

ME-ViT: A Single-Load Memory-Efficient FPGA Accelerator for Vision Transformers

Kyle Marino
University of Southern California
Los Angeles, USA
kmarino@usc.edu

Pengmiao Zhang
University of Southern California
Los Angeles, USA
pengmiao@usc.edu

Viktor K. Prasanna
University of Southern California
Los Angeles, USA
prasanna@usc.edu

Abstract—Vision Transformers (ViTs) have emerged as a state-of-the-art solution for object classification tasks. However, their computational demands and high parameter count make them unsuitable for real-time inference, prompting the need for efficient hardware implementations. Existing hardware accelerators for ViTs suffer from frequent off-chip memory access, restricting the achievable throughput by memory bandwidth. In devices with a high compute-to-communication ratio (e.g., edge FPGAs with limited bandwidth), off-chip memory access imposes a severe bottleneck on overall throughput. This work proposes ME-ViT, a novel Memory Efficient FPGA accelerator for ViT inference that minimizes memory traffic. We propose a *single-load policy* in designing ME-ViT: model parameters are only loaded once, intermediate results are stored on-chip, and all operations are implemented in a single processing element. To achieve this goal, we design a memory-efficient processing element (ME-PE), which processes multiple key operations of ViT inference on the same architecture through the reuse of *multi-purpose buffers*. We also integrate the Softmax and LayerNorm functions into the ME-PE, minimizing stalls between matrix multiplications. We evaluate ME-ViT on systolic array sizes of 32 and 16, achieving up to a $9.22\times$ and $17.89\times$ overall improvement in memory bandwidth, and a $2.16\times$ improvement in throughput per DSP for both designs over state-of-the-art ViT accelerators on FPGA. ME-ViT achieves a power efficiency improvement of up to $4.00\times$ ($1.03\times$) over a GPU (FPGA) baseline. ME-ViT enables up to 5 ME-PE instantiations on a Xilinx Alveo U200, achieving a $5.10\times$ improvement in throughput over the state-of-the-art FPGA baseline, and a $5.85\times$ ($1.51\times$) improvement in power efficiency over the GPU (FPGA) baseline.

Index Terms—Vision Transformer, FPGA Accelerator, Memory Bandwidth

I. INTRODUCTION

The self-attention based model of the Transformer [1] has led to significant advancements in machine learning, impacting a diverse range of applications [2]–[6]. Originally gaining prominence due to its remarkable success in natural language processing [7], the Transformer has been adapted to the domain of computer vision via Vision Transformers (ViTs) [8], achieving superior performance over convolutional networks [9], [10]. Despite the substantial achievements of ViTs, their implementation on real-time image data poses considerable computational and memory challenges due to the immense parameter counts associated with these models [11].

Significant effort has been put into accelerating the inference of ViTs, including model size reduction [12], [13], weight quantization [14], and algorithm optimization [15]. However,




Fig. 1: Roofline model of state-of-the-art (SOTA) architectures and ME-ViT for various models. Vertical axis is in log scale. ME-ViT optimizes memory bandwidth, enabling nearly peak performance in GOPS (Giga Operations per Second). SOTA implementations are bottlenecked by memory bandwidth. Four ViT variants are shown: ViT-Base model from [8] (ViT-B), and three models from [10] (DeiT-B, DeiT-S, and DeiT-T).

these methods do not directly address the main performance bottleneck of ViT inference on modern hardware: **memory bandwidth**. Computing capabilities of hardware such as GPUs and TPUs have outpaced memory bandwidth improvements, resulting in a poor Compute-to-Communication (C2C) ratio and limited model performance [16], [17]. Various works [18]–[20] have focused on algorithmic approaches to reducing memory bandwidth, but are still constrained from the overall architectural limitations imposed by the GPU.

Field Programmable Gate Arrays (FPGAs) provide a good platform for ViT acceleration due to the high computational parallelism and custom architectures that can be designed [21]–[23]. However, like GPUs, the high memory bandwidth needs of Transformers significantly limits their implementation on FPGAs. As shown in Figure 1, ViT and DeiT (a commonly-used variant of ViT) [10] models are severely bottlenecked by memory bandwidth without specialized optimization. The shortcomings of modern computing devices for memory-bound computing tasks such as ViT inference prompt the need for efficient and model-specific architectures suited for these tasks. Custom FPGA architectures can meet the computational and memory demands of ViT inference more

effectively than the general-purpose architecture of a GPU, while also consuming less power.

In this work, we aim to minimize the memory bandwidth for ViT inference on an FPGA. The development of such an optimization presents multiple unique challenges. **1) Avoiding write-backs and reloads in block matrix multiplications.** The large matrix multiplications present in ViTs require a block matrix multiplication (BMM) approach, which divides the large matrices into smaller blocks to meet the limited resources of an FPGA. This results in constant block write-backs and reads to off-chip memory, leading to high usage of memory bandwidth. To minimize the memory traffic, buffering all data on the FPGA is ideal. However, excessively large buffers can hinder effective utilization of available DSPs [24]. This is because too much buffering per systolic array will exhaust BRAM before all DSPs can be utilized. Thus, it is crucial to strategically reuse on-chip buffers, achieving a balance between minimal memory traffic and optimal FPGA DSP utilization. **2) Buffering intermediate results for the residual connections in ViT.** The residual connections add the values from the previous layer to the computed result of the current layer, which necessitates either the buffering or loading of a previous layer [25]. Buffering a layer uses more FPGA resources but reduces the memory traffic that comes with layer loading. Therefore, to minimize memory traffic and to efficiently handle these residual connections, the design of reusable buffers becomes a crucial aspect of the overall system architecture. **3) Reducing communications between FPGA accelerator and the host CPU.** Often, Softmax and LayerNorm operations in all layers of a ViT are performed on the host CPU, with only the matrix multiplications being offloaded to the accelerator. This is due to the computational insignificance of Softmax and LayerNorm compared to matrix multiplication [23] [22]. However, this causes frequent write-backs multiple times per layer, as well as a significant delay for the round trip to the CPU. Constant data transfer becomes a significant performance bottleneck, especially for larger FPGAs with higher processing capabilities.

We propose ME-ViT, a novel Memory-Efficient ViT accelerator on FPGA that addresses the above challenges. As shown in Figure 1, ME-ViT optimizes memory bandwidth and enables nearly peak performance for ViTs. ME-ViT is developed through two key optimizations. 1) A **Single-load policy** for model parameters loaded from off-chip memory. 2) **Multi-purpose buffers** for three key operations in ViT.

We propose a *single-load policy* as the key approach for minimizing memory traffic for ViT accelerators. It consists of three objectives. First, parameters loaded to the FPGA are only loaded once from off-chip memory (e.g., DRAM). If the value needs to be reused it is strategically buffered in the FPGA’s on-chip memory (e.g., BRAM). Second, intermediate data is not written back to off-chip memory between layers. This ensures that the absolute minimum number of memory transfers are used per inference. Third, all operations—including BMM, LayerNorm, Softmax, and activations—are performed within a single processing element to eliminate external data traffic.

The single-load policy is different than a block matrix multiplication approach which still requires reloading of the same weights between different blocks.

Multi-purpose buffers are designed to achieve the goals of the single-load policy while addressing its requirement of large BRAM allocations. To minimize resource usage, we design ME-PE, a single Memory Efficient Processing Element that executes the three key operations in ViT inference: Linear Projection (LP), Multi-headed Self-Attention (MSA), and Multi-Layer Perceptron (MLP). By designing a custom PE that conforms specifically to the MSA operations, we can strategically order the computations to obtain a minimal resource utilization. The minimal architecture needed for MSA can be repurposed for MLP calculation without requiring additional BRAM resources. BRAM is efficiently packed and repurposed for different stages of calculation. As a result, we are able to design a flexible PE that performs all operations for ViT inference while using minimal BRAM to ensure model parameters are only loaded once and all intermediate data write-back is avoided. Meanwhile, the reuse of resources within our design enables the implementation of multiple ME-PEs in an FPGA, further improving the throughput of ME-ViT.

Our main contributions are summarized as follows:

- We propose ME-ViT, a novel memory-efficient Vision Transformer accelerator on FPGA, which optimizes memory bandwidth and achieves nearly peak performance in operations per second.
- We propose a single-load policy as the core principle for minimizing memory access by only loading data once from DRAM, buffering intermediate results, and implementing all operations in a single PE.
- We design an ME-PE, a memory-efficient processing element with reusable multi-purpose buffers. This novel PE enables three key operations of ViT inference to be processed on the same architecture, minimizing resource usage and retaining intermediate data between operations.
- We integrate LayerNorm and Pseudo-Softmax (a hardware-optimized Softmax) in the ME-PE to avoid off-chip computation and reduce data traffic. These functions are pipelined with matrix multiplication to reduce computational stalls.
- We evaluate ME-ViT on a Xilinx Alveo U200. Using systolic array sizes of 32 and 16, ME-ViT achieves up to a $9.22\times$ and $17.89\times$ overall improvement in memory bandwidth, and up to a $2.16\times$ improvement in throughput per DSP over state-of-the-art ViT accelerators on FPGA. ME-ViT achieves up to $4.00\times$ ($1.03\times$) power efficiency over the GPU (FPGA) baseline. The ME-ViT with 5 ME-PEs implemented on board achieves a $5.1\times$ improvement in throughput over the state-of-the-art FPGA baseline, and a $5.85\times$ ($1.51\times$) improvement in power efficiency over the GPU (FPGA) baseline.

To the best of our knowledge, ME-ViT is the first Vision Transformer architecture that optimizes memory traffic and strictly enforces a minimal memory access policy.




Fig. 2: Vision Transformer architecture.

II. BACKGROUND

A. Vision Transformer

Vision Transformer (ViT) is a state-of-the-art deep learning architecture that utilizes the Transformer model for computer vision tasks. By leveraging self-attention mechanisms, ViTs have achieved significant advancements in image classification and object detection, revolutionizing the field of visual recognition. The ViT architecture is shown in Figure 2. The input image is broken up into patches and fed into the Transformer Encoder as a sequence. The Transformer Encoder is mainly composed of a multi-headed self-attention block (MSA), a multi-layer perceptron block (MLP), and layer normalization blocks (LN).

Multi-Headed Self-Attention. Self-attention takes the embedding of items as input, converts them to three matrices through linear projection, then feeds them into a scaled dot-product attention. The self-attention function is defined as:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{D_k}} \right) V \quad (1)$$

where Q is queries, K is keys, V is values, D is the model dimension, and D_k is the dimension of K . Considering one self-attention operation as one "head," Multi-headed Self-Attention (MSA) operation is shown as:

$$\begin{aligned} \text{MSA}(Q, K, V) &= \text{Concat}(\text{head}_1, \dots, \text{head}_h) W^O \\ \text{head}_i &= \text{Attention} \left(QW_i^Q, KW_i^K, VW_i^V \right) \end{aligned} \quad (2)$$

where the projection matrices $W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{D \times D_h}$, h is the total number of heads, i is the index of heads, and $D_h = D/h$ is the dimension of each head.

Multi-Layer Perceptron. The MLP block consists of two linear layers with an activation function:

$$\text{MLP}(x) = \text{GeLU}(xW^H + B^H)W^O + B^O \quad (3)$$

where W^H is the hidden layer weights, B^H is the hidden layer bias, W^O is the output layer weights, and B^O is the output layer bias, GeLU is an activation function [26].




Fig. 3: Systolic array architecture of size 4×4 . Dark gray boxes indicate DSPs, each split by a dashed line to illustrate DSP packing.

Transformer Encoder. For the input image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$, the formal expression of transformer encoder layers are shown in Equation 4.

$$\begin{aligned} \mathbf{z}_0 &= [\mathbf{x}_{\text{class}}; \mathbf{x}_p^1 \mathbf{E}; \mathbf{x}_p^2 \mathbf{E}; \dots; \mathbf{x}_p^N \mathbf{E}] + \mathbf{E}_{\text{pos}} \\ \mathbf{z}'_\ell &= \text{MSA}(\text{LN}(\mathbf{z}_{\ell-1})) + \mathbf{z}_{\ell-1} \\ \mathbf{z}_\ell &= \text{MLP}(\text{LN}(\mathbf{z}'_\ell)) + \mathbf{z}'_\ell \\ \mathbf{y} &= \text{LN}(\mathbf{z}_L) \end{aligned} \quad (4)$$


where $\mathbf{x}_{\text{class}}$ is a class token, $\mathbf{x}_p \in \mathbb{R}^{N \times (P^2 \times C)}$ is the image segmented to N patches, $\mathbf{E} \in \mathbb{R}^{(P^2 \cdot C) \times D}$ is the input embedding, $\mathbf{E}_{\text{pos}} \in \mathbb{R}^{(N+1) \times D}$ is the position embedding, $\ell = 1 \dots L$ is the index of layers.

In this work, we design a ViT accelerator by implementing the above functions on FPGA and optimizing the memory accesses in model inference.

B. ViT Accelerators on FPGA

Field Programmable Gate Arrays (FPGAs) have been extensively used for accelerating machine learning tasks [27] [28]. FPGAs consist of a programmable interconnect of logic gates and on-chip memories (BRAMs, URAMs, LUTRAMs), allowing custom hardware architectures to be designed. This flexibility enables hardware optimizations to meet the specific requirements of various machine learning tasks.

There have been several proposed architectures for accelerating ViT and, more generally, Transformer inference on FPGAs [23], [25], [29], [30]. These architectures typically quantize weights and activations to 8 bits to reduce model size and computation requirements [22], [23]. Some architectures compute Softmax and LayerNorm on the FPGA [25], [30], while others perform this on the host CPU [23]. Computing these functions on the CPU reduces the complexity of the FPGA design, but increases the memory overhead. Wang et al. [25] propose ViA, a ViT accelerator which performs the full calculation per layer on an FPGA but requires write-backs between layers and does not implement the full-size ViT models. Nag et al. [30] target an edge FPGA device and also incur frequent memory reads and write-backs between layers. Sun et al. [23] propose an accelerator design that requires the CPU to perform the Softmax and LayerNorm functions and also necessitates loading and unloading for each matrix multiplication. Lu et al. [29] propose an accelerator



(a) MLP Matrix Multiplication

(b) Partial Sum Method

Fig. 4: MLP matrix multiplication and partial sum method.

for the MSA and MLP layers of a Transformer with on-chip computation of Softmax and LayerNorm, but require memory access between matrix multiplications.

All of the above model architectures write back intermediate layer calculations and frequently load weights for matrix multiplication, which incur a large memory overhead [23], [25], [29], [30]. As these architectures scale to larger sizes, they become limited by memory bandwidth constraints. ME-ViT aims to address these memory constraints by eliminating intermediate write backs between layers and only loading weights once into the design, ensuring the minimal amount of data is transferred. This policy enables our design to easily scale without being constrained by memory limitations.

C. Matrix Multiplication on FPGA

1) *DSP Packing*: DSP packing is utilized to perform two simultaneous multiplies per DSP as described in [31]. Each DSP contains an 18-bit \times 27-bit multiplier, which can be used to perform $A \times B$ and $A \times C$ simultaneously by assigning A to the 18-bit operand, and $(B \ll 18 + C)$ to the 27-bit operand. The resulting multiply leaves the two distinct 16-bit products separated in the 45-bit DSP output. A systolic array of size $P_{SYS} \times P_{SYS}$ shown in Figure 3 is used to efficiently multiply matrices. With DSP packing, the systolic array can perform $P_{SYS} \times (2 \cdot P_{SYS})$ multiplies. Datapaths in Figures 8, 10, and 11 with double-arrows indicate the double-wide data paths, which are necessary for accommodating the combined bit width of the two separate products.

2) *Matrix Multiplication for MLP*: The MLP calculation involves multiplication with weight matrices W^H and W^O , which are too large to be buffered or computed fully within the PE. Matrix multiplication is split into row blocks and column blocks, where A_i refers to the i th row or column block of




Fig. 5: Overview of the ME-PE architecture.




Fig. 6: Scheduling of ME-PE modes across Transformer encoder for ME-ViT inference.

matrix A . A sub-block $B_{i,j}$ refers to the matrix block at the i th row block and j th column block of matrix B .

To perform the calculation without reloading weights, we implement the MLP method from [30] shown in Figure 4. A single sub-block of the intermediate result is calculated, which then is passed point-wise through an activation function. The result $M_{i,j}$ is multiplied by each sub-block in W^O . The result of each block multiplication is added to the previous value in the staged result. At the start of the cycle, the staged result contains the output layer bias, which handles bias addition implicitly. This performs the partial sum method shown in Figure 4b, which allows the M row block and W^O column block product to be broken up into separate iterations. For each iteration in Figure 4a, a column block from W^H and a row block from W^O is loaded. This approach allows the weights in the MLP layer to be fully used from a single load without requiring larger buffer sizes.

III. APPROACH

We introduce ME-ViT, a novel memory-efficient FPGA accelerator for Vision Transformers. ME-ViT minimizes memory traffic, which is accomplished through a versatile architecture of a memory-efficient processing element (ME-PE, Section III-A), the scheduling of ME-PE modes for the key operations in a Transformer Encoder (LP mode in Section III-B, MSA mode in Section III-C, and MLP mode in Section III-D), and the integration of LayerNorm and Softmax operations (Section III-E). Due to the high reuse of resources, we design an ME-ViT architecture with multiple ME-PEs working in parallel (Section III-F) that highly improves the throughput of ViT inference.

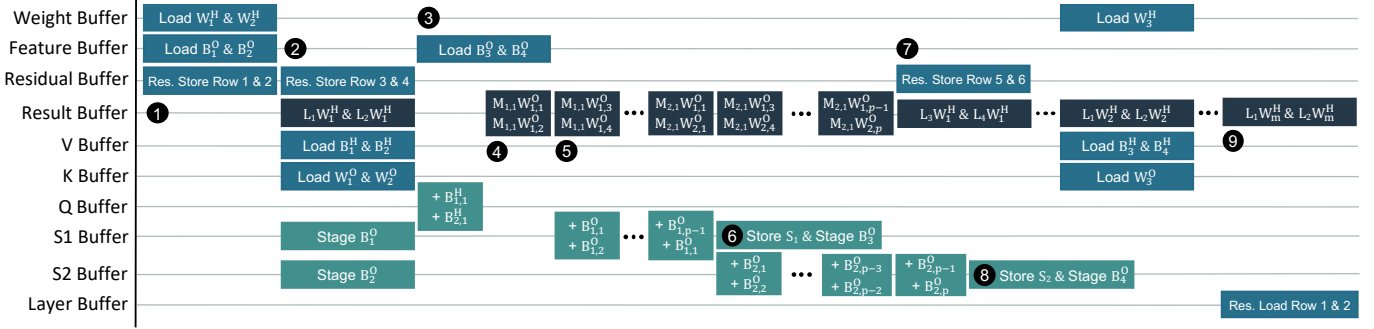


Fig. 12: Scheduling for MLP Mode.




Fig. 13: Design of LayerNorm and Pseudo-Sigmoid Modules.

are concurrently performed and stored in the Result Buffer. The column blocks B_1^H and B_2^H are loaded to the V Buffer and row blocks B_1^H and B_2^H are loaded into the K buffer. The row blocks B_1^H and B_2^H are staged by being moved to the S Buffer for future addition. At timestamp 3, the row blocks B_1^H and B_2^H are added to the two blocks stored in the Result Buffer. Instead of GeLU, the ReLU activation function is used for hardware simplicity. Various hardware approximations of GeLU exist and would not affect the scheduling or timing of this design. The activation function results are stored in the Q buffer. This timestamp results in a stall of the systolic array, however it is unavoidable since the result is needed immediately for the next multiply. This only incurs a delay of P_{SYS} clocks which is insignificant. The next pair of output biases are also loaded. At timestamp 4, the first two output layer block multiplications are computed. At timestamp 5, the results computed in timestamp 4 are added to the staged values in the S_1 buffer, and the next two output layer blocks are calculated. This process repeats until all sub-blocks involving $M_{1,1}$ are calculated. At timestamp 6, the S_1 buffer containing the staged values is stored back into the Feature Buffer, and multiplication repeats with $M_{2,1}$. At timestamp 7, the next pair of Layer Buffer row blocks is processed. This continues until

all row blocks in the Layer Buffer have been multiplied by W_1^H . After all layer row blocks have been processed, the cycle repeats with the next W^H column block. At timestamp 8, S_2 is stored and the next row block from the Feature Buffer is loaded to S_2 . On the last cycle in timestamp 9, the residual data stored in the Weight Buffer is transferred back to the Layer Buffer, replacing the existing values as they are no longer needed.

Two separate S Buffers are needed to reduce stalls for matrix multiplication. After one S buffer is filled it gets written back to Feature Buffer BRAM. While this takes place, the second S buffer gets written to. Otherwise, matrix multiplication would need to be stalled while the S buffer gets unloaded and reloaded with new values. This adds an additional resource overhead to MLP Mode that is not present for other modes. However, since the S buffer is relatively small and implemented in LUTRAM, this incurs minimal penalty.

E. Integrating LayerNorm and Pseudo-Sigmoid

1) *LayerNorm Module*: We implement LayerNorm in the ME-PE using a two-pass approach that is performed in parallel with matrix multiplication. LayerNorm is calculated in both the MSA and MLP blocks with:

$$LayerNorm_{i,j} = \frac{X_{i,j} - \mu_i}{\sqrt{\sigma_i^2 + \epsilon}} \cdot \gamma_j + \beta_j \quad (5)$$

With the mean and variance per i th row of matrix X as:

$$\mu_i = \frac{1}{n} \sum_{j=1}^n X_{i,j} \quad (6)$$

$$\sigma_i^2 = \frac{1}{n} \sum_{j=1}^n (X_{i,j} - \mu_i)^2 \quad (7)$$

The mean and variance can be calculated in parallel using the variance form:

$$\sigma_i^2 = \frac{1}{n} \sum_{j=1}^n X_{i,j}^2 - \left(\frac{1}{n} \sum_{j=1}^n X_{i,j} \right)^2 \quad (8)$$

This reduces the full LayerNorm calculation to two passes; The first pass accumulates a sum and squared sum of each row i . After the row sums are calculated, the $RowMean_i$ and $1/\sqrt{RowVar_i}$ are calculated using fixed-point arithmetic




Fig. 14: Multi-PE ME-ViT architecture on the Alveo U200.

functions. These constants are fed into the LayerNorm Module shown in Figure 13a. By passing each row element j through the LayerNorm Module in a pipelined manner, the final layer-normalized value is efficiently computed.

2) *Pseudo-Softmax Module*: We implement the Pseudo-Softmax function, a hardware-friendly alternative to the Softmax function proposed in [32]. We utilize a two-pass approach for computation which is parallelized with matrix multiplication. The Pseudo-Softmax uses base 2 instead of e to leverage floating point number properties to evaluate exponentiation.

$$\tilde{p}_i = \frac{2^{x_i}}{\sum_{k=1}^N 2^{x_k}} \quad (9)$$

Let a_i be a floating-point number with exponent x_i . This removes the need to calculate exponentiation in hardware since it is implicitly handled by the float representation.

$$\tilde{p}_i = \frac{a_i}{\sum_{k=1}^N a_k} \quad (10)$$

The summation term can be expressed as a single floating point number. In this representation, exp_{sum} and mant_{sum} denote the exponent and mantissa of the resulting number.

$$\sum_{k=1}^N a_k = 2^{\text{exp}_{\text{sum}}} \cdot \text{mant}_{\text{sum}} \quad (11)$$

The Pseudo-Softmax \tilde{p}_i for element x_i , is calculated as:

$$\tilde{p}_i = 2^{x_i - \text{exp}_{\text{sum}}} \cdot \frac{1}{1 \cdot \text{mant}_{\text{sum}}}. \quad (12)$$

This requires two passes over an input vector x . The first pass is to calculate the values exp_{sum} and mant_{sum} and the second pass is to calculate \tilde{p}_i . The reciprocal of the sum mantissa is determined after the floating point sum is computed. Since the Softmax value ranges from 0 to 1, the final result will be the reciprocal bit-shifted by $\text{exp}_{\text{sum}} - x_i + 1$ as per Equation 12. 127 is added to the score row to convert to a floating point exponent which is stored unsigned. A 1 is prepended to the mantissa since it is implicitly present in the floating point format. The result is stored in fixed-point format with only fractional bits to maximize accuracy. The upper 8 bits after the bit-shift contain the fixed-point representation of the Pseudo-Softmax function.

F. Multiple ME-PE Architecture for ME-ViT

A Multiple ME-PE architecture (Multi-PE) is proposed that contains parallel instantiations of the ME-PE along with a scheduler to coordinate data traffic between them. The Alveo

U200 contains 3 Super Logic Regions (SLRs), with SLRs 0 and 2 having 2275 DSPs each and SLR 1 only having 1317. With a $P_{\text{SYS}} = 32$, up to 5 PEs can fit as shown in Figure 14. A smaller P_{SYS} can fit more SAs in the FPGA, however the BRAM needed for such a design is the same as $P_{\text{SYS}} = 32$ since buffering requirements are unchanged. Therefore, only the $P_{\text{SYS}} = 32$ design is implemented to maximally utilize the available DSPs. Resource utilization is discussed in more detail in Section IV-B. The Multi-PE architecture can achieve remarkably higher throughput, but an increase in data traffic causes a bottleneck due to the limited 77 GB/s bandwidth to the FPGA DRAM. These results are discussed in Section IV-G.

IV. EVALUATION

A. Experimental Setup

We implement ME-ViT on the Xilinx Alveo U200 platform, consisting of 5867 DSPs, 1766 36k BRAMs, 892K LUTs, and 1831K FFs. It has 4 channels of DDR memory, and a total bandwidth of 77 GB/s. Experimental results are evaluated independently for each ME-ViT mode, and theoretical values are presented which remove extra latencies added from Vitis synthesis and Place and Route (P&R). All implementations are designed for 300 MHz, and 150 MHz figures are provided to compare with other designs. The ME-PE is designed and evaluated using Vitis HLS 2023.1. Power estimates are calculated using AMD Power Design Manager 2023.1.1.

A single ME-PE is analyzed on four common ViT model sizes shown in Table I. ME-PEs with systolic array size $P_{\text{SYS}} = 32$ and $P_{\text{SYS}} = 16$ are analyzed to provide insight into the performance scalability across FPGAs of varying DSP resources. As the size of the systolic array decreases, there is a corresponding reduction in total FPS (frames per second). Memory bandwidth also reduces despite smaller systolic arrays requiring more frequent data transfers. This relationship between scale and memory bandwidth is explored in Section IV-E. Finally, Multi-PE results are calculated based on single $P_{\text{SYS}} = 32$ ME-PE performance to maximally utilize the available DSPs.

TABLE I: Model Variants

Model	Image Size	Model Dimension	Num Heads	Layers	Parameter Count
ViT-B	256 ²	768	12	12	86M
DeiT-B	224 ²	768	12	12	86M
DeiT-S	224 ²	384	6	12	22M
DeiT-T	224 ²	192	3	12	6M

ViT variants shown in Table I are evaluated on the ME-ViT architecture. ViT-B refers to the base ViT model presented in [8] but with 256 input image resolution. DeiT-B, DeiT-S, and DeiT-T refer to model sizes presented in [10], all on 224 input image resolution. The key difference between DeiT models lies in their respective model dimensions: 768, 384, and 192.

B. Results on Hardware

Results for hardware utilization are shown in Table II. The three ME-ViT modes are implemented separately, and through-

put values in Table III are derived from the latency per mode. Since each mode largely utilizes the same resources but with different control logic, the unified design’s resources would marginally exceed that of the largest mode (MLP Mode). Resources are shown for the $P_{SYS} = 32$ ME-PE for the ViT-B model. Resource consumption is unchanged for DeiT-B, and BRAM usage drops to 176 and 144 for DeiT-S and DeiT-T respectively. For $P_{SYS} = 16$, 256 DSPs are used, with other values remaining the same as buffering requirements do not change.

TABLE II: Hardware Resource Utilization

Hardware Configuration	DSP	BRAM36	LUT (K)	FF (K)
LP Mode	1024	288	159	93
MSA Mode	1024	288	166	107
MLP Mode	1024	288	192	132
Auto Vit Acc [22]	2066	–	128	–

C. Performance Comparison

Performance comparison values for ME-ViT are shown in Table III. Theoretical values are presented which are calculated by leveraging extra parallelism which cannot be achieved through Vitis HLS. These figures are achievable with Verilog synthesis. All synthesized designs achieve an operating frequency of 300 MHz, however, competing solutions (Auto Vit Acc [22] and ViTA [30]) are implemented at lower frequencies. To accurately compare design metrics, 150 MHz figures are provided.

We compare ME-ViT performance against four baseline platforms: (1) CPU only platform (Intel i7-9800X), (2) high-power GPU accelerated platform (Nvidia Titan RTX), (3) low-power GPU accelerated platform (Nvidia Jetson TX2), (4) FPGA ViT accelerator presented in Auto Vit Acc [22], and (5) Edge FPGA ViT accelerator presented in ViTA [30].

A single ME-PE achieves a theoretical throughput of 26.4 FPS, a $1.04\times$ improvement over the Auto Vit Acc implementation. Notably, Auto Vit Acc is implemented at 150 MHz which outperforms a single ME-PE at the same clock frequency due to higher DSP usage. The ME-PE has a similar FPS/DSP efficiency as Auto Vit Acc at 150 MHz, and sees a $2.16\times$ improvement at 300 MHz. The ME-PE has a high power efficiency of 2.83 FPS/Watt, outperforming all other platforms except for ViTA. In comparison to a GPU with a similar power usage (TX2), ME-ViT has a $4.42\times$ improvement in FPS/Watt, demonstrating the high efficiency of the custom architecture approach.

D. Overall Throughput and Latency

Overall throughput measured in frames per second (FPS) for the 4 models is shown in Table IV. FPS improves as model sizes get smaller, and $P_{SYS} = 16$ experiences roughly $0.25\times$ the throughput of corresponding $P_{SYS} = 32$ designs. Latencies per ME-ViT mode are presented in Figure 15. MLP Mode exhibits the longest duration out of the three modes, taking approximately 60 percent of execution time across all models.




Fig. 15: Breakdown of latency per mode for ME-ViT on various models. Vitis implementation (HLS) introduces small extra latency overhead to the theoretical performance.




Fig. 16: Breakdown of bandwidth per mode for ME-ViT on various models, lower is better.

A reduction of input image size from 256 to 224 results in an average $1.17\times$ improvement for $P_{SYS} = 32$ and an average $1.08\times$ improvement for $P_{SYS} = 16$. For both systolic array sizes, a reduction of model dimension in half results in an average improvement of $3.7\times$.

E. Memory Bandwidth Comparison

As there exist no published results on memory bandwidth for similar FPGA architectures, a non optimized approach is calculated with the following characteristics that are common in various designs [21], [23], [30]: 1) Each BMM loads two input matrices. If an input block matrix was used for the previous multiply, it remains loaded. 2) All calculated matrix blocks are written back to DRAM. 3) Softmax and LayerNorm [33] are calculated on the CPU and are implicitly included in intermediate write-backs.

Figure 16 shows memory bandwidth figures for ME-ViT on all four models for both $P_{SYS} = 32$ and $P_{SYS} = 16$. Total

TABLE III: Platform Performance Comparison

Platform	CPU i7-9800X [23]	GPU Titan RTX [23]	GPU Jetson TX2 [22]	Auto Vit Acc [22]	ViTA [30]	ME-ViT		ME-ViT Theoretical		Multi-PE	
						150 MHz	300 MHz	150 MHz	300 MHz	150 MHz	300 MHz
						Latency (ms)	65.35	5.45	127	38.61	363.64
FPS	15.3	183.4	7.87	25.9	2.75	11.99	23.98	13.20	26.40	66.02	132.04
Power (W)	100	260	12.28	9.4	0.88	6.5	9.3	6.5	9.3	17.8	31.8
FPS/Watt	0.15	0.71	0.64	2.76	3.13	1.84	2.57	2.03	2.83	3.71	4.15
FPS/DSP	-	-	-	0.012	-	0.012	0.023	0.013	0.026	0.013	0.026

TABLE IV: Model Performance (FPS)

Model	HLS	Theoretical	HLS	Theoretical
	($P_{SYS}=32$)	($P_{SYS}=32$)	($P_{SYS}=16$)	($P_{SYS}=16$)
ViT-B	20.64	22.38	5.40	6.08
DeiT-B	23.98	26.40	5.81	6.64
DeiT-S	87.64	98.25	22.13	25.53
DeiT-T	298.52	352.27	78.55	94.13

TABLE V: Memory Bandwidth Improvement

Model	Total	Peak	Total	Peak
	($P_{SYS}=32$)	($P_{SYS}=32$)	($P_{SYS}=16$)	($P_{SYS}=16$)
ViT-B	9.22	13.07	17.14	25.58
DeiT-B	8.25	11.29	16.62	23.79
DeiT-S	7.06	14.60	17.53	27.60
DeiT-T	8.77	21.28	17.89	35.29

and peak bandwidth improvements are shown in Table V. Peak improvement occurs in the MSA Mode as this has the most back-and-forth traffic in the unoptimized case. Despite fewer model parameters to transfer, a higher improvement is seen on smaller models since less time is spent on computation and therefore a larger proportion of data movement needs to occur in the same time. There is an approximate $3.8\times$ reduction in latency between $P_{SYS} = 32$ and $P_{SYS} = 16$, but an approximate $2\times$ reduction in data transferred. This results in a larger reduction in memory bandwidth for $P_{SYS} = 16$.

F. Systolic Array Size Comparison

The systolic array has a large impact on the computational efficiency, defined as the total computation performed divided by the minimum computation required. Depending on how P_{SYS} divides both the model dimension and the layer height, computational efficiency can greatly vary. When P_{SYS} poorly matches these dimensions, the BMMs at the right and bottom boundaries fill only a small portion of the systolic array, leading to wasted computation.

As shown in Figure 17, there are periodic cycles in efficiency ranging from 0.95 to 0.5, with notable peaks occurring at 11, 17, 33, 50, and 66. In addition, simply increasing P_{SYS} leads to an overall downward trend in efficiency. These results underscore the importance of tuning hardware to fit the model dimensions.

G. Multi-PE ME-ViT Performance

We analyze the performance of the Multi-PE to determine the effectiveness of ME-ViT when scaled to larger FPGAs.



Fig. 17: Computational Efficiency vs. P_{SYS} for all models.

Fig. 18: Throughput (left) and performance (right) comparison between Multi-PE and the unoptimized design, higher is better. Throughput vertical axis is in log scale.

Multi-PE performance results are shown in Figure 18. The unoptimized design for ViT-B and DeiT-B can only support 3 ME-PEs before performance is limited by memory bandwidth. For DeiT-S and DeiT-T, the unoptimized design can only support 2 ME-PEs. ME-ViT allows 5 ME-PEs to be supported for all models, resulting in a $1.66\times$ improvement in both FPS and GOPS (Giga Operations per Second) for ViT-B and DeiT-B, and a $2.5\times$ improvement for DeiT-S and DeiT-T. The theoretical maximum GOPS for 5 ME-PEs is 3072, yet a maximum of 2682 is achieved due to inefficiencies with irregularly-sized matrix multiplication. Matrix blocks that do not completely fill the systolic array result in unused DSPs, leading to an overall reduction in GOPS.

For larger ViT models such as ViT-Large and ViT-Huge ($D = 1024$ and 1280), total BRAM usage will increase to 384 and 608 respectively with $P_{sys} = 32$. ME-PEs of this size can still fit inside a large FPGA like the Alveo U200, but fewer total will fit in the Multi-PE design. This results in a large number of unused DSPs, reducing the total GOPs from the theoretical maximum.

V. CONCLUSION

In this paper, we proposed ME-ViT, a novel ViT hardware accelerator that mitigates the high-bandwidth needs of ViT inference. ME-ViT minimizes the memory traffic for a ViT accelerator on an FPGA through a single-load policy and multi-purpose buffers within a memory-efficient processing element (ME-PE). ME-ViT achieves up to a $17.89\times$ overall improvement in memory bandwidth, and up to a $2.16\times$ improvement in throughput per DSP over state-of-the-art ViT accelerators on FPGA. ME-ViT enables implementation of up to 5 ME-PEs on a Xilinx Alveo U200, achieving a $5.10\times$ improvement in throughput over the FPGA baseline. Future research will focus on extending the ideas in this paper beyond ViTs to Large Language Models which inhibit a single-load policy due to limited on-chip memory.

ACKNOWLEDGMENT

This work has been supported by the U.S. National Science Foundation under grant numbers CNS-2009057 and SaTC-2104264. We are grateful to Bingyi Zhang for his insightful perspectives in the development of this paper.

REFERENCES

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [2] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, "Vivit: A video vision transformer," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 6836–6846.
- [3] H. Chang, H. Zhang, L. Jiang, C. Liu, and W. T. Freeman, "Maskgit: Masked generative image transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 11 315–11 325.
- [4] D. A. Hudson and L. Zitnick, "Generative adversarial transformers," in *International conference on machine learning*. PMLR, 2021, pp. 4487–4499.
- [5] P. Zhang, A. Srivastava, A. V. Nori, R. Kannan, and V. K. Prasanna, "Fine-grained address segmentation for attention-based variable-degree prefetching," in *Proceedings of the 19th ACM International Conference on Computing Frontiers*, 2022, pp. 103–112.
- [6] P. Zhang, R. Kannan, X. Tong, A. V. Nori, and V. K. Prasanna, "Sharp: Software hint-assisted memory access prediction for graph analytics," in *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, 2022, pp. 1–8.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [8] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [9] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.
- [10] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *International conference on machine learning*. PMLR, 2021, pp. 10 347–10 357.
- [11] Y. Li, G. Yuan, Y. Wen, J. Hu, G. Evangelidis, S. Tulyakov, Y. Wang, and J. Ren, "Efficientformer: Vision transformers at mobilenet speed," *Advances in Neural Information Processing Systems*, vol. 35, pp. 12 934–12 949, 2022.
- [12] S. Mehta and M. Rastegari, "Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer," *arXiv preprint arXiv:2110.02178*, 2021.
- [13] K. Wu, J. Zhang, H. Peng, M. Liu, B. Xiao, J. Fu, and L. Yuan, "Tinyvit: Fast pretraining distillation for small vision transformers," in *European Conference on Computer Vision*. Springer, 2022, pp. 68–85.
- [14] A. Bhandare, V. Sripathi, D. Karkada, V. Menon, S. Choi, K. Datta, and V. Salefore, "Efficient 8-bit quantization of transformer neural machine language translation model," *arXiv preprint arXiv:1906.00532*, 2019.
- [15] H. You, Z. Sun, H. Shi, Z. Yu, Y. Zhao, Y. Zhang, C. Li, B. Li, and Y. Lin, "Vitcod: Vision transformer acceleration via dedicated algorithm and accelerator co-design," in *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2023, pp. 273–286.
- [16] A. Ivanov, N. Dryden, T. Ben-Nun, S. Li, and T. Hoefler, "Data movement is all you need: A case study on optimizing transformers," *Proceedings of Machine Learning and Systems*, vol. 3, pp. 711–732, 2021.
- [17] A. Khan, A. K. Paul, C. Zimmer, S. Oral, S. Dash, S. Atchley, and F. Wang, "Hvac: Removing i/o bottleneck for large-scale deep learning applications," in *2022 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2022, pp. 324–335.
- [18] T. Dao, D. Fu, S. Ermon, A. Rudra, and C. Ré, "Flashattention: Fast and memory-efficient exact attention with io-awareness," *Advances in Neural Information Processing Systems*, vol. 35, pp. 16 344–16 359, 2022.
- [19] H. Tabani, A. Balasubramaniam, S. Marzban, E. Arani, and B. Zonooz, "Improving the efficiency of transformers for resource-constrained devices," in *2021 24th Euromicro Conference on Digital System Design (DSD)*. IEEE, 2021, pp. 449–456.
- [20] F. Busato, O. Green, N. Bombieri, and D. A. Bader, "Hornet: An efficient data structure for dynamic sparse graphs and matrices on gpus," in *2018 IEEE High Performance Extreme Computing Conference (HPEC)*, 2018, pp. 1–7.
- [21] W. Hu, D. Xu, Z. Fan, F. Liu, and Y. He, "Vis-top: Visual transformer overlay processor," *arXiv preprint arXiv:2110.10957*, 2021.
- [22] Z. Lit, M. Sun, A. Lu, H. Ma, G. Yuan, Y. Xie, H. Tang, Y. Li, M. Leeser, Z. Wang *et al.*, "Auto-vit-acc: An fpga-aware automatic acceleration framework for vision transformer with mixed-scheme quantization," in *2022 32nd International Conference on Field-Programmable Logic and Applications (FPL)*. IEEE, 2022, pp. 109–116.
- [23] M. Sun, H. Ma, G. Kang, Y. Jiang, T. Chen, X. Ma, Z. Wang, and Y. Wang, "Vaqf: Fully automatic software-hardware co-design framework for low-bit vision transformer," *arXiv preprint arXiv:2201.06618*, 2022.
- [24] A. Rahman, J. Lee, and K. Choi, "Efficient fpga acceleration of convolutional neural networks using logical-3d compute array," in *2016 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2016, pp. 1393–1398.
- [25] T. Wang, L. Gong, C. Wang, Y. Yang, Y. Gao, X. Zhou, and H. Chen, "Via: A novel vision-transformer accelerator based on fpga," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 11, pp. 4088–4099, 2022.
- [26] D. Hendrycks and K. Gimpel, "Gaussian error linear units (gelus)," *arXiv preprint arXiv:1606.08415*, 2016.
- [27] G. R. Nair, H.-S. Suh, M. Halappanavar, F. Liu, J.-s. Seo, and Y. Cao, "Fpga acceleration of gcn in light of the symmetry of graph adjacency matrix," in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2023, pp. 1–6.
- [28] V. Iskandar, M. A. A. E. Ghany, and D. Goehringer, "Near-memory computing on fpgas with 3d-stacked memories: Applications, architectures, and optimizations," *ACM Transactions on Reconfigurable Technology and Systems*, vol. 16, no. 1, pp. 1–32, 2022.
- [29] S. Lu, M. Wang, S. Liang, J. Lin, and Z. Wang, "Hardware accelerator for multi-head attention and position-wise feed-forward in the transformer," in *2020 IEEE 33rd International System-on-Chip Conference (SOCC)*. IEEE, 2020, pp. 84–89.
- [30] S. Nag, G. Datta, S. Kundu, N. Chandrathoodan, and P. A. Beerel, "Vita: A vision transformer inference accelerator for edge applications," *arXiv preprint arXiv:2302.09108*, 2023.
- [31] Xilinx. (2017) Deep learning with int8 optimization on xilinx devices. [Online]. Available: <https://docs.xilinx.com/v/u/en-US/wp486-deep-learning-int8>
- [32] G. Cardarilli, L. Di Nunzio, R. Fazzolari *et al.*, "A pseudo-softmax function for hardware-based high speed image classification," *Scientific Reports*, vol. 11, p. 15307, 2021.
- [33] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.