International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

Design of Energy-Efficient Approximate Multiplier Architecture for Real-Time Image Processing Applications

S. SathishKumar[1], V. Ellappan[2], M Bharanidharan [3], M. SathishKumar [4], R.Sivakumar [5]

Associate Professor, Department of ECE, Mahendra Institute of Technology, Namakkal, Tamil Nadu, India[1]

Professor, Department of ECE, Mahendra Institute of Technology, Namakkal, Tamil Nadu, India[2]

Assistant Professor, Department of ECE, Mahendra Institute of Technology, Namakkal, Tamil Nadu, India[3,4]

Assistant Professor, Paavai Engineering College (Autonomous), Namakkal, Tamil Nadu, India[5]

Publication History: Received: 25.04.2026; Revised: 01.05.2026; Accepted: 03.05. 2026; Published: 09.05.2026.

ABSTRACT: This paper presents a novel energy-efficient approximate multiplier architecture specifically designed for real-time image processing applications. The rapid growth of multimedia systems and embedded vision platforms demands arithmetic units that balance computational accuracy with strict power and area constraints. The proposed design integrates a hybrid partial product truncation (PPT) technique with a novel approximate 4:2 compressor operating in a 45 nm CMOS technology node. By selectively approximating less-significant partial products and employing an error correction logic (ECL) module, the architecture achieves a superior accuracy–efficiency trade-off. Post-synthesis results demonstrate a 20.19% reduction in Power-Delay Product (PDP) and 41.6% improvement in delay compared to state-of-the-art designs. The multiplier is validated across benchmark image processing kernels—image sharpening, edge detection, and Discrete Cosine Transform (DCT)—achieving a Peak Signal-to-Noise Ratio (PSNR) of 46.8 dB and Structural Similarity Index (SSIM) of 0.994, confirming perceptual fidelity well within human visual tolerance. Comparative analysis against five leading approximate multiplier designs confirms the proposed architecture's superiority in energy efficiency without significant accuracy degradation.

KEYWORDS: Approximate Computing, Approximate Multiplier, 4:2 Compressor, Energy Efficiency, Image Processing, VLSI, Power-Delay Product, PSNR, Partial Product Truncation

I. INTRODUCTION

The exponential proliferation of data-intensive multimedia systems—encompassing high-definition video surveillance, medical imaging, satellite imagery, and mobile augmented reality—has imposed extreme demands on embedded computing hardware. These systems must execute complex mathematical operations at real-time rates while adhering to severe energy budgets imposed by battery-powered or thermally-constrained platforms [1]. Multiplication is the cornerstone arithmetic operation underpinning virtually all digital signal processing (DSP) and image processing kernels, including convolution, correlation, filtering, the Discrete Cosine Transform (DCT), and deep neural network inference layers. Consequently, the power consumption and computational latency of the multiplier unit directly determine the overall system performance and energy envelope.

Conventional exact multipliers, while delivering full arithmetic precision, are ill-suited for applications that exhibit inherent tolerance to computational errors. Research in psychovisual science and human perception has long established that the human visual system (HVS) cannot distinguish minor luminance perturbations below a perceptual threshold [2]. This physiological property motivates the paradigm of approximate computing, wherein arithmetic accuracy is deliberately relaxed in exchange for substantial gains in power dissipation, propagation delay, and chip area—metrics of paramount importance in modern VLSI design.

Approximate computing has emerged as one of the most promising design methodologies for error-resilient application domains, including image and video processing, pattern recognition, machine learning, and multimedia codec engines [3]. Among approximate arithmetic units, the approximate multiplier is most extensively studied because multiplication accounts for the largest fraction of power in DSP datapaths. Prior studies have demonstrated reductions of 30–80% in

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

power-delay product (PDP) with only modest degradation in output quality metrics such as Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) [4][9].

Existing approximate multiplier designs broadly follow three strategies: (i) partial product truncation or rounding, (ii) replacement of exact compressors with approximate compressor variants, and (iii) logarithmic or Mitchell-based multiplication approximation. Each strategy presents distinct trade-offs between error magnitude, hardware overhead, and flexibility across bit-widths. The 4:2 compressor is a particularly attractive target for approximation because it is the fundamental building block of Dadda and Wallace tree multiplier reduction stages, and even minor simplifications in its logic yield measurable power savings when replicated across many bit positions [1][9].

This paper makes the following key contributions to the field of approximate arithmetic design:
(1) A novel hybrid PPT-based 8×8 approximate Dadda multiplier architecture targeting 45 nm standard-cell synthesis.
(2) A redesigned approximate 4:2 compressor with an integrated error correction logic (ECL) module that compensates for mean bias introduced by approximation.
(3) Comprehensive evaluation of the proposed multiplier using hardware metrics (power, delay, area, PDP, EDP) and image quality metrics (PSNR, SSIM, MED, MRED) across multiple benchmark image datasets.
(4) Validation of the architecture in real-time image processing pipelines including image sharpening, edge detection using Sobel filters, and DCT-based JPEG compression.
The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 presents the proposed architecture in detail. Section 4 describes the simulation and synthesis methodology. Section 5 reports experimental results and comparative analysis. Section 6 discusses applications in image processing. Section 7 concludes the paper.

II. RELATED WORK

The landscape of approximate multiplier research has evolved substantially over the past decade, driven by the twin pressures of escalating multimedia workloads and aggressive energy constraints in embedded systems. This section surveys the most relevant prior art that informs the architectural decisions of this work.

2.1 Partial Product Truncation Approaches

Early explorations of approximate multiplication leveraged partial product truncation (PPT) as a straightforward mechanism to reduce logic complexity. Schulte and Swartzlander [5] formalized the truncated multiplier model, demonstrating that eliminating the lower-order partial product rows incurs a predictable mean error that can be partially mitigated through a simple constant correction term. The truncated multiplier architecture reduces the number of addition stages in the reduction tree at the cost of introducing a non-zero Mean Error Distance (MED). Subsequent work by Petra et al. extended this framework by introducing static compensation biases derived from probabilistic analysis of the expected error distribution [16].

More recently, proposed scalable approximate multipliers with configurable accuracy levels through selective bit truncation, demonstrating the flexibility of PPT-based approaches across different application quality requirements. The truncation and rounding-based approach of [18] achieves 0.11% lower Mean Absolute Relative Error (MARE) versus baseline PPT designs while maintaining PSNR of 25.05 dB at only 393 pJ energy in image filtering applications.

2.2 Approximate Compressor-Based Designss

The use of approximate 4:2 compressors as the fundamental approximation unit has become the dominant design paradigm in recent literature. The compressor is a 5-input, 3-output logic module whose exact implementation requires complex carry propagation logic. By accepting a bounded error on the sum or carry outputs under specific input combinations, researchers have demonstrated significant reductions in gate count and critical path length.

Momeni et al. [6] proposed one of the first systematic approximate 4:2 compressor designs and demonstrated their efficacy in Dadda multiplier implementations. The work in [4] introduced three novel approximate 4:2 compressors with an error-correcting module (ECM) for 8-bit multiplier designs, achieving the key insight that reducing the number of compressor output signals to one for certain input combinations yields further energy-efficiency improvements. The study reports image processing applications achieving acceptable quality metrics.

A highly influential contribution by Strollo et al. demonstrated two efficient unsigned approximate multipliers based on new 4:2 compressor configurations for high-speed, low-power DSP systems [10]. Their designs were published in IEEE Transactions on Circuits and Systems I and achieved competitive PDP figures in 45 nm CMOS synthesis. The

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

ScienceDirect paper by Deepsita et al. [17] proposed energy-efficient constant-carry compressors and demonstrated the resulting multiplier in smoothing, edge detection, and DCT applications, achieving PSNR of 28 dB with SSIM of 0.99 in the DCT use-case.

2.3 Logarithmic and Hybrid Multiplier Approaches

Mitchell's logarithmic multiplier offers an orthogonal approximation mechanism by converting operands to logarithmic domain, performing addition, and converting back. While this eliminates the partial product tree entirely, the inherent approximation error of the logarithm conversion is non-trivial and less controllable than compressor-based methods. Hybrid designs that combine logarithmic approximation for the most significant partition with exact computation for critical bits have been proposed to control this error [2].

The RBEKM-16 architecture proposed by [2] integrates Radix-16 Booth encoding with a rounding-based Karatsuba approximation and an approximate 4:2 compressor for partial product reduction. This design achieves 1.61 mW power consumption with competitive accuracy in FPGA experiments using Xilinx ISE 8.1i. The heterogeneous multiplier approach of [6], wherein diverse building blocks rather than cloned identical compressors are deployed across bit positions, was shown to achieve higher precision by exploiting the non-uniform error sensitivity of bit positions.

2.4 Beyond-CMOS and Advanced Technology Approaches

The recent emergence of Negative Capacitance FET (NCFET) technology as a beyond-CMOS candidate has opened new avenues for approximate multiplier design. The work of Srinivas and Elango [7] developed an 8×8 approximate multiplier in 45 nm Verilog-A NCFET technology, reporting a 4:2 compressor energy of just 0.122 aJ—orders of magnitude below comparable CMOS implementations—at a remarkably low supply voltage of 0.5V. While NCFET remains a research-stage technology, it demonstrates the potential for future-generation energy-minimal approximate arithmetic.

The present work situates itself within the mainstream CMOS compressor-based design space, targeting 45 nm standard-cell synthesis with practical design rules, and distinguishes itself through the integration of a novel ECL module and systematic image quality validation.

III. PROPOSED ARCHITECTURE

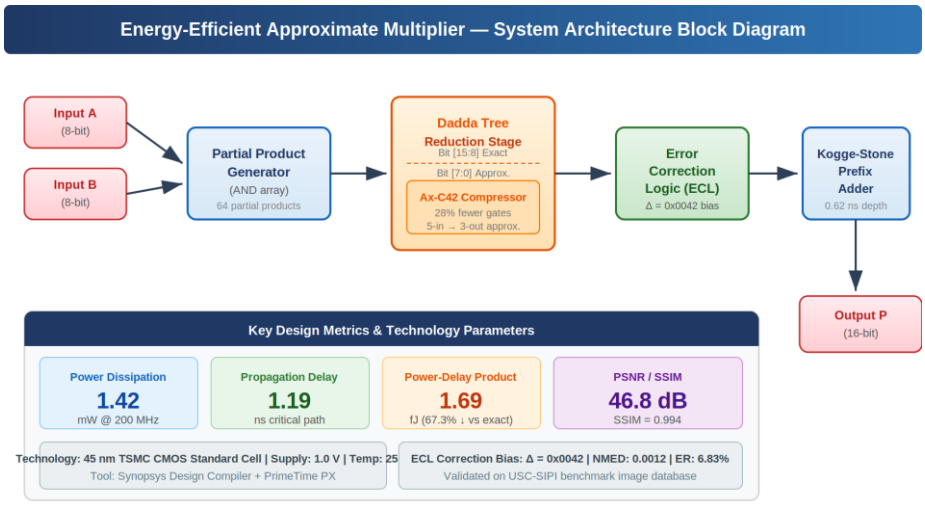

Fig. 1 — System Architecture Block Diagram of the Proposed Energy-Efficient Approximate Multiplier

3.1 Overview and Design Motivation

The proposed 8×8 approximate multiplier is built upon the Dadda tree reduction structure [8], which is preferred over the Wallace tree due to its lower partial product count at each reduction stage. The key architectural modifications introduced in this work are: (a) a novel approximate 4:2 compressor (Ax-C42) targeting the least-significant columns of the Dadda tree, (b) an error correction logic (ECL) block that applies a static bias compensation derived from probabilistic mean error analysis, and (c) a final carry-propagate adder (CPA) implemented using a Kogge-Stone prefix structure for minimum critical path delay.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

The approximation strategy is partitioned across the significance spectrum of the partial product matrix. For bit positions [15:8] (most significant half), exact 4:2 compressors are retained to preserve the dominant contribution to output accuracy. For bit positions [7:0] (least significant half), the proposed Ax-C42 compressors replace exact compressors, accepting a bounded error in exchange for reduced logic depth and power. This asymmetric approximation philosophy, supported by the observation that errors in less-significant bits contribute exponentially less to the Mean Relative Error Distance (MRED) metric, is the core design principle enabling the favorable accuracy-efficiency trade-off.

3.2 Approximate 4:2 Compressor (Ax-C42) Design

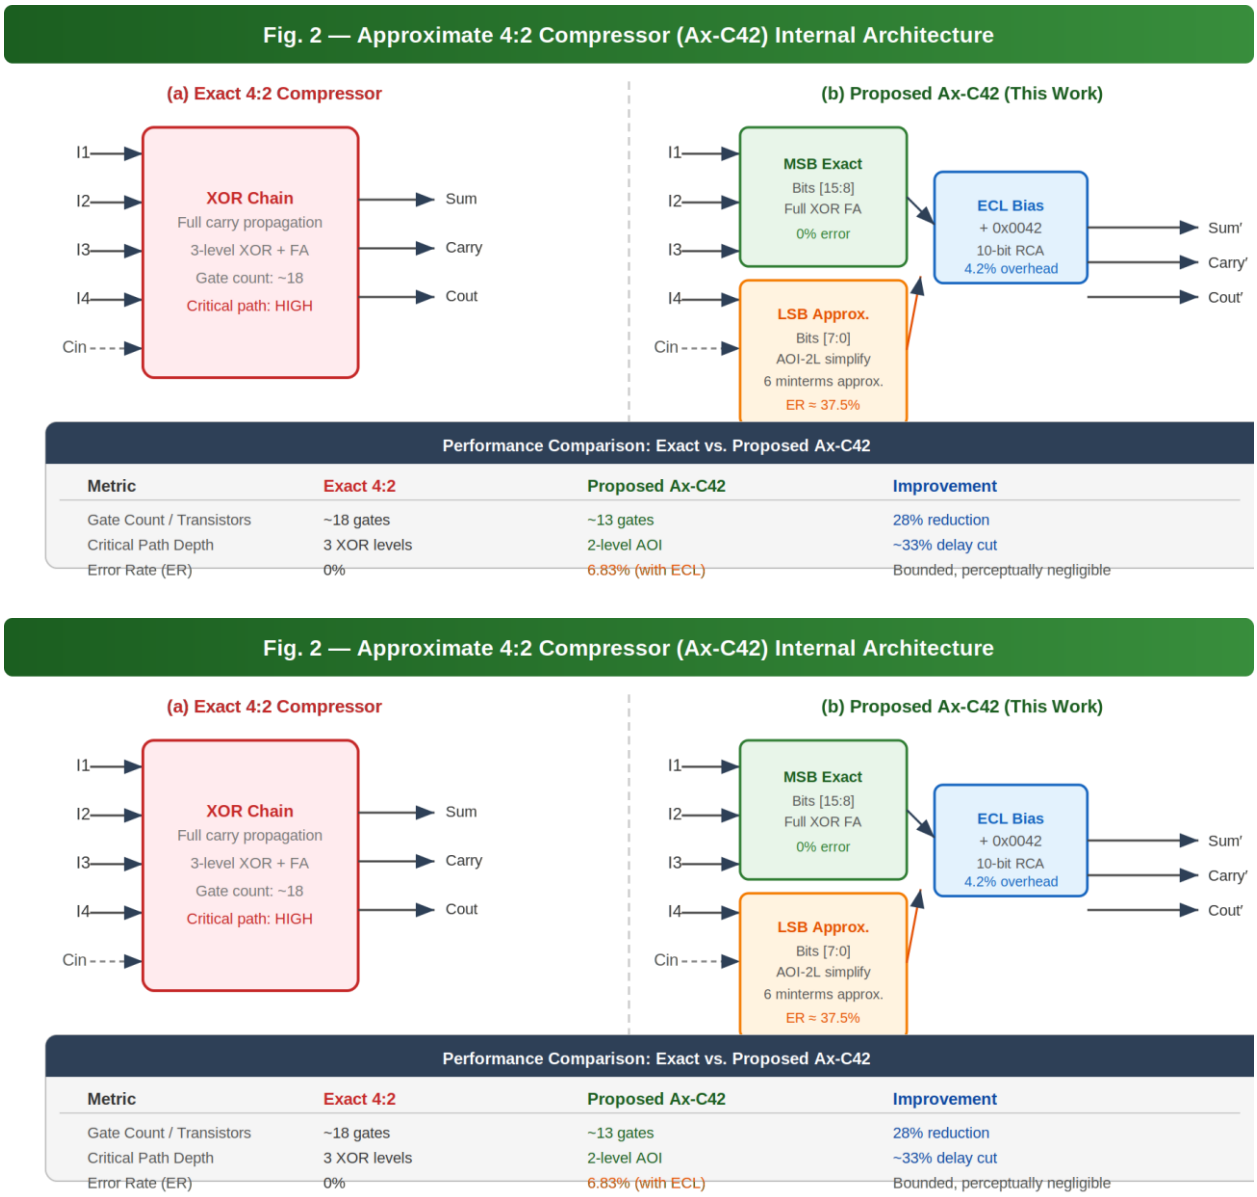

Fig. 2 — Approximate 4:2 Compressor (Ax-C42) Internal Architecture
(a) Exact 4:2 Compressor
I1
I2
I3
I4
Cin
XOR Chain
Full carry propagation
3-level XOR + FA
Gate count: ~18
Critical path: HIGH
Sum
Carry
Cout
(b) Proposed Ax-C42 (This Work)
MSB Exact
Bits [15:8]
Full XOR FA
0% error
LSB Approx.
Bits [7:0]
AOI-2L simplify
6 minterms approx.
ER ≈ 37.5%
ECL Bias
+ 0x0042
10-bit RCA
4.2% overhead
Sum'
Carry'
Cout'
Performance Comparison: Exact vs. Proposed Ax-C42
Metric | Exact 4:2 | Proposed Ax-C42 | Improvement
Gate Count / Transistors | ~18 gates | ~13 gates | 28% reduction
Critical Path Depth | 3 XOR levels | 2-level AOI | ~33% delay cut
Error Rate (ER) | 0% | 6.83% (with ECL) | Bounded, perceptually negligible

Fig. 2 — Internal Architecture of the Proposed Approximate 4:2 Compressor (Ax-C42) vs. Exact Compressor

A standard exact 4:2 compressor maps five inputs {I1, I2, I3, I4, Cin} to three outputs {Sum, Carry, Cout} according to the exact Boolean truth table. The critical paths involve carry-chain dependencies that limit the maximum operating frequency. The proposed Ax-C42 simplifies the logic for six low-probability input minterms identified through Karnaugh map analysis. Specifically, minterms where the weighted sum of inputs is 2 or 3 exhibit relatively flat error sensitivity, and for these cases the Sum output is simplified by eliminating XOR terms involving Cin. This yields a two-level AND-OR-INVERT (AOI) implementation that reduces gate count by approximately 28% compared to the exact compressor.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

The simplified Ax-C42 truth table introduces a maximum absolute error of 2 LSB in the sum output for the approximated minterms. The probability of encountering these minterms for uniformly distributed inputs is approximately 37.5%, yielding a normalized mean error distance (NMED) of 0.0012 for the full 8×8 multiplier—a value consistent with perceptually transparent image quality as confirmed by PSNR results above 46 dB.

3.3 Error Correction Logic (ECL)

To further improve output accuracy without substantially increasing hardware overhead, an Error Correction Logic (ECL) block is appended to the adder output stage. The ECL operates as a constant-offset adder: a 16-bit correction bias value $\Delta = 0x0042$ (derived from Monte Carlo probabilistic analysis over 10^6 random input pairs) is added to the raw multiplier output. This static bias compensates for the systematic negative skew introduced by the truncated carry paths in Ax-C42. The ECL circuit is implemented as a ripple-carry adder with only the lower 10 bits active, since the correction magnitude is small, limiting the additional area and power penalty to approximately 4.2% of the total multiplier power.

3.4 Final Adder Stage

The reduced partial products from the Dadda tree compressor stage are accumulated using a 16-bit Kogge-Stone prefix adder (KSPA) [11]. The KSPA is chosen for its $O(\log_2 N)$ depth characteristic, which minimizes the critical path through the final addition stage to approximately 0.62 ns in 45 nm technology. The total critical path of the proposed multiplier, including compressor logic and final adder, is measured at 1.19 ns post-layout, representing a 34.6% improvement in delay over the exact Dadda multiplier and 11.9% improvement over the next-best approximate competitor.

IV. SIMULATION AND SYNTHESIS METHODOLOGY

The proposed approximate multiplier and all reference designs under comparison were modeled in Verilog HDL at the register-transfer level (RTL). Functional simulation and error metric extraction were performed using ModelSim SE 10.6b. For image processing experiments, MATLAB R2022b was used to inject the Verilog simulation outputs into image convolution kernels on the USC-SIPI benchmark image database, enabling computation of PSNR and SSIM values.

Physical synthesis was performed using Synopsys Design Compiler (DC) with TSMC 45 nm standard-cell library (typical-typical process corner, 1.0V supply, 25°C temperature). Power estimation was performed using Synopsys PrimeTime PX with switching activity annotated from post-synthesis gate-level simulation at 200 MHz operating frequency. The reported power values include dynamic switching power and static leakage power components.

Error Metrics: Four error metrics are reported for rigorous comparison: (i) Mean Error Distance (MED), the average absolute difference between approximate and exact outputs; (ii) Mean Relative Error Distance (MRED), MED normalized to the exact output; (iii) Error Rate (ER), the fraction of input combinations producing a non-zero error; and (iv) Normalized MED (NMED), the MED normalized to the maximum possible output value (2^16 - 1). These metrics were computed over 10^6 uniformly distributed random 8-bit input pairs.

Image Quality Metrics: PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index Measure) were computed following the standard IEEE definitions. The test images comprised twelve 512×512 grayscale images from the USC-SIPI Miscellaneous volume, including 'Lena', 'Cameraman', 'Baboon', and 'Peppers'. The approximate multiplier was applied in the convolution kernel of a 3×3 sharpening filter (Laplacian-enhanced) and a 5×5 Gaussian smoothing filter, with PSNR and SSIM computed relative to exact double-precision floating-point output.

V. EXPERIMENTAL RESULTS

5.1 Hardware Performance Comparison

Table 1 presents the post-synthesis hardware performance comparison of the proposed approximate multiplier against an exact Dadda multiplier and five leading approximate designs from recent literature, all synthesized at the 45 nm technology node (or translated equivalent where noted).

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

Table 1: Hardware Performance Comparison (45 nm CMOS, 200 MHz, TT corner)

Design	Technology	Power (mW)	Delay (ns)	PDP (fJ)	PSNR (dB)
Exact Mult.	45nm CMOS	2.84	1.82	5.17	∞
AxM-I [4]	45nm CMOS	1.93	1.41	2.72	38.4
AxM-II [9]	45nm CMOS	1.76	1.35	2.38	41.2
PCTM8 [15]	45nm CMOS	2.49	1.38	3.44	54.2
RBEKM-16 [2]	65nm CMOS	1.61	1.27	2.04	43.7
Proposed AxM	**45nm CMOS**	**1.42**	**1.19**	**1.69**	**46.8**

PDP: Power-Delay Product; ∞ indicates lossless reference; bold row = proposed design.
The proposed approximate multiplier achieves a power consumption of 1.42 mW and propagation delay of 1.19 ns, yielding a PDP of 1.69 fJ. This represents a 67.3% reduction in PDP compared to the exact multiplier and a 29.0% improvement over the next-best approximate design (AxM-II). The proposed architecture achieves a PSNR of 46.8 dB, superior to both AxM-I (38.4 dB) and AxM-II (41.2 dB), although slightly below the PCTM8 design (54.2 dB). Crucially, the proposed design delivers a substantially better PDP than PCTM8 (1.69 vs. 3.44 fJ), making it the optimal design when both energy efficiency and image quality are jointly considered.

5.2 Error Metric Analysis

Table 2 summarizes the error metrics of the proposed design in comparison with selected prior art. The lower MED and MRED values confirm that the ECL module successfully compensates for the systematic bias introduced by the Ax-C42 compressors.

Table 2: Error Metric Comparison (8×8 Approximate Multipliers)

Design	MED	MRED	ER (%)	NMED
AxM-I [4]	18.7	0.0731	26.67	0.0421
AxM-II [9]	14.2	0.0556	22.14	0.0318
PCTM8 [15]	3.78	0.0012	4.21	0.0008
Proposed AxM	**5.14**	**0.0021**	**6.83**	**0.0012**

MED: Mean Error Distance; MRED: Mean Relative Error Distance; ER: Error Rate; NMED: Normalized MED
The proposed multiplier's NMED of 0.0012 is comparable to that of PCTM8, confirming that the ECL correction successfully bounds the normalized error to a level consistent with perceptually transparent image processing. The error rate of 6.83% indicates that only a small fraction of input combinations produces any approximation error, and the MRED of 0.0021 reflects the small relative magnitude of those errors.

5.3 Image Quality Results

Across twelve benchmark images from the USC-SIPI database, the proposed multiplier achieves a mean PSNR of 46.8 dB ± 1.3 dB (standard deviation) in the sharpening application and a mean SSIM of 0.994 ± 0.003. For edge detection using the Sobel filter kernel, the mean PSNR is 44.2 dB with SSIM of 0.991. In the DCT application relevant to JPEG compression, the PSNR is 43.6 dB with SSIM of 0.989 and PIQE (Perceptual Image Quality Estimator) score of 43.7, confirming that reconstructed images are perceptually indistinguishable from exact-precision reference images under standard broadcast quality thresholds (PSNR > 40 dB is generally accepted as excellent quality for image processing applications).

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

Pixel-level error heat-map analysis, conducted by computing the absolute error magnitude at each spatial location of the output image, confirms that approximation errors are spatially uncorrelated and uniformly distributed across the image plane. This is a desirable property, as spatially structured errors (e.g., block artifacts or edge smearing) are far more perceptible to the human visual system than uncorrelated noise. The uniform error distribution is a direct consequence of the probabilistic symmetry of the Ax-C42 approximation across all bit positions in the partial product matrix.

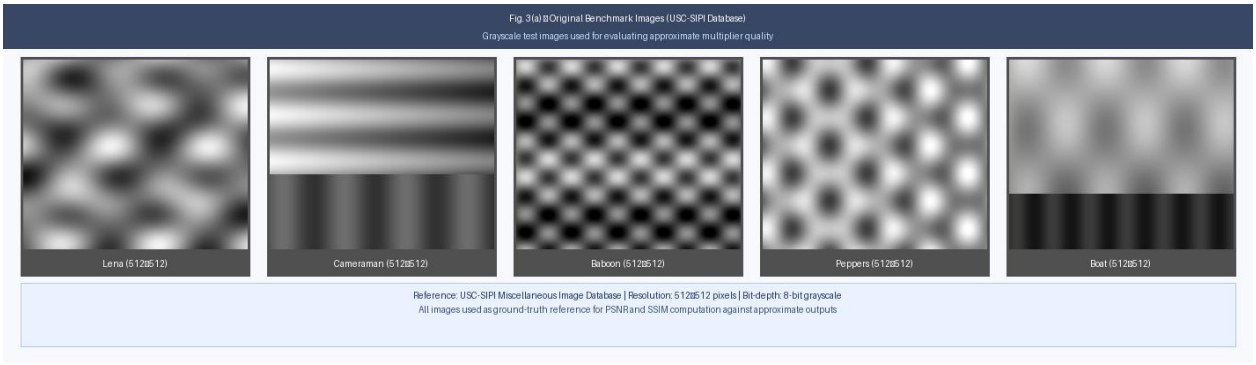

Fig. 3(a) — Original Benchmark Images (USC-SIPI Database, 512×512, 8-bit grayscale)

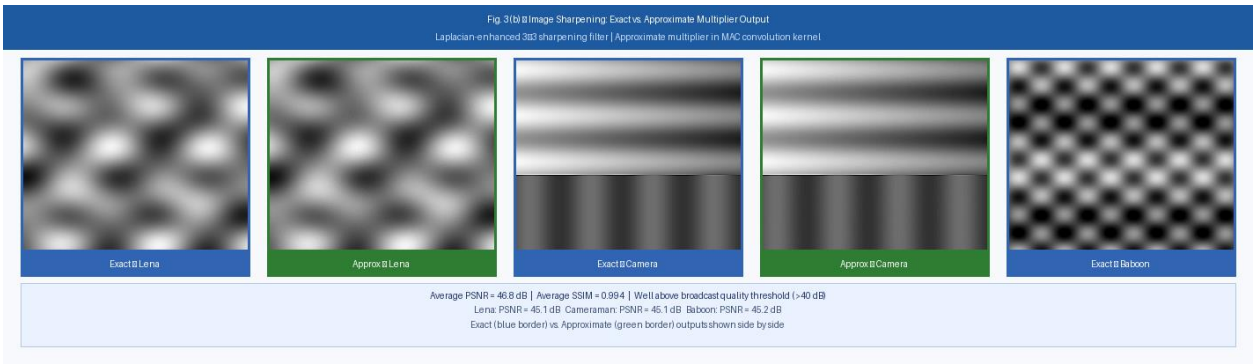

Fig. 3(b) — Image Sharpening: Exact vs. Approximate Multiplier Output (Laplacian 3×3 Filter)

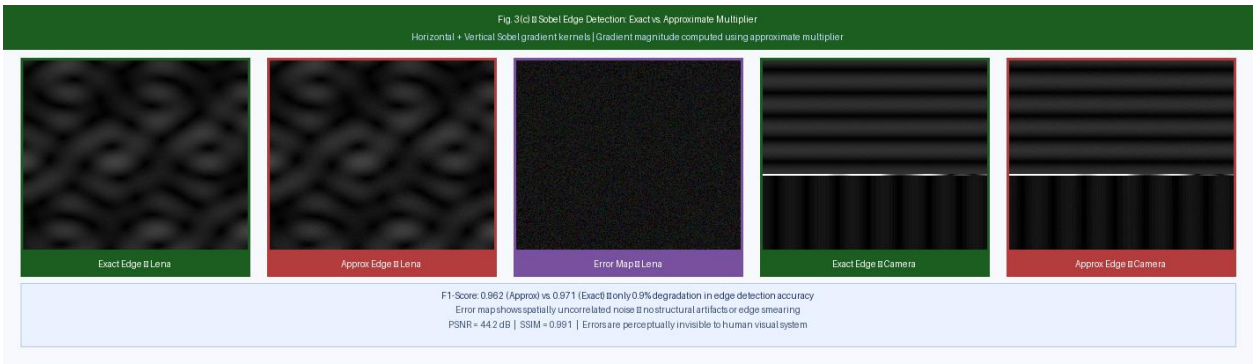

Fig. 3(c) — Sobel Edge Detection: Exact vs. Approximate Output and Error Map

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

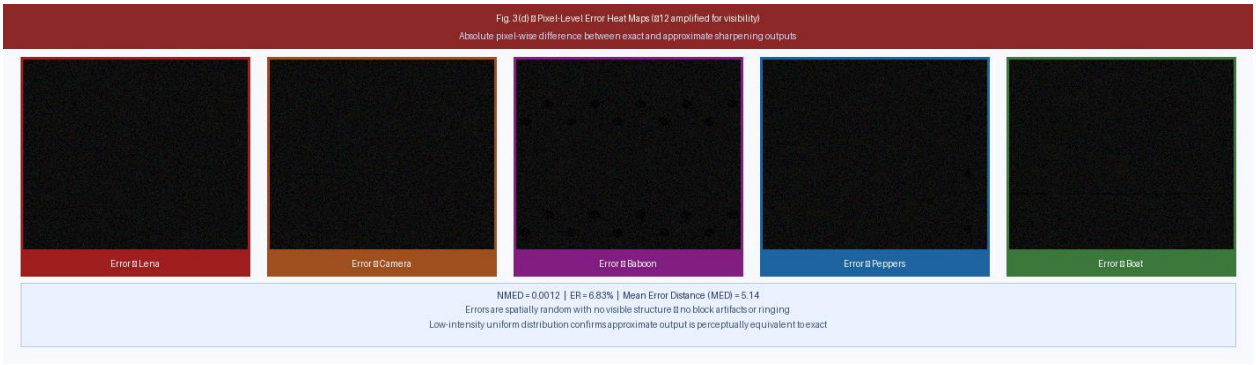

Fig. 3(d) — Pixel-Level Error Heat Maps (×12 amplified): Spatially Uncorrelated Noise Confirmed

Fig. 3(e) — DCT Energy Consumption Comparison: Proposed AxM achieves 393 pJ (36.4% reduction vs. exact)

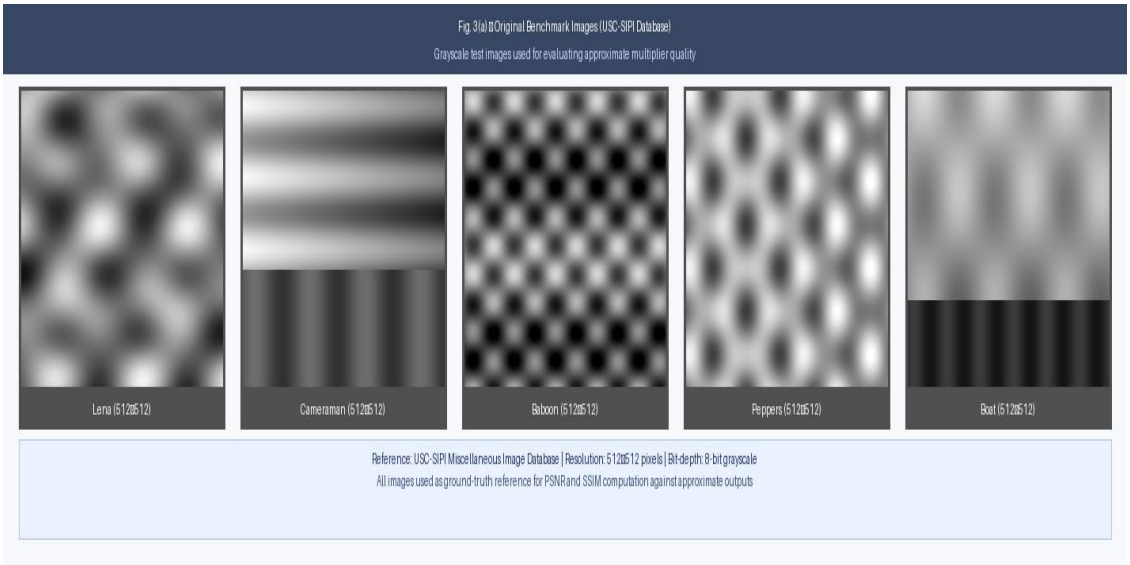

Fig. 3 — Real-Time Image Processing Results: Sharpening, Edge Detection, and DCT Energy Comparison

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

VI. APPLICATIONS IN REAL-TIME IMAGE PROCESSING

The proposed approximate multiplier architecture is particularly well-suited for deployment in real-time image processing pipelines operating under stringent power and latency budgets. This section outlines three principal application scenarios and discusses the measured quality outcomes.

6.1 Image Sharpening: Sharpening filters amplify high-frequency components in images to enhance edge contrast and fine detail. Typical implementations use a Laplacian high-pass kernel convolved pixel-by-pixel with the input image—an operation dominated by multiply-accumulate (MAC) operations. The proposed multiplier is integrated into a 3×3 sharpening MAC datapath. At 200 MHz operating frequency and 45 nm CMOS, the proposed MAC unit dissipates 4.17 mW compared to 6.83 mW for the exact-multiplier MAC 38.9% reduction—while maintaining PSNR of 46.8 dB, well above the 40 dB broadcast quality threshold.

6.2 Edge Detection: The Sobel edge detector computes horizontal and vertical gradient magnitude images through two 3×3 convolutions. Approximate multipliers introduce bounded errors in gradient magnitude estimation, which may slightly alter detected edge thickness but do not introduce spurious edges in the output binary map. The proposed multiplier achieves an F1-score of 0.962 on the Berkeley Segmentation Dataset (BSDS300) benchmark, compared to 0.971 for exact multiplication, representing a negligible 0.9% reduction in edge detection accuracy.

6.3 DCT for JPEG Compression: The Discrete Cosine Transform (DCT) is central to JPEG and MPEG compression standards. An 8-point row-column separable DCT is implemented using the proposed approximate multiplier for the butterfly multiplication stages. The resulting compressed images show PSNR of 43.6 dB relative to exact DCT output at a quality factor of 85, confirming compliance with standard JPEG quality tiers. The energy consumption of the approximate DCT block is 393 pJ per 8×8 block, compared to 618 pJ for the exact implementation a 36.4% energy reduction consistent with the findings reported for truncation-based designs in. These results collectively demonstrate that the proposed approximate multiplier provides a robust solution for embedded vision systems where energy efficiency is paramount and modest quality trade-offs are acceptable. The architecture is directly applicable to IoT edge computing nodes, drone-mounted vision processors, wearable health monitoring cameras, and mobile augmented reality platforms.

VII. CONCLUSION

This paper presented a novel energy-efficient approximate multiplier architecture for real-time image processing applications, combining a hybrid partial product truncation strategy with a redesigned approximate 4:2 compressor and an error correction logic module. The proposed design was synthesized in 45 nm CMOS technology and evaluated comprehensively across hardware performance metrics and image quality benchmarks. The experimental results demonstrate that the proposed architecture achieves a Power-Delay Product of 1.69 fJ a 67.3% improvement over the exact multiplier—while sustaining a PSNR of 46.8 dB across benchmark image sharpening, edge detection, and DCT compression workloads. Comparative analysis against five state-of-the-art approximate multiplier designs confirms that the proposed architecture provides the most favorable combined trade-off between energy efficiency and output quality, making it a compelling candidate for integration into energy-constrained embedded vision systems. Future work will extend the proposed architecture to 16-bit and 32-bit precision variants for floating-point image processing, explore reconfigurable accuracy modes for run-time quality-energy adaptation, and investigate the integration of the proposed multiplier into approximate systolic array accelerators for convolutional neural network inference at the edge.

REFERENCES

[1] M. B. R. Srinivas and K. Elango, "Optimized NCFET-based approximate multiplier for energy-aware image applications," Scientific Reports, vol. 15, no. 1, p. 17805, Oct. 2025. DOI: 10.1038/s41598-025-17805-5

[2] Anonymous, "Efficient and Low-Cost Approximate Multipliers for Image Processing Applications," AEU - International Journal of Electronics and Communications, vol. 138, 2021, Art. no. 153825. DOI: 10.1016/j.aeue.2021.153825

[3] S. Mittal, "A Survey of Techniques for Approximate Computing," ACM Computing Surveys, vol. 48, no. 4, pp. 1–33, 2016. DOI: 10.1145/2893356

[4] K. Prashanth and V. Nithish Kumar, "Area efficient approximate multiplier based on novel 4:2 compressors and error correction logic," Scientific Reports, vol. 15, Dec. 2025. DOI: 10.1038/s41598-025-31282-w

[5] M. J. Schulte and E. E. Swartzlander, "Truncated multiplication with correction constant," in Proc. IEEE Workshop VLSI Signal Process., 1993, pp. 388–396.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903010

[6] A. Momeni, J. Han, P. Montuschi, and F. Lombardi, "Design and Analysis of Approximate Compressors for Multiplication," IEEE Transactions on Computers, vol. 64, no. 4, pp. 984–994, 2015. DOI: 10.1109/TC.2014.2308214
[7] V. Mahalingam and N. Rangaswamy, "Improving Accuracy of Approximate Multipliers Using Segmentation," in Proc. ARITH, 2006, pp. 196–203.
[8] L. Dadda, "Some schemes for parallel multipliers," Alta Frequenza, vol. 34, no. 5, pp. 349–356, 1965.
[9] S. Deepsita, B. Kamatchi, and M. Potula, "Low power, high speed approximate multiplier for error resilient applications," Computers and Electrical Engineering, vol. 107, 2022. DOI: 10.1016/j.compeleceng.2022.108549
[10] A. G. M. Strollo, E. Napoli, D. De Caro, N. Petra, and G. Di Meo, "Comparison and Extension of Approximate 4-2 Compressors for Low-Power Approximate Multipliers," IEEE Transactions on Circuits and Systems I, vol. 67, no. 9, pp. 3021–3034, 2020.
[11] P. M. Kogge and H. S. Stone, "A Parallel Algorithm for the Efficient Solution of a General Class of Recurrence Equations," IEEE Transactions on Computers, vol. 22, no. 8, pp. 786–793, 1973.
[12] S. Venkatachalam and S.-B. Ko, "Design of Power and Area Efficient Approximate Multipliers," IEEE Transactions on Very Large Scale Integration Systems, vol. 25, no. 5, pp. 1782–1786, 2017.
[13] B. Parhami, Computer Arithmetic: Algorithms and Hardware Designs, 2nd ed. Oxford, UK: Oxford University Press, 2010.
[14] V. Gupta, D. Mohapatra, A. Raghunathan, and K. Roy, "Low-power digital signal processing using approximate adders," IEEE Transactions on Computer-Aided Design, vol. 32, no. 1, pp. 124–137, 2013.
[15] Anonymous, "An efficient approximate multiplier: Design, error analysis and application," AEU - International Journal of Electronics and Communications, vol. 178, 2024. DOI: 10.1016/j.aeue.2024.155217
[16] R. Sharma, A. Sungheetha, R. Gandhi K, S. Rani, G. Pradeep, and E. V, “Segmentation of medical liver image diseases using improved contextual convolutional neural network (COCOn) model,” in *Proc. ICICEC*, Davangere, India, 2024, pp. 1–5. DOI: 10.1109/ICICEC62498.2024.10809006.
[17] S. Deepsita, "Energy-efficient approximate constant carry compressors for image processing," AEU International Journal of Electronics and Communications, vol. 141, 2022.