{M \times N \times sizeof(output)}{\frac{PLIO_width}{8}} \quad (4)$$

$$\gamma = \frac{\text{Compute_cycles}}{\max(\text{Comm_A}, \text{Comm_B}, \text{Comm_C})} \quad (5)$$

The internal memory of the AIE presents a constraint on the size of M, K and N as shown by Equation 6. Since ping-pong buffering is used to overlap computation in AIE and the AIE-PL communication, the memory requirement is doubled.

$$\begin{aligned} &M \times K \times \text{sizeof}(\text{input}) + \\ &K \times N \times \text{sizeof}(\text{input}) + \\ &M \times N \times \text{sizeof}(\text{output}) \times 2 \leq 64 \text{ KB} \end{aligned} \quad (6)$$

We perform an exhaustive search for M, K and N that satisfies the constraints in Equation 6 and has high γ .

API sizing: AMD’s API library [1] provides an API for matrix multiplication called MMUL, which performs blocked matrix multiplication $C = A \times B$. This API is defined as `aie::mmul<M_Elems, K_Elems, N_Elems, TypeA, TypeB, AccumTag>` where `M_Elems`, `K_Elems`, `N_Elems` define the size. MMUL supports up to 9 sizes such as 4x4x4, 8x8x4 etc. We use the sizes M, K, and N resulting from the constraints mentioned in the previous sub-section and then evaluate all available MMUL API sizes for the specified precision by performing a sweep. We then select the MMUL API size that gives the best performance.

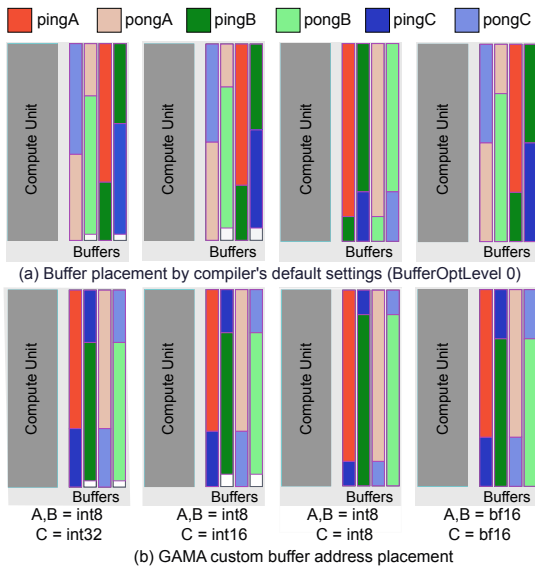


Fig. 2. GAMA’s custom buffer address placement within an AI engine minimizes stalls and improves performance.

Buffer placement: Each AIE has 64 KB of AIE memory divided into 4 banks of 16 KB each. Buffer placement or allocation assigns an address to a buffer inside the AIE memory’s banks. Optimal buffer placement in the AIE memory is crucial to prevent bank conflicts and subsequent memory stalls, thus improving kernel performance. The AIE programming interface provides users with APIs to control buffer placement. APIs can constrain the buffers to a certain AIE kernel location using `location<buffer>(buf_A) = location<kernel>(AIE_kernel_loc);` or can place

the buffers at a specific address in the AIE’s memory using `location<buffer>(buf_A) = {address(row, col, <addr>), address(row, col, <addr>)};` providing a more finer control over the placement. We refer to the former as “buffer location placement” and the latter as “buffer address placement”. AMD also provides users with various compiler optimization flags to optimize buffer placement automatically. Flags such as `BufferOptLevel` [4] are used to optimize buffer allocation. The `BufferOptLevel` flag has 10 levels (0 to 9); increasing the level progressively increases optimizations made by the compiler to place the buffers to minimize memory stalls. These flags are effective without buffer placement constraints and with small buffers.

Every AIE’s kernel can directly access buffers located in its neighboring AIEs, enabling buffer placement in the neighboring AIEs. Constraining buffers using buffer location placement makes it easy to scale the design to the AIE array. With the buffer size that maximally utilizes the AIE memory and the above constraint, `BufferOptLevel 0` needs to be used for successful compilation (all other `BufferOptLevel` levels give placement error). But this leads to several memory stalls.

We perform a custom buffer address placement using a simple algorithm that maximizes AIE memory utilization while minimizing stalls. Our custom buffer address placement is guided by the set of rules that determine the buffer location based on the matrix type (input and output) and buffer category (Ping or Pong). Our rules are as follows: (a) Never assign ping and pong buffers of the same matrix to the same bank. (b) Never assign ping and pong buffers of the same matrix to the adjacent bank. (c) Always assign the buffer of matrix A and B to a different bank. Algorithm 1 encapsulates these rules and automates buffer address placement for various matrix sizes and precisions. It performs a simple exhaustive search and generates buffer addresses that satisfy the rules mentioned above. It takes the dimensions of matrix M, K, N, and the input and output precision as inputs and first calculates the size of buffers A, B, and C (lines 1-5). If the sum of the sizes of the buffer exceeds 64 KB then it exits (Line 6). A buffer list is initialized for all the buffers that need address allocation (Line 8). The order of the list is important. We iterate over all banks of to check if the corresponding buffer is suitable for allocation (Lines 10-32). Banks have two spots for buffers, which means that max two buffers can fit in one bank. Matrix A and B are only allocated when there are two spots available in the bank and the adjacent banks do not contain the corresponding double buffers (Lines 13-14). If the constraint is satisfied, the buffer is placed in the bank, and a start address is allocated (lines 16-19). Matrix C can be placed as the second buffer in the banks that contain matrix A or B (Line 21). The matrix C buffer is placed and a start address is assigned to the buffer (Lines 22-26). The sum of both buffers can exceed the capacity (16 KB) of the bank, thus in this scenario the start address of the other buffer in the next bank is shifted by the offset (Lines 28-30). Once finished, the algorithm returns the starting addresses of all six buffers.

Figure 2 (a) shows the compiler’s default buffer placement

with *BufferOptLevel* 0 and (b) shows a result of our custom buffer address placement algorithm. Our custom buffer address placement results in an average 12% fewer cycles compared to automated buffer placement. We use the same algorithm for all buffer placements for the rest of the paper.

Algorithm 1 AIE buffer address placement algorithm

```

1: ip: M, K, N, ip_p, op_p
2: op: buffer_location
3: Initialize: aie_mem.size = 65536 aie_mem.bnks = 4
4: buf_A, buf_B  $\leftarrow M \times K \times ip_p, K \times N \times ip_p$ 
5: buf_C  $\leftarrow M \times N \times op_p$ 
6: if CHK_OFL(M, K, N, ip_p, op_p, aie_mem.size) then
7:   exit
8: buf_list[:]  $\leftarrow$  ping_A, pong_A, ping_B, pong_B, ping_C, pong_C
9: bnks  $\leftarrow$  aie_mem.bnks
10: for buf in buf_list do
11:   for b in range(bnks) do
12:     if buf  $\in$  (Mat_A or Mat_B) then
13:       if IS_ADJACENT(buf, b, aie_mem) OR b.free_spots  $\leq$  1 then
14:         continue
15:       else
16:         bnks[b].buffer  $\leftarrow$  buf
17:         bnks[b].free_space  $- =$  buf.size
18:         bnks[b].free_spots  $- =$ 
19:         buf.start_addr  $\leftarrow$  bnks[b].start_addr
20:     else if buf  $\in$  (Mat_C) then
21:       if bnks[b].free_spot > 0 then
22:         bnks[b].buffer  $\leftarrow$  buf
23:         bnks[b].free_space  $- =$  buf.size
24:         if bnks[b].free_spots = 2 then
25:           buf.s_addr  $\leftarrow$  bnks[b].s_addr
26:         if bnks[b].free_spots = 1 then
27:           buf.s_addr  $\leftarrow$  bnks[b].s_addr + bnks[b].buffer[0].size
28:         if bnks[b].free_space < 0 then
29:           offset  $\leftarrow$  (bnks[b].buffer[0].size +
bnks[b].buffer[1].size) - bnks[b].size
30:           bnks[b+1].buffer.s_addr  $\leftarrow$  bnks[b+1].s_addr + offset
31:           bnks[b].free_spots  $- =$ 
32:         break
33: return buf_list.addr

```

B. Chaining the AIEs together




Fig. 3. Pack of 4 AIEs computing a larger GEMM by using the cascade interface for reductions. Final workload size for the pack in terms of M, K and N of a single AIE kernel $\rightarrow M=M, K=4*K$ and $N=N$

AIE to AIE connectivity: AIEs are connected in series, as shown in Figure 3, to form a pack, allowing larger GEMMs to be implemented with increased parallelism and reuse. There are two PLIOs for input to each AIE in the pack and one PLIO for output from the pack. Reductions are performed inside the pack, and only final results are written back to the PL. This reduces the output PLIOs in the design, since only the last AIE has to write data back to the PL. Cascade, via-switch connection, and buffer are the various methods for partial sum communication from one AIE to the next AIE.




Fig. 4. Efficiency drops as pack size increases. Pack size of 4 has the maximum efficiency among the sizes that can scale well to the whole AIE array (unhatched portion).




Fig. 5. Buffer address placement for a pack of four AIEs that perform a GEMM operation together. Note the output buffers are placed in the 3rd AIE.

CHARM [20], ARIES [22] and RSN-XNN [15] use a cascade interface, while MaxEVA [14] and AMA [8] use buffer. Buffer communication does not stream data from one AIE to another; it uses buffers in the neighboring AIE’s memory. These buffers consume memory. Via-switch and cascade connections stream data from one AIE to another avoiding this memory overhead. Via-switch interface is only 32-bit and using it increases communication latency significantly. AIE architecture has a 512-bit wide cascade interface compared to AIE’s 384-bit. The cascade with the larger bit width and no memory overhead becomes the best candidate for AIE to AIE communication. We use cascade interface for all partial sum communication. Cascade connections make the pack run in a data-flow fashion. **Different kernel types:** The pack includes three distinct kernels. The first kernel (used in the leftmost AIE in the pack) takes Matrix A and Matrix B as inputs from the PLIO and output’s partial sum through cascade interface. The second kernel (used in the middle AIEs in the pack) has two input PLIO ports, one input cascade port, and one output cascade port. The third kernel (used in the last AIE in the pack) receives two inputs from PLIO, one input from the cascade port, and one output PLIO port.

Pack size: Due to the flexibility of the design, any number of AIEs can be chained as long as PLIO resources are available to support the design. We calculate the kernel compute efficiency (KCE) for the pack to gauge the effect of cascade on the design. The average the kernel compute cycles across all the AIE’s in the pack is used to calculate the KCE. Figure 4 shows how the kernel compute efficiency varies as we change the pack size from 2 to 38. The max pack size is 38 in our case because VE2802 (the device we use) only has 38 columns in the AIE grid, and we use horizontal cascade connections. Designs for all of these pack sizes compile and run, but not all of them scale well. The hatched portion in the figure




Fig. 6. Our scaling and mapping strategy shown for hyperparameter settings of $Y=2$, $G=4$, $X=2$. Final workload size in terms of M , K and N of a single AIE kernel $\rightarrow M=Y*M$, $K=G*K$ and $N=X*N$

shows the pack sizes that do not scale to the complete array, as they run out of input PLIO or output PLIO resources (PLIO requirement increases with pack size). Small packs (2-3) are bound by output PLIOs, and larger packs (11-38) are bound by input PLIOs. Thus, only pack sizes of 3-10 are scalable. Although the cascade interface does not have memory overhead and has a large bit width, it incurs cascade stalls. Cascade stalls occur when the producer kernel tries to write data elements that exceed the bit width of the cascade interface. These stalls accumulate as we chain more AIEs, leading to a reduction in kernel compute efficiency as the pack size increases. We select a pack size of 4 as it has the highest kernel compute efficiency and is scalable to the AIE array.

AIEs connected in a pack perform a GEMM operation with matrix sizes larger than the matrix sizes for a single AIE. For example, if the kernel sizes are $M=48$, $K=240$ and $N=48$ then the native size of the pack becomes $M=48$, $K=960$ and $N=48$. **Buffer placement:** We use cascade interface to transfer the data from one AIE to the next AIE in a pack as seen in the previous section. Cascade avoids the need for output buffers. Thus, only the last AIE in the pack generates results for the overall GEMM operation performed by the pack. Figure 5 shows the placement of the input and output buffers for a pack size of four. In the pack of four, the 4th AIE generates the output, but we place the output buffers of the 4th AIE into the 3rd AIE. Doing this further avoids memory stalls. The three AIEs (1st, 2nd, and 4th) have only 4 input buffers each, so they do not need custom buffer address placement. However, the 3rd AIE will have all six buffers (4 input buffers and 2 output buffers) and therefore needs the custom buffer address placement from Section IV-A. All buffers are contained within the pack, simplifying scaling to the AIE array.

C. Scaling to the complete AIE array

To maximize performance, the design of a pack is scaled across the AIE array by replicating the pack. To improve data reuse, we leverage the PLIO broadcast mechanism. This minimizes redundant data transfers and optimizes overall performance. Broadcasting also saves a lot of PLIOs by enabling reuse, which is crucial since PLIO interfaces are limited.

Scaling hyperparameters and constraints: The scaling in our design is governed by three hyperparameters. Y : Replicates the pack along the Y-axis of the AIE array. G : Defines the size of each pack. X : Replicates the pack along the X-axis of the AIE array. Figure 6 shows the scaling for the case ($Y = 2, G = 4, X = 2$). PLIOs are shown using different colors that show broadcast and reuse. Matrix A and matrix B are assigned to 16 AIE as shown in Figure 6.




Fig. 7. A staggered placement is used while replicating pack of kernels across the whole array to avoid routing congestion.

Scaling to the complete AIE array must satisfy two main constraints: (1) The first constraint, shown in Equation 7, ensures compatibility with the geometry of the VE2802 device. (2) The second constraint in Equation 8 ensures the total PLIOs and AIEs do not surpass available resources.

$$\begin{aligned}
 Y &\leq \# \text{of rows in the device} \quad (8) \\
 G \times X &\leq \# \text{of columns in the device} \quad (38) \\
 Y \times G \times X &\leq \text{AIE_cores} \\
 Y \times G + G \times X &\leq \text{PLIO_in} \\
 Y \times X &\leq \text{PLIO_out} \quad (8)
 \end{aligned}$$

Kernel placement: Scaling the design to the entire array by satisfying the constraints stated above leads to compilation failure due to PLIO routing congestion. This happens because the third AIE utilizes three distinct PLIOs (two read, one write), leading to congestion in the switch network's vertical lane when extended across all rows. To avoid this congestion, GAMA uses a zigzag or skewed kernel placement strategy across the array, as shown in Figure 7. We skew the kernel placement by two since skewing by one also leads to congestion, and skewing by three reduces the AIE array utilization. The first two AIEs in each alternate rows of AIEs are not used. The pattern alternates the third AIE's row position, easing vertical switch lane pressure. This placement successfully avoids congestion, leading to a scalable and compilable design.

Note that this manual kernel placement and the custom buffer address placement from the previous subsection reduces the compiler's efforts, lowering the compilation time by 6x.

V. EVALUATION RESULTS

A. Experimental setup

Our experiments are carried out using the VE2802 Versal AIE-ML device on the VEK280 board using Vitis version 2024.1. We utilize the `aiesimulator` tool offered by AMD

TABLE III
SINGLE AIE KERNEL COMPUTE CYCLES (KCC) AND KERNEL COMPUTE EFFICIENCY (KCE) FOR DIFFERENT BUFFER PLACEMENT STRATEGIES

Configuration					KCC (Theoretical)	Unconstrained buff ^{a,\$}		Buffer location placement ^{b,\$}		Buffer address placement ^{c,\$}		
Precision (ip-op)	M	K	N	MMUL API		KCC (Measured)	KCE (%)	KCC (Measured)	KCE (%)	KCC (Measured)	KCE (%)	%Perf recovered
int8-int32	48	240	48	4x8x8	2160	2426	89%	3076	70%	2590	83%	13%
int8-int16	64	184	64	4x8x8	2944	3141	94%	3923	75%	3345	88%	13%
int8-int8	64	224	64	4x8x8	3584	3686	97%	4340	83%	3831	94%	11%
bf16-bf16	64	96	64	8x8x4	3072	3135	98%	3598	85%	3255	94%	9%

^a Unconstrained buffer placement with *BufferOptLevel* 9.

^b Buffers constrained to same AIE using buffer location placement and *BufferOptLevel* 0.

^c Buffers constrained to same AIE using buffer address placement and *BufferOptLevel* 0.

^{\$} Per-AIE memory utilization is shown in Table II

TABLE II
SINGLE AIE COMPUTE TO PLIO COMMUNICATION RATIO (γ) AND AIE MEMORY UTILIZATION

Precision (ip-op)	M	K	N	γ	AIE Mem Usage ^a	AIE Mem UCB ^b (%)	AIE Mem CB ^c (%)
int8-int32	48	240	48	0.72	64512	63%	98%
int8-int16	64	184	64	0.96	63488	25%	97%
int8-int8	64	224	64	0.96	65536	12%	100%
bf16-bf16	64	96	64	0.96	65536	25%	100%

^a AIE Memory usage for the M, K and N size matrix

^b Per-AIE Memory utilization for unconstrained buffer placement

^c Per-AIE Memory utilization for constrained buffer location and address placement

to simulate our designs and measure performance. AIEs operate at a frequency of 1.25 GHz, while the PL frequency is set to 300 MHz. A 128-bit PLIO configuration is used for both input and output throughout the experiments. AIE-to-AIE communication is done using the 512-bit cascade interface.

B. Performance of single AIE kernels

At the single AIE level, the two most important aspects are the MMUL API size and the single AIE kernel size (MxKxN). As discussed in Section IV-A, the size of a single AIE kernel depends on factors such as the compute-communication ratio γ , AIE memory usage and input/output precision.

We do an exhaustive search for M, K and N sizes, and then select the ones with the highest value of γ while satisfying the memory requirements stated in Equation 6. Different precisions result in different sizes for M,K and N.

Table II presents the results of this exhaustive search. For each precision, we show the kernel size (M, K, and N) along with the compute-to-communication ratio γ and the internal memory usage. Our constrained buffer address placement achieves 98%, 97%, 100%, and 100% AIE memory utilization for int8-int32 (ip-op), int8-int16, int8-int8, and bf16-bf16 precisions, respectively. GAMA achieves the highest memory utilization compared to all other prior work [8], [14], [15], [20]. All precisions except int8-int32 achieve a 0.96 compute-to-communication ratio, as int8-int32's larger output size increases communication latency.

To calculate the efficiency of the kernel, we measure the kernel compute cycles (KCC) and then take the ratio of theoretical KCC to measured KCC and call it kernel compute efficiency (KCE). We use KCE as a metric to find the best MMUL API

size. We sweep all MMUL sizes to find the optimal M, K, and N sizes, as noted in Section IV-A. Table III shows the KCE for the best respective MMUL size for each precision. The table compares the KCC and KCE for three scenarios: unconstrained buffers (*BufferOptLevel* 9), buffers constrained to the same AIE using buffer location placement (*BufferOptLevel* 0), and then buffer constrained to the same AIE using our buffer address placement. Unconstrained single AIE buffers show the best performance with the highest average KCE of 94%. However, as mentioned earlier, this is not scalable because the buffers can be allocated in neighboring AIEs by the compiler. It significantly reduces AIE memory utilization as shown in Table II. We use this best-case performance as a baseline for further comparison. Buffers constrained to the same AIE using buffer location placement and *BufferOptLevel* 0 shows an average 16% drop in KCE. The observed drop occurs due to the random fragmentation of buffers across the AIE banks. However, buffers constrained to the same AIE using our custom buffer address placement recover on average 12% of this loss, bringing the average KCE to 90%, while maximizing AIE memory utilization and ensuring scalability.

C. Performance of a pack of AIEs

Figure 3 shows the AIE connected as a pack. Each AIE in the pack uses the kernel size from the previous section, optimizing compute to communication ratio γ , AIE memory utilization, and KCE. We use a pack size of 4 as discussed in Section IV-B. Table IV shows the effect of different placement constraints on the pack. Unconstrained buffers is considered as the baseline with the highest average KCE of 88%. Buffers constrained to the pack using buffer location placement show 14% loss in KCE. Our custom buffer placement recovers 12% of the loss, bringing the average KCE to 86%.

Table IV also shows the efficiency loss incurred as a result of the cascade interface. We use the KCC of unconstrained single AIE buffer from Table III as baseline to compare with the KCC of unconstrained buffer cycles of the pack. Since unconstrained buffer placement avoids any memory stalls, the efficiency loss can be characterized as a loss due to cascade. On average, we see a 7% reduction in KCE due to cascade stalls, compared to single AIE case, across all precisions.

D. Scaling to the complete array

With custom buffer address placement, resulting in buffers being placed within a pack, we are able to achieve easy and

TABLE IV
BUFFER PLACEMENT’S EFFECT ON PERFORMANCE FOR A PACK OF 4 AIEs

Configuration				Unconstrained buff ^a			Buffer location placement ^b		Buffer address placement ^c		
Precision (ip-op)	M	K	N	KCC*	KCE (%)	% Cascade stalls	KCC*	KCE (%)	KCC*	KCE (%)	% performance recovered
int8-int32	48	960	48	2665	81%	9%	3198	68%	2711	80%	12%
int8-int16	64	736	64	3326	89%	6%	4126	71%	3419	86%	15%
int8-int8	64	896	64	3980	90%	7%	4273	84%	4009	89%	6%
bf16-bf16	64	384	64	3361	91%	7%	4340	71%	3404	90%	19%

^a Unconstrained buffer placement with *BufferOptLevel* 9.

^b Buffers constrained to same AIE using buffer location placement and *BufferOptLevel* 0.

^c Buffers constrained to same AIE using our buffer address placement with *BufferOptLevel* 0.

* All KCC values are measured cycles averaged across 4 AIEs. (M, K, N) denotes the final size of the GEMM operation on the pack of 4 AIEs

TABLE V
PERFORMANCE SCALED TO THE WHOLE AIE ARRAY.

Precision (ip-op)	M	K	N	Throughput (TOPs/TBFLOPs)	TE
int8-int32	384	960	432	133	69%
int8-int16	512	736	576	159	82%
int8-int8	512	896	576	165	85%
bf16-bf16	512	384	576	83	86%

(M, K, N) denotes the final size of the GEMM.

AIEs = 288 (94.7%), Memory Banks = 2304 (94.7%), PLIOs In = 68 (60.7%), PLIOs out = 72 (85.7%), Y=8, G=4, X=9.

TABLE VI
RELATIVE PERFORMANCE COMPARISON WITH PRIOR WORK, COMPARING THE THROUGHPUT EFFICIENCY (TE)

Precision (ip-op)	GAMA TE*	Best prior work	Prior work TE+	% Improvement
int8-int32	69%	MAXEVA [14]	60%	9%
int8-int16	82%	AMA [8]	73.3%	8.7%
int8-int8	85%	AutoMM [24]	31.3%	53.6%
int8-int8	85%	ARIES [22]	45.9%	39%
bf16-bf16	86%	-	-	-

* AMD Versal VE2802 (VEK280 Board)

+ AMD Versal VC1902 (VCK190 Board)

efficient scaling across the whole AIE array. The hyperparameter values of 8, 4 and 9 for Y, G and X, respectively, show the best scaling on VE2802. Table V shows the final throughput achieved for the complete AIE array for different precisions. Our design achieves 94.7% (288/304) AIE array utilization, 94.7% of memory bank utilization, 60% PLIO in utilization and 85.7% PLIO out utilization. We calculate execution time from timestamps generated by `aiesimulator` in the output file. This throughput measurement technique was introduced by MaxEVA [14] and is also used by AMA [8], keeping the results consistent. The highest throughput of 165 TOPs is observed for int8-int8 due to low communication latency due to 8-bit precision. Similarly, int8-int32 has the lowest throughput of 133 TOPs, the larger output precision occupies more memory and increases communication latency.

E. Comparing with other state-of-the-art designs

Our work is the first work to map GEMM operations on full AIE-ML array devices. Prior work [8], [14], [15], [20]

use AIE. Hence, we only compare the relative achieved performance with prior work. Relative comparison is done using Throughput Efficiency (TE) i.e. ratio of achieved throughput to peak throughput. Moreover, different prior works use different precisions - some use int8-int8 [20]–[22], [24], int8-int16 [8], int8-int32 [14]. No other prior work use all the 4 precisions like GAMA. Hence, we can only compare a part of our results with each prior work. Table VI compares our approach with [8], [20], [22], [24] different works on Versal VC1902.

We compare int8 input and int32 output with MAXEVA [14], since that is the only work that writes 32-bit output back to the PL. GAMA achieves 133 TOPs (69%) TE. MaxEVA achieves 77 TOPs (60%) TE. GAMA extracts 9% more TE. MAXEVA’s efficiency is capped at 80% as only 80% of AIEs perform matrix multiplication, with 20% used for addition.

AMA [8] uses a similar kernel design compared to ours. Their matrix multiplication kernels also perform the reduction. They communicate partial sums to the next engine using buffer connections. AMA only writes 16 bit outputs compared to MAXEVA’s 32 bit output. Comparing the TE of the chip, we are 8.7% higher than AMA for the 16-bit output.

Lastly, we compare GAMA with CHARM [24] and ARIES [22] for int8-int8 precision. Our design outperforms CHARM and ARIES by 53.6% and 39%, respectively, in terms of TE. RSN-XNN [15] only demonstrates performance for fp32. Since AIE-ML lacks native support for fp32, a direct comparison between GAMA and RSN-XNN cannot be made.

VI. CONCLUSION

In this work, we introduce the GAMA framework that accelerates GEMM on AMD Versal AIE-ML architecture. Our framework maximizes AIE memory utilization, provides a custom buffer placement algorithm to minimize memory stalls, and a staggered kernel placement to facilitate scaling over the AIE array. GAMA framework provides the highest utilization of peak throughput for Versal compared to all other state-of-the-art frameworks. Though our improvements focus on AIE-ML, optimizing memory usage and prioritizing scalability can also benefit other accelerator architectures.

Future work will incorporate the memory tiles present in the AIE-ML architecture. Maximizing utilization of these memories is critical to improving the overall performance due to Versal’s limited peak off-chip DRAM bandwidth.

REFERENCES

- [1] “AI Engine API User Guide.” 2022.
- [2] AMD, “AMD Versal ACAP,” <https://www.amd.com/en/products/adaptive-socs-and-fpgas/versal.html>, 2024, [Online; accessed 02-may-2024].
- [3] AMD/Xilinx, “Versal ACAP AI Engine Architecture Manual (AM009).” 2021.
- [4] —, “AI Engine Kernel and Graph Programming Guide (UG1079).” 2022.
- [5] —, “Versal Adaptive SoC AIE-ML Architecture Manual (AM020).” 2023.
- [6] P. Chatarasi, S. Neuendorffer, S. Bayliss, K. Vissers, and V. Sarkar, “Vyasa: A High-Performance Vectorizing Compiler for Tensor Convolutions on the Xilinx AI Engine,” in *2020 IEEE High Performance Extreme Computing Conference (HPEC)*, 2020, pp. 1–10.
- [7] P. Chen, P. Manjunath, S. Wijeratne, B. Zhang, and V. Prasanna, “Exploiting On-Chip Heterogeneity of Versal Architecture for GNN Inference Acceleration,” in *2023 33rd International Conference on Field-Programmable Logic and Applications (FPL)*. Gothenburg, Sweden: IEEE, Sep. 2023, pp. 219–227. [Online]. Available: <https://ieeexplore.ieee.org/document/10296434/>
- [8] X. Deng, S. Wang, T. Gao, J. Liu, L. Liu, and N. Zheng, “AMA: An Analytical Approach to Maximizing the Efficiency of Deep Learning on Versal AI Engine,” in *2024 34th International Conference on Field-Programmable Logic and Applications (FPL)*, 2024, pp. 227–235.
- [9] P. Dong, J. Zhuang, Z. Yang, S. Ji, Y. Li, D. Xu, H. Huang, J. Hu, A. K. Jones, Y. Shi, Y. Wang, and P. Zhou, “EQ-ViT: Algorithm-Hardware Co-Design for End-to-End Acceleration of Real-Time Vision Transformer Inference on Versal ACAP Architecture,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 43, no. 11, pp. 3949–3960, 2024.
- [10] C. Lattner, M. Amini, U. Bondhugula, A. Cohen, A. Davis, J. Pienaar, R. Riddle, T. Shpeisman, N. Vasilache, and O. Zinenko, “MLIR: A Compiler Infrastructure for the End of Moore’s Law,” 2020. [Online]. Available: <https://arxiv.org/abs/2002.11054>
- [11] N. Perryman, C. Wilson, and A. George, “Evaluation of Xilinx Versal Architecture for Next-Gen Edge Computing in Space,” in *2023 IEEE Aerospace Conference*, Mar. 2023, pp. 1–11, iSSN: 1095-323X.
- [12] J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, and S. Amarasinghe, “Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines,” in *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI ’13. New York, NY, USA: Association for Computing Machinery, 2013, p. 519–530. [Online]. Available: <https://doi.org/10.1145/2491956.2462176>
- [13] G. Singh, A. Khodamoradi, K. Denolf, J. Lo, J. Gómez-Luna, J. Melber, A. Bisca, H. Corporaal, and O. Mutlu, “SPARTA: Spatial Acceleration for Efficient and Scalable Horizontal Diffusion Weather Stencil Computation,” in *ICS*, 2023.
- [14] E. Taka, A. Arora, K.-C. Wu, and D. Marculescu, “MaxEVA: Maximizing the Efficiency of Matrix Multiplication on Versal AI Engine.” Nov. 2023, arXiv:2311.04980 [cs]. [Online]. Available: <http://arxiv.org/abs/2311.04980>
- [15] C. Wang, X. Zhang, J. Cong, and J. C. Hoe, “Reconfigurable Stream Network Architecture,” 2025. [Online]. Available: <https://arxiv.org/abs/2411.17966>
- [16] M. Wierse, “Evaluation of Xilinx Versal Device,” Bachelor Thesis, ETH Zurich, Zurich, 2023-02.
- [17] Z. Yang, J. Zhuang, J. Yin, C. Yu, A. K. Jones, and P. Zhou, “AIM: Accelerating Arbitrary-Precision Integer Multiplication on Heterogeneous Reconfigurable Computing Platform Versal ACAP,” in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. San Francisco, CA, USA: IEEE, Oct. 2023, pp. 1–9. [Online]. Available: <https://ieeexplore.ieee.org/document/10323754/>
- [18] A. Yemme and S. S. Garani, “A Scalable GPT-2 Inference Hardware Architecture on FPGA,” in *2023 International Joint Conference on Neural Networks (IJCNN)*, Jun. 2023, pp. 1–8, iSSN: 2161-4407.
- [19] C. Zhang, T. Geng, A. Guo, J. Tian, M. Herbordt, A. Li, and D. Tao, “H-GCN: A Graph Convolutional Network Accelerator on Versal ACAP Architecture,” in *2022 32nd International Conference on Field-Programmable Logic and Applications (FPL)*, Aug. 2022, pp. 200–208, iSSN: 1946-1488. [Online]. Available: <https://ieeexplore.ieee.org/document/10035160>
- [20] J. Zhuang, J. Lau, H. Ye, Z. Yang, Y. Du, J. Lo, K. Denolf, S. Neuendorffer, A. Jones, J. Hu, D. Chen, J. Cong, and P. Zhou, “CHARM: Composing Heterogeneous Accelerators for Matrix Multiply on Versal ACAP Architecture,” in *Proceedings of the 2023 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*. Monterey CA USA: ACM, Feb. 2023, pp. 153–164. [Online]. Available: <https://dl.acm.org/doi/10.1145/3543622.3573210>
- [21] J. Zhuang, J. Lau, H. Ye, Z. Yang, S. Ji, J. Lo, K. Denolf, S. Neuendorffer, A. Jones, J. Hu, Y. Shi, D. Chen, J. Cong, and P. Zhou, “CHARM 2.0: Composing Heterogeneous Accelerators for Deep Learning on Versal ACAP Architecture,” *ACM Trans. Reconfigurable Technol. Syst.*, vol. 17, no. 3, Sep. 2024. [Online]. Available: <https://doi.org/10.1145/3686163>
- [22] J. Zhuang, S. Xiang, H. Chen, N. Zhang, Z. Yang, T. Mao, Z. Zhang, and P. Zhou, “ARIES: An Agile MLIR-Based Compilation Flow for Reconfigurable Devices with AI Engines,” in *Proceedings of the 2025 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, ser. FPGA ’25. New York, NY, USA: Association for Computing Machinery, 2025, p. 92–102. [Online]. Available: <https://doi.org/10.1145/3706628.3708870>
- [23] J. Zhuang, Z. Yang, S. Ji, H. Huang, A. K. Jones, J. Hu, Y. Shi, and P. Zhou, “SSR: Spatial Sequential Hybrid Architecture for Latency Throughput Tradeoff in Transformer Acceleration,” Feb. 2024, arXiv:2401.10417 [cs]. [Online]. Available: <http://arxiv.org/abs/2401.10417>
- [24] J. Zhuang, Z. Yang, and P. Zhou, “AutoMM: Energy-Efficient Multi-Data-Type Matrix Multiply Design on Heterogeneous Programmable System-on-Chip,” May 2023, arXiv:2305.18698 [cs]. [Online]. Available: <http://arxiv.org/abs/2305.18698>