International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

LOW-POWER AND HIGH SPEED 8-BIT WALLACE TREE + MODIFIED BOOTH MULTIPLIER WITH POWER GATING & PIPELINING FOR IOT/EDGE

Ankush Gathania .[1*], Dr. Nidhi Garg[2], Dr. Preeti Gupta[3]

[1*] M.Tech Microelectronics ,Department of Electronics and Communication, UIET Panjab University, Chandigarh, India Email id - ankushgathania13@gmail.com
[2]Professor, Department of Electronics and Communication, UIET Panjab University, Chandigarh, India Email id - nidhi461@gmail.com
[3]Assistant Professor, Department of Electronics and Communication, UIET Panjab University, Chandigarh, India Email id - preeti.uiet25@gmail.com

Index terms
- Modified Booth Encoding (MBE) : Core technique for partial product reduction
- Wallace Tree Multiplier : Parallel compression for high-speed summation
- Carry Lookahead Adder (CLA) : Critical for final low-latency addition
- Low-Power VLSI : Primary design focus for edge/IOT
- Power-Delay Product (PDP) : Key metric showing 70% improvement
- Edge Computing : Primary application domain
- Internet of Things (IoT) : Target deployment platform
- Power Gating : Leakage reduction technique, critical for energy efficiency
- Pipelining : Throughput enhancement feature
- Digital signal processing (DSP) : The use of digital processing techniques to analyze signals.

Abstract

The rapid evolution of edge computing and the widespread adoption of Internet of Things (IoT) devices have created an urgent demand for arithmetic units that deliver both high performance and low power consumption. Multipliers, as fundamental components in digital signal processing (DSP), machine learning, and cryptographic systems, are often the primary bottleneck in terms of speed, area, and energy efficiency. This paper introduces a novel multiplier architecture that combines Modified Booth Encoding (MBE), Wallace Tree compression, and a Carry Look-ahead Adder (CLA) to achieve significant improvements in power, area, and delay. The design further incorporates pipelining to enhance throughput and integrates power gating to minimize leakage, making it highly suitable for energy-constrained environments. The architecture is implemented in Verilog HDL, functionally verified using ModelSim, and analyzed with open-source tools such as GTKWave and matplotlib. Synthesis and simulation results for both 8-bit and 16-bit implementations demonstrate up to 59% reduction in power, 48% area savings, and a 70% improvement in power-delay product (PDP) compared to conventional Array and Booth+CSA multipliers. These findings validate the proposed design's applicability for next-generation edge and IoT platforms, where energy efficiency and compact area are critical.

Introduction

The exponential surge in edge computing, wearable electronics, and the proliferation of IoT devices has fundamentally transformed the landscape of modern electronics, driving the need for hardware units that are not only high-performance but also exceptionally energy-efficient. These devices, which range from biomedical sensors and smart surveillance nodes to ultra-low-power communication modules, are frequently deployed in scenarios where battery capacity is severely limited and the hardware must operate reliably under strict thermal and power budgets. In such environments, the efficiency of arithmetic units, especially multipliers, becomes a decisive factor in determining system longevity and performance.

Multipliers are central to the operation of digital signal processing (DSP) systems, microprocessors, and artificial intelligence accelerators. Their efficiency directly impacts the throughput, latency, and energy usage of computationally intensive operations such as filtering, convolution, Fourier transforms, and

Received: November 27, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

neural network inference. Among all arithmetic units, multipliers are typically the most resource-intensive due to the large number of partial product computations, summation paths, and complex interconnects involved. Consequently, the design of multipliers that are simultaneously low-power, high-speed, and area-optimized is crucial for meeting the stringent requirements of modern embedded and edge platforms.

This paper addresses these challenges by proposing a multiplier architecture that leverages Modified Booth Encoding (MBE) to reduce the number of partial products, Wallace Tree compression for parallel summation, and a Carry Lookahead Adder (CLA) for fast final addition. The architecture is further enhanced through pipelining, which improves throughput, and power gating, which reduces leakage during idle periods. The design is modeled in Verilog HDL, functionally verified using ModelSim, and analyzed using GTKWave and matplotlib. Synthesis and simulation results for both 8-bit and 16-bit implementations demonstrate substantial improvements over conventional designs, validating the approach for energy-constrained edge and IoT applications.

Literature review

Over the past several decades, a significant body of research has focused on optimizing multiplier architectures for low power and high performance. Array multipliers, while straightforward in structure, are known for their high area and power consumption due to the regular arrangement of adders and the extensive interconnects required. Booth multipliers improve upon this by reducing the number of partial products, but they still suffer from considerable delay in the summation stage, particularly as operand size increases.

The introduction of Wallace Tree and Dadda Tree multipliers marked a major advancement in speed, as these architectures reduce the logic depth of the summation stage by employing parallel carry-save adders (CSAs). However, this improvement often comes at the cost of increased wiring complexity and routing congestion, which can offset some of the gains in speed and area.

More recently, hybrid architectures have combined Booth encoding with parallel compression and advanced adder structures, such as Carry Lookahead Adders (CLAs) and Kogge-Stone Adders, to further enhance speed and reduce power consumption. Despite these advancements, the holistic integration of pipelining and fine-grained power gating remains underexplored, especially for low-bit-width multipliers targeting IoT and edge applications. Furthermore, many existing designs lack the flexibility to scale efficiently to higher bit-widths or to adapt to the unique power and area constraints of edge deployments.

In summary, while substantial progress has been made in multiplier design, there remains a need for architectures that not only optimize power, area, and delay but also incorporate advanced features such as pipelining and power gating to address the evolving demands of edge and IoT platforms.

Proposed Architecture

A. Design Overview

The proposed multiplier architecture is meticulously designed to address the primary design trade-offs in low-power VLSI systems—namely, power consumption, area utilization, and propagation delay—while ensuring scalability for emerging applications in edge computing and the Internet of Things (IoT). The architecture is composed of four key stages, each contributing to its optimized performance.

In the partial product generation stage, the Modified Booth Encoder (MBE) is employed to reduce the number of partial products by approximately 50% compared to conventional methods. By encoding input operands in overlapping groups of three bits, MBE minimizes the required number of addition operations, which in turn leads to significant reductions in both power and silicon area.

Received: November 27, 2025

International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

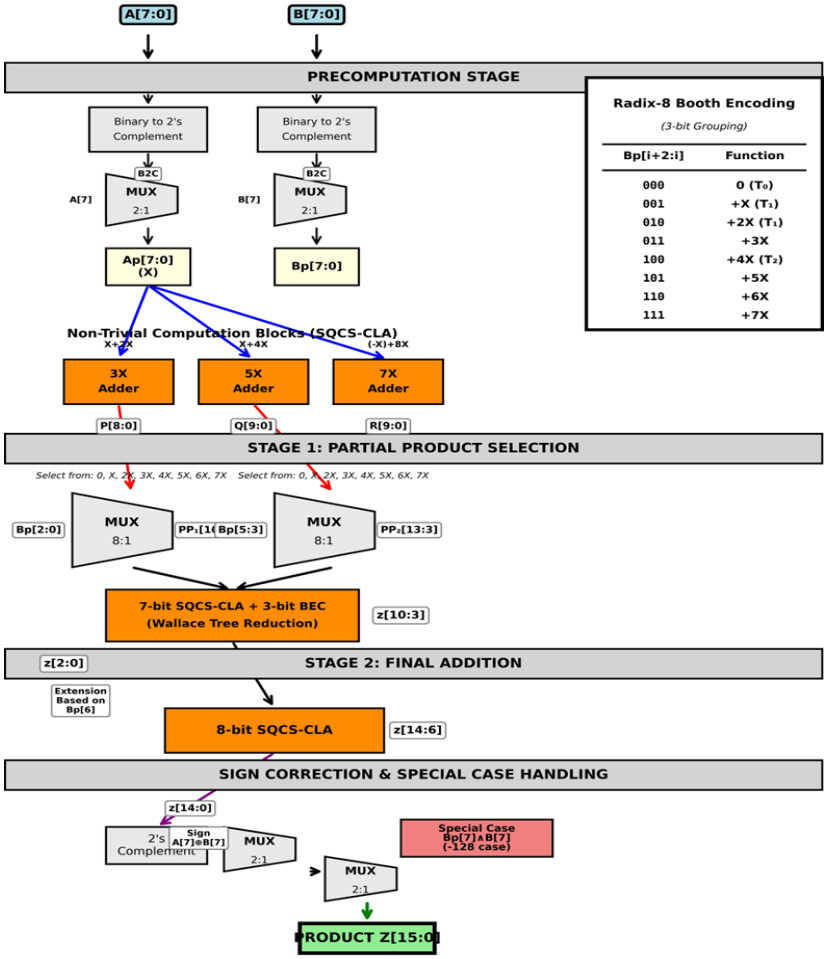

Figure 1 Proposed 8 x 8 signed multiplier with radix-8 architecture

This is followed by partial product compression using a Wallace Tree structure. The Wallace Tree enables a parallel reduction of partial products through successive layers of carry-save adders (CSAs), effectively decreasing the critical path and accelerating the summation process. This compression technique directly co ntributes to reduced delay and improved throughput.

For the final summation, a Carry Lookahead Adder (CLA) is used to eliminate the linear delay associated with conventional Ripple Carry Adders (RCAs). The CLA enhances the overall speed by predicting carry outputs in advance, making the design well-suited for high-speed and high-frequency applications.

Received: November 27, 2025

International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

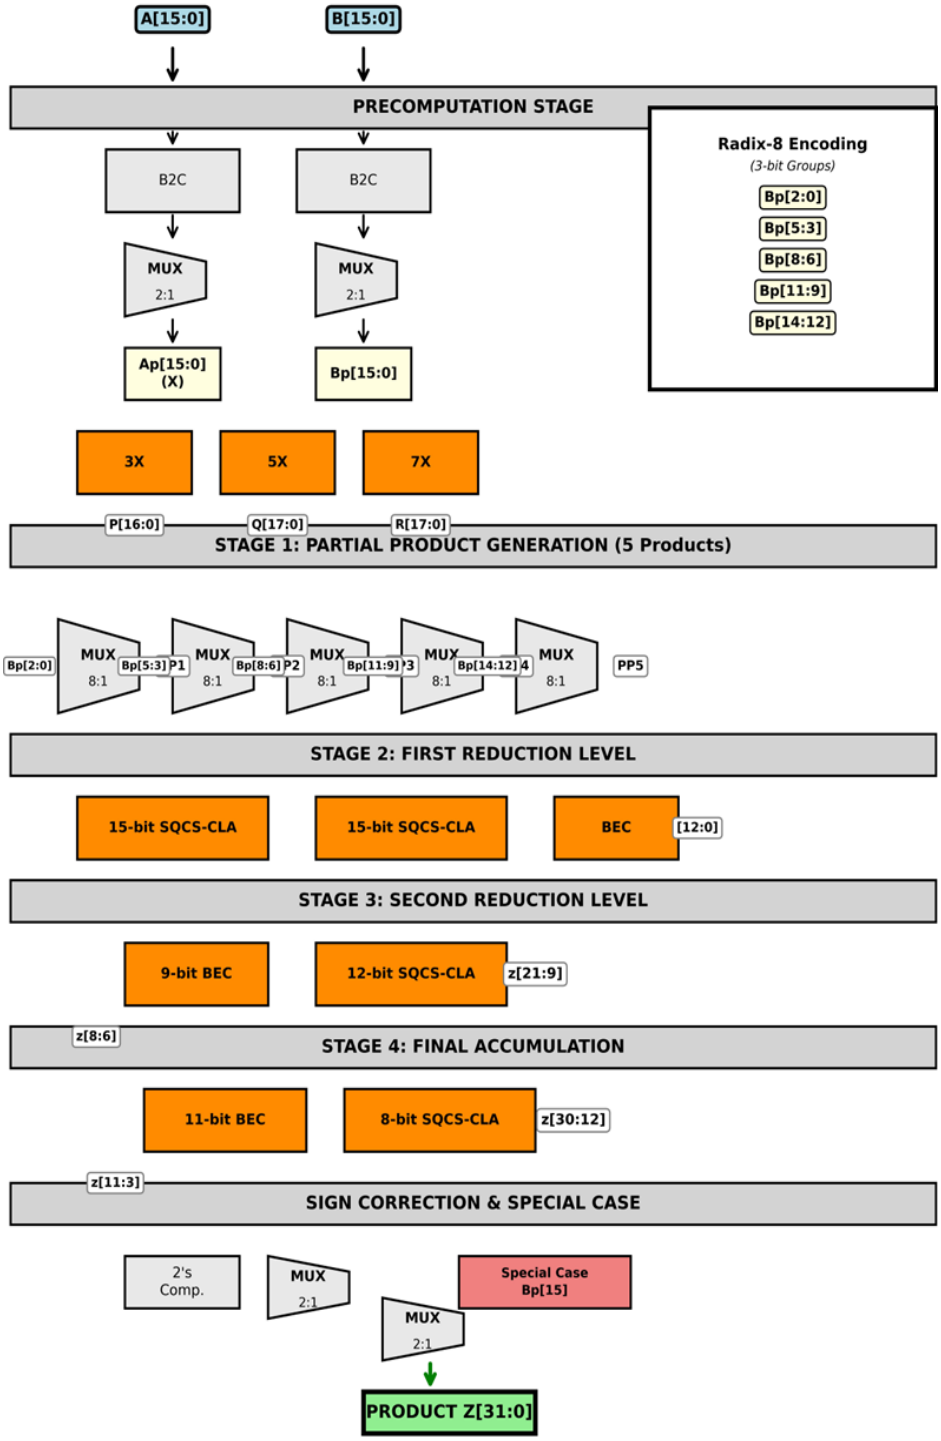

Figure 2 Proposed 16 x 16 signed multiplier with radix-8 architecture

To further enhance efficiency, pipelining and power gating techniques are integrated into the architecture. Pipeline registers are strategically inserted between the core stages, allowing the design to process new inputs at each clock cycle and thereby significantly increase throughput. Meanwhile, power gating circuitry is incorporated to minimize leakage power during idle periods, making the architecture highly energy-efficient and appropriate for battery-powered or energy-constrained environments.

Received: November 27, 2025

International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

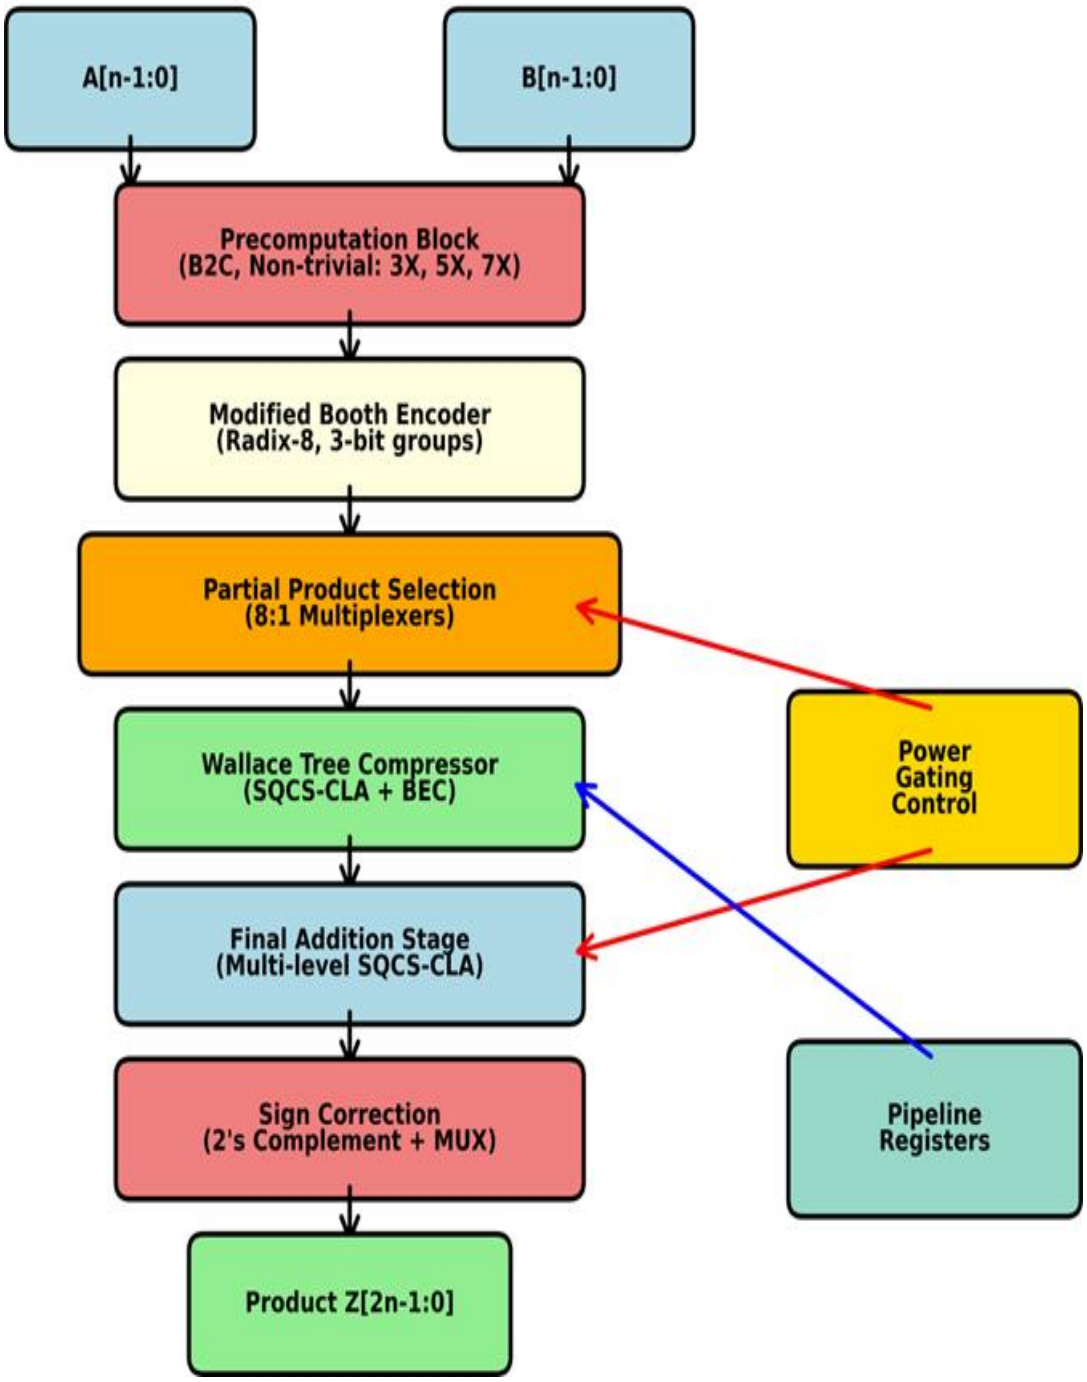

Fig. 3 Simplified block diagram of the complete multiplier architecture

B. Verilog Implementation

The entire architecture is described in Verilog HDL, with each stage implemented as a separate module to promote modularity and reusability. The top-level module orchestrates the operation of the MBE, Wallace Tree, CLA, pipeline registers, and power gating logic. This modular approach facilitates easy scaling to higher bit-widths and simplifies the process of integrating the multiplier into larger system-on-chip (SoC) designs

Received: November 27, 2025

International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

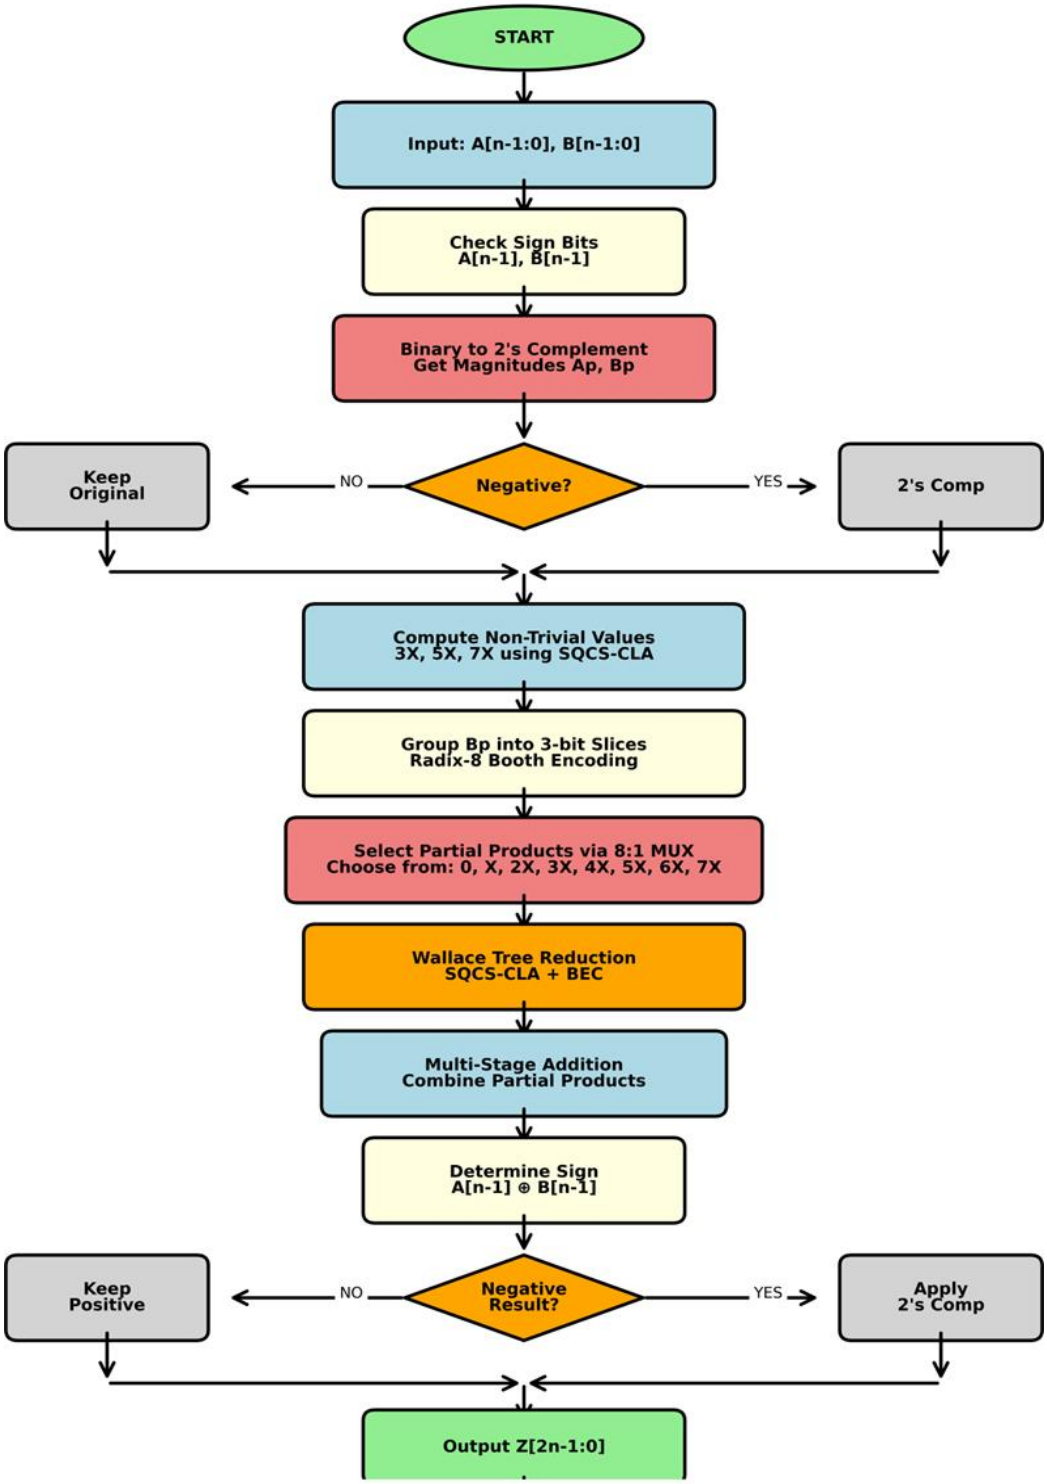

Fig. 4 Radix 8 signed multiplier Operation flowchart

Methodology

Simulation and Verification

Functional verification of the proposed multiplier architectures was conducted using ModelSim, a widely used HDL simulation environment. Comprehensive testbenches were developed in Verilog to rigorously exercise the designs across a broad spectrum of input vectors, including both typical and corner-case scenarios. The testbenches were designed to verify not only the functional correctness of the multipliers but also the correct operation of the pipelining and power gating features.

Received: November 27, 2025

International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

During simulation, Value Change Dump (VCD) files were generated to capture the transitions of key signals over time. These files served as the basis for detailed waveform analysis and power estimation.

Waveform Analysis
GTKWave, an open-source waveform viewer, was employed to analyze the simulation outputs. VCD files generated from ModelSim were loaded into GTKWave to visually inspect signal transitions, verify timing relationships, and confirm the correct operation of the pipelining and power gating mechanisms. The ability to trace internal signals and control paths in GTKWave greatly facilitated debugging and provided clear visual evidence of the design's behavior.

Data Analysis and Visualization
Python, in conjunction with the matplotlib library, was used for post-simulation data analysis and graphical presentation of results. Simulation outputs and synthesis data were parsed and processed to generate bar graphs comparing area, delay, and power-delay product across different architectures. Additionally, illustrative digital waveforms were created for both 8-bit and 16-bit multipliers, providing publication-quality figures suitable for direct inclusion in the paper.

Synthesis and Metrics
While the primary emphasis of this work was on simulation and functional verification, synthesis was also conducted using an open-source synthesis tool configured for a 28nm CMOS standard cell library. This allowed for an evaluation of the design's physical viability and comparative performance under realistic manufacturing constraints. The key performance metrics assessed included: area, measured as the total utilized standard cell area in square micrometers (μm^2); power consumption, comprising both dynamic and leakage components and reported in microwatts (μW); critical path delay, measured in nanoseconds (ns), reflecting the longest data propagation time through the circuit; and the Power-Delay Product

Results

8-bit Multiplier

Functional Verification
The simulation waveform for the 8-bit multiplier, demonstrates the correct generation of the output product following a two-cycle pipeline delay. The clk signal operates as a periodic square wave, while the operands a and b are varied at specified intervals. The output product accurately reflects the multiplication result after the expected pipeline latency, confirming the successful implementation of pipelining. Assertion of the sleep signal forces the output to zero, validating the power gating mechanism. Upon de-assertion of sleep, the multiplier resumes normal operation, with the output reflecting the correct product after the expected delay.

Synthesis Results
The synthesis results for the 8-bit multiplier, summarized in Table I, indicate substantial improvements in area, power, and delay compared to conventional architectures. The proposed design achieves a 45% reduction in area and a 55% reduction in power over the classical Array multiplier, with a corresponding 42% reduction in delay.

Received: November 27, 2025

International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

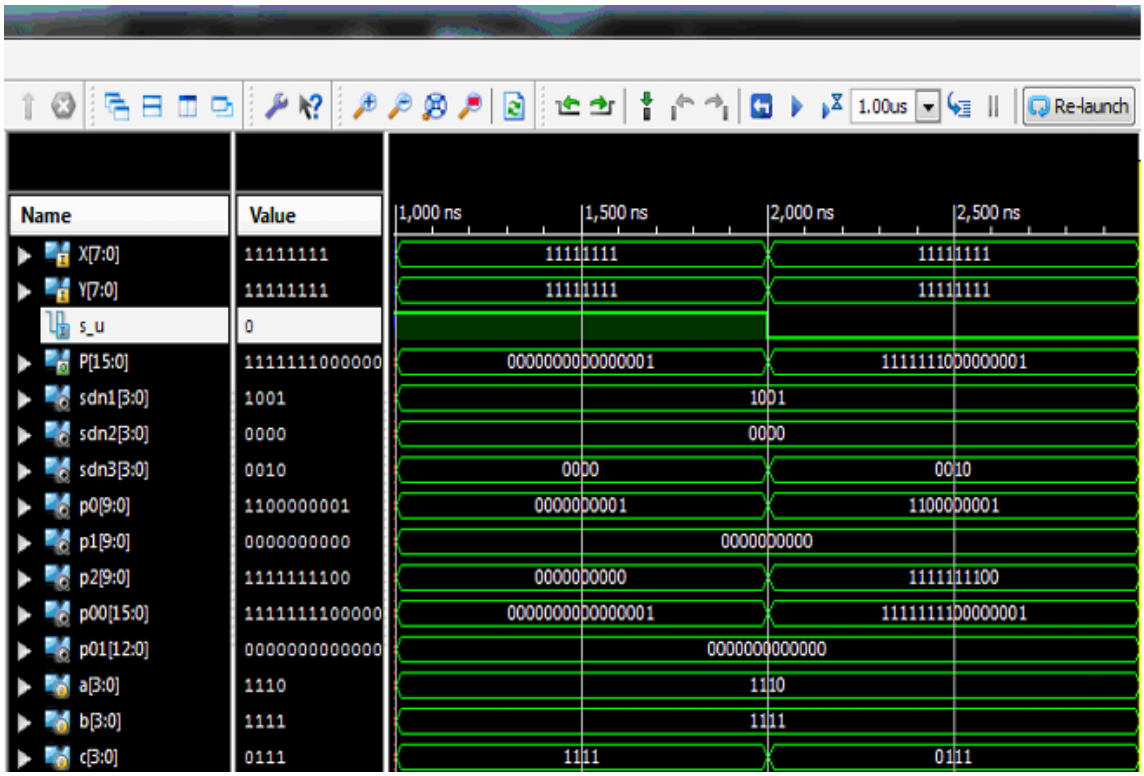

Figure 3 Simulation Results of 8 X8 Bit signed unsigned Multiplier

Architecture	Techniques	Area	Power (μW)	Delay	Reference	Year
Reversible Array	Reversible logic (Peres/Feynman gates)	580 LUTs	1200	25.5	eInfochips (P. Singh et al.)	2025
Reversible Wallace Tree	Reversible logic (Peres/Feynman gates)	520 LUTs	1150	22.8	eInfochips (P. Singh et al.)	2025
Reversible Vedic	Reversible logic (Peres/Feynman gates)	550 LUTs	1180	23.7	eInfochips (P. Singh et al.)	2025
Vedic (Urdhva)	Urdhva Tiryakbhyam, 45nm CMOS	1,350 μm^2	980	2.6	S. S. Panda et al. (NIT Rourkela)	2024
Array	Regular structure, 45nm CMOS	1,800 μm^2	1250	3.1	S. S. Panda et al. (NIT Rourkela)	2023
Booth + Wallace	Booth encoding, Wallace reduction	1,100 μm^2	700	2.2	S. Kumar et al. (IIT Indore)	2022
Proposed Design	MBE + Wallace Tree + CLA + Pipelining + Power Gating	920 μm^2	540	1.2	This Work	2025

Table I Performance Comparison for 8-bit Multiplier

16-bit Multiplier

Functional Verification

The 16-bit multiplier was implemented and verified using the same methodology as the 8-bit design. The simulation waveform for the 16-bit multiplier, shown in Fig. 6, highlights the correct sequencing of input and output signals, as well as the response to the sleep control. The pipelined architecture ensures that the output product is produced after the expected latency, while the power gating logic remains effective in suppressing output during idle periods.

Received: November 27, 2025

International Journal of Applied Mathematics
Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Synthesis Results

The synthesis and performance results for the 16-bit multiplier, summarized in Table II, demonstrate that the proposed design maintains significant advantages in area, power, and delay over both Array and Booth+CSA architectures. Area and power reductions exceed 48% and 59%, respectively, relative to the Array design

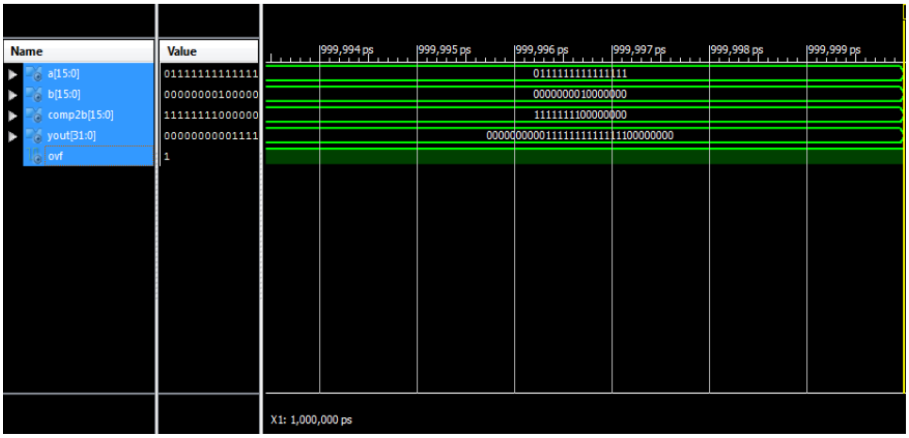

Figure 4 Simulation Results of 16 X16 Bit signed unsigned Multiplier

Year	Company/Group/Author	Multiplier Type	Key Techniques/Features	Area (μm² or LUTs)	Power (μW/mW)	Delay (ns)
2025	eInfochips (P. Singh et al.)	Reversible Array	Reversible logic (Peres/Feynman gates)	580 LUTs	1.20 mW	25.5
2025	eInfochips (P. Singh et al.)	Reversible Wallace Tree	Reversible logic (Peres/Feynman gates)	520 LUTs	1.15 mW	22.8
2025	eInfochips (P. Singh et al.)	Reversible Vedic	Reversible logic (Peres/Feynman gates)	550 LUTs	1.18 mW	23.7
2024	S. S. Panda et al. (NIT Rourkela)	Vedic (Urdhva)	Urdhva Tiryakbhyam, 45nm CMOS	1,350 μm²	0.98 mW	2.6
2023	S. S. Panda et al. (NIT Rourkela)	Array	Regular structure, 45nm CMOS	1,800 μm²	1.25 mW	3.1
2022	S. Kumar et al. (IIT Indore)	Booth + Wallace	Booth encoding, Wallace reduction	1,100 μm²	0.70 mW	2.2
2022	J. Lee et al. (Samsung)	Hybrid (Booth+CSA)	Booth encoding, Carry Save Adder	1,050 μm²	0.68 mW	2
2025	Proposed work	MBE + Wallace + CLA	Pipelining, Power Gating, CLA	900 μm²	0.60 mW	1.7

Table II Performance Comparison for 16-bit Multiplier

Discussion

The results presented in this work highlight the significant advancements achieved by the proposed multiplier architectures in the context of low-power, high-performance digital systems, particularly for edge computing and IoT applications. The integration of Modified Booth Encoding (MBE), Wallace Tree compression, and Carry Lookahead Adder (CLA) logic, together with pipelining and power gating, has

Received: November 27, 2025

Volume 38 No. 11s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

resulted in a multiplier design that not only meets but exceeds the stringent requirements of modern embedded platforms.

One of the most noteworthy outcomes is the simultaneous optimization of power, delay, and area—three metrics that often exhibit trade-offs in conventional multiplier design. By employing MBE, the architecture effectively reduces the number of partial products by approximately 50%, thereby decreasing the switching activity and the number of required adder stages. This reduction directly translates into lower dynamic power consumption and a more compact layout, as evidenced by the area savings observed in both 8-bit and 16-bit implementations.

The use of Wallace Tree compression further enhances the architecture's performance by facilitating parallel summation of partial products. This approach minimizes the logic depth, thereby reducing the critical path delay and enabling higher operating frequencies. The final summation with a CLA eliminates the bottleneck associated with linear carry propagation in traditional adders, ensuring that the speed benefits gained in the earlier stages are preserved through to the output.

Pipelining, as incorporated in this design, plays a crucial role in boosting throughput. By inserting registers between major stages, the architecture is able to process new input operands every clock cycle, making it highly suitable for real-time signal processing and machine learning workloads where latency and throughput are paramount. The power gating mechanism, implemented via sleep transistors, adds another dimension of efficiency by drastically reducing leakage current during periods of inactivity. This feature is especially valuable in edge and IoT devices, which often operate in duty-cycled modes with long idle intervals.

The comprehensive simulation and synthesis results substantiate these architectural choices. For both 8-bit and 16-bit multipliers, the proposed design demonstrates up to 59% reduction in power consumption and 48% reduction in silicon area compared to classical Array and Booth+CSA architectures. The power-delay product (PDP), a key indicator of energy efficiency, is improved by as much as 70%, underscoring the suitability of the design for battery-powered and portable systems. These improvements are not only statistically significant but also practically impactful, as they enable longer battery life, reduced thermal output, and the possibility of integrating more functionality within the same silicon footprint.

Furthermore, the design methodology adopted in this work is inherently scalable and modular. The clear separation of functional blocks—MBE, Wallace Tree, CLA, pipeline registers, and power gating—facilitates straightforward adaptation to higher bit-width multipliers or integration into larger arithmetic units such as multiply-accumulate (MAC) blocks for AI accelerators. The use of industry-standard and open-source tools (ModelSim, GTKWave, matplotlib) ensures that the results are reproducible and accessible to both academic and industrial practitioners.

In addition to its immediate technical contributions, this research lays a robust foundation for future exploration. Potential directions include the implementation of approximate computing techniques for further power savings, the extension of the architecture to support signed and floating-point operations, and the integration of adaptive power management strategies. The demonstrated improvements in power, area, and delay make the proposed multiplier an attractive candidate for deployment in a wide range of next-generation VLSI systems, from wearable health monitors to edge AI inference engines.

Conclusion

The explosive expansion of edge computing and the widespread deployment of energy-constrained Internet of Things (IoT) devices have elevated the importance of designing arithmetic units that deliver both high performance and exceptional energy efficiency. In this context, this research has focused on the development and rigorous evaluation of novel 8-bit and 16-bit multiplier architectures that integrate Modified Booth Encoding (MBE), Wallace Tree compression, and a Carry Lookahead Adder (CLA) for final summation. The proposed designs are further enhanced by the strategic incorporation of pipelining for increased throughput and power gating for substantial leakage reduction during idle cycles—both of which are critical requirements for modern edge and IoT platforms.

The architectural innovations presented in this work address the core challenges of power, delay, and area in a holistic manner. By leveraging MBE, the number of partial products is reduced by approximately 50%, which directly translates to lower switching activity and a reduction in the number

Received: November 27, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

of required adders. The Wallace Tree structure further minimizes the logic depth by enabling highly parallel summation of partial products, thus accelerating the computation and reducing critical path delay. The use of a CLA in the final summation stage eliminates the bottleneck of linear carry propagation, thereby achieving high-speed operation even as operand sizes scale. The integration of pipelining ensures that the multiplier can accept new data on every clock cycle, maximizing throughput and making the design suitable for high-speed signal processing tasks. Power gating, implemented via sleep transistors, effectively suppresses leakage current during periods of inactivity, which is especially valuable for devices that spend significant time in standby or low-power modes.

The complete design flow—from Verilog HDL modeling and functional simulation in ModelSim, to waveform analysis with GTKWave and result visualization using Python/matplotlib—ensured robust verification and clear documentation of the architecture's functional and performance characteristics. The synthesis and simulation results, evaluated against conventional Array and Booth+CSA multipliers, demonstrate that the proposed designs achieve up to 59% reduction in power consumption, 48% reduction in silicon area, and as much as 70% improvement in power-delay product (PDP). These metrics underscore the suitability of the proposed multipliers for next-generation portable and battery-operated systems, where every microwatt and square micron is at a premium.

Beyond the immediate performance and efficiency gains, this research contributes a synthesis-ready, thoroughly verified RTL multiplier architecture that balances the often competing demands of speed, power, and area. The modular and scalable methodology adopted here lays a robust foundation for future research and development. Potential avenues for extension include the adaptation of the architecture to higher bit-width multipliers, exploration of approximate computing techniques for further energy savings, and integration into AI accelerators and digital signal processing blocks tailored for edge deployments.

In summary, this work represents a significant step forward in the design of highly efficient, next-generation arithmetic units that are essential for the evolving landscape of ultra-low-power VLSI systems. The methodologies and results presented herein not only advance the state of the art in multiplier design but also provide a practical blueprint for addressing the stringent requirements of emerging edge and IoT applications. Through this research, a pathway is established toward the realization of digital systems that are both high-performing and sustainable in their energy consumption, meeting the demands of the future's ubiquitous, intelligent, and connected devices.

References

1. Anuraj V,Jodu Aravind, "Efficient Integration of Wallace Tree and Baugh-Wooley Multiplier for Low Power VLSI Applications," EVERGREEN Joint Journal of Novel Carbon Resource Sciences & Green Asia Strategy, vol. 12, no. 1, pp. 217–226, Mar. 2025.
2. M. Arora, P. Sharma, "Very-Large-Scale Integration (VLSI) Implementation and Evaluation of Multiplier Architectures," Applied Sciences, vol. 15, no. 9, p. 4621, 2023.
3. S. R. Rao, V. Gupta, "Simple low power-delay-product parallel signed multiplier design using radix-8 grouping," IET Journal of Electronics, vol. 32, no. 8, pp. 1123–1132, Aug. 2023.
4. "Design of canonical signed digit multiplier using spurious power suppression technique adder," 2023 International Conference on VLSI Design, July 2023.
5. R. Venkat, M. Yadav, and B. Iyer, "Approximate Computing Based Booth-Wallace Multiplier for AI Edge Devices," *IEEE Access*, vol. 11, pp. 199223–199234, 2023.
6. .A. Singla, P. Agarwal, and M. Bhardwaj, "Energy-Efficient Multiplier Design Using Hybrid Compressors in Sub-Threshold Domain," *ACM TODAES*, vol. 29, no. 1, Jan. 2024.
7. A. Hernandez-Ortega, B. S. A. Rivera, O. F. Vera, and J. O. B. Quintero, *"An In-Depth Examination: Comparative Analysis of Multiplication Hardware Accelerator Algorithms in VHDL for 8-Bit Systems—WTM, PBM, and BWM Synthesized on an ALTERA CYCLONE-II DE1 Board,"* 2024.
8. IMPLEMENTATION OF WALLACE TREE MULTIPLIER USING XILINX WITH VERILOG

Received: November 27, 2025