PEANO-ViT: Power-Efficient Approximations of Non-Linearities in Vision Transformers

Mohammad Erfan Sadeghi
University of Southern California
Los Angeles, California, USA
sadeghim@usc.edu

Arash Fayyazi
University of Southern California
Los Angeles, California, USA
fayyazi@usc.edu

Seyedarmin Azizi
University of Southern California
Los Angeles, California, USA
seyedarm@usc.edu

Massoud Pedram
University of Southern California
Los Angeles, California, USA
pedram@usc.edu

ABSTRACT

The deployment of Vision Transformers (ViTs) on hardware platforms, specially Field-Programmable Gate Arrays (FPGAs), presents many challenges, which are mainly due to the substantial computational and power requirements of their non-linear functions, notably layer normalization, softmax, and Gaussian Error Linear Unit (GELU). These critical functions pose significant obstacles to efficient hardware implementation due to their complex mathematical operations and the inherent resource count and architectural limitations of FPGAs. PEANO-ViT offers a novel approach to streamlining the implementation of the layer normalization layer by introducing a division-free technique that simultaneously approximates the division and square root function. Additionally, PEANO-ViT provides a multi-scale division strategy to eliminate division operations in the softmax layer, aided by a Padé-based approximation for the exponential function. Finally, PEANO-ViT introduces a piece-wise linear approximation for the GELU function, carefully designed to bypass the computationally intensive operations associated with GELU. In our comprehensive evaluations, PEANO-ViT exhibits minimal accuracy degradation ($\leq$ 0.5% for DeiT-B) while significantly enhancing power efficiency, achieving improvements of 1.91×, 1.39×, and 8.01× for layer normalization, softmax, and GELU, respectively. This improvement is achieved through substantial reductions in DSP, LUT, and register counts for these non-linear operations. Consequently, PEANO-ViT enables efficient deployment of Vision Transformers on resource- and power-constrained FPGA platforms.

CCS CONCEPTS

• **Hardware** → **High-level and register-transfer level synthesis**; • **Computing methodologies** → **Computer vision**.

KEYWORDS

Vision Transformers, FPGA Implementation, Deep Learning Efficiency, Hardware Acceleration

1 INTRODUCTION

The landscape of computer vision has been fundamentally transformed with the advent of deep learning, among which Vision Transformers (ViTs) [4, 8, 12] have emerged as a particularly promising approach. Unlike traditional convolutional neural networks (CNNs) that rely on local receptive fields, ViTs leverage the power of self-attention mechanisms to capture global dependencies within an image, enabling a more comprehensive understanding of visual data. This capability has placed ViTs at the forefront of research, demonstrating state-of-the-art performance across a wide range of tasks in computer vision. Overall, deep learning has revolutionized various domains by providing robust algorithms capable of learning complex patterns from large datasets, thus enabling unprecedented advancements in the application of artificial intelligence across numerous fields, from healthcare [10] to recommendation systems to scientific research.

ViTs rely on a series of identical encoder blocks to progressively extract complex features from an image. These encoder blocks consist of two principal components: Multi-headed Attention (MHA) and Feed-Forward Network (FFN), each prefaced with a layer normalization block. Embedded within MHA and FFN are linear layers, GELU, and softmax, integrated via two residual connections that bookend the normalization stages. The output of the final encoder block goes through a classifier to obtain the class predictions.

Despite their exceptional performance, ViTs face significant challenges for practical deployment due to their extensive parameter count and considerable computational demands. A wide range of methods has been explored to improve the efficiency of ViTs, including approaches like quantization [7], pruning [16], and low-rank approximations [1]. However, the deployment of ViTs in practical applications, especially on hardware platforms such as Field-Programmable Gate Arrays (FPGAs), presents fundamental challenges. Among these, the non-linear layers—layer normalization, softmax, and GELU—integral to the architecture of ViTs, stand out. While crucial for the network's ability to model complex patterns, these functions are computationally intensive and thus present a critical challenge for the efficient implementation on FPGAs.

Our research delivers two key contributions. Firstly, we introduce PEANO-ViT, a novel approach that utilizes hardware-optimized approximation techniques for the non-linear functions within ViTs. Our approach in PEANO-ViT offers a comprehensive solution to the challenges posed by implementing key functions in ViTs on

FPGA platforms. By leveraging innovative techniques such as the Padé-based approximation for the exponential function and incorporating bit manipulation operations for efficient division in the softmax layer, we strive for a well-balanced and resource-efficient implementation that prioritizes performance and resource conservation. The layer normalization implementation effectively tackles computational challenges by approximating the reciprocal of the square root, $\frac{1}{\sqrt{x}}$, in a novel manner. Furthermore, our adoption of a piece-wise linear approximation for GELU not only minimizes resource usage but also closely preserves the original function's behavior. Secondly, we demonstrate through comprehensive experiments that PEANO-ViT enables the efficient execution of ViTs on FPGAs, with minimal impact on accuracy and significant improvements in computational efficiency and power consumption.

2 RELATED WORK

Transformers [13], originally developed for tackling long sequences in natural language processing tasks, served as the inspiration behind ViT [4] for computer vision applications. ViTs achieve impressive results by processing images as sequences of tokens and leveraging the power of self-attention. However, while crucial for performance, the core non-linear functions in ViTs – softmax, GELU, and layer normalization – are computationally expensive and hinder efficient hardware implementation. Several studies have explored hardware-efficient strategies for these layers, presenting various approximation techniques that balance approximation accuracy with computational cost. Their characteristic in comparison to PEANO-ViT is summarized in Table 1. The calculations for basic layer normalization, softmax, and GELU are depicted in equations (1-3), respectively. In equation 1, γ and β are learnable parameters while μ and σ represents the average and variance of input data of the layer normalization function.

$$\text{LayerNorm}(x_i) = \frac{x_i - \mu}{\sigma} * \gamma + \beta \tag{1}$$

$$\text{Softmax}(x_i) = \frac{e^{x_i}}{\sum e^{x_i}} \tag{2}$$

$$\text{GELU}(x) \approx 0.5x \left(1 + \tanh\left[\sqrt{\frac{2}{\pi}} \left(x + 0.044715x^3 \right) \right] \right) \tag{3}$$

2.1 Softmax Implementations

The implementation of the softmax layer has emerged as a focal point of research, with numerous studies dedicated to optimizing its efficiency through various approximation techniques. The main challenges for an efficient implementation of softmax on hardware platforms arise from the non-linear function of e^x and the final division operation for normalizing the output values. Previous research efforts, such as those by [11] targeted the efficient calculation of exponential function and but were hindered by the costly division operation. In contrast, studies by [5], [14], and [6] adopted bit manipulation techniques to simplify the exponential function approximation and eliminate the need for division. Although these methods are beneficial for reducing computational demands and are well-suited for hardware implementation, they still have a high

computational complexity due to their inherently iterative nature, causing increased inference latency.

2.2 Layer Normalization Implementations

For hardware implementation of layer normalization, significant hurdles include the efficient approximation of the square root function and managing division operations. the approach introduced in [14] tackles the division operation issue but continues to employ the exact yet resource-intensive formula of square root, resulting in lower throughput.

2.3 GELU Implementations

Beyond layer normalization and softmax, the GELU function's approximation also poses a significant challenge in the hardware deployment of ViTs. This is due to its intricate non-linear nature, which necessitates the execution of the $tanh(x)$ function alongside polynomial calculations. Authors of [6] have explored the approximation of the GELU function by simplifying the non-linear 2^x function using bit manipulation operations. Additionally, [9] has presented an innovative method that leverages existing softmax hardware to facilitate GELU computations. While these approaches are designed to be hardware-efficient and minimize resource consumption, the computational latency remains a concern. This is due to the iterative nature of some of the bit manipulation operations in [6], and the use of non-optimized hardware for GELU in [9].

3 METHODOLOGY

In this section, we describe the techniques utilized to approximate the layer normalization, softmax, and GELU functions. Our emphasis was on developing methods that avoid divisions and ensure compatibility with hardware implementations while also aiming to preserve the accuracy of the model as much as possible.

3.1 Layer Normalization

As described in subsection 2.2, the main challenges of implementing layer normalization on hardware platforms such as FPGAs are the non-linear square root function and the costly division operation. Inspired by SOLE [14], we propose a method to directly approximate $\frac{1}{\sqrt{X}}$. We start with the following identities:

$$\frac{1}{\sqrt{X}} = 2^{\log_2 \frac{1}{\sqrt{X}}}, \quad \log_2 \frac{1}{\sqrt{X}} = \frac{-1}{2} \log_2 X \tag{4}$$

Based on [14], we use equations (5-6) to approximate $\log_2 X$, in which k_x is the leading '1' bit of X and $x \in [0, 1)$:

$$X = \sum_{i=0}^{n-1} 2^i b_i = 2^{k_x} + \sum_{i=0}^{k_x-1} 2^i b_i = 2^{k_x}(1 + x) \tag{5}$$

$$\log_2 X \approx k_x + x \tag{6}$$

Therefore, we can have the following approximation:

$$\frac{1}{\sqrt{X}} \approx 2^{\frac{-(k_x+x)}{2}} \tag{7}$$

Calculating the $2^{\frac{-(k_x+x)}{2}}$ term is the only step remaining. We note that $2^\alpha = 2^u * 2^v$ in which u is an integer number and $v \in [0, 1)$.

Table 1: Comparison of state-of-the-art methods for implementation of non-linear layers

Approach	Layer normalization approximation	Softmax approximation	GELU approximation	All division-free approximations	Accuracy and resource aware flexible approximations
Softermax [11]	✗	✓	✗	✗	✗
Koca et al.[5]	✗	✓	✗	✓	✗
Peltekis et al.[9]	✗	✓	✓	✓	✗
SOLE [14]	✓	✓	✗	✗	✓
Li et al.[6]	✗	✓	✓	✓	✗
LTrans-OPU [2]	✓	✓	✓	✓	✗
PEANO-ViT (Ours)	✓	✓	✓	✓	✓

Algorithm 1 PEANO Layer Normalization

Input: $x_1, \ldots, x_n, \gamma, \beta, fracPow2[2^m] = \{2^{(0.0\ldots0)_2}, \ldots, 2^{(0.1\ldots1)_2}\}$
Output: $y_1, \ldots, y_n$
1: $Avg = \frac{1}{n} \sum_{i=1}^{n} x_i$ //average of inputs
2: $AvgSQ = \frac{1}{n} \sum_{i=1}^{n} x_i^2$ //average of inputs squared
3: $Var = AvgSQ - Avg^2$ //variance of inputs
4: $k_{Var} = LeadingOne(Var)$ //leading '1' bit of variance
5: $x_{Var} = Var[k_{Var} - 1 : 0]$ //contains the bits after k_{Var}
6: $log_2Approx = -(k_{Var} + x_{Var}) >> 1$
7: $u = \lfloor log_2Approx \rfloor$
8: $v = u - log_2Approx$
9: $\tilde{v} = fracBits(v, m)$ //$\tilde{v}$ keeps the top m fractional bits in v
10: $recipSqrt = fracPow2[\tilde{v}] << u$ //approximation of $\frac{1}{\sqrt{Var}}$
11: **for** $i = 1$ **to** n **do**
12:     $y_i = (x_i - Avg) * recipSqrt * \gamma + \beta$
13: **end for**
14: **return** $y_1, y_2, \ldots, y_n$

To avoid calculating 2^v, we keep the top m bits of v's binary representation as $v \approx \tilde{v} = (0.v_{-1} \ldots v_{-m})_2$ and pre-store $2^{(0.0\ldots0)_2}$ up to $2^{(0.1\ldots1)_2}$. Since u is an integer number, 2^u can be implemented using the shift operation. Thus, the approximation of $\frac{1}{\sqrt{X}}$ can be obtained from two equations below:

$$2^{\frac{-(k_x + x)}{2}} = 2^u \cdot 2^v \tag{8}$$

$$\frac{1}{\sqrt{X}} \approx 2^{\tilde{v}} << u \tag{9}$$

Figure 1b shows the $\frac{1}{\sqrt{X}}$ compared to our approximation and the overall layer normalization method is described in algorithm 1. Using these approximations, we have simultaneously tackled the two problems of efficient implementation of the square root function and approximating the division operation. It is important to highlight that m, an adjustable integer parameter, enables a trade-off between the precision of the approximation and the on-chip memory requirements for storing $2^{\tilde{v}}$. Increasing m improves the approximation accuracy at the cost of demanding more on-chip memory. This flexibility will be discussed in detail in Section 4.3.

3.2 Softmax

Our method for softmax approximation includes two steps. First, we introduce a Padé-based approximation for the exponential function.

In the second step, we eliminate the division operations by proposing a multi-scale reciprocal approximation (MSR-approx) method. The Padé approximation $Pade_{[m,n]}(x) = \frac{a_0 + a_1 x + \ldots + a_{m-1} x^m}{b_0 + b_1 x + \ldots + b_{n-1} x^n}$ of a function $f(x)$ is the ratio of 2 polynomial functions. It represents a better approximation of an arbitrary nonlinear function compared to pure polynomial approximations of the same degree. For approximating the e^x term, we have set $m = n = 2$ to get a Padé approximation as follows:

$$e^x \approx \frac{12 + 6x + x^2}{12 - 6x + x^2} \tag{10}$$

To compute the $Pade_{[2,2]}(x)$ approximation of e^x, we only need to compute $x^2 = x \cdot x$ and $6x = x << 2 + x << 1$ thanks to the numerator and denominator having similar functional forms. Figure 1a illustrates the Pade-based approximation of the function compared to e^x. As can be seen, the proposed approximation is very accurate for $x \in [-3, 2]$. This observation motivated us first to add 2 to all inputs (after subtracting the maximum value) and then set e^x to 0 for the values of less than -3 after the first step's calculations. Our final approximation of the exponential function is thus as follows:

$$PEANOexp(\tilde{x}) = \begin{cases} 0 & \text{if } \tilde{x} < -3 \\ \frac{12 + 6\tilde{x} + \tilde{x}^2}{12 - 6\tilde{x} + \tilde{x}^2} & \text{if } \tilde{x} \geq -3 \end{cases} \tag{11}$$

Where $\tilde{x} = x - max(x_i) + 2$. The above approximation adds another division operation to the main calculation of softmax. The first division is for the computation of $PEANOexp(x)$ while the second division is needed for the softmax's output normalization. Since $\tilde{x} \in [-3, 2]$, values of the $PEANOexp(x)$'s denominator lie in the interval of $[4, 39]$. This motivated us to pre-store some of $\frac{1}{x}$ values and subsequently use them to approximate the reciprocal function. However, unlike the denominator of $PEANOexp(x)$, the denominator of the second division has a huge range of values. Therefore, pre-storing values to approximate the second division is not feasible (unless a very large lookup table is used, which would result in high memory usage.)

To solve the aforesaid problem, we propose a multi-scale reciprocal approximation (MSR-approx) scheme for both division operations in the softmax. First we replace X (the denominator) with $\tilde{X}$ using the equation below:

$$\tilde{X} = Scale \cdot \lfloor \frac{X}{Scale} \rfloor \tag{12}$$

And the reciprocal function approximation is described as,

$$\frac{1}{X} \approx \frac{1}{\tilde{X}} = \frac{1}{Scale} \cdot \frac{1}{\lfloor \frac{X}{Scale} \rfloor} \qquad (13)$$

Next, we force $Scale = 2^\alpha$ to be an integer power of 2 so that $\frac{1}{Scale}$ can be implemented by using a right shift by α. This constraint also helps with the calculation of $\lfloor \frac{X}{Scale} \rfloor$ since it simply means dropping out the α right bits of X. The only thing we need to do is to pre-store $\lfloor \frac{X}{Scale} \rfloor$ values, which is still problematic due to the fact that the range of X can be extremely wide for the second division operation. This arises from the assumption of fixed α for all X values while using a dynamic value of α will solve the problem of X's large variable range as described in algorithm 2.

Algorithm 2 shows the multi-scale approximation of the reciprocal function, which uses an adjustable integer threshold α^* and pre-stored values of $\{\frac{1}{1}, \ldots, \frac{1}{2^{\alpha^*+1}-1}\}$. The MSR-approx maps all values of X into the interval of $[1, 2^{\alpha^*+1}-1]$ via defining a flexible Scale value, which solves the problem of the dynamic range of X. For instance, if $\alpha^* = 4$ then for $X \in [1, 31]$ then $\lfloor \frac{X}{Scale} \rfloor \in \{1, \ldots, 31\}$, and $\lfloor \frac{X}{Scale} \rfloor \in \{16, \ldots, 31\}$ for the other values of X. Hence, we only need to pre-store $\{\frac{1}{1}, \ldots, \frac{1}{31}\}$. Figure 1c illustrates our MSR-approx method compared to original reciprocal function for $\alpha^* = 4$. Choosing α^* is a trade-off between the accuracy of MSR-approx and the memory required for pre-storing values (see Section 4.3). Larger α^* proposes a more accurate approximation of reciprocal function while requiring larger memory for pre-stored values. The softmax using the MSR-approx scheme is presented in algorithm 3.

An alternative approach for improving the accuracy of the multi-scale division is to use linear interpolation between pre-stored points (instead of directly using any of these points.) For instance, if $X = 59$ and $\alpha^* = 4$, the scale is equal to 2, so in the basic MSR-approx method, we approximate $\frac{1}{59}$ using $\frac{1}{\lfloor \frac{59}{2} \rfloor} = \frac{1}{29}$. Instead, we can do linear interpolation between $\frac{1}{29}$ and $\frac{1}{30}$ to have a more accurate approximation of $\frac{1}{59}$. The MSR approximation enhanced with linear interpolation (called LMSR-approx) attains superior accuracy at the expense of a slight increased resource consumption and computational cycles, illustrating a clear trade-off between accuracy and resource efficiency.

3.3 GELU

PEANO-ViT uses a piece-wise linear approach to approximate the Gaussian Error Linear Unit (GELU). Unlike ViT's other non-linear functions, such as the square root and exponential functions, GELU exhibits a predominantly linear behavior across both the lower and upper extremes of its domain. Additionally, the GELU activation function maintains a narrow range of values within its non-linear region. These characteristics motivate the adoption of a piece-wise linear approximation as a highly suitable method for replicating the functionality of the GELU function.

Our method employs six breakpoints for GELU computations, resulting in seven linear segments. The initial breakpoints are set at $x = -3$ and $x = 3$, chosen to emulate the GELU's linear behavior as x approaches $\pm\infty$. Importantly, like many established activation functions (e.g., ReLU, PReLU, GELU, SiLU), our approximation ensures that the activation function intersects the origin,
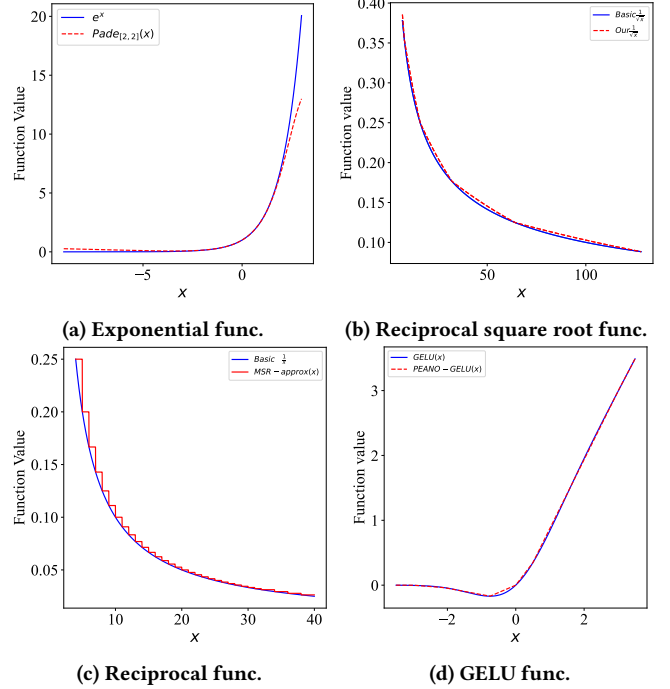

(a) Exponential func.   **(b) Reciprocal square root func.**

(c) Reciprocal func.   **(d) GELU func.**

Figure 1: Comparison of standard functions with our approximations.

introducing a third breakpoint at $x = 0$. To capture GELU's capability for generating negative outputs, a breakpoint at $x = -0.75$ approximates its minimum value, enhancing the fidelity of our approximation. To optimize the representation of GELU's transitional non-linear behavior within the intervals $[-3, -0.75]$ and $[0, 3]$, additional breakpoints at $x = -2.1$ and $x = 0.5$ are introduced. These points were determined through the minimization of the mean square error, ensuring a more accurate approximation in the specified ranges. With the mentioned breakpoints, figure 1d visualizes our final approximation which is described in the equation below:

$$PEANO-GELU(x) = \begin{cases} 0 & \text{if } x < -3 \\ -0.0414(x+3) & \text{if } -3 \le x < -2.1 \\ -0.0982(x+2.1) - 0.0373 & \text{if } -2.1 \le x < -0.75 \\ 0.2266(x+0.75) - 0.17 & \text{if } -0.75 \le x < 0 \\ 0.6914x & \text{if } 0 \le x < 0.5 \\ 1.0617(x-0.5) + 0.3457 & \text{if } 0.5 \le x < 3 \\ x & \text{if } x \ge 3 \end{cases}$$

3.4 FPGA Implementation

The overall FPGA implementation of PEANO-ViT's non-linear layers is illustrated in Figure 2. Notably, each non-linear function processes N elements concurrently, enabling an approximate $N - fold$ reduction in computation time. To enhance processing speed further, FIFO queues have been integrated between the reading, storing, and computing stages across all three implementations. Distinct from GELU, both layer normalization and softmax necessitate dual readings of input data—the initial for preliminary calculations and the subsequent for the normalization phase. Integrating an extra FIFO in parallel to the primary data stream notably decreases the latency for both the layer normalization and softmax modules by eliminating the requirement to temporarily store input values for a second calculation phase. Increasing the parameter N accelerates the processing of non-linear functions at the cost of more FPGA

Algorithm 2 Multi-Scale Reciprocal approximation (MSR-approx)

Input: $x, \alpha^*, StoredRecip[2^{\alpha^*+1} - 1] = \{\frac{1}{1}, \ldots, \frac{1}{2^{\alpha^*+1}-1}\}$
Output: y //approximation of $\frac{1}{x}$
  1:   $logInterval = LeadingOne(x)$
  2:   **if** $logInterval \leq \alpha^*$ **then**
  3:      $\alpha = 0$
  4:   **else**
  5:      $\alpha = logInterval - \alpha^*$
  6:   **end if**
  7:   $Scale = 2^{\alpha}$
  8:   $y = (StoredRecip[\lfloor x >> \alpha \rfloor]) >> \alpha$
  9:   **return** y

Algorithm 3 PEANO Softmax

Input: $x_1, \ldots, x_n$
Output: $y_1, \ldots, y_n$
  1:   $MaxInput = max(x_i)$ //Maximum of inputs
  2:   $\tilde{x}_i = x_i - MaxInput + 2$ //Shifting inputs by 2 - MaxInput
  3:   **for** $i = 1$ **to** n **do**
  4:      **if** $\tilde{x}_i < -3$ **then**
  5:          $PEANOexp_i = 0$
  6:      **else**
  7:          $PEANOexp_i = (12 + 6\tilde{x}_i + \tilde{x}_i^2)$
  8:               $\times$ MSR-approx$(12 - 6\tilde{x}_i + 6\tilde{x}_i^2)$
  9:      **end if**
10:   **end for**
11:   $Sum = \sum_{i=1}^{n} PEANOexp_i$ //Summation of exponential terms
12:   **for** $i = 1$ **to** n **do**
13:      $y_i = PEANOexp_i \times$MSR-approx(Sum)
14:   **end for**
15:   **return** $y_1, \ldots, y_n$

Table 2: Accuracy Loss of approximations on ImageNet-1K benchmark. The results of [14] and [6], if available, are directly sourced from the papers. FP32 and FiP16 stand for 32-bit floating-point and 16-bit fixed-point, respectively.

Model	Approach	Approximations	Accuracy
DeiT-S	Baseline(FP32)	-	79.85%
	SOLE [14](FP32)	Layer normalization + softmax	79.27%
	PEANO-ViT(Ours)(FP32)	Layer normalization + softmax	**79.36%**
	PEANO-ViT(Ours)(FiP16)	All non-linearities	**79.13%**
DeiT-B	Baseline(FP32)	-	81.85%
	SOLE [14](FP32)	Layer normalization + softmax	81.60%
	PEANO-ViT(Ours)(FP32)	Layer normalization + softmax	81.55%
	PEANO-ViT(Ours)(FiP16)	All non-linearities	81.35%
	PEANO-ViT(Ours) W LMSR-approx(FiP16)	All non-linearities	**81.65%**
Swin-B	Baseline(FP32)	-	83.60%
	SOLE [14](FP32)	Layer normalization + softmax	83.05%
	PEANO-ViT(Ours)(FP32)	Layer normalization + softmax	**83.60%**
	PEANO-ViT(Ours)(FiP16)	All non-linearities	**83.56%**
ViT-L	Baseline(FP32)	-	85.15%
	Li et al.[6](FiP16)	Softmax + GELU	84.78%
	PEANO-ViT(Ours)(FiP16)	Softmax + GELU	**85.03%**
	PEANO-ViT(Ours)(FiP16)	All non-linearities	**84.83%**

resource consumption. Consequently, PEANO-ViT becomes a configurable hardware framework alongside its software flexibilities.

4 RESULTS AND DISCUSSIONS

In this study, the PEANO-ViT model was implemented on a Xilinx UltraScale+ VU9P board running at a frequency of 250 MHz. We utilized the Vivado power report from Xilinx to evaluate the power consumption of each design. To evaluate the performance of PEANO-ViT, we employed the publicly available ImageNet-1K dataset [3] and three different model architectures, namely ViT [4], DeiT [12] and Swin [8], across various sizes (small, base, and large). It is important to point out that our experimental setup does not require extensive retraining. Instead, we conducted only two epochs of fine-tuning after integrating each approximation into the model. We utilized pre-trained models from the TIMM library [15] as our starting point and implemented our approximations using PyTorch.

4.1 ImageNet Classification

Table 2 provides a comparison of accuracy losses for four ViT-based models utilizing the PEANO-ViT approximations against techniques proposed by [14] and [6] implemented on FPGA and GPU platforms, respectively. In our analysis, we set the layer normalization parameter $m = 4$ and the MSR-approximation parameter $\alpha^* = 4$ without any linear interpolation. The superior performance of PEANO-ViT compared to [6] and [14] stems from its independent approximations of the softmax, GELU, and layer normalization functions, while [6] focuses solely on softmax and GELU, and [14] on layer normalization and softmax. The results of Table 2 indicate that PEANO-ViT exhibits minimal accuracy degradation when applying approximations to all non-linear blocks. Furthermore, when using a similar approximation approach, PEANO-ViT achieves lower accuracy reduction across DeiT-S, Swin-B, and ViT-L models compared to the methods outlined in [6] and SOLE [14]. For the DeiT-B model, PEANO-ViT shows reduced accuracy degradation compared to SOLE [14] when switching from MSR-approximation to LMSR-approximation. Notably, PEANO-ViT offers the ability to further minimize accuracy loss by adjusting m and α^* and by incorporating linear interpolation in the MSR approximation (LMSR-approx).

4.2 Hardware Cost

Table 3 details the power efficiency gain and reduction in resource usage achieved by implementing PEANO-ViT. By utilizing the rapid and hardware-compatible approximations introduced by PEANO-ViT, the significant power consumption and resource usage associated with hardware-intensive and costly iterative methods for exact non-linear implementation have been greatly diminished. Furthermore, Table 3 provides the resource utilization breakdown for each non-linear layer of PEANO-ViT. In processing layers such as normalization, softmax, and GELU, we simultaneously handle 16 elements, resulting in a Level of Parallelism (LoP) of 16 to enable a fair comparison with LTrans-OPU. This LoP can be adjusted to align with resource availability and latency objectives, making PEANO-ViT a versatile framework for enhancing the speed of machine learning tasks. Increasing the LoP enhances processing speed but may lead to higher resource consumption and power usage.
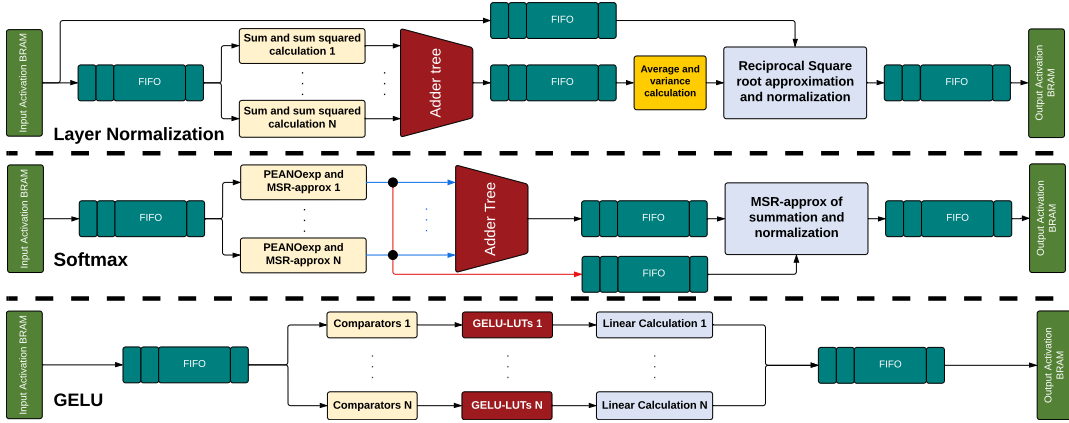
Figure 2: Overall FPGA implementation of PEANO-ViT

Table 3: Hardware metrics for DeiT-B Implementation

Non-linear layer	Approach	DSP	DSP (Reduction)	LUT	LUT (Reduction)	Register	Register (Reduction)	Power efficiency
Layer normalization	Standard layer normalization	51	-	24609	-	29831	-	1×
	LTrans-OPU [2]	0	**100%**	60902	-147.4%	7850	**73.6%**	0.99×
	PEANO layer normalization (Ours)	52	-1.9%	8157	**66.8%**	8621	71.1%	**1.91×**
Softmax	Standard softmax	64	-	9745	-	10648	-	1×
	LTrans-OPU [2]	0	**100%**	238569	-2348.1%	13837	-29.9%	0.19×
	PEANO softmax W MSR-approx (Ours)	48	25%	5595	**42.5%**	3831	**64%**	**1.39×**
	PEANO softmax W LMSR-approx (Ours)	49	23.4%	5741	41.1%	3876	63.6%	1.38×
GELU	Standard GELU	128	-	101267	-	88293	-	1×
	LTrans-OPU [2]	0	**100%**	11314	88.8%	2499	**97.1%**	6.76×
	PEANO GELU (Ours)	16	87.5%	2940	**97.1%**	2951	96.6%	**8.01×**

Table 4: Effect of PEANO-ViT parameters on approximations accuracy

Fuction	Test input interval	Changed parameter	MSE
Reciprocal square root	$[1, 128]$	$m = 3$	4.93×10^{-5}
		$m = 4$	9.56×10^{-6}
		$m = 5$	7.86×10^{-6}
Reciprocal	$[8, 64]$	$\alpha^* = 4, MSR$	4.19×10^{-6}
		$\alpha^* = 5, MSR$	4.03×10^{-6}
		$\alpha^* = 4, LMSR$	3.63×10^{-9}
		$\alpha^* = 5, LMSR$	3.58×10^{-9}
GELU	$[-4, 4]$	7 segments	2.65×10^{-4}
		10 segments	8.31×10^{-5}

4.3 Flexibility of PEANO-ViT

PEANO-ViT is a highly versatile framework that can be tailored to meet specific accuracy goals, hardware resource limitations, and power consumption requirements. This adaptability is achieved through the adjustment of key parameters such as m for layer normalization, α^* for softmax, and the selection between MSR or LMSR approximations for softmax. Furthermore, the framework offers flexibility in determining the number of linear segments for approximating the GELU function. Table 4 illustrates the impact of different configurations on the mean square error accuracy of approximated functions. Increasing the values of m and α^*, expanding the number of linear segments in GELU, and choosing LMSR over MSR result in improved accuracy but also consume higher hardware resources, resulting in increased power consumption.

5 CONCLUSION

PEANO-ViT optimizes ViT models by approximating non-linear blocks and eliminating division operations, maintaining high accuracy with minimal reduction. This approach enhances power efficiency and resource savings, setting a new benchmark for sustainable deep learning. Its flexibility allows for customized adjustments in accuracy, hardware resources, and power consumption, ensuring it meets specific performance requirements without sacrificing efficiency or accuracy.

REFERENCES

[1] Seyedarmin Azizi, Mahdi Nazemi, and Massoud Pedram. 2024. Memory-Efficient Vision Transformers: An Activation-Aware Mixed-Rank Compression Strategy. arXiv:2402.06004 [cs.CV]

[2] Yueyin Bai et al. 2023. LTrans-OPU: A Low-Latency FPGA-Based Overlay Processor for Transformer Networks. In *33rd International Conference on Field-Programmable Logic and Applications, FPL 2023*. IEEE, 283–287.

[3] Jia Deng et al. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.

[4] Alexey Dosovitskiy et al. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *9th International Conference on Learning Representations*.

[5] Nazim Altar Koca et al. 2023. Hardware-efficient Softmax Approximation for Self-Attention Networks. In *IEEE International Symposium on Circuits and Systems*.

[6] Tianyang Li et al. 2023. A high speed reconfigurable architecture for softmax and GELU in vision transformer. *Electronics Letters* 59, 5 (2023), e12751.

[7] Zhenhua Liu et al. 2021. Post-Training Quantization for Vision Transformer. In *Annual Conference on Neural Information Processing Systems 2021*.

[8] Ze Liu et al. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *2021 IEEE/CVF International Conference on Computer Vision*.

[9] Christodoulos Peltekis et al. 2024. Reusing Softmax Hardware Unit for GELU Computation in Transformers. (2024). arXiv:2402.10118

[10] Parsa Razmara, Tina Khezresmaeilzadeh, and B. Keith Jenkins. 2024. Fever Detection with Infrared Thermography: Enhancing Accuracy through Machine Learning Techniques. arXiv:2407.15302 [cs.LG] https://arxiv.org/abs/2407.15302

[11] Jacob R. Stevens et al. 2021. Softermax: Hardware/Software Co-Design of an Efficient Softmax for Transformers. In *58th ACM/IEEE Design Automation Conf.*

[12] Hugo Touvron et al. 2021. Training data-efficient image transformers & distillation through attention. In *Proceedings of the 38th Int. Conf. on Machine Learning.*

[13] Ashish Vaswani et al. 2017. Attention is All you Need. In *Annual Conference on Neural Information Processing Systems 2017.*

[14] Wenxun Wang et al. 2023. SOLE: Hardware-Software Co-design of Softmax and LayerNorm for Efficient Transformer Inference. In *IEEE/ACM International Conference on Computer Aided Design.*

[15] Ross Wightman. 2019. PyTorch Image Models. https://github.com/rwightman/pytorch-image-models. https://doi.org/10.5281/zenodo.4414861

[16] Fang Yu et al. 2022. Width & Depth Pruning for Vision Transformers. In *Thirty-Sixth AAAI Conference on Artificial Intelligence.*