

SIMDive: Approximate SIMD Soft Multiplier-Divider for FPGAs with Tunable Accuracy

Zahra Ebrahimi, Salim Ullah, Akash Kumar

Chair For Processor Design, Center for Advancing Electronics Dresden (CfaED),
Technische Universität Dresden, Germany
zahra.ebrahimi_mamaghani@tu-dresden.de

ABSTRACT

The ever-increasing quest for data-level parallelism and variable precision in ubiquitous multimedia and *Deep Neural Network* (DNN) applications has motivated the use of *Single Instruction, Multiple Data* (SIMD) architectures. To alleviate energy as their main resource constraint, approximate computing has re-emerged, albeit mainly specialized for their *Application-Specific Integrated Circuit* (ASIC) implementations. This paper, presents for the first time, an SIMD architecture based on novel multiplier and divider with tunable accuracy, targeted for *Field-Programmable Gate Arrays* (FPGAs). The proposed hybrid architecture implements Mitchell's algorithms and supports precision variability from 8 to 32 bits. Experimental results obtained from Vivado, multimedia and DNN applications indicate superiority of proposed architecture (both in SISD and SIMD) over accurate and state-of-the-art approximate counterparts. In particular, the proposed SISD divider outperforms the accurate *Intellectual Property* (IP) divider provided by Xilinx with 4× higher speed and 4.6× less energy and tolerating only <0.8% error. Moreover, the proposed SIMD multiplier-divider supersedes accurate SIMD multiplier by achieving up to 26%, 45%, 36%, and 56% improvement in area, throughput, power, and energy, respectively. The library of implementations will be open-sourced and available at https://Address_blinded_for_review.edu, to springboard future research on approximate multipliers/dividers for approximate computing and FPGA communities.

KEYWORDS

Field-Programmable Gate Arrays, Approximate Hybrid Multiplier-Divider, Single Instruction Multiple Data, Mitchell's Algorithm, Energy-Efficiency, Area-Optimization.

1 INTRODUCTION

The computationally-intensive nature of upcoming *Internet of Things* (IoT) and multimedia, has dictated the demand to feature energy-efficient, multi-precision *Single Instruction, Multiple Data* (SIMD) architectures. Approximate computing paradigm has shown to serve as a viable energy-efficient solution for these applications after the strives to prevent breakdown of Moore's law with the cease of Dennard scaling era [11]. This technique has also become pronounced in machine learning domain, as it can trade stringent resource budget with a tolerable quality relaxation.

Multipliers and dividers are the most highly-used/resource-hungry arithmetic units in the kernel of these applications (dominating 99% of computational energy [11]) which carves out a prominent niche for their approximation. In particular, long latency of divider limits the overall speed of applications. Approximation of division has recently gained attention as, although less frequent, this operation is still inevitable in these applications. For instance, it is used within K-means in unsupervised clustering, Discrete Cosine Transform in JPEG compression and AlexNet Convolutional Neural Network (up to half a million times in three layers) [15]. However, the required numerical-precision for these operations is not fixed among all or even within an application (e.g., not only the precision varies during learning and retrieval phase, but also among layers of a neural network [13]).

Field-Programmable Gate Arrays (FPGAs), rewarded by a high degree of parallelism to accelerate these applications, have been

augmented with hard-wired DSP blocks to excel multiplications. Nevertheless, in spite of their advantages, hosting off-the-shelf fixed-precision DSP blocks falls short on fulfilling design requirements in a variety of domains. Beside being unable to perform division, some shortcomings that testify on their inefficiency are: 1) their fixed locations in FPGAs impose routing complexity and often results in degraded performance of some circuits [16] (and Viterbi decoder, Reed-Solomon and JPEG encoders discussed in [31]); 2) unable to be efficiently-utilized for multiplication precision below 18-bit [4, 18] (the comparable performance and better energy-efficiency of small-scale LUT-based multipliers over DSP blocks further encourages their deployment in various applications, e.g., neural networks) 3) their limited ratio versus LUTs (<0.001) in multiplication-intensive applications or concurrently executing programs. This forces designers to utilize soft- *Intellectual Property* (IP) versions of multipliers and dividers provided by major FPGA vendors such as Xilinx and Intel [38, 37].

Most of the architectures in the literature which are mainly customized for ASIC, have focused on approximating fixed word-length multipliers or dividers, or SIMD multipliers [28, 29, 23]. These techniques are not generic since approximation principles (as defined for ASIC) neglect differences in the underlying reconfigurable infrastructure and yield insignificant improvements when directly synthesized and ported to FPGAs [32]. Few designs have targeted FPGAs which are either *approximate SISD* [31, 32] or *accurate SIMD* multipliers [24, 25, 17, 14, 26]. Moreover, lack of support for division in their SIMD architecture imposes substantial overhead on the target design. This highlights the need for exploring novel avenues to provide a roadmap enabling approximate SIMD-fashion multipliers/dividers, specifically for FPGAs.

This paper presents for the first time an *SIMD approximate soft Multiplier-Divider targeted for FPGAs with Tunable accuracy* – SIMDive. The proposed hybrid architecture not only eliminates the need of reconfiguration and its associated overheads but can also support both integer multiplication and division, for the first time, in an integrated hardware with the word-lengths of 8-, 16-, and 32-bits. We build our architecture based upon Mitchell's algorithm which translates multiplication into addition. This translation results in the simplest linearly-approximated logarithmic multiplier [21]. In addition, by altering only the additions to subtractions, division can be derived with the same steps. These translations enable resource-saving, and perfectly fit FPGAs as they are already equipped with fast carry chains hardened to accelerate addition and subtraction. As such, the novel contributions of this paper are outlined as follows:

- **First integrated approximate multiplier-divider architecture.** We first propose a novel multiplier and a divider based on Mitchell's algorithms, each of which has smaller resource and higher speed than its accurate counterpart (specially divider with 4.6× less energy and 4× higher speed). Afterwards, we architect a hybrid design that can be used in either functionality without the need of reconfiguration with still less energy and delay than accurate multiplier.
- **An SIMD architecture for proposed multiplier-divider, customized for FPGAs.** To speedup original Mitchell's algorithm and adapt SIMD approach, we propose a *Leading-One Detector* (LOD), used in the first step of this algorithm. The proposed LOD detects the position of leading one in parallel in

Table 1: Summary of SISD (approximate) and SIMD (accurate/approximate) multipliers and dividers in the literature

SIMD	Approx	Mul/Div	Description of work	Platform	Improvement	Avg Rel Error (%)
X	✓	✓/X	4x4 and 8x8 with approximate partial products [32]	FPGA	{Area, power, latency}+	1.6
X	✓	✓/X	4-, 8-, and 16-bit with approximate PPs using 4x2 instances [31]	FPGA	{Area, energy}+	0.3
X	✓	✓/X	Improving accuracy of Mitchell's logarithmic multiplication [28]	ASIC	{Area, power}++	2.7
X	✓	X/✓	Improving accuracy of Mitchell's division (same approach in [28] is applied for division) [29]	ASIC	{Area, power}++, Delay+++	2.9
X	✓	X/✓	MSB division based on position of leading-one ([12], [7])	ASIC	{Delay, power, energy}++	0.8 to 3, 13.4
X	✓	X/✓	Multiplying truncated/rounded dividend with reciprocal of divisor: [34], [33], [39], [3]	ASIC	{Area, power, delay}+++	4.2, 2.8, 2.4, 4.9
X	✓	X/✓	Inexact subtractors in array divider [19], Approximating FP mantissa division to subtraction [9]	ASIC	{Area, power, delay}+	1.8, 0.6 for FP
✓	X	✓/X	Accurate variable-precision multiplier (8x8 to 32x32) [25]	FPGA	Throughput ++	Accurate
✓	X	✓/X	8- to 32-bit Add/Sub/Mul for FPGA-based multimedia processors [17], [26]	FPGA	Throughput ++	Accurate
✓	✓	✓/X	Variable-precision multiplier based on 8-bit truncated instances [23]	ASIC	{Area, Energy} +	1.2
✓	✓	✓/✓	First approximate hybrid multiplier/divider (SISD and SIMD) with tunable accuracy	FPGA	Area ++, {Throughput, energy}+++	0.7

each 4-bit segment of inputs solely with two 6-LUTs instead of large priority encoders, used in previous studies. Moreover, by adding one controlling signal, our proposed design also successfully implements an approximate SIMD divider, altogether smaller than accurate multiplier.

- **Tunable accuracy using novel light-weight error-reduction scheme.** We propose error-reduction strategy applicable to both multiplication and division. We use solely one 6-LUT for determining each bit of 64 error-reduction terms and minimal extra circuitry for their addition to the baseline multiplier/divider. In fact, our error-coefficients are added with the same LUTs and their associated fast carry chains that are already used for the addition step of Mitchell's algorithm¹. This addresses the prolonged critical path in cutting-edge error-refinement approaches [28, 29, 20]. Hence, we are able to increase accuracy by one-bit and limit error to a desirable bound using one more LUT (99.2% accuracy with eight LUTs). Moreover, our error-refinement scheme is easily scalable and can be coupled with multiplier/divider of any size.

Evaluation of SIMDiv against established approximate multipliers and dividers in the literature demonstrates that it outperforms cutting-edge designs (both SISD and SIMD) w.r.t the accuracy and performance metrics. To allow extensible designs and accelerate research in FPGA and approximate computing communities, the RTL and behavioral models of SIMDiv will be open-sourced and available at https://Address_blinded_for_review.edu.

2 RELATED WORK

The summary of state-of-the-art (SoA) studies in the literature (Table 1), can be classified in two major categories:

2.1 SISD Approximate Multiplier/Divider

SoA studies in this category, [29, 29, 3, 28, 33, 39, 34, 7, 12, 9, 20] for ASIC and [32, 31] for FPGAs use four approaches: 1) LSBs truncation which impose higher error (>4%) in divider compared to multiplier. 2) Hierarchical structure of inexact multipliers. This has weak-scalability as Karnaugh map simplifications were optimized for a specific-size (hence, error can drastically accumulate when transported to a larger design). 3) Use approximate adders (subtractors) for multiplication (division) which offers limited resource saving. 4) Multiplying rounded dividend with reciprocal of divisor in which achieving low error (<1%) comes with the cost of overhead at least in one of resource metrics compared to exact divider and/or SIMDiv 5) Exploit approximate multiplication/division algorithms based on leading one position. In addition, approaches implemented in ASIC-platforms have provided smaller gains in FPGAs. Ultimately, penalty of separate resources for implementing multiplier and divider still exists. Specifically, shortcoming attributed to Mitchell-based designs are: 1) approximating log of inputs individually (by piece-wise segmentation of each power-of-two-interval), neglects magnitude of error after multiplication [20]. 2) Lengthened critical path of due to selection of error-coefficient depends on the intermediate result of Mitchell's algorithm. 3) Many overflow cases after adding the

¹The ternary addition is devised based on directly configuring LUT primitives from Xilinx UNISIM library [36].

error-reduction term.

2.2 SIMD Accurate/Approximate Multiplier

Authors in [18, 4] have shown performance/energy improvements in FPGA-based DNNs by modifying ASIC-based DSP block to perform double approximate multiplications with a common operand. In a recent study, [23] has proposed an approximate SIMD design (using 8x8 truncated multipliers) for ASIC platforms. Targeting FPGAs, few works have presented SIMD soft multipliers [25, 24, 26, 17] that implement accurate multiplication with precision 8- to 32- bit.

By providing better resource gains and accuracy when compared to SoAs, our work differs from [28, 29] (which have similar and untunable approximation schemes, separately for multiplication and division), as: 1) we design a light-weight accuracy knob that can bound error to < 0.8%, customized specifically for FPGAs. 2) Through our error-reduction approach, the primary weakness of these works, i.e., the prolonged critical path is addressed. 3) We introduce the first hybrid multiplier/divider, smaller than an accurate multiplier, without the need for reconfiguration or changing architecture of FPGAs, 4) With a customized LOD designed for FPGAs, we enable precision-variability of proposed multiplier/divider targeted for SIMD architectures. SIMDiv can speed up execution of many applications featuring data parallelism while also reducing the share of energy expenditure by coalescing multiple memory accesses, both of which increase computational efficiency.

3 PROPOSED ARCHITECTURE

3.1 Preliminaries: Mitchell's Algorithms

Consider the binary representation of N -bit unsigned integer A (Eq. 1), in which k reveals the position of leading one. The rest of the bits (starting from position $k - 1$ to 0) are considered as the fractional part (falls in the range $0 \leq x < 1$).

$$A = 2^k + \sum_{i=0}^{k-1} 2^i b_i = 2^k(1+x) \xrightarrow{e.g.} 43 = 2^5(1+0.01011)_2, 10 = 2^3(1+0.01)_2 \quad (1)$$

In linear mathematics, $\log_2(1+x)$ is approximated to x for this range; therefore, the approximate log value of input A is:

$$\widetilde{\text{Log}}_2(A) \approx k + x \Rightarrow \text{Log}_2(43) \approx (101.01011)_2, \text{Log}_2(10) \approx (11.01)_2 \quad (2)$$

In the same manner for second input, summation (subtraction) of two parts is obtained in Eq. 3 (Eq. 4).

$$\widetilde{\text{Log}}_2(\tilde{P}) = (k_1+k_2)+(x_1+x_2) \Rightarrow K_s = (1000)_2, X_s = (0.10011)_2 \quad (3)$$

$$\widetilde{\text{Log}}_2(\tilde{D}) = (k_1-k_2)+(x_1-x_2) \Rightarrow K_s = (10)_2, X_s = (0.00011)_2 \quad (4)$$

Finally, by applying anti-log, binary representation of approximate product (quotient) are derived by Eq. 5 (Eq. 6):

$$\tilde{P} = \begin{cases} 2^{k_1+k_2}(1+x_1+x_2), & x_1+x_2 < 1 \\ 2^{k_1+k_2+1}(x_1+x_2), & x_1+x_2 \geq 1 \end{cases} \quad (5)$$

$$\Rightarrow \tilde{P} = (110011000)_2 = 408, P_{\text{accurate}} = 430$$

$$\tilde{D} = \begin{cases} 2^{k_1-k_2-1}(2+x_1-x_2), & x_1-x_2 < 0 \\ 2^{k_1-k_2}(1+x_1-x_2), & x_1-x_2 \geq 0 \end{cases} \quad (6)$$

$$\Rightarrow \tilde{D} = (100)_2 = 4, \quad D_{\text{accurate}} = 4$$

3.2 SIMDive: Approximate SIMD Multiplier-Divider

The overall structure of proposed SIMD multiplier-divider is illustrated in Fig. 2. Controlling signals *precision* and *Mul/Div mode* shown in Fig. 2 (a), serve to establish diverse sub-word size and functionalities of each module, respectively. We used the default one-hot encoding preferred by FPGA manufacturers as it is proven to be more resource-efficient than binary encoding for reconfigurable fabrics [2]. One-hot encoding scheme also enables two further advantages: 1) it allows us support *mixed-precision* and *mixed-functionality* for the first time: the proposed SIMD multiplier-divider can either operate as a single 32x32 unit or be decomposed into a twin 16x16, one 16x16 and two 8x8, or quad 8x8 units each of which can act separately as a integer multiplier or a divider². Supporting dynamic mixed-functionality in our design eliminates the need of separate resources/the overhead of reconfiguration and accommodates both operations in a single module. This feature which makes it stand out from previous works is of great interest, especially in multiplication-intensive workloads with fewer division (appealing for both DNN and multimedia). Moreover, in case of sub-word parallel processing, a complete 32-bit SIMDive unit is not occupied for a division and it can also orchestrate multiplication if needed. 2) The separate *data-size* signals can also be used to power-gate each sub-unit, individually, in case of idleness for a long period (can be investigated as a future track, as it is not in the scope of this paper). In the following, we describe the structure of each 8-bit Mul/Div unit based on Mitchell's algorithm.

Referring to Fig. 2 (b), detection of leading one is orchestrated in parallel in our proposed approximate log calculator for each 4-bit segment of the input by using two 6-LUTs: the first LUT acts as a zero-detection flag and detects whether the four bits are zero. The second 6-LUT (used as two 5-LUTs) is directly configured in such a way to reveal the position of leading-one (0 to 3). Afterwards, depending on the required operand word-length, integer and fractional parts are determined based on the most significant non-zero segment. It is worth underlining that segmentation granularity enables a trade-off in latency-energy which could be a determining parameter for a designer in relatively-large multipliers. Based on our analysis, 4-bit segmentation rendered the best energy-delay product for 32-bit SIMD architecture. In the next step, each 4-bit addition for integer and fractional parts shown in Fig. 2 (b) are fulfilled by a Virtex-7 slice. As shown in part (c) of this figure, each slice includes four 6-LUTs and its associated fast carry chains, together implement a CLA. Extending the 8-bit addition to 16- and 32-bits in our SIMD architecture is also easily achieved by connecting the C_{out} from previous adder to the C_{in} of next adder (handled by yellow multiplexers in part (a) of this figure). Integer division is also performed by altering additions to subtractions (Eq. 5 and Eq. 6). We support this operation by a 2's complement module which calculates the negative of the second input before feeding it to the adder.

3.3 Proposed Light-Weight Error-Reduction Scheme

Mitchell's error for multiplier and divider (8-bit) is plotted by a heat map provided in Fig. 1, through which four points can be observed:

$$E_P = P - \tilde{P} = \begin{cases} 2^{k_1+k_2}(x_1x_2), & x_1+x_2 < 1 \\ 2^{k_1+k_2}(1-x_1-x_2+x_1x_2), & x_1+x_2 \geq 1 \end{cases} \quad (7)$$

²Naturally in case of divider, second operand has 16, 8, and 4-bit, respectively.

$$E_D = D - \tilde{D} = \begin{cases} 2^{k_1-k_2} \frac{(x_1(x_2-1)+x_2-(x_2)^2)}{2(1+x_2)}, & x_1-x_2 < 0 \\ 2^{k_1-k_2} \frac{(x_1x_2-(x_2)^2)}{1+x_2}, & x_1-x_2 \geq 0 \end{cases} \quad (8)$$

- Different error magnitude in each power-of-two interval (Fig. 1 (a), (d)) demonstrates that adding a single correction term to the output cannot fit for all multiplier/divider sizes.
- Eq. 7 and Eq. 8 prove proportional replication of error in each power-of-two. Which is to say, irrespective of k_1 and k_2 , unique refinement schemes for each of multiplier/divider may fit all sizes and they can be added to fractional part before scaling to save more resources.
- Fig. 1 (b), (e) exhibit the non-uniform, but symmetrical error distribution: errors tends to be the same at the beginning and end of each power-of-two interval, encouraging the same reduction approach for all multiplier or divider sizes.
- Finally, Fig. 1 (c) shows diverse distribution of relative error. This means employing a single error-coefficient to the whole interval (as proposed in SoA MBM [28] and INZeD [29]), is not efficient and results in many output overflow cases.

Coalescing the insights from above points incentivize using multiple error-reduction terms appropriately opted based on summation of fractional parts to cope with overflow and yet enduring minimal latency overhead. In our analysis, we attempted to optimize two factors: 1) *error magnitude* \times *error distribution* in each region (can be estimated as the integral of error-magnitude of the region). 2) partitioning overhead which depends on the number of MSBs checked in fractional parts (e.g., checking up to 11 bits as proposed in [20] would thwart the gain of approximation). Elaborately building our architecture upon these observations and benefiting from the underlying FPGA structure, we have proposed a novel error-reduction scheme based on efficiently utilizing 6-LUTs: we assign 3 MSBs of each fractional part to LUT inputs which is responsible for calculating one bit of the error-coefficient. As illustrated in Fig. 1 (b) and (e), the squarish region for all combination of inputs is subdivided to 64 sub-regions by merely using 3 MSBs of inputs. We assign a distinctive coefficient to each of these regions, representing their average error for each sub-interval. 64 output entries of i^{th} LUT determine i^{th} bit of these 64 coefficients in binary representation. Therefore, using only one LUT, we can efficiently determine one bit of 64 error-coefficient terms. Having 64 coefficients appropriately calculated based on the combination of both operands addresses both drawbacks discussed in Section 2 (neglecting magnitude of error due to separately approximating each operand and overflow cases). It has to be emphasized that employing a single coefficient also results in several cases having equal or near error same as peak error in original Mitchell. These cases neither are measured, nor handled in [28, 29]. In our scheme, such errors are alleviated by assigning an appropriate coefficient to each pair of input.

As discussed, LUTs and their associated fast carry chain in Xilinx UNISIM library [36] can be configured to implement a ternary adder. This perfectly suits our error-healing approach as *we are able to combine the process of adding error-reduction term with fractional parts within the same resources in a single step*. Regardless of adder size, only one more bit at MSB is needed in ternary addition (compared to binary version), since $\text{frac}1_i + \text{frac}2_i + \text{error_term}_i + C_{in}$ (C_{out} of from previous bit) may result in 3 bits which necessitates one more LUT at the end of the chain. Moreover, the delay of FPGA primitives is fixed and adding error-reduction term at the same time when fractional parts are added keeps the overall delay of the circuit nearly untouched (addresses lengthened critical path discussed in Section 2).




Figure 1: Mitchell’s error: (a) integer multiplier, (b): top view of multiplier relative error in each power of two interval (same for any size of multiplier) (c): relative error distribution in multiplier, (d): integer divider, (e): top view of divider relative error in each power of two interval (same for any size of divider)

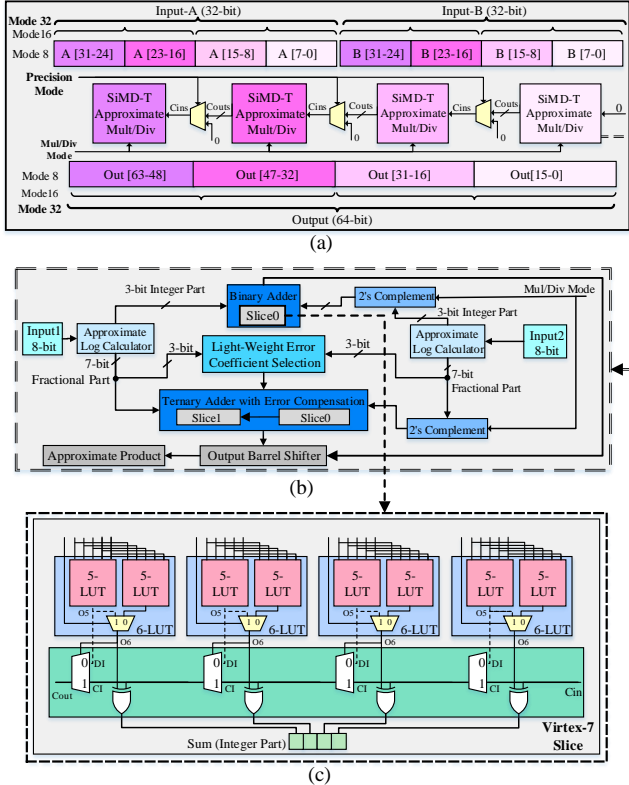


Figure 2: (a) SIMDiv Structure, (b) Proposed 8-bit multiplier/divider, (c) Virtex-7 slice (used for addition/subtraction of integer/fractional parts)

3.4 Applicability to other FPGAs

Our approach is easily applicable to various FPGAs without the need for architectural modification. It can achieve even better accuracy in higher-bit LUTs (e.g., 8-bit ALMs in Intel’s Stratix and Arria series): considering 4-bits of each fractional parts will provide 256 sub-regions. Therefore, solely one LUT can enable 256 error-reduction coefficients which can significantly improve accuracy³. Our error-analysis reveals that the proposed light-weight error-reduction approach in this paper *can significantly reduce average relative error to < 0.1% using 8-LUTs*.

4 RESULTS AND DISCUSSION

4.1 Experimental Setup

Here, we have evaluated SIMDiv against five SIMD and SISD accurate and approximate cutting-edge multipliers and dividers. They include: performance-optimized accurate IPs of multiplier (LogiCORE v12) [38] and divider (LogiCORE v5.1) [37] with default configuration, i.e., single-cycle radix-2, non-blocking,

³Ternary addition also can be implemented efficiently in Stratix and Arria series with higher frequency [10]

provided by Xilinx Vivado, Mitchell [21], SoAs MBM [28], INZeD [29] and AAXD dividers [29] as they have the best resource-error trade-off when compared to the rest of designs ([12, 19, 34, 7, 8]), CA [31] (based on approximate 4x4 multipliers) customized for FPGAs, and truncated multiplier (with 7x7 or 15x7 as the basic multiplier, the more accurate one is also exploited in SIMD structure). Note, hierarchical SIMD divider is not mathematically feasible by decomposing large one to small instances [5]. Even implementing division in reciprocal-multiplication approach would be approximate and user still needs divider IP to generate reciprocal function. However, by exploiting Mitchell’s algorithm in our proposed SIMDiv, we have made its approximated SIMD mode possible (as division is translated to subtracting two groups of numbers and then a simple shifting).

All designs are implemented from scratch in 16-bit SISD mode to provide insight about their individual resource footprints and accuracy. Afterwards, they have been exploited in 32-bit SIMD architecture. Each circuit is coded in VHDL and synthesized and implemented by Vivado 17.4 for Virtex-7 VC707 FPGA. Area, throughput, and power are reported from Vivado and Power Analyzer simulations over 10^6 for SISD and 10^9 inputs for SIMD mode uniformly distributed in a random order in the whole 16- and 32-bit interval, respectively. For precise estimations, throughput and energy dissipation are calculated based on the total execution time and power consumed for all the inputs fed to the SIMDiv. We used the cost function defined by [1], i.e., $Area \times Energy \times Delay / (1 - NED)$, Where Normalized Error Distance (NED) is the error distance for all inputs divided by maximum error. Design metrics are also reported separately since: 1) the weighted product of quality-resource metrics may not always be an appropriate figure of merit and lacks distinctiveness [27]. 2) Depending on designer/application preference, each metric can have more importance over the other. The behavioral models of multipliers/dividers are also developed in MATLAB, C++, and Python to calculate average absolute relative error and peak absolute relative error (referred to as relative and peak errors, respectively) for all possible multiplier inputs. We have also deployed SIMDiv during the inference phase of a lightweight Artificial Neural Network (ANN), Gaussian Image Smoothing, and Multiply-based Image blending applications to test the applicability of SIMDiv in real-world applications.

4.2 Simulation and Synthesis Results

Tables 2 and 3 summarize design metrics and error analysis. Following conclusions are notable referring to these tables:

- *SIMDiv corroborates its superiority by improving resource consumption*: Comparing SIMDiv with truncated and hierarchical-based counterparts, designed upon incorporating smaller instances, justifies following points:

1) all resource footprints are improved in Mitchell-based designs compared to the accurate counterparts. In particular, delay and energy are improved by 4x and 4.6x, respectively, in our proposed divider in SISD mode, as compared to accurate counterpart. In contrast, CA [31] with hierarchical

Table 2: Design metrics in SISD multipliers (16x16) and dividers (16/8)

	SISD Circuit	Area (6-LUT)	Delay (nS)	Power (mW)	Energy (μ J)	ARE ¹ (%)	PRE ² (%)	CF ³	
Multiplier	Accurate IP [38]	287	6.4	47.8	306	-	-	1	
	CA [31]	245	6.8	46.6	317	0.3	19.04	0.68	
	Trunc (four 7x7)	200	5.6	40.5	227	2.35	100	0.46	
	Trunc (two 15x7)	202	5.8	42	245	1.19	100	0.51	
	Mitchell [21]	174	4.7	35.5	167	3.85	11.11	0.43	
	MBM [28]	186	5.4	36.3	196	2.63	8.81	0.41	
	Proposed	211	4.8	37.7	178	0.82	4.9	0.34	
Divider	Accurate IP [37]	168	21.4	24.6	526	-	-	1	
	AAXD (12/6) [12]	216	18.5	22.8	412	0.74	100	0.96	
	AAXD (8/4) [12]	148	11.9	21	263	2.99	100	0.3	
	Mitchell [21]	119	5.3	20.3	107	4.11	13	0.17	
	INZeD [29]	160	5.9	22.5	137	2.93	9.5	0.11	
		Proposed	140	5.4	21.2	114	0.77	5.24	0.06
		Proposed Integrated Mul-Div	268	5.8	44	254	0.82	5.24	0.22

¹Average Relative Error, ²Peak Relative Error, ³ Cost Function (normalized to accurate IP)

Table 3: Design metrics of proposed approximate 32-bit SIMD multiplier-divider and SoA multipliers/dividers implemented in SIMD fashion

	SIMD Basic Block	Area (LUT)	Throughput (μ S)	Power (mW)	Energy (mJ)
Multiplier	Accurate Multiplier [24]	1125	561	121	862
	CA [31]	1071	412	106	1028
	Truncated (using 31x7)	970	645	98	608
	Accurate Divider (32-bit, SISD)	592	22	38	954
Multiplier & Divider	Mitchell Mul-Div [21]	782	851	72	339
	MBM-INZeD [28]-[29]	910	694	89.4	515
	Proposed SIMDive	834	817	77.5	379

implementation approach dissipates even more energy with lower throughput than accurate multiplier.

2) Approximation applied on hierarchical multipliers is rewarding in accuracy-resource trade-off only when it is done from scratch for each size. Otherwise by integrating smaller instances, error can be significantly accumulated as truncated bits are also placed in upper bit-positions. This means that error of CA multiplier would increase by multiple times in 32-bit SISD or SIMD mode by using smaller instances. This will drastically sacrifice output accuracy to achieve resource saving in a larger design. In contrast, instead of implementing hierarchical approach by connecting approximate modules, we have exploited resource reuse in Mitchell algorithm as much as possible: e.g., 8-bit accurate adders are connected to make a 32-bit instance, or detection of leading one is performed in parallel in each 4-bit segment of inputs. Therefore, no additional error is incurred in our design when used in 16- or 32-bit.

3) Through novel light-weight error-reduction scheme specifically customized for FPGAs, SIMDive has achieved significant improvement in terms of delay, energy, and accuracy, specially compared to SoA INZeD [29] (mitchell-based) and AAXD [12] (dynamically truncation of division operands).

4) Transforming our proposed integrated multiplier-divider from 16-bit SISD to 32-bit SIMD has increased the area by factor of 3. This factor is 4 for hierarchical designs which are still either SIMD multiplier or SISD divider (the increase in delay is also more pronounced in hierarchical designs). The reason behind this is two-fold: a) as discussed before, our proposed LOD detects position of leading one in each 4-bit segment. This enables resource reuse and imposes small overhead when 32-bit LOD is converted to 8-bit. b) Mitchell's algorithm is also inherently more suited for SIMD architecture, as converting SISD to SIMD in addition steps of the algorithm poses small overhead (modifying 32-bit SISD adder in fractional part to four 8-bit adders in SIMD mode). Overall, when input size is multiplied by a factor of x, resource footprint grows quadratic (x^2) in hierarchical multiplication approach, while it increases less aggressively by utilizing our logarithmic designs.

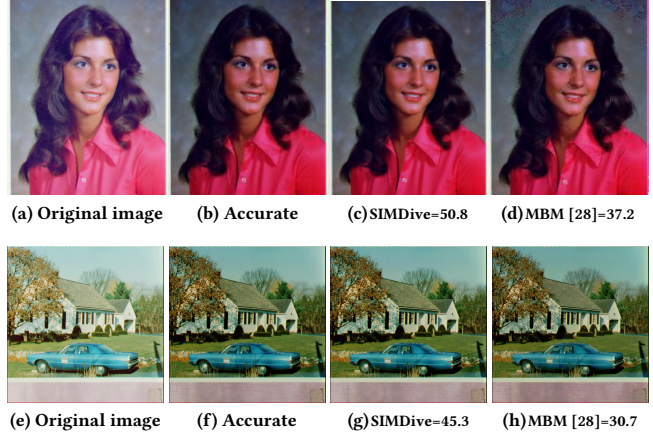


Figure 3: Image blending application with various multipliers (PSNR values are w.r.t. the accurate multiplier-based filter)

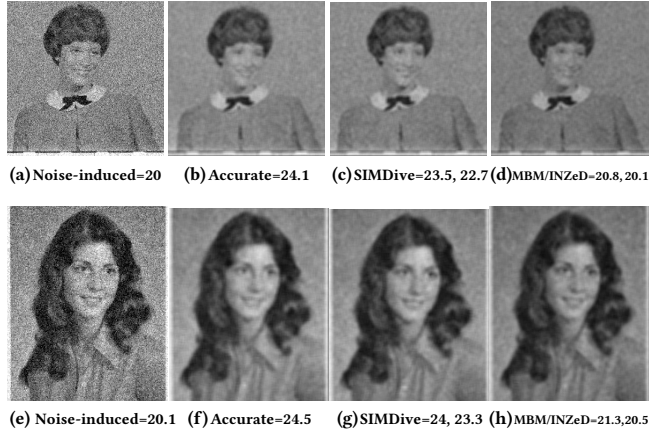


Figure 4: Gaussian noise removal filter (divider and hybrid mul/div modes, PSNR values are w.r.t the original noise-free image)

- The proposed error-refinement approach surpasses SoAs: By augmenting Mitchell's algorithm with our novel error-reduction schemes independent from the input size, we successfully achieved the lowest peak error among approximate designs (up to 20x). Additionally, average relative error is also limited to < 0.8%. Although exhaustive 32x32 test is prohibitively time consuming [23], average error in Mitchell-based designs will not significantly change in larger bit-width. Moreover, our proposed error-reduction scheme is independent from input-width and as shown by the result, it outperforms MBM design with respect to both error metrics (5x less average and 2.5x less peak error). Finally, boosting precision in our approach comes with a negligible cost: one more LUT increases the precision of error-coefficient by one bit (as discussed in Section 3.C).

4.3 High-Level ANN & Image Processing Applications

For further quality assessment of SIMDive in high-level applications, we considered Image Blending and Gaussian Image Smoothing applications for *Miscellaneous* test images provided by USC-SIPI Database [30]. Peak signal-to-noise ratio (PSNR) metric is used to measure quality of images. For Image Blending application, all multipliers in the application have been replaced with approximate versions provided by SIMDive and MBM [28]. The average PSNR value produced by SIMDive-based application for five random images from the database is 46.6, whereas it is

Table 4: ANN classification accuracy with accurate/approximate Mul/Div

MNIST Dataset	No. of Hidden layers	Nodes in each Hidden layer	Inference Accuracy %			
			Double Precision	8-bit Fixed Precision		
			Accurate	Accurate	SIMDive	MBM/INZeD [28, 29]
Digits	2	100	97.09	96.67	96.68	96.62
Digits	3	100	96.56	96.23	96.22	96.17
Fashion	2	100	85.18	84.07	84.09	84.08
Fashion	3	100	85.29	84.26	84.27	84.39
Area (normalized to 8-bit accurate)			-	1	0.78	0.74
Energy (normalized to 8-bit accurate)			-	1	0.62	0.72

limited to 32.1 for the MBM counterpart. Fig. 3 presents visual qualities of two processed images. For Gaussian Image Smoothing filter, we modified the application for two distinct modes: 1) only approximate version of divider is deployed. 2) Hybrid: both multiplication and division are replaced with approximate alternatives. For the former case, SIMDive-based Gaussian filter is capable of producing Superior PSNR (24.5), over INZeD counterpart (20.9). Interestingly in the latter case, not only the PSNR value of hybrid SIMDive (23.3) surpasses the hybrid MBM/INZeD (21.3), the values have not changed significantly compared to their "division only" approximation mode. This further motivates the deployment of hybrid SIMDive mul/div. The visual quality of two processed images are shown in Fig. 4.

We also utilized our proposed SIMDive-based approximate multiplier in an ANN, provided by [22], for testing its efficacy on the classification accuracy of MNIST handwritten digits [6] and MNIST fashion [35] datasets. The MNIST datasets provide 28×28 grayscale, 60,000 and 10,000 training and testing images, respectively. These images are categorized into ten different classes. The ANN under consideration exploits fully connected layers with 100 nodes for classification. Two different network configurations implemented: (a) two fully connected hidden layers and (b) three fully connected hidden layers. In both configurations, the input and output layers have 784 and 10 neurons, respectively. For both datasets and both configurations, the network was trained with 60,000 images using floating-point numbers; while during the inference phase for 10,000 testing images, we quantized the network parameters and activations to 8-bit fixed point precision and evaluated the classification accuracy with accurate and approximate multipliers. The classification accuracy results are described in Table 4. Due to the inherent error resilience of ANNs, many of the quantization induced errors have been partially healed for fashion dataset. Interestingly, SIMDive-based ANN not only achieves same or higher classification accuracy compared to accurate and MBM/INZeD counterparts, it also outperforms accurate design and provides 22% and 38% improvement in terms of area and energy, respectively.

5 FUTURE WORKS AND CONCLUSION

We proposed for the first time approximate SISD and SIMD soft multiplier/divider with better throughput and energy than the cutting-edge SIMD/SISD counterparts. In addition we proposed an accuracy control knob with our light-weight error-reduction scheme for the hybrid architecture which tunes error to a desirable bound and achieve high accuracy ($> 99.2\%$). We intend to utilize the proposed coalesced multiplier/divider in other domains, e.g. floating point units (mantissa multiplication and division). Moreover, to evaluate the potential gains of precision variability in more details, a customized SIMD-architecture of SIMDive-based ANN will be investigated in future tracks. Last but not least, considering the orthogonal contribution of approximate adder/subtractor in the literature, they can be employed to

add/subtract fractional part LSBs in tandem with accurate ones for MSBs without imposing high level of inaccuracy.

REFERENCES

- [1] Omid Akbari et al. "RAP-CLA: A Reconfigurable Approximate Carry Look-Ahead Adder". In: *TCAS-II* (2018).
- [2] Peter J Ashenden. *The designer's guide to VHDL*. Morgan Kaufmann, 2010.
- [3] Setareh Behroozi et al. "SAADI: A Scalable Accuracy Approximate Divider for Dynamic Energy-quality Scaling". In: *ASP-DAC*. 2019.
- [4] Andrew Boutros et al. "Embracing diversity: Enhanced DSP blocks for low-precision deep learning on FPGAs". In: *FPL*. 2018.
- [5] Linbin Chen et al. "On the Design of Approximate Restoring Dividers for Error-Tolerant Applications". In: *TC* (2016).
- [6] Li Deng. "The MNIST Database of Handwritten Digit Images for Machine Learning Research [Best of the Web]". In: *ISPREG* (2012).
- [7] Soheil Hashemi et al. "A low-power dynamic divider for approximate applications". In: *DAC*. 2016.
- [8] Soheil Hashemi et al. "DRUM: A Dynamic Range Unbiased Multiplier for Approximate Applications". In: *ICCAD*. 2015.
- [9] Mohsen Imani et al. "CADE: Configurable Approximate Divider for Energy Efficiency". In: *DATE*. 2019.
- [10] Intel. *ALM Ternary adder implementation, 2019*. <https://www.intel.com/content/www/us/en/programmable/documentation/jbr1444752564689.html>.
- [11] Shubham Jain et al. "Compensated-DNN: energy efficient low-precision deep neural networks by compensating quantization errors". In: *DAC*. 2018.
- [12] Honglan Jiang et al. "Adaptive approximation in arithmetic circuits: A low-power unsigned divider design". In: *DATE*. 2018.
- [13] Patrick Judd et al. "Stripes: Bit-serial deep neural network computing". In: *Micro*. 2016.
- [14] Shafiqat Khan et al. "Reconfigurable SWP operator for multimedia processing". In: *ASAP*. 2009.
- [15] Alex Krizhevsky et al. "ImageNet classification with deep convolutional neural networks". In: *NIPS*. 2012.
- [16] Ian Kuo and Jonathan Rose. "Measuring the gap between FPGAs and ASICs". In: *TCAD* (2007).
- [17] Marco Lanuzza et al. "Low-cost fully reconfigurable data-path for FPGA-based multimedia processor". In: *FPL*. 2005.
- [18] Sugil Lee et al. "Double MAC on a DSP: Boosting the performance of convolutional neural networks on FPGAs". In: *TCAD* (2018).
- [19] Weiqiang Liu et al. "Combining Restoring Array and Logarithmic Dividers into an Approximate Hybrid Design". In: *ARITH*. 2018.
- [20] Joshua Yung Lih Low and Ching Chuen Jong. "Unified Mitchell-based approximation for efficient logarithmic conversion circuit". In: *TC* (2015).
- [21] John N Mitchell. "Computer multiplication and division using binary logarithms". In: *IRETEC* (1962).
- [22] *MNIST-CNN, 2016*. <https://github.com/integeruser/MNIST-CNN>.
- [23] Roberto R Osorio and Gabriel Rodriguez. "Truncated SIMD Multiplier Architecture for Approximate Computing in Low-Power Programmable Processors". In: *IEEE Access* (2019).
- [24] Stefania Perri et al. "SIMD Multipliers for Accelerating Embedded Processors in FPGAs". In: *JCS* (2006).
- [25] Stefania Perri et al. "Variable precision arithmetic circuits for FPGA-based multimedia processors". In: *TVLSI* (2004).
- [26] Sohan Purohit et al. "Power-efficient high throughput reconfigurable datapath design for portable multimedia devices". In: *ReConFig*. 2008.
- [27] Semeen Rehman et al. "Architectural-space exploration of approximate multipliers". In: *ICCAD*. 2016.
- [28] Hassaan Saadat et al. "Minimally Biased Multipliers for Approximate Integer and Floating-Point Multiplication". In: *TCAD* (2018).
- [29] Hassan Saadat et al. "Approximate Integer and Floating-Point Dividers with Near-Zero Error Bias". In: *DAC*. 2019.
- [30] *SIFI Image Database, 2019*. <http://sifi.usc.edu/database/database.php>.
- [31] Salim Ullah et al. "Area-optimized low-latency approximate multipliers for FPGA-based hardware accelerators". In: *DAC*. 2018.
- [32] Salim Ullah et al. "SMApplib: library of FPGA-based approximate multipliers". In: *DAC*. 2018.
- [33] Marzieh Vaeztourshizi et al. "An Energy-Efficient, Yet Highly-Accurate, Approximate Non-Iterative Divider". In: *ISLPED*. 2018.
- [34] Shaghayegh Vahdat et al. "TruncApp: A truncation-based approximate divider for energy efficient DSP applications". In: *DATE*. 2017.
- [35] Han Xiao et al. "Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms". In: *CoRR* (2017).
- [36] Xilinx. *7 Series FPGA Programmable Guide for HDL Designs, 2013*. https://www.xilinx.com/support/documentation/sw_manuals/xilinx14_1/7series_hdl.pdf.
- [37] Xilinx. *LogiCORE IP Divider v5.1, 2016*. https://www.xilinx.com/support/documentation/ip_documentation/div_gen/v5_1/pg151-div-gen.pdf.
- [38] Xilinx. *LogiCORE IP multiplier v12.0, 2015*. https://www.xilinx.com/support/documentation/ip_documentation/multi_gen/v12_0/pg108-mult-gen.pdf.
- [39] Reza Zandegani et al. "SEERAD: A high speed yet energy-efficient rounding-based approximate divider". In: *DATE*. 2016.