

A Length Adaptive Algorithm-Hardware Co-design of Transformer on FPGA Through Sparse Attention and Dynamic Pipelining

Hongwu Peng^{1,+}, Shaoyi Huang^{1,+}, Shiyang Chen², Bingbing Li¹, Tong Geng³, Ang Li³, Weiwen Jiang⁴, Wujie Wen⁵, Jinbo Bi¹, Hang Liu² and Caiwen Ding¹

⁺These authors contributed equally.

¹University of Connecticut, USA. ²Stevens Institute of Technology, USA. ³Pacific Northwest National Laboratory, USA. ⁴George Mason University, USA. ⁵Lehigh University, USA.

¹{hongwu.peng, shaoyi.huang, bingbing.li, jinbo.bi, caiwen.ding}@uconn.edu, ²{schen94, hliu77}@stevens.edu,

³{tong.geng, ang.li}@pnl.gov, ⁴wjiang8@gmu.edu, ⁵wuw219@lehigh.edu

ABSTRACT

Transformers are considered one of the most important deep learning models since 2018, in part because it establishes state-of-the-art (SOTA) records and could potentially replace existing Deep Neural Networks (DNNs). Despite the remarkable triumphs, the prolonged turnaround time of Transformer models is a widely recognized roadblock. The variety of sequence lengths imposes additional computing overhead where inputs need to be zero-padded to the maximum sentence length in the batch to accommodate the parallel computing platforms. This paper targets the field-programmable gate array (FPGA) and proposes a coherent sequence length adaptive algorithm-hardware co-design for Transformer acceleration. Particularly, we develop a hardware-friendly sparse attention operator and a length-aware hardware resource scheduling algorithm. The proposed sparse attention operator brings the complexity of attention-based models down to linear complexity and alleviates the off-chip memory traffic. The proposed length-aware resource hardware scheduling algorithm dynamically allocates the hardware resources to fill up the pipeline slots and eliminates bubbles for NLP tasks. Experiments show that our design has very small accuracy loss and has $80.2 \times$ and $2.6 \times$ speedup compared to CPU and GPU implementation, and $4 \times$ higher energy efficiency than state-of-the-art GPU accelerator optimized via CUBLAS GEMM.

KEYWORDS

Transformer, Attention, BERT, Length adaptive, FPGA

1 INTRODUCTION

Transformers are considered as one of the most important deep learning models since 2018 [1], in part because it could potentially replace existing Deep Neural Networks (DNNs), such as Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) [2]. By leveraging self-attention [3], Transformers have established state-of-the-art (SOTA) records (beyond human level) in various fields. Using computer vision as an example, CNNs were the first choice previously [4]; but nowadays, Transformer is gradually becoming the potential alternative, both from theoretical demonstration and empirical explorations, e.g., [2, 5–7].

Despite the remarkable triumphs, the prolonged turnaround time of Transformer models is a widely recognized roadblock that concerns real-world applications. Fig.s 1(a) and (b) depict the architecture of one encoder a.k.a. the building block for Transformers.

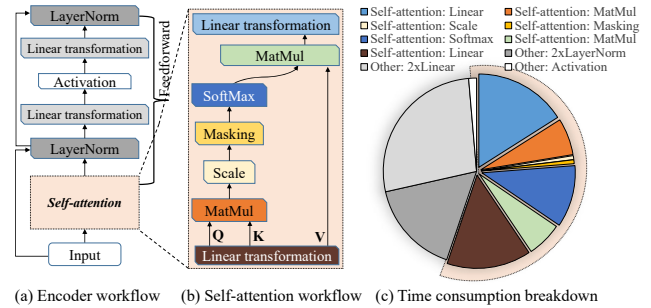


Figure 1: The architecture of a four-head encoder. The time consumption is measured on for TensorRT [15] on WikiText-2 dataset [16], where the input sequence has 128 tokens.

Briefly, one encoder takes the word embeddings of a sequence as input. These embeddings pass through the self-attention mechanism to produce an attention matrix. This matrix is fed through layer normalization, linear transformation, and activation operations to derive the output. Various Transformer model variants often stack different numbers of encoders and decoders together [3]. The bad news is that the time consumption of a single encoder in Fig. 1 could easily reach 100s of μ s, which is $\sim 10 \times$ slower than a typical CNN model. Of such a long latency, around 60% of the time is spent in the self-attention workflow. According to our preliminary study, the time consumption ratio of self-attention is projected to climb if the number of tokens in the input sequence increases, especially in the NLP field [8] further. There exist several initial attention-aware optimization attempts, such as sparse attention [9] and attention approximation [10–14], to leverage runtime approximations or domain knowledge, i.e., tokens only attend their nearby tokens & a few sampled tokens as the summary of the sentence instead of attending all tokens. However, these explorations, unfortunately, fall short by either lacking generality or high computation overheads.

To make it worse, Transformer brings the challenge of a wide variety of input lengths, where inputs need to be zero-padded to the maximum sentence length in the batch to accommodate the parallel computing platforms such as GPU and Field-Programmable Gate Array (FPGA) [17]. By nature, RNN-based models, e.g., GRU and LSTM, process the inputs sequentially. Thus, the inputs could be divided into unified fixed-length sub-inputs and processed independently. Transformers leverages parallel processing and therefore cannot enjoy the benefit of fixed-length sub-inputs. Existing works

on sequence length standardization fall into two categories. The first one is padding or truncation, which forces the sequence length to be the same. It leads to enormous computation overhead due to the unnecessary computation of the padding part. The second category divides a sequence batch into micro-batches (padding within the micro-batch) to mitigate the computation overhead. However, the various and irregular sequence length undermines overall performance and throughput at the inter micro-batch level. Together with the prolonged turnaround time, achieving fast and efficient Transformer models becomes a grand challenge.

Across all the popular hardware, e.g., CPUs, GPUs, FPGAs, and Application-Specific Integrated Circuits (ASICs), FPGAs strike an effective balance among massive parallelism, high energy efficiency and short development cycle, hence lend themselves as the top choice to expedite the Transformer architecture. In this paper, we believe that the ideal Transformers acceleration should have a coherent algorithm–hardware co-design. We believe that **(i) Transformers should have their dedicated efficient algorithm designs**. Since self-attention cares more about the value relativity of all the attention scores than the absolute value of any specific attention score, we propose an efficient scheme to exploit two different self-attention approximations adaptively. Note, our approximation mechanisms are quantization-based designs that are not only computation-efficient but also hardware-friendly. We also think that **(ii) Transformers should efficiently support various sequence length inputs**. For instance, SQuAD v2.0 [8] has an average and maximum sequence length of 171 and 975, respectively. When padding the sequence with 975, it causes 5.7× computational and memory bandwidth overhead on average. The inputs are sorted and processed according to the order of length. Compared to existing works, we achieve 4 × higher energy efficiency than GPU accelerator optimized through CUBLAS GEMM routine [18, 19] with small accuracy loss, and comparable energy efficiency compared to ASIC accelerator designs [12, 13]. Our contributions are:

- We propose sparse attention which is computation-efficient and hardware-friendly to reduce the need for computational resources and memory bandwidth.
- We propose a sequence length adaptive design to allocate coarse pipeline stages dynamically to eliminate pipeline bubbles and achieve the highest possible throughput under different sequence length inputs.
- Transformer exhibits a highly skewed distribution of computation complexity among the operators. We further develop a loop fusion to orchestrate the multiple attention operators and re-arrange various Transformer computations to enhance temporal locality and efficient hardware design with finer granularity.

2 RELATED WORK

Attention-aware Optimization. We also notice several recent attention-aware optimization attempts, such as *sparse attention* [9] and *attention approximation* [10–14], which unfortunately fall short by either lack of generality or high computation overheads. Particularly, *sparse attention mechanisms* [9] leverage domain knowledge, i.e., tokens only attend their nearby tokens and a few sampled tokens as the summary of the sentence instead of attending all

tokens, to reduce computation and memory consumption during self-attention computation. Such design requires a pre-determined attention mask that lacks generality.

Approximation-based attention leverages runtime approximations to derive sparse attention which faces non-trivial overheads. BP-Transformer [10] and Reformer [11] convert the self-attention computation into a nearest neighbor search problem and use either tree-based search, Locality Sensitive Hashing (LSH), or low-rank approximation to find similar tokens for attention. A³ [12] embraces architecture innovation to estimate the closeness between tokens. However, the estimation process still requires full access to original matrices and does not alleviate memory bottleneck for attention computation according to SpAtten [13]. ELSA [14] uses LSH distance for attention rank approximation, but it again suffers significant overheads for LSH.

Sequence length standardization. TensorRT [15] utilized the padding and truncation [17] to standardize the sequence length for parallel computing. It makes the hardware design regularized but leads to enormous computation overhead due to the unnecessary computation of the zero-padding part. TurboTransformer [20] divided a batch into micro-batches with similar lengths; however, within the micro-batch, it still required maximum sequence length padding. To make things worse, when we implement this method on FPGA, it introduces significant pipeline bubbles.

3 SPARSE ATTENTION ALGORITHM

3.1 Overview

Approximation-based attention leverages run-time approximations to derive sparse attention which faces non-trivial overheads. BP-Transformer [10] and Reformer [11] convert the self-attention computation into a nearest neighbor search problem and use either tree-based search, Locality Sensitive Hashing (LSH), or low-rank approximation to find similar tokens for attention. A³ [12] embraces architecture innovation to estimate the closeness between tokens. However, the estimation process still requires full access to original matrices and does not alleviate memory bottleneck for attention computation according to SpAtten [13].

We propose to compute sparse self-attention via rapid attention rank approximation and sparse attention computation. We firstly quantize the full precision \mathbf{Q} and \mathbf{K} into low bits representation, i.e., 1 bit or 4 bits. Then we conduct matrix multiplication on quantized value and get the Top- k candidates index. At last, we conduct full precision sparse attention computation based on Top- k candidates. The algorithm **reduces the attention complexity from $O(n^2)$ to $O(n)$** , where n is sequence length.

3.2 Sparse Attention Via Q & K Quantization

In self-attention, the input is transformed into three matrices \mathbf{Q} , \mathbf{K} and \mathbf{V} . Then \mathbf{Q} and \mathbf{K} are multiplied to arrive at $\mathbf{S} = \mathbf{Q} \cdot \mathbf{K}^T$, where each element in the resultant matrix is an attention score. A row of attention scores in \mathbf{S} represent the dot-product between a row vector in \mathbf{Q} and all row vectors in \mathbf{K} respectively. Step ① in Fig. 3 illustrates one row of \mathbf{Q} multiplying with \mathbf{K} , where \mathbf{q} is one row in matrix \mathbf{Q} . Subsequently, we perform a softmax operation on \mathbf{S} , which is $\mathbf{S}_i = \frac{\exp(\mathbf{qk}_j)}{\sum_{j=1}^4 \exp(\mathbf{qk}_j)}$, i.e., step ② in Fig. 3. The key observation is: since softmax is a normalization method, *it is the*

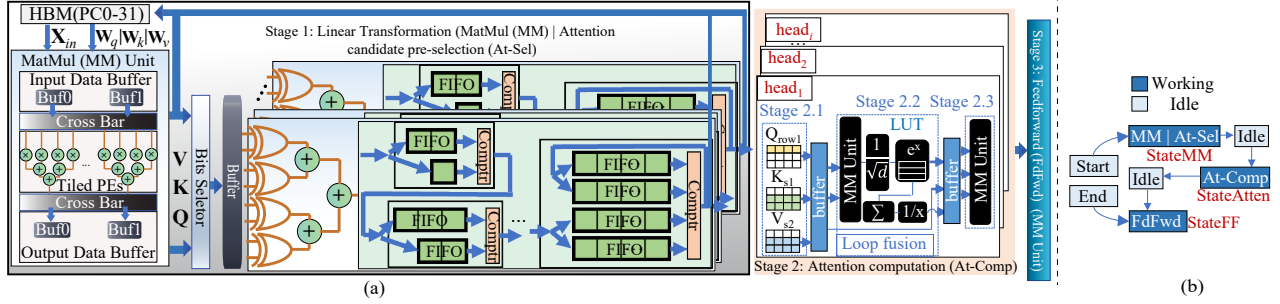


Figure 2: (a) Sparse attention on FPGA; (b) State machine.

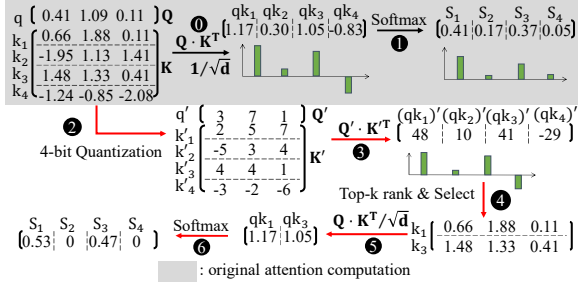


Figure 3: Candidate selection from quantized result.

value relativity of all the attention scores, as opposed to the absolute value of any specific attention score, that matters.

We propose to quantize Q and K from the full-precision representation (usually 32-bit floating-point) into a low-precision integer representation. Because both quantization and exponential operations used in softmax are monotonically increasing operators, the quantized results maintain the order of attention scores. We use our fast quantized matrix multiplication to extract dominant attention values. Afterward, we perform accurate attention computation only for dominant attention scores. The design is depicted in Fig. 3.

Particularly, we first find out the suitable scaling factor M for the given tensor to quantization, then perform $x' = \text{round}(\frac{2^3-1}{M}x)$, which casts all the floating point values into a desired integer. For example, the scaling factor M of K in Figure 3 is 0.77, so each element is multiplied with $\frac{2^3-1}{0.77}$ and rounded to the nearest integer. We follow a similar procedure to quantize q into q' . Subsequently, we again use a look-up table to perform the multiplication. For instance, if we multiply two 4-bit integers, the look-up table only needs 256 entries. We can easily estimate the multiplied value. At the end of step ②, we derive the $Q' \cdot K'^T$. As the examples indicate, the quantized results keep the same rank and distribution compared with their full-precision counterpart.

We conduct Top- k sort and select the Top- k ranked attention scores for exact matrix multiplication, which derives more accurate softmax values. This is faster than the original design because we only need to compute Top- k attention scores. In step ④ in Fig. 3, we select Top-2 element k_1 and k_3 to perform matrix multiplication and softmax, which is used as an approximation of the result of self-attention. Subsequently, we will perform full-precision $Q \cdot K^T$ for the selected attention scores at ⑤ and final softmax at step ⑥.

4 HARDWARE ACCELERATOR DESIGN AND SCHEDULING ALGORITHM

The proposed quantization-based sparse attention system design is more fit on FPGAs than general-purpose processors because the latter are instruction-driven architecture. At the same time, FPGAs are data-driven architecture that avoids instruction fetch and related memory access. When compared to the popular GPU accelerators, FPGAs excel for the following reasons: (1) The on-chip memory capacity of FPGAs is much higher (360 \times) than that of GPUs (i.e., 35 MB in Xilinx Alveo U200 vs. 96 KB in V100). The FPGA on-chip memory features its high memory bandwidth (31 TB/s) and low access latency (single clock cycle), enabling higher throughput and lower latency design [21–28]. With more on-chip memory size, we can achieve a better computation to communication (CTC) ratio for the same operations, i.e., matrix multiply and matrix add. (2) FPGA provides more design opportunities on fine-grained and coarse-grained pipelining and loop fusion techniques. We can have better data locality optimization, and design space freedom on FPGA through polyhedral analysis and proper loop reschedule. (3) FPGAs are more power-efficient, and therefore more suitable for resource-constrained scenarios.

The FPGA platform enables better intra-attention coarse grain pipelining design and leaves more freedom on FPGA resource allocation. The commonly used way for NLP tasks with variable length inputs is to unify input sequence to a fixed length through padding and cutting. However, length padding introduces unnecessary overhead, and length cutting leads to information loss. To accelerate the NLP tasks with variable sequence length, we propose sequence length adaptive Transformer hardware design on the FPGA platform and the corresponding hardware scheduling algorithms to optimize the coarse-grained stage throughput. Compared to commonly used padding and cutting methods, our proposed method has higher hardware throughput and less information loss.

4.1 Sparse Attention Accelerator Design

We break down the original single Transformer Encoder stage pipeline into three coarse-grained pipelines and overlap their execution time by inserting buffers for each concatenated pipeline pair, as shown in Fig. 2(a). Stage 1 contains the linear transformation (using MatMul (MM)) and quick attention approximations (a.k.a, pre-selection) (At-Sel) hardware. The MM result is directly fed into the bits selector hardware for ultra-low bit quantization. The result is stored in the on-chip buffer for candidate pre-selection computation (utilizing LUT hardware for approximate distance calculation). The approximate distance output and address are then fed to the

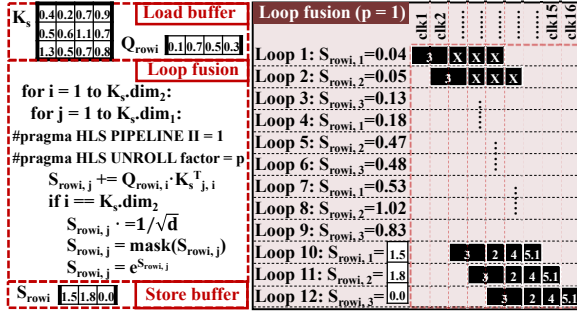


Figure 4: An attention kernel fusion example.

merge sort hardware for high throughput ($\Pi=1$) scalable Top- k sort [29]. The Top- k results (e.g., index and value pairs) are stored back to HBM for inter-stage buffering. Stage 2 is attention computation (At-Comp) and Stage 3 is feedforward (FdFwd). Stage 2 is divided into three sub-stages and implemented with the intra-layer coarse-grained pipeline to enhance hardware utilization. Stage 2.1, the data loading stage, utilizes the Top- k results from stage 1 to choose the attention candidates. As discussed in Section. 3.2, each Q_{rowi} have selected candidates K_{si} and V_{si} matrix for attention computation. Stage 2.2 is implemented with loop fusion and is composed of operators ④ to ⑥ in Fig. 3. Stage 2.3 is composed of $Z_i = S_i \cdot V / \text{sum}(S_i)$ operation for each row i . The double-buffers added between stages buffer the data produced/consumed by the previous/current stage for coarse-grained pipelining. Stage 2.3 conduct the final MM operation between Z_i and V Stage 3 (feedforward) is composed of MM, add, layer normalization, and GELU unit.

In stage 2, We divide the softmax ⑥ in Fig. 3 into two operations: exponent calculation and normalization. We leverage the FPGA’s fine-grained pipelining characteristic to fuse multiple attention operators (i.e., ⑤ to ⑥.1) into a single loop. An example is given in Fig. 4, the scaling, mask, and exponential operations are conducted at the last loop iterations. Q_{rowi} is the processed row and K_s is the corresponding selected candidate for the attention. $K_s.dim_1$ and $K_s.dim_2$ are the size of 1st and 2nd dimension of K_s . S_{rowi} is the processed attention score result. Thanks to the reconfigurable architecture of FPGAs, we enable fusing the loops with different iteration trip counts while GPU designs (e.g., TensorRT) only can fuse specific loops (e.g., scale + matrix multiply). Both Stage 1 and 3 leverage on-chip memory and proper loop allocation to mitigate communication bottlenecks between the on and off-chip memory.

4.2 Length-aware Scheduling Algorithm

Different length sequences tasks consume different computational resources and have different latency, which introduces pipeline bubbles due to irregular and unpredictable dataflow. Because all operators have $O(n)$ complexity where n is sequence length, we proposed a novel length-aware coarse-grained pipeline algorithm to dynamically adjust the hardware resource allocation to illuminate the pipeline bubble. We leverage FPGA characteristics to adjust the resource allocation according to stages’ computation complexity, eliminate redundant computation and achieve ultra-high throughput. We propose the techniques detailed below.

Resource Scheduling Algorithm. The slowest stage constrains the throughput of the coarse-grained pipeline. We first develop a Encoder coarse-grained stage allocation algorithm (Algorithm 1) to schedule operators efficiently. For original operator graph $G =$

Algorithm 1: Encoder coarse-grained Stage Allocation.

Input: Encoder operator graph $G = (V, E)$, Encoder operator weight set $W(V, s_{avg})$, and Encoder priority set $P(V, s_{avg})$;
Output: operator subgraph of each Encoder computation stage $G_k = (V_k, E_k)$;
 Traverse $G = (V, E)$ and compute priority set $P(V, s_{avg})$ for the Encoder;
 $k \leftarrow 1, N(V) \leftarrow \{1\}, G_1 \leftarrow v_1$; // add a operator to a new stage
foreach $v_j \in V$ in decreasing order of $P(v, s_{avg})$ **do**
 foreach $N'(v_j) \in G_k$ **do**
 $N'(v_j) \leftarrow N(v_j) \cdot \lceil \frac{W(v_j, s_{avg})}{W(v_i, s_{avg})} \rceil$;
 end
 if resource constraints are satisfied then
 $G_j \leftarrow v_j$; // add a operator to current stage
 $N(V) \leftarrow N'(V)$; // update operator parallelisms
 else
 $k, K \leftarrow k + 1$;
 $G_k \leftarrow v_j$; // add a operator to a new stage
 end
end
return $\{G_1, G_2, \dots, G_K\}$;

(V, E) , each vertex $v_i \in V$ represents an operator and the edge e_{ij} represents the data dependency between v_i and v_j . Each vertex v_i has a weight $W(v_i, s_{avg})$ which is the associated arithmetic computational complexity. It takes the G and Encoder operator weight set $W(V, s_{avg})$ at average sequence length s_{avg} , and Encoder operator priority set at average sequence length $P(V, s_{avg})$ as input and outputs operator subgraph of each Encoder computation stage $G_k = (V_k, E_k)$ shown in Eq. 1, where $W(v_i, s_{avg})$ and $P(v_i, s_{avg})$ denote their value at vertex (v_i, s_{avg}) . To fully utilize the resources of a certain FPGA chip for sequence length adaptive design, we further adjust the operator parallelism $N(v_i, s_i)$ for intra coarse-grained pipeline stages and enumerate pipeline replication factor $R(G_k, s_i)$ to obtain the optimal setting with the help of analytical performance and resource models.

$$P(v_i, s_{avg}) = \begin{cases} W(v_i, s_{avg}) + \max_{v_j \in \text{Succ}(v_i)} P(v_j, s_{avg}), & v_i \neq v_{\text{sink}} \\ W(v_{\text{sink}}, s), & \text{otherwise} \end{cases} \quad (1)$$

Length-aware Coarse-grained Pipeline Algorithm. We then develop a length adaptive resource scheduling method to dynamically patch the pipeline bubbles for a batch of tasks with different sequence lengths. The effectiveness of the proposed scheduling method relies on the fact that **all operators have $O(n)$ complexity**. The batch inputs are sorted and processed according to the decreasing order of length, under the control of a dedicated state machine (three states shown in Fig. 2(b): *StateMM*, *StateAtten* and *StateFF*) during its Encoder activation period. The state machines dynamically allocate hardware resources (stages and buffers) to eliminate pipeline bubbles and ensure a high hardware utilization for batch sentences with varying lengths. Each stage has almost 100% utilization, and there is no pipeline bubble. We could significantly reduce the latency (denoted as "saved").

A length-aware scheduling timing diagram example is given in Fig. 5, where the batch size is 5, and the input sequence length varies from 72 to 140. The batch inputs are sorted and fed into the Encoder coarse-grained stages in decreasing order of sequence length. Fig. 5(a) shows how each of the Encoder coarse-grained stages processes each sequence input. Most attention-based models have multiple Encoder layers, so the batch input is processed by the layer order. Fig. 5(b) shows the hardware resource occupation of Encoder coarse-grained stages. With the state machine-based

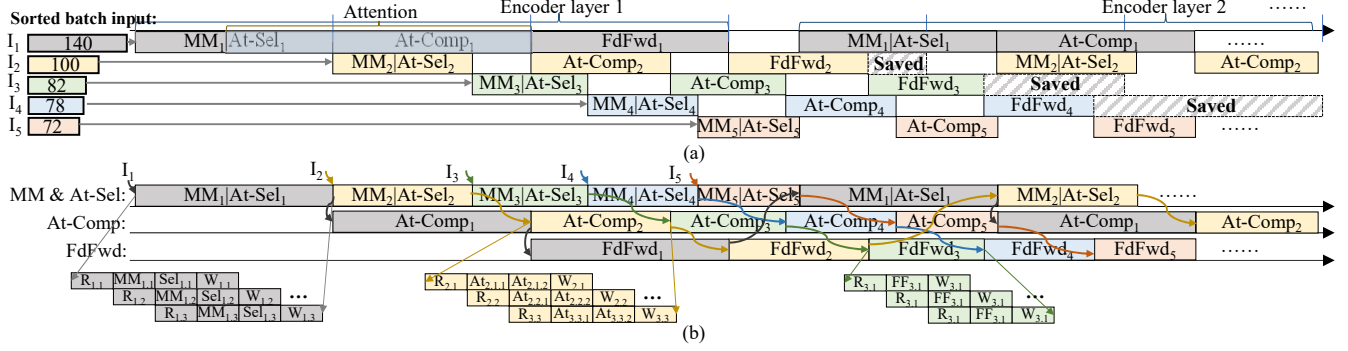


Figure 5: Length-aware coarse-grained dynamic pipeline algorithm example: (a) timing diagram; (b) hardware utilization.

scheduling algorithm implemented, the pipeline stages of different sequence length inputs and different Encoder layers are patched together without pipeline bubble, so both stages have almost 100% hardware utilization. The intra-layer coarse-grained pipeline is implemented in each stage to exploit the trade-off between spatial & temporal data locality and hardware resource occupation. The communication and computation are overlapped with each other through coarse-grained pipeline and data prefetching.

5 EXPERIMENT

We evaluate several well-known self-attention centric models to demonstrate the algorithm & hardware design performance. For NLP models, we choose four of the most popular ones: BERT-base [30], BERT-large, and DistilBERT [31], and RoBERTa [32]. Model configurations are given in Table 1. These models are self-attention-centric and have a similar structure. For BERT-base, DistilBERT and RoBERTa, we run 3 representative datasets to evaluate the performance: SQuAD v1.1 [33], RTE [34], and MRPC [35]. For BERT-large, we run SQuAD v1.1 for evaluation. The minimum sequence length, average sequence, and maximum sequence of those datasets are also given. Max/Avg ratio also corresponds to the computational overhead introduced through padding.

Table 1: Model & evaluation dataset.

Model	Layers	Hidden dim	Num. of Heads
DistilBERT	6	768	12
BERT-base, RoBERTa	12	768	12
BERT-large	24	1024	16
Evaluation dataset	Avg	Max	Max/Avg
SQuAD v1.1	177	821	4.6
RTE	68	253	3.7
MRPC	53	86	1.6

The FPGA hardware design and evaluation are conducted on the Alveo U280 platform. We also evaluate the hardware performance on CPU, edge GPU, and GPU server platforms for cross-platform comparison: Intel(R) Xeon(R) Gold 5218 CPU, Jetson TX2, and Quadro RTX 6000. The FPGA design is conducted on software version Vivado 2020.1, whereas GPU and CPU design is completed on Pytorch 1.10.0 and Transformers 4.13.0.dev0.

5.1 Sparse Attention Accuracy Evaluation

We evaluate the models mentioned above and datasets on the Top- k sparse attention algorithm. The state-of-the-art models are quantized into 8 bits fixed-point representation without accuracy drop [36]. The Q & K quantization is conducted based on 1-bit quantization, which is a sign function. For the sparse attention algorithm, the

quantized models are directly used without model fine-tuning. For SQuAD v1.1 and MRPC datasets, we use the F1 score as our accuracy measure. For the RTE dataset, the raw accuracy is reported. Experiments are conducted on the corresponding validation dataset.

Fig. 6 shows the accuracy test of evaluated models and datasets. We choose k value from 10 to 50 to assess the effectiveness of sparse attention, where the k value determines the degree of approximation for sparse attention computation. Generally, smaller k indicates aggressive approximation and leads to a higher accuracy drop. For most of the evaluation, Top-10 sparse attention lead to non-negligible performance degradation. Top-30 provides a good trade-off between the accuracy drop and sparsity ratio, whereas all evaluations have less than 2% accuracy drop. With a Top-30 sparse attention, the attention computation complexity can be reduced by more than 80% in average.

5.2 Cross-platform Throughput Evaluation

Based on the model accuracy evaluation, we choose a sweet point $k = 30$ for Top- k pre-selection, then we mapped models into FPGA hardware coarse-grained stages through Algorithm 1. We exploit the design space to maximize the hardware throughput and CTC ratio for the hardware design. The attainable FPGA design frequency is 200 MHz, and most of the hardware resources (BRAM, FF, LUT) consumption are congested inside the SLR0 of the Alveo U280 board since the only SLR0 is connected to HBM channels. HBM channels provide a maximum of 460 GB/s bandwidth. The batch size is set as 16 to maximize the hardware utilization.

For the FPGA platform, 8 bits fixed-point number multiply & accumulate consumes 1 DSP unit. And there are 3000 DSP units within the SLR0 in total. So the maximum attainable computation throughput of the FPGA platform is 1.2 TFLOPS. Thanks to the reconfiguration structure, the FPGA design can efficiently map variable sequence length computation into hardware through the scheduling algorithm discussed in Sec. 4.2. The FPGA design can surpass GPU server performance in all the models & tasks evaluation through efficient scheduling. BERT-base on SQuAD v1.1, RTE, and MRPC, and BERT-large model on SQuAD v1.1 are used for hardware design demonstration. DistilBERT and RoBERTa has similar structure and thus similar hardware performance to BERT-base.

The end-to-end hardware throughput comparison is given in Fig. 7a. FPGA baseline indicates the FPGA design without length-aware scheduling and sparse attention algorithm implemented. The sequence length is padded to the maximum sequence length for the CPU and GPU design to evaluate the tasks. The geomean speedup of the FPGA length-aware sparse attention design is $80.2 \times, 41.3 \times, 2.6$





Figure 6: Accuracy evaluation of Top-k sparse attention.



(a) End to end cross platform throughput comparison

(b) Cross platform attention throughput comparison

Figure 7: Cross platform hardware evaluation.

\times , $3.1 \times$ than CPU, edge GPU, GPU server, and FPGA baseline design. As for the self-attention computation, the hardware throughput is also recorded during the evaluation, and the corresponding speedup is given in Fig. 7b. The FPGA sparse attention hardware achieves a geomean speedup of $1073 \times$, $550 \times$, $35 \times$, $41 \times$ than CPU, edge GPU, GPU server, and FPGA baseline design. Our FPGA design achieves an equivalent hardware throughput of 3.6 TFLOPS on 8 bits fixed-point operations with length-aware scheduling and sparse attention algorithm implemented.

5.3 Cross-work Energy Efficiency Comparison

We further conduct cross-platform and cross-work energy efficiency comparison between GPU baseline, GPU design optimized through CUBLAS GEMM routine [18], FPGA [37], and ASIC implementations [12, 13] of Transformer accelerator. The result is shown in Table 2. Our FPGA surpasses GPU baseline implementation and existing state-of-the-art FPGA design in terms of both throughput and energy efficiency with an acceptable accuracy drop. With the length-aware scheduling and sparse attention algorithms implemented, our work has a comparable energy efficiency to existing ASIC designs dedicated to transformer acceleration.

Table 2: Energy efficiency & throughput comparison.

Work/platform	Throughput (GOPS)	Energy eff. (GOP/J)	Accuracy drop (average)(%)
GPU RTX 6000	1380	8	1.8
GPU V100: E.T. [18]	7550	25	2.1
Ours FPGA	3600	102	1.8
FPGA design [37]	76	N/A	3.8
ASIC: A ³ [12]	221	269	1.6
ASIC: SpAtten [13]	360	382	1.1

6 CONCLUSION

In this paper, we propose a hardware-friendly sparse attention algorithm through query and key values quantization, where we bring down the complexity of self-attention from $O(n^2)$ to $O(n)$. We develop a length-aware hardware scheduling algorithm to accommodate variable sequence length computation into coarse-grained pipeline stages without pipeline bubbles. To alleviate the off-chip

memory access, we further develop an attention kernel fusion to process the attention computation. We exploit temporal data locality through the on-chip buffer to enhance the CTC ratio and push the hardware design to the computation roof. Experimental results show that we achieve $80.2 \times$ and $2.6 \times$ speedup compared to CPU and GPU implementation with 1.8% accuracy loss. Our FPGA design has more than $4 \times$ higher energy efficiency than GPU accelerator optimized through CUBLAS GEMM routine and has comparable energy efficiency compared to state-of-the-art ASIC designs.

ACKNOWLEDGEMENT

This work was in part supported by the NSF CRII Award No. 2000722, CAREER Award No. 2011236, No. 2006748, No. 2046102, and DOE Award No. 66150. Any opinions, findings and conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- [1] Nikolas Adaloglou. Why Multi-head Self Attention Works: Math, Intuitions and 10+1 Hidden Insights. <https://theaisummer.com/self-attention/>, 2021. [Online; accessed August 30, 2021].
- [2] Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. On the Relationship Between Self-Attention and Convolutional Layers. In *ICLR*, 2019.
- [3] Ashish Vaswani et al. Attention is All You Need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016.
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiahua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16×16 Words: Transformers for Image Recognition at Scale. In *ICLR*, 2021.
- [6] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *CVPR*, 2021.
- [7] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-End Human Pose and Mesh Reconstruction with Transformers. In *CVPR*, 2021.
- [8] Pranav Rajpurkar, Robin Jia, and Percy Liang. Know What You Don’t Know: Unanswerable Questions for SQuAD. In *ACL*, pages 784–789, 2018.
- [9] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big Bird: Transformers for Longer Sequences. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *NeurIPS*, volume 33, pages 17283–17297. Curran Associates, Inc., 2020.
- [10] Zihao Ye, Qipeng Guo, Quan Gan, Xipeng Qiu, and Zheng Zhang. BP-Transformer: Modelling Long-range Context via Binary Partitioning. *arXiv preprint arXiv:1911.04070*, 2019.
- [11] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The Efficient Transformer. In *ICLR*, 2019.

- [12] Tae Jun Ham, Sung Jun Jung, Seonghak Kim, Young H Oh, Yeonhong Park, Yoonho Song, Jung-Hun Park, Sanghee Lee, Kyoung Park, Jae W Lee, and Deog-Kyoon Jeong. A³: Accelerating Attention Mechanisms in Neural Networks with Approximation. In *2020 HPCA*, pages 328–341. IEEE, 2020.
- [13] Hanrui Wang, Zhekai Zhang, and Song Han. SpAtten: Efficient Sparse Attention Architecture with Cascade Token and Head Pruning. In *2021 HPCA*, pages 97–110. IEEE, 2021.
- [14] Tae Jun Ham, Yejin Lee, Seong Hoon Seo, Soosung Kim, Hyunji Choi, Sung Jun Jung, and Jae W Lee. ELSA: Hardware-Software Co-design for Efficient, Lightweight Self-Attention Mechanism in Neural Networks. In *ISCA*. IEEE, 2021.
- [15] NVIDIA. TensorRT. Retrieved from <https://developer.nvidia.com/tensorrt>. Online; accessed: October 6, 2021.
- [16] Merity Stephen, Xiong Caiming, Bradbury James, and Socher Richard. Pointer Sentinel Mixture Models. ICLR, 2016.
- [17] Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. Transformers: State-of-the-art natural language processing. In *EMNLP*, pages 38–45, 2020.
- [18] Shiyang Chen, Shaoyi Huang, Santosh Pandey, Bingbing Li, Guang R Gao, Long Zheng, Caiwen Ding, and Hang Liu. Et: re-thinking self-attention for transformer models on gpus. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–18, 2021.
- [19] Shaoyi Huang, Shiyang Chen, Hongwu Peng, Daniel Manu, Zhenglun Kong, Geng Yuan, Lei Yang, Shusen Wang, Hang Liu, and Caiwen Ding. Hmc-tran: A tensor-core inspired hierarchical model compression for transformer-based dnns on gpu. In *Proceedings of the 2021 on Great Lakes Symposium on VLSI*, pages 169–174, 2021.
- [20] Jiarui Fang, Yang Yu, Chengduo Zhao, and Jie Zhou. Turbotransformers: an efficient gpu serving system for transformer models. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 389–402, 2021.
- [21] Hongwu Peng, Balaji Narayanasamy, Asif Imran Emon, Zhao Yuan, Rongxuan Zhang, and Fang Luo. Selective digital active emi filtering using resonant controller. In *2020 IEEE International Symposium on Electromagnetic Compatibility & Signal/Power Integrity (EMCSI)*, pages 632–639. IEEE, 2020.
- [22] Zeke Wang, Hongjing Huang, Jie Zhang, and Gustavo Alonso. Benchmarking High Bandwidth Memory on FPGAs. *arXiv preprint arXiv:2005.04324*, 2020.
- [23] Panjie Qi, Yuhong Song, Hongwu Peng, Shaoyi Huang, Qingfeng Zhuge, and Edwin Hsing-Mean Sha. Accommodating transformer onto fpga: Coupling the balanced model compression and fpga-implementation optimization. In *Proceedings of the 2021 on Great Lakes Symposium on VLSI*, pages 163–168, 2021.
- [24] Hongwu Peng, Shanglin Zhou, Scott Weitze, Jiaxin Li, Sahidul Islam, Tong Geng, Ang Li, Wei Zhang, Minghu Song, Mimi Xie, et al. Binary complex neural network acceleration on fpga. In *2021 IEEE 32nd International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, pages 85–92. IEEE, 2021.
- [25] Hongwu Peng, Shaoyi Huang, Tong Geng, Ang Li, Weiwen Jiang, Hang Liu, Shusen Wang, and Caiwen Ding. Accelerating transformer-based deep learning models on fpgas using column balanced block pruning. In *2021 22nd International Symposium on Quality Electronic Design (ISQED)*, pages 142–148. IEEE, 2021.
- [26] Geng Yuan, Zhiheng Liao, Xiaolong Ma, Yuxuan Cai, Zhenglun Kong, Xuan Shen, Jingyan Fu, Zhengang Li, Chengming Zhang, Hongwu Peng, et al. Improving dnn fault tolerance using weight pruning and differential crossbar mapping for reram-based edge ai. In *2021 22nd International Symposium on Quality Electronic Design (ISQED)*, pages 135–141. IEEE, 2021.
- [27] Shaoyi Huang, Ning Liu, Yueying Liang, Hongwu Peng, Hongjia Li, Dongkuan Xu, Mimi Xie, and Caiwen Ding. An automatic and efficient bert pruning for edge ai systems. In *2022 23rd International Symposium on Quality Electronic Design (ISQED)*, pages 1–6. IEEE, 2022.
- [28] Hongwu Peng, Balaji Narayanasamy, Asif Imran Emon, Zhao Yuan, Mustafeez Ul Hassan, and Fang Luo. Design and implementation of selective active emi filter with digital resonant controller. In *2020 IEEE Energy Conversion Congress and Exposition (ECCE)*, pages 5855–5861. IEEE, 2020.
- [29] Hongwu Peng et al. Optimizing FPGA-based Accelerator Design for Large-Scale Molecular Similarity Search. In *ICCAD*, 2021.
- [30] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *2019 ACL*, pages 4171–4186, 2019.
- [31] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [32] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [33] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100, 000+ questions for machine comprehension of text. In *EMNLP*, 2016.
- [34] Ido Dagan, Bill Dolan, Bernardo Magnini, and Dan Roth. Recognizing textual entailment: Rational, evaluation and approaches—erratum. *Natural Language Engineering*, 16(1):105–105, 2010.
- [35] Bill Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Third International Workshop on Paraphrasing*. AFNLP, 2005.
- [36] Wei Zhang, Lu Hou, Yichun Yin, Lifeng Shang, Xiao Chen, Xin Jiang, and Qun Liu. TernaryBERT: Distillation-aware Ultra-low Bit BERT. In *2020 EMNLP*, pages 509–521, 2020.
- [37] Panjie Qi, Edwin Hsing-Mean Sha, Qingfeng Zhuge, Hongwu Peng, Shaoyi Huang, Zhenglun Kong, Yuhong Song, and Bingbing Li. Accelerating framework of transformer by hardware design and model compression co-optimization. In *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, pages 1–9. IEEE, 2021.