NeuroMAX: A High Throughput, Multi-Threaded, Log-Based Accelerator for Convolutional Neural Networks

Mahmood Azhar Qureshi
Kansas State University
Manhattan, KS, US
mahmood102@ksu.edu

Arslan Munir
Kansas State University
Manhattan, KS, US
amunir@ksu.edu

ABSTRACT

Convolutional neural networks (CNNs) require high throughput hardware accelerators for real time applications owing to their huge computational cost. Most traditional CNN accelerators rely on single core, linear processing elements (PEs) in conjunction with 1D dataflows for accelerating convolution operations. This limits the maximum achievable ratio of peak throughput per PE count to unity. Most of the past works optimize their dataflows to attain close to a 100% hardware utilization to reach this ratio. In this paper, we introduce a high throughput, multi-threaded, log-based PE core. The designed core provides a 200% increase in peak throughput per PE count while only incurring a 6% increase in area overhead compared to a single, linear multiplier PE core with same output bit precision. We also present a 2D weight broadcast dataflow which exploits the multi-threaded nature of the PE cores to achieve a high hardware utilization per layer for various CNNs. The entire architecture, which we refer to as NeuroMAX, is implemented on Xilinx Zynq 7020 SoC at 200 MHz processing clock. Detailed analysis is performed on throughput, hardware utilization, area and power breakdown, and latency to show performance improvement compared to previous FPGA and ASIC designs.

KEYWORDS

Convolutional neural networks (CNNs), hardware accelerator, multi-threaded, throughput, hardware utilization

ACM Reference Format:

Mahmood Azhar Qureshi and Arslan Munir. 2020. NeuroMAX: A High Throughput, Multi-Threaded, Log-Based Accelerator for Convolutional Neural Networks. In *Proceedings of IEEE/ACM International Conference on Computer-Aided Design (ICCAD '20)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3400302.3415638>

1 INTRODUCTION

Convolutional neural networks (CNNs) enable embedding of AI into devices for vision-based applications with an unprecedented accuracy. The early proposed high accuracy CNNs [1–3] required tens of millions of parameters and computations for one inference pass. This computational complexity along with high memory requirements greatly hampered their deployment on low energy, resource constrained devices. In addition to this, many CNN architectures used varying kernel sizes which results in reconfigurability requirement as well as low hardware utilization in accelerator designs.

Separable convolution for CNNs was introduced the first time in mobilenets [4, 5] to reduce the number of multiply and accumulates (MACs). In addition, many modern CNNs use kernels of size 3×3 to promote ease of accelerator design with high hardware utilization and throughput.

Design of an efficient dataflow for scheduling data into the accelerator is equally important. An inefficient dataflow results in reduced hardware utilization which causes a decrease in throughput. Dataflow should also promote the reusability of data since, in most cases, the same kernels are being applied on the entire input feature map. It has been shown previously that the movement of data to/from DDR memory is $200\times$ more costly in terms of energy consumption than a standard MAC operation [6]. Thus, the dataflow design should not only optimize the throughput and area, but also the data movement, in order to ensure reduced energy expenditure. Log-based accelerators have recently gained quite a lot of traction because of their simpler structure as compared to traditional accelerators with linear processing elements (PEs). Each PE in traditional CNN accelerator cores is essentially responsible for one multiplication in convolution operation. Log PEs replace the bulky multiplier cores with low cost barrel shifters without incurring a significant loss in accuracy. We clarify that cost here primarily refers to the area cost, which is determined by the number of LookUp Tables (LUTs) for field-programmable gate arrays (FPGAs) and gate count for application-specific integrated circuits (ASICs). This area cost is important because there are limited resources on-chip and thus this area cost also translates to monetary cost of system-on-chip (SoC). Many past approaches have designed log-based PE elements but have not exploited the low cost and overhead of such PEs. They instead rely on already established spatial architectures and 1D dataflows used for linear PEs. Our proposed **NeuroMAX** accelerator core comprises of 108 PEs arranged in a $6 \times 3 \times 6$, 3D spatial grid. The presented accelerator optimizes the most commonly used 3×3 and 1×1 kernel sizes to achieve high throughput and utilization. It can also be used for larger kernel sizes because of its grid structure and configurable 2D dataflow.

Our main contributions are as follows:

- We design a multi-threaded, low cost, log-based PE core. Using this core, we generate a spatial grid of 108 PEs, capable of performing a wide variety of convolution operations with high hardware utilization.
- We develop a 2D dataflow which exploits the thread based PE design to maximize the throughput and enhance data reuse to minimize the DDR memory access.
- We implement the entire NeuroMAX architecture in an FPGA and show improved performance in terms of area, throughput, hardware utilization, latency and power efficiency compared to past approaches.

This work is licensed under a Creative Commons Attribution International 4.0 License.

ICCAD '20, November 2–5, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8026-3/20/11.

<https://doi.org/10.1145/3400302.3415638>

Figure 1: Linear vs. Log Quantization (a) 1.5 bits linear VGG16 net (b) 5.0 bits log VGG16 (c) 5.1 bits log VGG16 (d) 1.5 bits linear SqueezeNet (e) 5.0 bits log SqueezeNet (f) 5.1 bits log SqueezeNet

2 RELATED WORK

Many hardware accelerators have been proposed recently and in the past. [7] proposed a non-systolic array, reconfigurable spatial architecture along with a new dataflow scheme called *row stationary* to maximize the data reuse. However, this design incurs high PE cost owing to local storage and control in PE. It also has low hardware utilization which results in low throughput per PE. [8] proposes an FPGA-based CNN accelerator with integrated depth-wise separable mode of operation. This accelerator, however, has low throughput because of the usage of 32-bit floating point format. [9] proposes an FPGA-based CNN accelerator having a dedicated matrix multiplication engine (MME) on Arria 10 SoC. It achieves a frame rate of 266 fps, however, its MME engine has a huge digital signal processing (DSP) cost of 1200+ DSP blocks. [10] is the improved version of [7] with higher hardware utilization and throughput. [11] introduced the concept of logarithmic data representation for neural network accelerator designs. It also gives accuracy comparison between linear and log quantization. [12] proposes an accelerator design using arbitrary log base. It, however, does not utilize the low hardware overhead of the log-based PE and instead rely on linear PE arrangements. [13] proposes a reconfigurable design for various convolution kernels. It uses a propagated input data flow scheme but incurs high latency and low hardware utilization. [14] proposes a rescheduled dataflow for convolution to optimize the energy efficiency. [15] proposes a vectorwise accelerator architecture with the goal of maximizing the hardware utilization. It supports various kernel sizes from 1×1 to 5×5 .

Although some of the recent designs achieve high hardware utilization, they are not able to increase the peak throughput per PE count beyond unity owing to the use of single core, linear PEs with high area cost. This paper overcomes the limitations of prior works by leveraging *log* PEs with multiple low cost threads within each log PE, and designing a 2D dataflow which promises high throughput by exploiting multi-level parallelism.

3 LOG MAPPING

Log mapping or log quantization maps an input value x to a logarithmically quantized value x' . Many trained neural nets have weights w and input activations a which are non-uniformly distributed. Mapping these 32-bit floating point (fp32), non-uniformly distributed values over fixed point, linearly quantized values introduces significant amount of quantization noise for small bit width.

Most hardware platforms use fixed point arithmetic for data manipulation where the fixed point number is represented in signed $Qm.n$ format. Here, m represents the integer part whereas n represents the fractional part. The range of values which can be represented are $range_{lin} = [-2^{m-1}, 2^{m-1} - \epsilon]$ where, $\epsilon = 2^{-n}$, is the step size.

A linear quantizer rounds the fp32 value to the nearest multiple of ϵ and then clips it as follows:

$$x_q = clip \left[\left(round \left(\frac{x}{\epsilon} \right) \right) \cdot \epsilon, -2^{m-1}, 2^{m-1} - \epsilon \right] \quad (1)$$

where,

$$clip(x, min, max) = \begin{cases} max, & x \geq max \\ x, & min < x < max \\ min, & otherwise \end{cases} \quad (2)$$

A log quantizer takes as input, x and the quantization parameters $< m, n, b >$, where b is the logarithmic base, and produces a log quantized value x' as output. The quantization process can be written as:

$$x' = clip \left[(round(log_b(|x|))) \cdot \epsilon, -2^{m-1}, 2^{m-1} - \epsilon \right] \quad (3)$$

$$x_q = \begin{cases} 0, & x = 0 \\ sign(x) \cdot b^{x'}, & otherwise \end{cases} \quad (4)$$

Figure 1 shows some of the quantization results for the first five convolution layers of VGG16 [2] and SqueezeNet [16]. Instead of using log base-2 for quantization, we use log base- $\sqrt{2}$ for more accurate mapping as shown in Figure 1(c) and (f). Infact, we observe that VGG16, pretrained on ImageNet dataset, with fp32 data, after base- $\sqrt{2}$ quantization, has top-1 accuracy decrement by only $\approx 3.5\%$ from 67.5% to 63.8%. This is opposed to log base-2 quantization which decreases the accuracy by $\approx 10\%$. This observation has also been verified in [11].

4 HARDWARE ARCHITECTURE

4.1 Top-Level

Figure 2 shows the top level hardware architecture of the proposed NeuroMAX CNN accelerator on Zynq-7020 SoC. The CONV core is the accelerator module containing a memory block, a state controller, PE grid, adder stages and post processing module. The memory block contains the weight, input and output SRAMs with a total cumulative size of 3.8Mb. The PE grid consists of 108 PEs arranged in $6 \times 3 \times 6$ 3D array. Figure 2 also shows the internal structure of the PE grid containing PE matrices numbered from 0 to 5.

Figure 2: NeuroMAX System Architecture

The PE matrices are all connected to their respective input, weight and output SRAM blocks. Each PE matrix processes independent channels in parallel for standard and separable convolutions for maximizing the throughput. The outputs from the PE matrices are provided to their respective adder nets within the PE grid. A total of six adder net 0s are present corresponding to six PE matrices. The configuration of these adder nets remain constant regardless of the type of convolution used or the filter size. The output from the adder net 0 is provided to six configurable two-stage adders whose input connections change based on the filter size and the convolution type. The first adder stage is referred to as adder net 1 and the second stage is the channel accumulation stage.

To perform a convolution operation, a tile of log quantized input fmap and weight data is loaded from the off-chip DDR memory into the SRAMs in the CONV core by AXI DMA and interconnect. The processor also sends the parameter information containing the values for filter size, input width, input height, output width, output height and total channels to the state controller inside the CONV core. The state controller modifies the configurable adders and determines the dataflow to be used for convolution operation. The linear convolution outputs are sent to the post processing block which performs ReLU operation and quantizes the results back into log values using pre-computed log table. These output log values are loaded into the output SRAMs and sent back to the off-chip DDR memory to be used for processing the next layer. No intermediate outputs or partial sums are stored in the DDR memory and all the intermediate processing is done within the CONV core to minimize the off-chip traffic.

4.2 PE Matrix

Figure 3 shows the hierarchical design of a single PE matrix (PE matrix 0) from left to right in a bottom up view. Each PE receives a 1D vector of weight values and one input value. It should be noted that both the input ($i0'$) and the weight values ($w0_{0-2}'$) are log quantized. The output vectors from PEs are provided to the adder net 0 which generates 18 psums (o1-o18). This adder net works by summing the same color coded values generated within a row of PEs as shown in Figure 4.

Figure 3(b) shows the internal structure of a single PE element ($PE_{0_0_0}$). There are three compute cores or threads, each processing a single weight data and the input value, and in turn, produce three outputs (p_{11}, p_{12}, p_{13}). The lowest level of the PE matrix is a thread within an individual PE as shown in Figure 3(a). Basic log-based multiplication operation is performed in a single thread. Assuming we have two log quantized values, w_q' and a_q' , representing the original weight (w_q) and the activation input (a_q) respectively, the multiplication of these values in log domain can be carried out as:

$$w_q a_q = \text{sign}(w_q) \cdot 2^{g_q'} \quad (5)$$

where,

$$g_q' = w_q' + a_q' \quad (6)$$

[12] showed a method to implement the exponential in equation (5) in hardware by decomposing the exponent into its integer and fractional part as:

$$w_q a_q = \text{sign}(w_q) \cdot 2^{\text{INT}(g_q')} \cdot 2^{\text{FRAC}(g_q')} \quad (7)$$

The integer part $2^{\text{INT}(g_q')}$ can be implemented by a shift operation, whereas, the fractional part can be pre-computed and stored within the thread. The total number of fractional computations depends on the total number of fractional bits (n) used. In our case, we have $n = 1$ and thus store $2^n = 2$ values in the thread memory. Equation (7) can now be rewritten as:

$$w_q a_q = \text{sign}(w_q) \cdot (\text{LUT}(\text{FRAC}(g_q')) \gg -\text{INT}(g_q')) \quad (8)$$

The hardware implementation of equation (8) is shown in Figure 3(a). Since weights can have negative values, which is not accounted for in the log computations, we use an additional bit to represent the log weight data with the most significant bit $w_q'[6]$ representing the sign of the weight before quantization. This is not required for the input fmap values since most modern CNNs use ReLU activations which eliminate the negative outputs.

5 DATA FLOW AND PROCESSING

The main idea behind designing an efficient dataflow is to minimize the data movement to/from the costly off-chip DDR memory. One MAC operation typically requires three memory reads corresponding to weight, ifmap, psum and one memory write corresponding to the updated psum. A neural net like AlexNet, with 724M MACs, will need $\approx 3000M$ DDR memory accesses. Many efficient dataflows have been presented in literature to minimize this data movement. Some of these include *output stationary*, *weight stationary* and *row stationary* [17]. Since convolution operation requires the reuse of filter weights, input and psums in successive operations, the dataflows are designed to optimize the re-usability without accessing the DDR memory. We introduce a 2D weight broadcast dataflow for maximizing the re-usability of the weights, input and psums.

5.1 3 x 3 Convolution

Figure 5 shows a 3×3 convolution example. Here, a 12×6 input is convolved with a 3×3 filter to produce a 10×4 output for stride 1 and a 6×3 output for stride 2. A total of 108 bits, corresponding to the 6×3 input tile, are received from the AXI4 interconnect and stored in the input SRAM. This input tile is modified by the state controller and provided to the PE matrix in a row shifted pattern as shown in Figure 6(a) and (c) for stride 1 and stride 2, respectively. We also acquire a 2D weight array and broadcast it to the PE matrix as shown in Figure 6(b). Figure 7 shows the dataflow and the operation of the PE matrix for the first 6×3 input tile and the weight matrix at time stamp $t = 1$. The entire input tile and

Figure 3: (a) Compute Thread (b) Collection of Threads to Make a PE (d) 6x3 PE Matrix 0 and Adder Net 0

Row0: o1 = $p_{11} + p_{14} + p_{17}$, o2 = $p_{12} + p_{15} + p_{18}$, o3 = $p_{13} + p_{16} + p_{19}$
Row1: o4 = $p_{21} + p_{24} + p_{27}$, o5 = $p_{22} + p_{25} + p_{28}$, o6 = $p_{23} + p_{26} + p_{29}$
Row2: o7 = $p_{31} + p_{34} + p_{37}$, o8 = $p_{32} + p_{35} + p_{38}$, o9 = $p_{33} + p_{36} + p_{39}$
Row3: o10 = $p_{41} + p_{44} + p_{47}$, o11 = $p_{42} + p_{45} + p_{48}$, o12 = $p_{43} + p_{46} + p_{49}$
Row4: o13 = $p_{51} + p_{54} + p_{57}$, o14 = $p_{52} + p_{55} + p_{58}$, o15 = $p_{53} + p_{56} + p_{59}$
Row5: o16 = $p_{61} + p_{64} + p_{67}$, o17 = $p_{62} + p_{65} + p_{68}$, o18 = $p_{63} + p_{66} + p_{69}$

Figure 4: Adder Net 0 psum Generation

Figure 5: A Convolution Example with 12x6 Input, 3x3 Filter for Stride 1 and 2

the 2D weight array is loaded into the PE matrix simultaneously. Because of the multi-threaded structure of PEs, each PE performs three multiplication operations using three threads and the outputs are row-wise summed to generate the psums (o1-o18) using adder net 0 (Figure 4). The dataflow chart and the processing of the entire 12×6 input is shown in Figure 8. The output $a_0wa_{0,2}$, in Figure 8, represents the three outputs a_0wa_0 , a_0wa_1 , a_0wa_2 generated by three threads within a PE. The adder net 0 computes the partial

Figure 6: State Controller Operation (a) Input, Stride 1 (b) Filter Weights (c) Input, Stride 2

Figure 7: 2D Weight Broadcast Dataflow

sum outputs (o1-o18) the same way as shown in Figure 4, where $p_{11} = a_0wa_0$, $p_{14} = b_0wb_0$ and $p_{17} = c_0wc_0$ are the same colored (green) outputs along the row.

t=1	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂	t=2	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂	t=3	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂	t=4	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂
c ₀ b ₀ a ₀	a ₀ wa _{0,1,2}	b ₀ wb _{0,1,2}	c ₀ wc _{0,1,2}	d ₀ c ₀ b ₀	b ₀ wa _{0,1,2}	c ₀ wb _{0,1,2}	d ₀ wc _{0,1,2}	e ₀ d ₀ c ₀	c ₀ wa _{0,1,2}	d ₀ wb _{0,1,2}	e ₀ wc _{0,1,2}	f ₀ e ₀ d ₀	d ₀ wa _{0,1,2}	e ₀ wb _{0,1,2}	f ₀ wc _{0,1,2}
c ₁ b ₁ a ₁	a ₁ wa _{0,1,2}	b ₁ wb _{0,1,2}	c ₁ wc _{0,1,2}	d ₁ c ₁ b ₁	b ₁ wa _{0,1,2}	c ₁ wb _{0,1,2}	d ₁ wc _{0,1,2}	e ₁ d ₁ c ₁	c ₁ wa _{0,1,2}	d ₁ wb _{0,1,2}	e ₁ wc _{0,1,2}	f ₁ e ₁ d ₁	d ₁ wa _{0,1,2}	e ₁ wb _{0,1,2}	f ₁ wc _{0,1,2}
c ₂ b ₂ a ₂	a ₂ wa _{0,1,2}	b ₂ wb _{0,1,2}	c ₂ wc _{0,1,2}	d ₂ c ₂ b ₂	b ₂ wa _{0,1,2}	c ₂ wb _{0,1,2}	d ₂ wc _{0,1,2}	e ₂ d ₂ c ₂	c ₂ wa _{0,1,2}	d ₂ wb _{0,1,2}	e ₂ wc _{0,1,2}	f ₂ e ₂ d ₂	d ₂ wa _{0,1,2}	e ₂ wb _{0,1,2}	f ₂ wc _{0,1,2}
c ₃ b ₃ a ₃	a ₃ wa _{0,1,2}	b ₃ wb _{0,1,2}	c ₃ wc _{0,1,2}	d ₃ c ₃ b ₃	b ₃ wa _{0,1,2}	c ₃ wb _{0,1,2}	d ₃ wc _{0,1,2}	e ₃ d ₃ c ₃	c ₃ wa _{0,1,2}	d ₃ wb _{0,1,2}	e ₃ wc _{0,1,2}	f ₃ e ₃ d ₃	d ₃ wa _{0,1,2}	e ₃ wb _{0,1,2}	f ₃ wc _{0,1,2}
c ₄ b ₄ a ₄	a ₄ wa _{0,1,2}	b ₄ wb _{0,1,2}	c ₄ wc _{0,1,2}	d ₄ c ₄ b ₄	b ₄ wa _{0,1,2}	c ₄ wb _{0,1,2}	d ₄ wc _{0,1,2}	e ₄ d ₄ c ₄	c ₄ wa _{0,1,2}	d ₄ wb _{0,1,2}	e ₄ wc _{0,1,2}	f ₄ e ₄ d ₄	d ₄ wa _{0,1,2}	e ₄ wb _{0,1,2}	f ₄ wc _{0,1,2}
c ₅ b ₅ a ₅	a ₅ wa _{0,1,2}	b ₅ wb _{0,1,2}	c ₅ wc _{0,1,2}	d ₅ c ₅ b ₅	b ₅ wa _{0,1,2}	c ₅ wb _{0,1,2}	d ₅ wc _{0,1,2}	e ₅ d ₅ c ₅	c ₅ wa _{0,1,2}	d ₅ wb _{0,1,2}	e ₅ wc _{0,1,2}	f ₅ e ₅ d ₅	d ₅ wa _{0,1,2}	e ₅ wb _{0,1,2}	f ₅ wc _{0,1,2}

t=5	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂	t=6	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂	t=7	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂	t=8	wa ₀₋₂	wb ₀₋₂	wc ₀₋₂
c ₆ b ₆ a ₆	a ₆ wa _{0,1,2}	b ₆ wb _{0,1,2}	c ₆ wc _{0,1,2}	d ₆ c ₆ b ₆	b ₆ wa _{0,1,2}	c ₆ wb _{0,1,2}	d ₆ wc _{0,1,2}	e ₆ d ₆ c ₆	c ₆ wa _{0,1,2}	d ₆ wb _{0,1,2}	e ₆ wc _{0,1,2}	f ₆ e ₆ d ₆	d ₆ wa _{0,1,2}	e ₆ wb _{0,1,2}	f ₆ wc _{0,1,2}
c ₇ b ₇ a ₇	a ₇ wa _{0,1,2}	b ₇ wb _{0,1,2}	c ₇ wc _{0,1,2}	d ₇ c ₇ b ₇	b ₇ wa _{0,1,2}	c ₇ wb _{0,1,2}	d ₇ wc _{0,1,2}	e ₇ d ₇ c ₇	c ₇ wa _{0,1,2}	d ₇ wb _{0,1,2}	e ₇ wc _{0,1,2}	f ₇ e ₇ d ₇	d ₇ wa _{0,1,2}	e ₇ wb _{0,1,2}	f ₇ wc _{0,1,2}
c ₈ b ₈ a ₈	a ₈ wa _{0,1,2}	b ₈ wb _{0,1,2}	c ₈ wc _{0,1,2}	d ₈ c ₈ b ₈	b ₈ wa _{0,1,2}	c ₈ wb _{0,1,2}	d ₈ wc _{0,1,2}	e ₈ d ₈ c ₈	c ₈ wa _{0,1,2}	d ₈ wb _{0,1,2}	e ₈ wc _{0,1,2}	f ₈ e ₈ d ₈	d ₈ wa _{0,1,2}	e ₈ wb _{0,1,2}	f ₈ wc _{0,1,2}
c ₉ b ₉ a ₉	a ₉ wa _{0,1,2}	b ₉ wb _{0,1,2}	c ₉ wc _{0,1,2}	d ₉ c ₉ b ₉	b ₉ wa _{0,1,2}	c ₉ wb _{0,1,2}	d ₉ wc _{0,1,2}	e ₉ d ₉ c ₉	c ₉ wa _{0,1,2}	d ₉ wb _{0,1,2}	e ₉ wc _{0,1,2}	f ₉ e ₉ d ₉	d ₉ wa _{0,1,2}	e ₉ wb _{0,1,2}	f ₉ wc _{0,1,2}
c ₁₀ b ₁₀ a ₁₀	a ₁₀ wa _{0,1,2}	b ₁₀ wb _{0,1,2}	c ₁₀ wc _{0,1,2}	d ₁₀ c ₁₀ b ₁₀	b ₁₀ wa _{0,1,2}	c ₁₀ wb _{0,1,2}	d ₁₀ wc _{0,1,2}	e ₁₀ d ₁₀ c ₁₀	c ₁₀ wa _{0,1,2}	d ₁₀ wb _{0,1,2}	e ₁₀ wc _{0,1,2}	f ₁₀ e ₁₀ d ₁₀	d ₁₀ wa _{0,1,2}	e ₁₀ wb _{0,1,2}	f ₁₀ wc _{0,1,2}
c ₁₁ b ₁₁ a ₁₁	a ₁₁ wa _{0,1,2}	b ₁₁ wb _{0,1,2}	c ₁₁ wc _{0,1,2}	d ₁₁ c ₁₁ b ₁₁	b ₁₁ wa _{0,1,2}	c ₁₁ wb _{0,1,2}	d ₁₁ wc _{0,1,2}	e ₁₁ d ₁₁ c ₁₁	c ₁₁ wa _{0,1,2}	d ₁₁ wb _{0,1,2}	e ₁₁ wc _{0,1,2}	f ₁₁ e ₁₁ d ₁₁	d ₁₁ wa _{0,1,2}	e ₁₁ wb _{0,1,2}	f ₁₁ wc _{0,1,2}

Figure 8: Dataflow Chart for 3×3 Stride 1 Convolution in Figure 5

Figure 9: Adder Net 1 Configuration (a) Stride 1 (b) Stride 2

The dark red outputs (row 5 and 6) in Figure 5 for stride 1 and green outputs (row 3) for stride 2 represent the boundary outputs. The boundary condition occurs when the filter overlaps two different column-wise input tile sectors. For clarity, we assume that the first input tile at $t = 1$ is processed by the PE matrix. This corresponds to the first six rows and the first three columns of the input. The PE matrix will process the last row-wise input tile at $t = 4$ which corresponds to the first six rows and the last three columns of the input as shown in Figure 6(a). The input tile will then jump to the next column-wise 6x3 input tile which corresponds to the last six rows and the first three columns at $t = 5$ as shown in Figure 6(a). However, it can be seen that the row 5 and 6 in the output are dependent on the overlapping results from the two concurrent column-wise input tile sectors (e.g. at $t = 1$ and $t = 5$, $t = 2$ and $t = 6$ and so on). To resolve this, the three dependent psums (o13, o17 and o16), generated from row 5 and row 6, of first column-wise tile sector of the 12x6 input, are passed through a variable length shift

Figure 10: 1×1 Convolution Example

register with the maximum length equal to the width of the input. These psums are subsequently utilized when the next column-wise 6x3 input tile (a_6 to a_{11}) is being processed. Thus, the rows 1 to 4 in the output are generated during the time intervals $t = 1$ to $t = 4$, whereas the rows 5 to 10 are generated during the time interval $t = 5$ to $t = 8$.

The output in Figure 5, for stride 1, is generated by alternate colored, column-wise summation of the psums in the adder net 1 as shown in Figure 9(a). Figure 9(b) shows the output generation for stride 2 case. The shift registers (VAR Len SR) for generating the boundary outputs are also shown in Figure 9. It can be observed that because of the optimized dataflow, only 2 out of 18 or 11% psums require local storage as opposed to >50% psums requiring storage (local or off-chip) in previously proposed dataflows. The throughput for the above example is 45 OPS/cycle (total OPS/total cycles = $360/8 = 45$), which results in an 83.3% overall thread utilization ($45/(3 \times 6 \times 3) \times 100$). We will simply use thread utilization as hardware utilization in this context.

5.2 1 x 1 Convolution

1×1 convolutions are very popular in modern CNNs. These convolutions, along with the depth-wise separable, are replacing the normal 2-D convolutions because of the less number of MAC operations [4]. The 1×1 CONV operation convolves $1 \times 1 \times C \times P$ filters with a $M \times N \times C$ input to produce $M \times N \times P$ outputs. Here, C is the number of channels, P is the number of filters, M is the input width and N is the input height.

Figure 10 shows a 1×1 CONV example where a 3×6×6 input is convolved with 6, 1×1×6 filters to produce a 3×6×6 output. Since this convolution generates the psums by channel accumulation, the outputs from the multiple PE matrices are utilized. For the example

Figure 11: State Controller Load Operation

Figure 12: Dataflow Chart for 1×1 Convolution in Figure 10 in Figure 10, the state controller data scheduling for PE matrix 0 and 1 is shown in Figure 11. It can be seen that the first three channels of the input are convolved with the first three channels of all the filters in PE matrix 0, whereas, the last three channels of the input are convolved with the last three channels of all the filters in PE matrix 1. The time stamps during specific processing of input and weights in PE matrices are also shown in Figure 11. It should be noted that for an input with more channels, the rest of the PE matrices will also be used. Thus, by using the dataflow in Figure 11, the architecture can process 18 channels concurrently by using the 6 PE matrices, with each PE matrix processing 3 input and filter channels.

The dataflow chart for the PE matrix 0 for the example in Figure 10 is shown in Figure 12. The same dataflow chart can also be generated for the PE matrix 1. As mentioned earlier, the psums in 1×1 convolution are calculated using channel-wise accumulation. The eighteen outputs (o_1 - o_{18}) generated by the individual PE matrices are summed in their respective adder net 1s. The input connections for adder net 1 (AN 1_0) and the channel accumulator (CA 0) of PE

Figure 13: (a) Channel-wise accumulation for PE matrix 0 (b) All channel-wise accumulations

Figure 14: (a) 5×5 Convolution Example (b) Input Load Operation (c) Weight Load Operation

matrix 0 are shown in Figure 13(a). Here, o_{10} is the psum output from the PE matrix 0 and o_{15} is the psum output from the PE matrix 5. Since the example in Figure 10 is small, it only requires the first two PE matrices and their outputs, that is, only o_{10} - o_{18} and o_{15} - o_{18} are active. The output in Figure 10 is generated by using all six adder net 1s and the channel accumulators as shown in Figure 13(b). The throughput for the above example is 108 OPS/cycle (total OPS/total cycles = $(6 \times 6 \times 3 \times 6/6 = 108)$, which results in a 100% overall thread utilization $(108/(3 \times 6 \times 3 \times 2)) \times 100$.

5.3 Higher Order Convolutions

The proposed NeuroMAX accelerator is designed to optimize 3×3 and 1×1 convolutions. It can, however, also be used to accelerate larger kernel sizes. [18] proposed a kernel decomposition method such that an additional support for 4×4 and 5×5 filter is needed to implement any filter size. Figure 14(a) gives an example of 5×5 convolution. As the size of the PE matrix is 6×3 , a filter of width greater than 3 and height greater than 6 needs multiple cycles to calculate the output value. This can be seen in Figure 14(b) and (c) where the last two columns of the input matrix and the weight matrix are loaded at time stamp $t = 2$. Figure 15 shows the dataflow

t=1	wa_{0-4}	wb_{0-4}	wc_{0-4}	t=2	wd_{0-4}	we_{0-4}	0
$c_0b_0a_0$	a_0wa_{0-4}	b_0wb_{0-4}	c_0wc_{0-4}	$0e_0d_0$	d_0wd_{0-4}	e_0we_{0-4}	0
$c_1b_1a_1$	a_1wa_{1-4}	b_1wb_{1-4}	c_1wc_{1-4}	$0e_1d_1$	d_1wd_{1-4}	e_1we_{1-4}	0
$c_2b_2a_2$	a_2wa_{2-4}	b_2wb_{2-4}	c_2wc_{2-4}	$0e_2d_2$	d_2wd_{2-4}	e_2we_{2-4}	0
$c_3b_3a_3$	a_3wa_{3-4}	b_3wb_{3-4}	c_3wc_{3-4}	$0e_3d_3$	d_3wd_{3-4}	e_3we_{3-4}	0
$c_4b_4a_4$	a_4wa_{4-4}	b_4wb_{4-4}	c_4wc_{4-4}	$0e_4d_4$	d_4wd_{4-4}	e_4we_{4-4}	0
$c_5b_5a_5$	a_5wa_{5-4}	b_5wb_{5-4}	c_5wc_{5-4}	$0e_5d_5$	d_5wd_{5-4}	e_5we_{5-4}	0
000	0	0	0	000	0	0	0

Figure 15: Dataflow Chart for 5×5 Convolution in Figure 14(a)

Figure 16: Adder Configuration for 5×5 Convolution

chart which accounts for this configuration. The generated psums (o1-o18) are provided to the adder net 1 as shown in Figure 16. For this convolution, the output values are calculated as:

$$Va_0, Va_2 = ((o1+o5+o9)+(o10+o14))_{old} + (o1+o5+o9)_{new} \quad (9)$$

$$Va_1 = ((o4+o8+o12)+(o13+o17))_{old} + (o4+o8+o12)_{new} \quad (10)$$

In equations (9) and (10), the *old* value corresponds to the convolution output from the first three columns of the input and the weight matrix at $t = 1$, whereas, the *new* value corresponds to the last two columns at $t = 2$. The adder net 1 and the channel accumulator configuration for this convolution is shown in Figure 16. A similar configuration and dataflow chart is used for implementing a 4×4 convolution. In addition to this, the CONV core can also perform pooling operation by choosing the appropriate stride and kernel.

6 IMPLEMENTATION AND RESULTS

This section discusses the implementation of the proposed NeuroMAX accelerator and presents the area cost, power consumption, performance, throughput, and hardware utilization results. The accelerator has been implemented on the PL side of Xilinx Zynq-7020 SoC operating at 200MHz. Figure 17 shows cost comparison between our multi-threaded log PE core and an area optimized linear multiplier core with equal output bit precision and latency. It can be seen that by choosing a thread count of 3 (shown as log (3) in Figure 17), the LUT and FF cost is only 1.05 $\times$ and 1.14 $\times$ that of the linear PE. Thus, a total of 108 linear PEs would be equivalent, in cost, to ≈ 122 multi-threaded log PEs. For fairness, we will use the cost adjusted PE number for performance comparison.

Figure 17: Linear vs Log PE LUT and FF Cost (16-bit)

Table 1: Resource Utilization

Property	Accelerator	Utilization
#LUTs	20680	38%
#FFs	17207	16%
#36kB BRAMs	108	77%
Power	2.727 W	NA

Figure 18: Breakdown of (a) LUT cost (b) FF cost (c) Power Consumption for NeuroMAX

Table 1 shows the resource utilization of the implemented accelerator core as well as the total power consumption (static + dynamic). Figure 18 (a), (b), (c) shows the breakdown of LUT cost, FF cost and power consumption among different modules of the accelerator. The PE grid and the adder net 0 combined have the highest LUT and FF count (81% and 91%, respectively). The post processing block consumes negligible resources. The processing system (ARM core) dominates the power consumption (57%), while the PE grid and adder net 0 have the second highest consumption (26%) of the total.

Figure 19 shows layer by layer hardware utilization for various CNN architectures. We achieve an average utilization of 95%, 84% and, 86% for VGG-16, MobileNet v1 and, ResNet-34, respectively. The dip in hardware utilization in some layers of mobilenet and ResNet-34 is because of stride 2 convolutions which utilize only 50% of the available PE cores. The low utilization in the first layer of VGG16 is because it only has 3 channels and since each PE matrix processes one channel, the last 3 PE matrices remain idle which gives an exact utilization of 50%.

Chang et al. [15] recently presented an accelerator design with 1D broadcast dataflow which promises higher utilization and throughput (GOPS) than all the previous designs. We, therefore, compare our design against [15] in Figure 20 for various CNNs. [15] uses a total of 168 PE cores and provides a utilization of 99% with throughput 166.32 GOPS, 93.4% with throughput 156.91 GOPS and, 90.2%

Table 2: Comparison of NeuroMAX with Previous Designs

Property	NeuroMAX	[7]	[8]	[9]	[10]	[12]	[15]
Technology	Zynq-7020 SoC	65nm	Zynq-7100	Arria 10 SoC	65nm	Virtex-7	40nm
Precision(bits)	6-bit log	16-bit	32fp	16-bit	8-20 bits	5-bit log	16-bit
PE number	122(adjusted)	168	1926	1278	192	256	168
Processing clock (MHz)	200	200	100	133	200	Unreported	500
Peak Throughput (GOPS)	324	84	17.11	170.6	153.6	Unreported	168
Peak Throughput/PE	2.7(adjusted)	0.5	0.008	0.13	0.8	Unreported	1
Cost (LUTs(a),gates(b))	20.6k(a)	1176k(b)	142k(a)	66k(a)	2695k(b)	29k(a)	266k(b)
Power (W)	2.72	0.278	4.083	Unreported	0.460	3.756	0.155

Figure 19: Hardware Utilization of NeuroMAX for (a)VGG-16 (b) MobileNet v1 (c) ResNet-34

Figure 20: PE count vs Utilization vs Throughput Comparison of NeuroMAX with [15]

with throughput 151.54 GOPS for VGG16, ResNet-34 and mobilenet, respectively. We use 122 PE cores (cost adjusted), a 28% decrease from [15], and provide a throughput of 307.8 GOPS, an 85% increase, 281.8 GOPS, a 79.4% increase, and 268.92 GOPS, a 77.4% increase, for the three CNNs, respectively. This increase in throughput with lower PE count is attributed towards our low cost, multi-threaded PE core design and an efficient 2D dataflow. We also achieve somewhat similar hardware utilization, that is, 94% for VGG16, 87.3% for ResNet-34 and, 83% for mobilenet. It should be noted that [15] implements the accelerator on an ASIC, whereas, we use an FPGA, thus, an accurate comparison in LUT count, FF count and, power consumption cannot be made. It is, however, evident that the design in [15] when ported into FPGA will have $\approx 31\%$ more LUTs and FFs owing to more number of PEs used. Table 2 shows the comparison of our accelerator with previous state of the art ASIC and FPGA designs. We see an improved performance in terms of PE number, peak throughput and peak throughput/PE ratio. Only [15] has a peak throughput/PE ratio equal to unity with average around 0.85. Our peak throughput/PE is 3 with average around 2.7 after cost adjustment. The power comparison reveals that the FPGA-based designs inherently consume more power compared to ASICs. We can,

however, see that NeuroMAX consumes significantly less power and has lower cost in terms of LUT count compared to other FPGA designs. Table 3 gives a layer-by-layer processing latency comparison for VGG16. Both [7] and [15] benchmark the latency of their accelerators on this CNN, therefore, we also evaluate and compare NeuroMAX's performance on VGG16. It should be noted however that [15] uses 500MHz processing clock in their design. For fair comparison, we make suitable adjustments in their reported values. Our proposed NeuroMAX accelerator has 93% and, 47% decrease in latency, when compared to [7] and [15], respectively, at 200 MHz clock.

Table 3: VGG16 Latency Comparison

Layer	NeuroMAX	[7]	[15]
CONV1_1(ms)	1.35	38.0	2.57
CONV1_2(ms)	28.9	810.6	55.04
CONV2_1(ms)	14.4	405.3	27.43
CONV2_2(ms)	29.26	810.8	55.7
CONV3_1(ms)	14.54	204	27.7
CONV3_2(ms)	28.6	408.1	54.5
CONV3_3(ms)	28.7	408.1	54.6
CONV4_1(ms)	14.4	105.1	27.42
CONV4_2(ms)	29	210.0	55.23
CONV4_3(ms)	29.5	210.0	56.19
CONV5_1(ms)	7.24	48.3	13.79
CONV5_2(ms)	7.23	48.5	13.77
CONV5_3(ms)	7.11	48.5	13.54
Total(ms)	240.23	3755.3	457.5

7 CONCLUSION

This paper proposes NeuroMAX, a high throughput accelerator using multi-threaded, log-based PE cores. The designed PE cores are capable of providing a 200% increase in peak throughput while only increasing the area overhead by 6%, when compared to a standard multiplier-based PE core. We also design an efficient 2D weight broadcast dataflow scheme which exploits the multi-level parallelism of our processing engine and enables hardware utilization close to 100%. The accelerator is capable of performing a wide variety of convolutions including standard and separable 3×3 stride 1 and $2, 4 \times 4, 5 \times 5$ and 1×1 depthwise, required in modern CNN architectures. We have implemented NeuroMAX on Xilinx Zynq-7020 SoC and have evaluated various performance parameters. The design can provide at least a throughput increase of 77.4% and a latency decrease of 47% with a 28% decrease in PE count against recently proposed accelerator designs for modern CNNs. NeuroMAX also provides at least a 27% and a 29% decrease in power consumption and LUT count, respectively, against prior FPGA-based CNN accelerators.

REFERENCES

- [1] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'12, page 1097–1105, Red Hook, NY, USA, 2012. Curran Associates Inc.
- [2] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2014.
- [3] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [4] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017.
- [5] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks, 2018.
- [6] Mark Horowitz. 1.1 computing's energy problem (and what we can do about it). volume 57, pages 10–14, 02 2014.
- [7] Y. Chen, T. Krishna, J. S. Emer, and V. Sze. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE Journal of Solid-State Circuits*, 52(1):127–138, 2017.
- [8] Bing Liu, Danyin Zou, Lei Feng, Shou Feng, Ping Fu, and Junbao Li. An fpga-based cnn accelerator integrating depthwise separable convolution. *Electronics*, 8(3):281, 2019.
- [9] L. Bai, Y. Zhao, and X. Huang. A cnn accelerator on fpga using depthwise separable convolution. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 65(10):1415–1419, 2018.
- [10] Yu-Hsin Chen, Tien-Ju Yang, Joel Emer, and Vivienne Sze. Eyeriss v2: A flexible accelerator for emerging deep neural networks on mobile devices, 2018.
- [11] Daisuke Miyashita, Edward H. Lee, and Boris Murmann. Convolutional neural networks using logarithmic data representation, 2016.
- [12] Sebastian Vogel, Mengyu Liang, Andre Guntoro, Walter Stechele, and Gerd Ascheid. Efficient hardware acceleration of cnns using logarithmic data representation with arbitrary log-base. In *Proceedings of the International Conference on Computer-Aided Design, ICCAD '18*, New York, NY, USA, 2018. Association for Computing Machinery.
- [13] Y. Huan, J. Xu, L. Zheng, H. Tenhunen, and Z. Zou. A 3d tiled low power accelerator for convolutional neural network. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2018.
- [14] J. Jo, S. Kim, and I. Park. Energy-efficient convolution architecture based on rescheduled dataflow. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 65(12):4196–4207, 2018.
- [15] K. Chang and T. Chang. Vwa: Hardware efficient vectorwise accelerator for convolutional neural network. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 67(1):145–154, 2020.
- [16] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size, 2016.
- [17] V. Sze, Y. Chen, T. Yang, and J. S. Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.
- [18] Y. Lin and T. S. Chang. Data and hardware efficient design for convolutional neural network. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 65(5):1642–1651, 2018.