

# FAMOUS: Flexible Accelerator for the Attention Mechanism of Transformer on UltraScale+ FPGAs

Ehsan Kabir\*, Md. Arafat Kabir\*, Austin R.J. Downey<sup>§</sup>, Jason D. Bakos<sup>†</sup>, David Andrews\*, Miaoqing Huang\*

\*Department of EECS, University of Arkansas, Fayetteville, Department of <sup>†</sup>CSE, <sup>§</sup>ME, University of South Carolina, USA  
 {ekabir, makabir, dandrews, mqhuang}@uark.edu, austindowney@sc.edu, jbakos@cse.sc.edu

**Abstract**—This paper proposes *FAMOUS*, a flexible hardware accelerator for dense multi-head attention (MHA) computation of Transformer neural networks (TNNs) on field-programmable gate arrays (FPGAs). It is optimized for high utilization of processing elements and on-chip memories to improve parallelism and reduce latency. An efficient tiling of large matrices has been employed to distribute memory and computing resources across different modules on various FPGA platforms. The design is evaluated on Xilinx Alveo U55C data center cards containing UltraScale+ FPGAs. Experimental results showed that it can attain a maximum throughput, the number of parallel attention heads, embedding dimension, and tile size of 328 (giga operations/second (GOPS)), 8, 768 and 64 respectively on the U55C. Furthermore, it is  $3.28\times$  and  $2.6\times$  faster than the Intel Xeon Gold 5220R CPU and NVIDIA V100 GPU respectively. It is also  $1.3\times$  faster than the fastest state-of-the-art FPGA-based accelerator.

**Index Terms**—FPGA, Transformer, Attention, High-Level Synthesis, Natural Language Processing, Accelerators.

## I. INTRODUCTION

Transformer neural networks have demonstrated significant advancements in natural language processing (NLP), machine translation, computer vision [1], [2], and other domains in recent years. They contain a remarkable feature named multi-headed attention (MHA) mechanism which is different from the traditional convolutional neural network (CNN), recurrent neural network (RNN), and long short term memory (LSTM) model. It enables a high level of computational parallelism for both the training and inference phases making it highly suitable for acceleration on hardware like GPUs and FPGAs, with FPGAs being particularly advantageous due to their high degree of parallelism, low latency, and energy efficiency [3]. Most of the FPGA or ASIC-based hardware accelerators for transformers [4] have specialized sparse architecture for a specific application. Thus, they lack the flexibility to be reconfigured for a different model during runtime. Most works [4] used high-level synthesis (HLS) tools, but it is challenging to write efficient HLS code that can effectively manage certain FPGA resources like DSPs for optimal performance [5]. Furthermore, MHA uses a large amount of the block RAMs (BRAM) [6]. Since FPGAs usually have limited BRAM, creating a good partitioning scheme that works well with the architecture is necessary and can be challenging.

To address these challenges, this paper makes the following contributions:

- An efficient tiling of weight matrices to accommodate large models in on-chip memory.
- A novel architecture ensuring high BRAM and DSP utilization for efficient parallel processing of the transformer’s attention mechanism with low latency.
- A parameterized HLS code that enables users to modify some parameters at design time from HLS tool.
- A runtime programmable feature that enables users to modify some parameters at runtime from software.

## II. BACKGROUND

There are several building blocks in transformers of which the multi-head attention (MHA) is described here. Fig. 1 illustrates the scaled dot product attention in each head, which is a crucial part of the MHA layer. The output of MHA can be represented as Equation 1 & 2. The input sequence  $X$  is linearly mapped into  $Q_i, K_i, V_i$  matrices using weights and biases. The parameter  $d_k = d_{model}/h$  is the 2<sup>nd</sup> dimension of  $Q_i$  and  $K_i$ .  $d_{model}$  is a hyperparameter called embedding dimension and  $h$  is number of heads. ‘i’ is the index for attention heads.




Fig. 1: Multihead Attention Layer.

$$Attention(Q_i, K_i, V_i) = softmax \left( Mask \left( \frac{Q_i K_i^T}{\sqrt{d_k}} \right) \right) V_i \quad (1)$$

$$Q_i = X \times W_q + B_q, K_i = X \times W_k + B_k, V_i = X \times W_v + B_v \quad (2)$$

## III. ACCELERATOR ARCHITECTURE

The core of the accelerator shown in Fig. 2 was designed in C language on Vitis high-level synthesis (HLS) 2022.2.1 tool. There are three main processing modules in it. They are denoted as  $QKV_{PM}$ ,  $QK_{PM}$  and  $SV_{PM}$  according to the output they produce. The number of instances for these modules depends on the number of attention heads ( $h$ ). Each module contains an array of processing elements (PE). A PE is comprised of a DSP48 performing multiplication and accumulation (MAC) operations. The number of PEs ( $t$ ) depends on the unrolling factor of the inner loop and the initiation interval of the pipelined outer loop.  $QKV_{PM}$  module generates the query, key, and value matrices. The

arrays used in this module are divided into subarrays using our tiling technique to fit into the BRAMs.  $QK_{PM}$  module performs the matrix-matrix multiplication operations between the Q and K matrices. As these matrices are relatively small, they are not tiled. The output from  $QK_{PM}$  module is transmitted to the  $SV_{PM}$  module after softmax operation, where it undergoes matrix-matrix multiplication operations with the value (V) matrix.




Fig. 2: Accelerator Architecture for Attention Mechanism

#### IV. TILING TECHNIQUE

Fig. 3 describes our unique tiling strategy. The weight matrices are tiled along the second dimension (column of the matrix) only because the first dimension (row of the matrix) is already reduced by the number of heads. Thus, they are




Fig. 3: Tiling Technique in Multihead Attention Layer.

loaded ( $\frac{d_{model}}{TS}$ ) times. Input buffers of each attention head are declared as a two-dimensional matrix of size ( $SL \times TS$ ). Therefore, tiling is applied along the column of the matrix, and they are also loaded ( $\frac{d_{model}}{TS}$ ) times. TS is tile\_size and SL is sequence\_length.

#### V. RUNTIME PROGRAMMABLE FEATURE

The parameters such as attention heads, embedding dimension, and sequence length were runtime programmable. These parameters can be sent to **FAMOUS** from the software using the steps shown in Fig. 4.




Fig. 4: Process for Incorporating Programmability.

#### VI. EVALUATION AND RESULTS

Table I illustrates the runtime programmable capability, resource utilization, and performance of **FAMOUS**. Tests 1, 2, and 3 examine the effect of varying the number of heads, tests 4 and 5 evaluate changes in embedding dimensions, and tests 6, 7, and 8 analyze variations in sequence length, all in relation to latency and throughput (GOPS (giga operations per second)). On Alveo U55C, the lowest latency of 0.94 ms

and the highest GOPS of 328 were achieved for 8 parallel heads when the tile size was 64. Table II compared **FAMOUS** with some GPUs and CPUs running approximately at 1.5GHz frequency. We achieved  $3.28\times$ ,  $2.6\times$ ,  $1.17\times$  speed up, and an increase in throughput compared to Intel Xeon Gold 5220R CPU, NVIDIA V100 GPU, and Intel E5 2698 v4 CPU respectively because of higher parallelism. Table III compared **FAMOUS** with other FPGA-based accelerators. Our latency is lower and GOPS is higher than all other works except for Calabash [6] because it excluded computation time for Q, K, and V calculations.

TABLE I: Overall Results for MHA Accelerator.

Test no.	Sequence Length	Embedding Dimension	Number of Heads	Tile Size	FPGA	Data Format	DSPs	BRAMs 18k	LUTs	FFs	Latency (ms)	GOPS
#1	64	768	8	64	Alveo U55C	8bit fixed	4157 (46%)	3148 (78%)	1284782 (98%)	661996 (25%)	0.94	328
#2			4								1.401	220
#3			2								2.281	135
#4	64	512	8	64	Alveo U55C	8bit fixed	4157 (46%)	3148 (78%)	1284782 (98%)	661996 (25%)	0.597	184
#5		256									0.352	312
#6	128	768	8	64	Alveo U55C	8bit fixed	4157 (46%)	3148 (78%)	1284782 (98%)	661996 (25%)	2	314
#7	32										0.534	285
#8	16										13	16

TABLE II: Comparison with Other Acceleration Platforms.

Platform	Intel E5 CPU [6]	NVIDIA V100 GPU [7]	Intel Xeon CPU [8]	NVIDIA P100 GPU [8]	<b>FAMOUS</b> (Alveo U55C FPGA)
Topologies	64, 768, 12	64, 512, 4	64, 512, 8	64, 512, 4	64, 768, 8   64, 512, 8
GOP	0.308	0.11	0.11	0.11	0.308
Latency (ms)	1.1	1.5578	1.96	0.496	0.94
GOPS	280	71	56	221	184

TABLE III: Comparison with FPGA Accelerators.

Works	Calabash [6]	Lu et al. [9]	Ye et al. [8]	Li et al. [7]	Peng et al. [4]	<b>FAMOUS</b>
FPGAs	Xilinx VU9P	Xilinx VU13P	Alveo U250	Xilinx VU37P	Alveo U200	Alveo U55C
Method	HDL	HDL	HDL	HLS	HLS	HLS
DSPs	4227	129	4189	1260	623	4157
BRAMs	640	498	1781	448	—	3148
GOPS	1288	128	171	72	97	623
Latency (ms)	0.239 <sup>a</sup>	0.8536 <sup>b</sup>	0.642	1.5264	1.706 <sup>c</sup>	0.494

<sup>a</sup> Q, K, V matrix computation time ignored.

<sup>b</sup> Time adjusted for 8 attention heads.

<sup>c</sup> Time extracted for attention mechanism from a full transformer.

#### VII. CONCLUSION

This research presents a flexible FPGA-based accelerator for the multi-head attention layer of transformer neural networks, designed using high-level synthesis. It supports runtime programmability for various topologies without requiring synthesis. Efficient tiling enables large models to fit on-chip while optimizing computation. The accelerator achieves 328 GOPS throughput, outperforming some CPUs and GPUs, with  $1.3\times$  lower latency than the fastest state-of-the-art FPGA-based solutions.

#### REFERENCES

- [1] A. Vaswani et al., "Attention is all you need," *NeurIPS*, 2017.
- [2] T. Wang et al., "ViA," *IEEE TCAD*, 2022.
- [3] K. Guo et al., "[dl] a survey of fpga-based neural network inference accelerators," *ACM TRET*, 2019.
- [4] H. Peng et al., "Accelerating Transformer-based Deep Learning Models on FPGAs using Column Balanced Block Pruning," in *ISQED*, 2021.
- [5] E. Kabir et al., "Accelerating lstm-based high-rate dynamic system models," in *FPL*, 2023.
- [6] Z. Luo et al., "Calabash," in *FPL*, 2023.
- [7] T. Li et al., "Unified Accelerator for Attention and Convolution in Inference Based on FPGA," in *ISCAS*, 2023.
- [8] W. Ye et al., "Accelerating Attention Mechanism on FPGAs based on Efficient Reconfigurable Systolic Array," *ACM TECS*, 2023.
- [9] S. Lu et al., "Hardware Accelerator for Multi-Head Attention and Position-Wise Feed-Forward in the Transformer," in *SOCC*, 2020.