IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

Verilog-Based FPGA Realization of KINA: A Karatsuba-Initiated Accelerator for Ring-Binary-LWE Post-Quantum Cryptography

Arukonda Suresh[1] and K Naresh[2]
PG Scholar, Dept. of ECE[1]
Assistant Professor, Dept. of ECE[2]
Sree Dattha Institute of Engineering and Science, Hyderabad, Telangana, India.
suresharukonda01@gmail.com[1], balu.482@gmail.com[2]

Abstract: *The transition toward post-quantum cryptography requires hardware architectures that can execute lattice-based polynomial arithmetic with low latency and moderate resource cost. This paper presents a Verilog-based realization and evaluation of the Karatsuba Initiated Novel Accelerator (KINA) concept for Ring-Binary Learning With Errors based encryption (RBLWE-ENC). The core computation is polynomial multiplication in Z_q/(x^n+1), where one operand is binary and the other is integer-valued. Conventional schoolbook multiplication incurs O(n^2) coefficient operations, while the RBLWE parameter setting is not naturally suited to direct Number Theoretic Transform acceleration. The implemented architecture therefore applies a one-level Karatsuba decomposition that forms three half-size products, T_L, T_H, and T_M, and recombines them through a linear-combination datapath. Circular shift registers support column-wise accumulation, binary point-wise products are realized with simple logic, and the two-bit middle operand is handled by a compact multiplexer-based multiplier. A finite-state controller coordinates clear, load, compute, output, and done phases. For the educational RTL prototype, n=8 and q=256 were used for simulation and synthesis. XSim verification completed successfully with a final visible cipher coefficient of 146 (0x92) and a decoded message bit of 1. RTL elaboration and synthesis confirmed full module connectivity and device mapping, while the reported post-synthesis on-chip power estimate was 0.089 W. The study also analyzes the scalable parallelism factor u, for which multiplication latency decreases from n/2 to n/(2u) cycles at the cost of additional datapath resources. The results demonstrate the functional feasibility of a compact Verilog KINA implementation and provide a practical path toward larger n=256 and n=512 FPGA realizations for lightweight post-quantum security.*

Keywords: post-quantum cryptography, RBLWE, Ring-Binary-LWE, Karatsuba multiplication, KINA, FPGA, Verilog HDL, lattice-based cryptography, polynomial multiplier, hardware accelerator.

I. INTRODUCTION

Public-key cryptography underpins authentication, key establishment, secure firmware updates, cloud access, encrypted communication, and digital signatures. Widely deployed systems based on RSA and elliptic-curve cryptography derive their security from integer factorization and discrete logarithm problems. Shor's quantum algorithm changes this security assumption because a sufficiently capable quantum computer can solve both problem classes in polynomial time [3]. This has motivated the development of post-quantum cryptography (PQC), in which algorithms are designed to resist known classical and quantum attacks [2], [4].

Lattice-based cryptography is a leading PQC family because it offers practical encryption, key encapsulation, and digital-signature constructions while supporting efficient implementation in software and hardware [5], [10], [11]. The

DOI: 10.48175/IJARSCT-38401

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

NIST standardization effort has further accelerated deployment: FIPS 203 standardizes ML-KEM, FIPS 204 standardizes ML-DSA, and FIPS 205 standardizes SLH-DSA [6]-[9]. Although the RBLWE scheme considered in this work is not itself one of the finalized NIST standards, it belongs to the broader research line of lattice-based cryptography and is attractive for hardware study because binary coefficients simplify the dominant arithmetic operations.

Ring-Binary Learning With Errors based encryption (RBLWE-ENC) is a lightweight ring variant in which binary secret/error structures and a small modulus can reduce arithmetic complexity. The reference parameter sets discussed in the source design include (n,q)=(256,256) and (512,256) [1]. The major computation in key generation, encryption, and decryption is polynomial multiplication over the quotient ring Z_q/(x^n+1). One multiplicand is binary, while the other contains integer coefficients. This asymmetry creates an opportunity for specialized hardware because multiplication by a binary coefficient can be realized as conditional selection rather than as a full general-purpose multiplier.

The conventional hardware baseline uses schoolbook polynomial multiplication. Its implementation is conceptually simple, but it requires O(n^2) coefficient interactions and therefore produces increasing latency and area-time cost as n grows. Transform-domain approaches such as the Number Theoretic Transform (NTT) are highly effective in several lattice schemes, but the RBLWE parameter choices considered here are not naturally favorable for a direct NTT implementation [1]. A hardware-oriented alternative is therefore needed that improves speed without introducing the modulus and transform overhead of an NTT datapath.

This paper develops a Verilog realization of the Karatsuba Initiated Novel Accelerator (KINA) concept originally introduced for RBLWE hardware acceleration [1]. The objective is not to claim invention of the KINA algorithm; rather, the contribution is a compact RTL realization, simulation-oriented architecture, verification flow, and implementation analysis that can be reproduced in an FPGA design environment. The design applies a controlled one-level Karatsuba decomposition, uses circular shift registers for column-wise accumulation, computes three sub-products in parallel, and performs ring-aware recombination through a linear-combination unit. The same arithmetic core supports the dominant operations of key generation, encryption, and decryption.

The principal contributions of this implementation study are: (i) a clear Verilog module hierarchy covering input processing, circular shift registers, point-wise multiplication, accumulation, linear combination, threshold decoding, and FSM control; (ii) a small n=8, q=256 proof-of-concept configuration suitable for waveform-level verification before scaling; (iii) functional verification in XSim with observable ciphertext and decoded-bit outputs; (iv) RTL elaboration, synthesis mapping, and power-estimation evidence from Vivado; and (v) an analysis of the parallelism parameter u that links resource cost to multiplication latency. The remainder of the paper presents background and related work, the RBLWE arithmetic model, the KINA derivation, the RTL architecture, implementation methodology, results, security considerations, and future scaling directions.

II. BACKGROUND AND RELATED WORK

A. Learning With Errors and Ring Variants

The Learning With Errors (LWE) problem introduced a powerful foundation for lattice-based cryptography by hiding a secret vector within noisy linear equations [10]. Ring-LWE moves the computation into polynomial quotient rings, reducing key sizes and providing efficient structured arithmetic [11]. Research on small and binary secret/error distributions has explored both hardness and implementation advantages [21]-[24]. In hardware, a binary coefficient is especially useful because a product b_i d_j, with b_i in {0,1}, can be implemented by gating or multiplexing d_j rather than invoking a conventional multiplier.

RBLWE extends this implementation-oriented direction by using binary structures in a ring setting. Its key generation, encryption, and decryption phases repeatedly invoke polynomial multiplication, polynomial addition, and threshold decoding. The arithmetic therefore maps naturally to FPGA logic elements, registers, adders, and multiplexers. The design challenge is to preserve this lightweight structure while reducing the O(n^2) latency of direct schoolbook multiplication.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

B. Hardware Acceleration of Lattice Cryptography

A substantial body of work has investigated lattice-based cryptographic processors on reconfigurable hardware. Early implementations emphasized practical public-key encryption on FPGA [15], low-cost and area-efficient accelerators [16], and compact ring-LWE cryptoprocessors [17]. High-speed polynomial multipliers were developed for ring-LWE and related homomorphic-encryption workloads [12], while later work studied standard-lattice encryption and key encapsulation on embedded devices [14], [18].

For binary ring-LWE and related lightweight constructions, research has addressed power side-channel countermeasures [26], edge-oriented cryptoprocessors [27], low-complexity hybrid-field multiplication [28], LUT-like finite-field arithmetic [29], area/power optimized architectures [30], and compact arithmetic engines [31]-[35]. These studies show that the hardware cost of lattice cryptography is strongly influenced by polynomial multiplication, coefficient representation, memory movement, and the selected degree of parallelism.

The KINA architecture [1] specifically targets the RBLWE multiplication structure. Rather than attempting a full recursive Karatsuba tree, it uses a one-level decomposition that trades three half-size products for the four half-size products that would arise from a direct split. This controlled use of Karatsuba is important because repeated recursion can enlarge intermediate coefficient widths and increase the number of adders, registers, and routing resources. The architecture therefore seeks a practical FPGA balance between algebraic reduction and datapath overhead.

Table I summarizes the most directly relevant hardware studies.

Reference	Primary contribution	Relevance to this work
Aysu et al. [26]	Binary ring-LWE hardware with power side-channel countermeasures	Demonstrates security-aware hardware implementation.
Ebrahimi et al. [27]	PQC cryptoprocessors for edge and resource-constrained IoT devices	Establishes relevance of lightweight PQC accelerators for IoT.
He et al. [28]	Low-complexity polynomial multiplication over hybrid fields	Reduces arithmetic cost for binary ring-LWE.
Xie et al. [29]	Lookup-table-like finite-field arithmetic	Explores multiplier simplification for binary ring-LWE.
Shahbazi and Ko [30]	Area- and power-efficient post-quantum cryptosystem	Emphasizes resource-constrained implementation.
Xie et al. [31]	Efficient finite-field arithmetic hardware	Improves binary ring-LWE arithmetic efficiency.
Lucas et al. [32]	Lightweight binary ring-LWE accelerator	Targets compact hardware acceleration.
Imaña et al. [33]	Inverted binary ring-LWE arithmetic	Develops alternative arithmetic organization.
Xu et al. [34]	Robust binary ring-LWE decryption hardware	Focuses on decryption correctness and robustness.
He et al. [35]	Compact RBLWE FPGA accelerators	Provides compact FPGA baselines.
He et al. [1]	KINA with Karatsuba-initiated multiplication	Introduces the architecture realized in Verilog in this study.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

III. RBLWE-ENC COMPUTATIONAL MODEL

A. Polynomial Ring

Let the computations take place in the quotient ring R_q=Z_q[x]/(x^n+1). For an n-coefficient binary polynomial B(x) and an n-coefficient integer polynomial D(x), the dominant operation is the negacyclic product

$$T(x) = B(x)\, D(x) \bmod (x^n + 1) \quad (1)$$

Because x^n=-1 in the quotient ring, terms whose degree is greater than or equal to n are folded back into the lower coefficient range with a sign inversion. This negacyclic reduction is a central part of the recombination hardware because the raw product of two degree-(n-1) polynomials can contain degrees up to 2n-2.

The schoolbook implementation evaluates all coefficient interactions and then applies the modular folding. Its arithmetic complexity grows quadratically with n. For n=256 or 512, this produces a sizeable number of multiply-accumulate operations. Although the binary nature of B reduces individual multiplication cost, the number of coefficient interactions still remains O(n^2).

B. Key Generation, Encryption, and Decryption

The RBLWE-ENC flow contains key generation, encryption, and decryption. Let a be a public polynomial parameter. During key generation, Alice generates random binary polynomials r_1 and r_2; r_2 is retained as the secret key and the public key polynomial p is formed as

$$p = r\_1 - a\, r\_2 \quad (2)$$

For encryption, a plaintext message m is encoded into a coefficient representation m~ and three binary error polynomials e_1, e_2, and e_3 are introduced. The two ciphertext components are

$$c\_1 = a\, e\_1 + e\_2 \quad (3)$$

$$c\_2 = p\, e\_1 + e\_3 + m\sim \quad (4)$$

Decryption combines c_1 with the secret key and c_2, then applies a coefficient threshold decoder:

$$m = decode(c\_1\, r\_2 + c\_2) \quad (5)$$

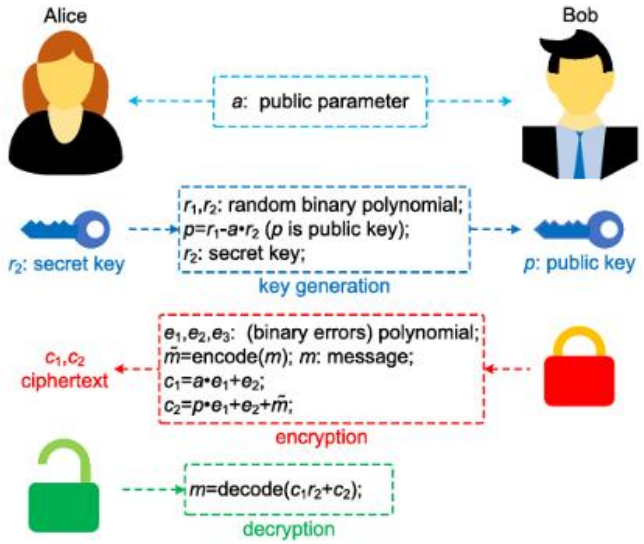

Fig. 1. Major phases of RBLWE-ENC.

Fig. 1. Major phases of RBLWE-ENC (adapted from the KINA architecture source [1]).

The exact encoding and thresholding details depend on the RBLWE formulation, but the hardware observation is consistent across phases: the same binary-by-integer polynomial multiplication engine can be reused, with phase-

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

specific additions and output handling. Figure 1 summarizes the cryptographic flow used as the basis of the RTL architecture.

IV. KARATSUBA-INITIATED MULTIPLICATION

A. One-Level Decomposition

The central acceleration strategy is a one-level Karatsuba decomposition. The even- and odd-indexed coefficients of each n-term polynomial are separated into lower and higher half-size polynomials. In compact notation, the operands can be written as

$$B = B\_L + x\, B\_H, \quad D = D\_L + x\, D\_H \tag{6}$$

Substitution into the product gives

$$BD = B\_L\, D\_L + x(B\_L\, D\_H + B\_H\, D\_L) + x^2\, B\_H\, D\_H \tag{7}$$

A direct split would require four half-size products. Karatsuba replaces the two cross-products by a single middle product. Define

$$T\_L = B\_L\ D\_L, \quad T\_H = B\_H\ D\_H, \quad T\_M = (B\_L + B\_H)(D\_L + D\_H) \tag{8}$$

Since T_M=T_L+T_H+B_LD_H+B_HD_L, the cross-term becomes T_M-T_L-T_H. After grouping terms, the recombination used by the KINA datapath is

$$T = (1 - x)T\_L + (x^2 - x)T\_H + x\, T\_M \tag{9}$$

The arithmetic significance is that one n-size multiplication is transformed into three n/2-size products. At the operation-count level, the multiplication work becomes approximately $3n^2/4$ instead of n^2, excluding the lower-cost additions, coefficient selection, and ring reduction. This is the reason KINA can improve the area-time balance without the transform setup required by an NTT implementation.

A fully recursive Karatsuba tree is deliberately avoided. In this RBLWE setting, the middle polynomial B_M=B_L+B_H contains values 0, 1, or 2 rather than only binary coefficients. Additional recursion would continue to increase intermediate widths and adder complexity. One level therefore provides a controlled reduction that remains hardware friendly.

B. Binary and Two-Bit Point-Wise Products

The decomposed terms T_L and T_H multiply binary coefficients by integer coefficients. Hardware can therefore realize each product with an AND-mask or conditional selection. The middle term is slightly more complex because B_M contains two-bit values. Instead of a general multiplier, the datapath selects one of three values: 0, d_M,i, or 2d_M,i. The factor two is generated by a left shift, so the middle path can still avoid a full multiplier. This property is one of the principal reasons the KINA decomposition remains efficient for Ring-Binary-LWE.

V. PROPOSED VERILOG KINA ARCHITECTURE

The Verilog architecture follows the three functional blocks of the source KINA design: input processing, parallel sub-product computation, and linear recombination. A controller, counters, output logic, and decoder complete the accelerator. Figure 2 shows the top-level dataflow.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

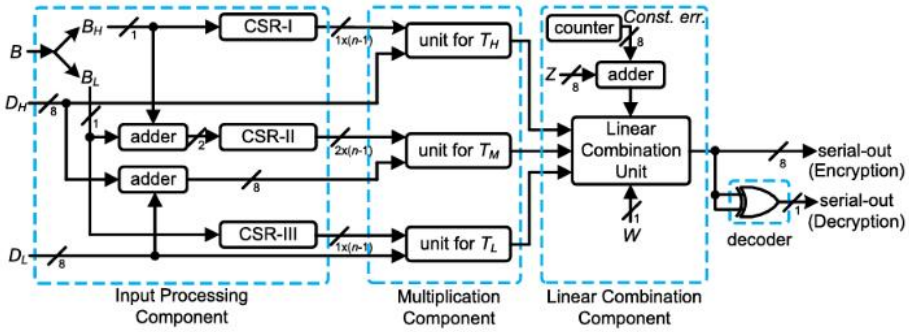

Fig. 2. Proposed RBLWE-ENC accelerator (KINA, $u = 1$), where Z and W denote the additional two polynomials (Z is an integer polynomial and W is a binary polynomial. CSR: circular shift-register. A constant error (Const. err.), delivered from a counter, is also needed when executing the addition with the integer polynomial Z, following the suggestion of [34].

Fig. 2. Top-level KINA accelerator with input processing, multiplication, and linear-combination components (adapted from [1]).

A. Input Processing and Circular Shift Registers

The input-processing block accepts the binary polynomial B and integer polynomial D and separates their even and odd coefficients into B_L, B_H, D_L, and D_H. Two adders form B_M=B_L+B_H and D_M=D_L+D_H. The binary lower and higher coefficient sets are stored in circular shift registers (CSRs), while the D-side coefficients are fed to the multiply-accumulate paths in the order required by column-wise accumulation.

The CSR organization avoids constructing every shifted version of the polynomial explicitly. After initial loading, the register loop is closed and the stored coefficients are rotated, generating successive columns of the corresponding circulant/negacyclic multiplication structure. For the basic u=1 architecture, one column is processed per cycle. The B_L and B_H registers hold binary values, while the B_M register requires two-bit storage because each element may be 0, 1, or 2.

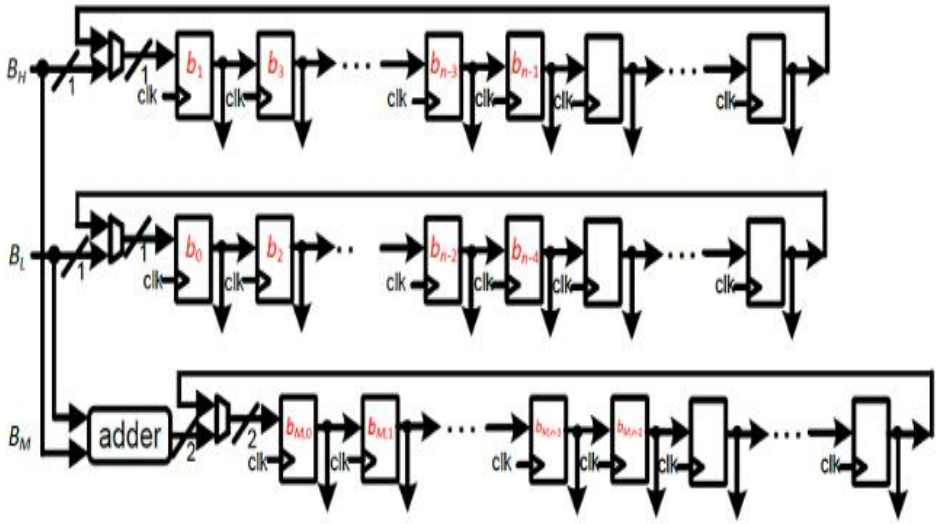

Fig. 3. Details of the CSRs for B_H, B_L, and B_M (from up), respectively.

Fig. 3. Circular shift-register arrangement for B_H, B_L, and B_M (adapted from [1]).

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

B. Parallel Multiplication-Accumulation Units

Three computation units operate in parallel. The T_L unit receives B_L and D_L, the T_H unit receives B_H and D_H, and the T_M unit receives B_M and D_M. In T_L and T_H, each point-wise binary product is produced by simple gating. In T_M, a compact selection network chooses 0, d, or 2d. Each point-wise product is added to a running register value, creating a multiply-accumulate structure whose feedback register holds the partial column sum.

After the required n/2 column cycles in the basic design, the sub-product coefficients are available for the recombination stage. The use of independent units allows T_L, T_H, and T_M to be accumulated concurrently, while the serial/column-wise organization keeps the structure smaller than a fully parallel coefficient array.

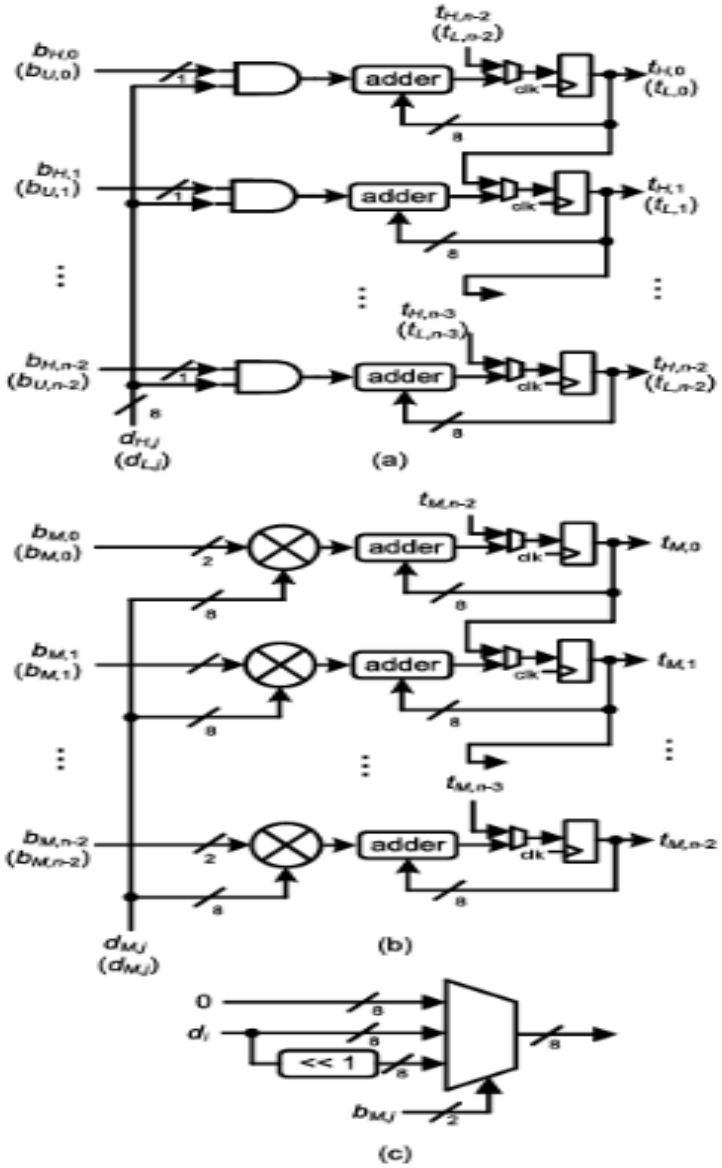

Fig. 4. (a) Processing unit for T_H or T_L. (b) Processing unit for T_M. (c) Point-wise multiplier for T_M.

Fig. 4. Processing units for T_L/T_H and T_M, including the compact point-wise multiplier (adapted from [1]).

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

C. Linear Combination and Ring Reduction

The linear-combination block implements (9) and enforces the quotient-ring relation x^n=-1. It receives serial coefficient streams from T_L, T_H, and T_M, applies sign inversions where subtraction is required, selects the appropriate coefficient positions with multiplexers, and adds the terms to produce the final serial coefficient T. Phase-specific quantities such as an error term, message term, or constant error contribution can be added at the output depending on whether the core is used for key generation, encryption, or decryption.

A serial output is advantageous for integration because it avoids an n-coefficient-wide interface. In encryption mode, the project uses an 8-bit coefficient output consistent with q=256. In decryption mode, the coefficient passes to a threshold decoder that emits a single message bit.

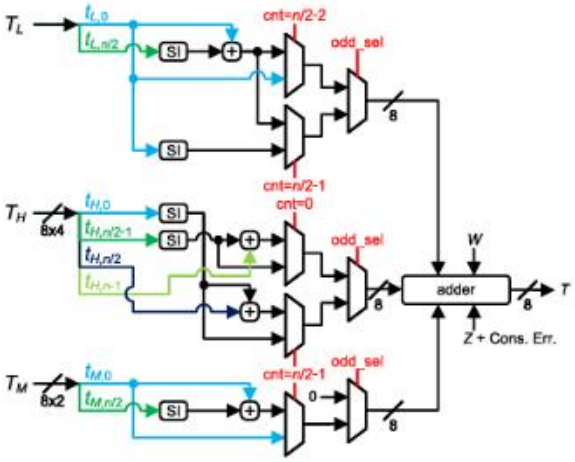

Fig. 5. Linear combination component, where SI denotes the sign inverter. Cons. Err.: constant error. The green and blue signals denote that they are connected with the corresponding registers in Fig. 4; while the red signals are control signals, e.g., "odd_sel" denotes that the MUX works in the lower channel when the odd order of output coefficient is selected to be produced when "odd_sel" = "1" (similar to the "cnt" control signals attached to the MUXes, e.g., "cnt = $N-1$" means that the MUX works in the lower channel when this counting signal is "1").

Fig. 5. Linear-combination component for recombining T_L, T_H, and T_M and handling ring wrap-around (adapted from [1]).

D. Control FSM

A finite-state machine coordinates register clearing, coefficient loading, accumulation, serial output, and completion signaling. The fixed schedule is independent of secret coefficient values, which gives the datapath a constant-cycle execution structure. Constant-cycle behavior is useful for reducing straightforward timing leakage, although it does not by itself protect against power, electromagnetic, or fault-injection side channels.

FSM state	Function
CLEAR / RESET	Clear accumulators, shift registers, counters, outputs, and control flags.
LOAD	Accept B/D coefficients, construct lower/high/middle terms, and fill the CSR structures.
COMPUTE	Enable point-wise multiplication and accumulation for T_L, T_H, and T_M.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

FSM state	Function
OUTPUT	Enable linear recombination, ring-aware coefficient selection, ciphertext output, and optional decoding.
DONE	Assert completion flag and hold/return to the idle condition for the next transaction.

E. Verilog Module Hierarchy

RTL module	Role in the accelerator
pointwise_multiplier	Binary or two-bit coefficient multiplication; T_M uses 0/d/2d selection.
accumulator_bank	Registers and adders for partial column sums.
mul_acc_unit	Combines point-wise multiplication and accumulation.
csr_bank	Stores and circularly rotates B_L, B_H, and B_M coefficient sequences.
input_processing	Performs decomposition and generates middle terms.
multiplication_component	Runs T_L, T_H, and T_M units concurrently.
linear_combination	Implements Karatsuba recombination and ring wrap-around.
decoder	Converts recovered coefficient values into message bits during decryption.
control_fsm	Sequences reset, load, compute, output, and done phases.
kina_top	Integrates datapath, controller, I/O, and phase-specific additions.

VI. HIGHER-SPEED KINA EXTENSION

The basic accelerator uses u=1, meaning one matrix/column contribution is processed per clock cycle. The KINA framework permits a tunable speed parameter u. When u columns are processed in parallel, the multiplication accumulation latency changes from n/2 cycles to approximately

$$L\_mul = n / (2u) \text{ cycles} \quad (10)$$

The throughput improvement is obtained by replicating point-wise multiplication lanes, extending the register output structure, and using adder trees to combine simultaneous partial products. Consequently, u is a direct architecture knob: small u minimizes area, while larger u reduces latency and increases throughput. This is a classic FPGA area-speed trade-off because replicated adders, registers, selection logic, and routing consume additional resources.

Figure 6 shows the shift-register modification for higher-speed processing. The register network advances by u positions rather than one position, and u integer coefficients are delivered in parallel. Figure 7 shows the layered multiplication structure in which multiple partial products feed adder trees.

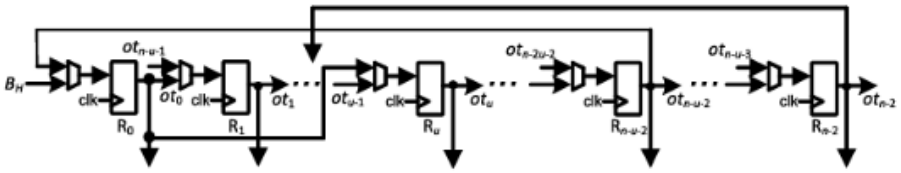

Fig. 6. Internal structure of the shift-register for B for the proposed higher speed KINA ("ot_i" denotes the output of the individual register).

Fig. 6. Higher-speed KINA shift-register structure for processing multiple columns per cycle (adapted from [1]).

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

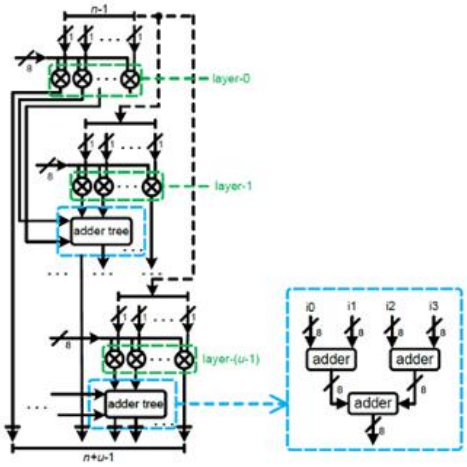

Fig. 7. Layered higher-speed computation unit with point-wise multipliers and adder trees (adapted from [1]).

Design option	Parallelism	Multiplication latency (cycles)	Typical design objective
Basic KINA	u=1	n/2	Area-conscious lightweight implementation
Balanced high-speed	u=2 or 4	n/(2u)	FPGA implementation balancing resource cost and latency
Throughput-oriented	u=8 or 16	n/(2u)	Server/gateway setting where throughput dominates area

VII. RTL IMPLEMENTATION AND VERIFICATION METHODOLOGY

A. Proof-of-Concept Parameter Set

The project implementation uses n=8 and q=256 as an educational verification configuration. This small value of n is not presented as a cryptographic security parameter; it is used to keep waveform inspection, testbench generation, and debugging manageable. The architectural data paths, control sequence, and decomposition can then be scaled toward the n=256 and n=512 parameter sets discussed in the reference KINA work [1].

Parameter	Proof-of-concept setting	Scalable setting	Implementation meaning
Polynomial degree n	8	256 or 512	n=8 used for RTL/waveform verification; larger values are the intended FPGA scaling targets.
Modulus q	256	256	Eight-bit coefficient path.
Binary operand	B	B	Enables gating/selection instead of full multiplication.
Parallelism u	1 in the basic proof of concept	1, 2, 4, 8, ...	Controls the area-latency trade-off.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

Parameter	Proof-of-concept setting	Scalable setting	Implementation meaning
Output	8-bit cipher coefficient / 1-bit message	Same logical roles	Serial coefficient output with optional threshold decoding.

B. Functional Verification Sequence

Verification proceeds from local arithmetic to full-system behavior. First, the T_L, T_H, and T_M units are checked individually against expected half-size products. Second, the linear-combination result is compared with reference polynomial multiplication modulo x^n+1. Third, the top-level FSM is exercised so that the waveform includes reset, start, load, compute, output, and done phases. The important observable signals in the supplied design include clk, rst, start, b_in, d_in, index_bit, z_coeff, w_coeff, const_err, cipher_out, message_bit, and done.

The project documentation records that the design was simulated and synthesized in Xilinx Vivado. The narrative does not state the exact Vivado release or the exact target device identifier, so those details are not invented in this paper and should be inserted from the actual project settings before formal submission. The available evidence consists of the behavioral waveform, XSim console output, RTL elaborated schematic, synthesized device view, and post-synthesis power-estimation report.

C. Synthesis Metrics of Interest

For a complete FPGA evaluation, synthesis and implementation should record lookup tables, flip-flops, slices, maximum operating frequency, cycle latency, critical-path delay, area-delay product, throughput, and power. The source KINA study reports these metrics for practical FPGA parameter sets and demonstrates that increasing u lowers latency while consuming more logic [1]. In the present n=8 proof-of-concept report, numerical LUT/FF/Fmax values are not stated in the textual results, so the paper reports only the measured/visible project outputs that are explicitly documented rather than fabricating missing figures.

VIII. RESULTS AND DISCUSSION

A. Behavioral Simulation

The behavioral simulation verifies the top-level control sequence and datapath activity. Figure 8 shows the XSim waveform with the relevant input, control, intermediate, and output signals. The design moves through the expected phases, updates the output during the output interval, and reaches the done condition. The final visible ciphertext coefficient is 0x92, corresponding to decimal 146. This confirms that the serial output path is active and that the top-level controller reaches a valid terminal state.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

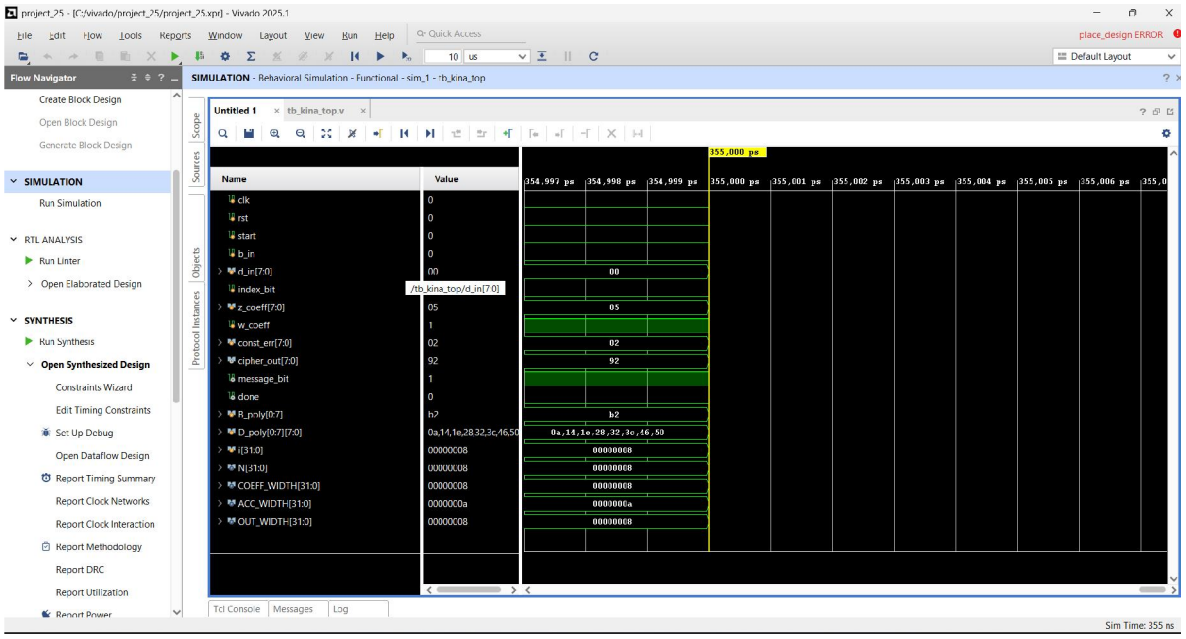

Fig. 8. Behavioral simulation waveform of the Verilog KINA top module.

B. XSim Console Verification

The XSim console provides a second verification point. The project output records successful completion of the simulation, a final cipher output of 146, and a final message bit of 1. Figure 9 reproduces the console evidence. The agreement between the waveform-visible coefficient and the printed result improves confidence that the observed value is not merely a transient cursor selection in the waveform viewer.

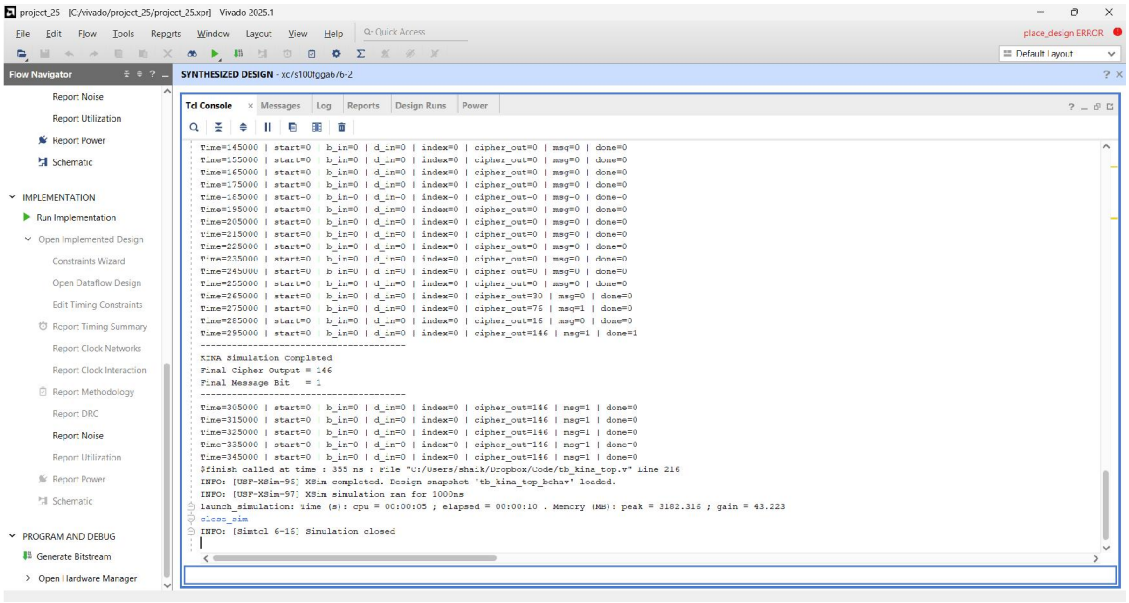

Fig. 9. XSim console output confirming final cipher output 146 and message bit 1.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

C. RTL Elaboration and Synthesis Mapping

RTL elaboration checks the connectivity of the implemented hierarchy before device mapping. Figure 10 shows the elaborated schematic, in which the controller, input-processing path, multiplication structures, linear-combination block, and decoder are connected through the top-level design. This verifies that the modular RTL has been integrated into a complete synthesizable network rather than remaining as disconnected test modules.

The synthesized device view in Figure 11 indicates that the Verilog design was accepted by synthesis and mapped to the selected FPGA target. The available project report does not provide a numerical utilization summary in text, so only the synthesis success and mapping evidence are claimed here.

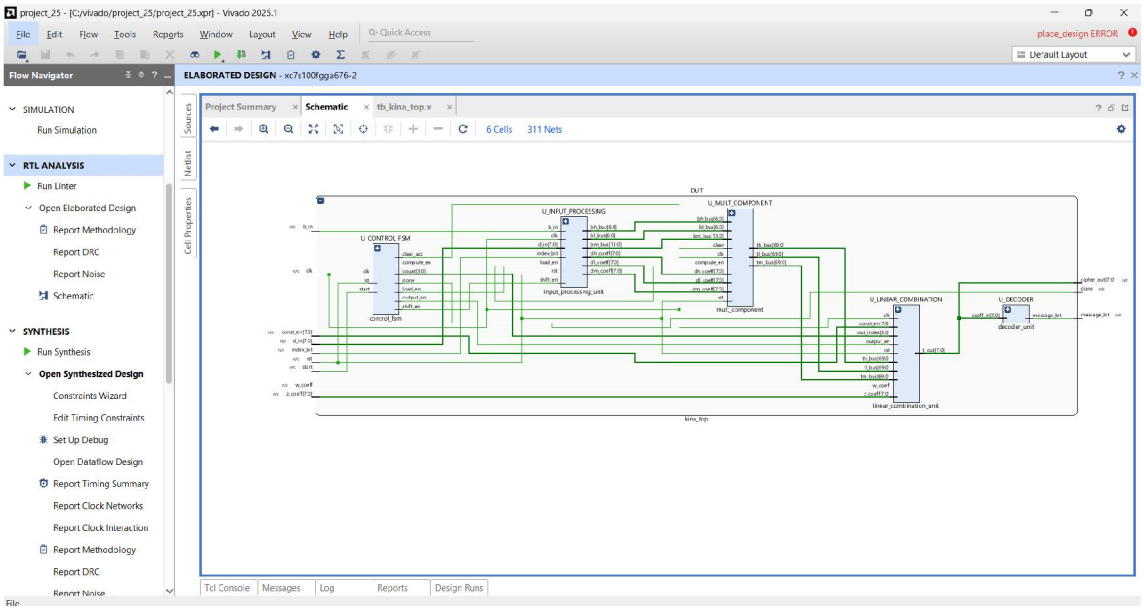

Fig. 10. RTL elaborated schematic showing top-level module connectivity.

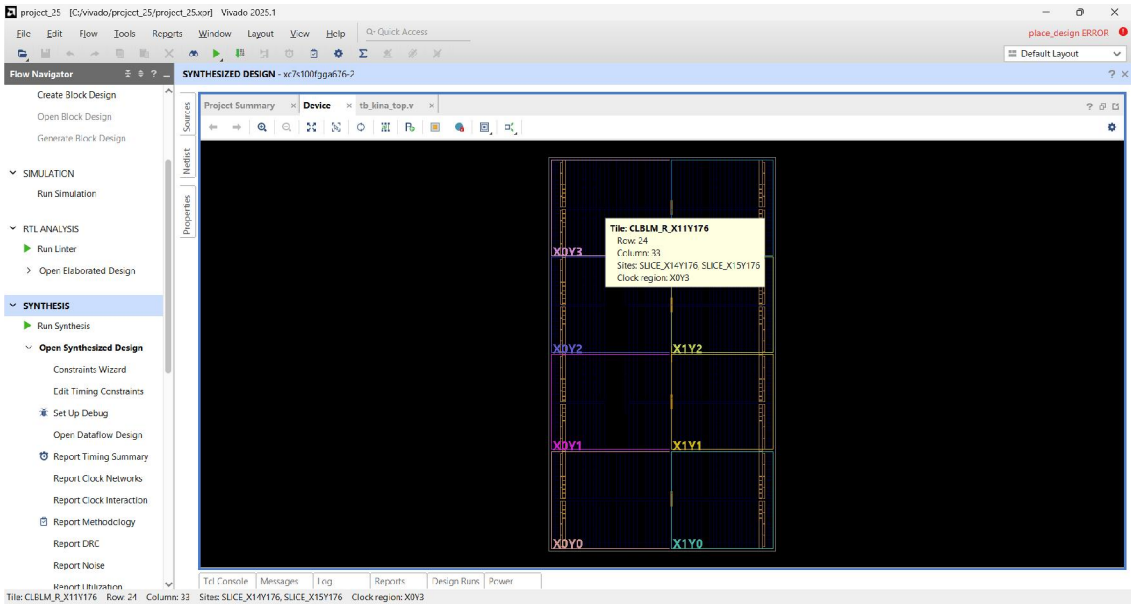

Fig. 11. Synthesized device view for the selected FPGA target.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 2, September 2026
Impact Factor: 8.2

D. Post-Synthesis Power Estimate

Figure 12 shows the available Vivado power-estimation report. The documented total on-chip power estimate is 0.089 W. This value should be interpreted as an early post-synthesis estimate rather than a final board-level measurement. The project documentation explicitly notes that the estimate should be updated after complete implementation with valid switching activity. Nevertheless, the result is useful as a baseline showing that the small n=8 RTL implementation is compatible with a low-power hardware realization.

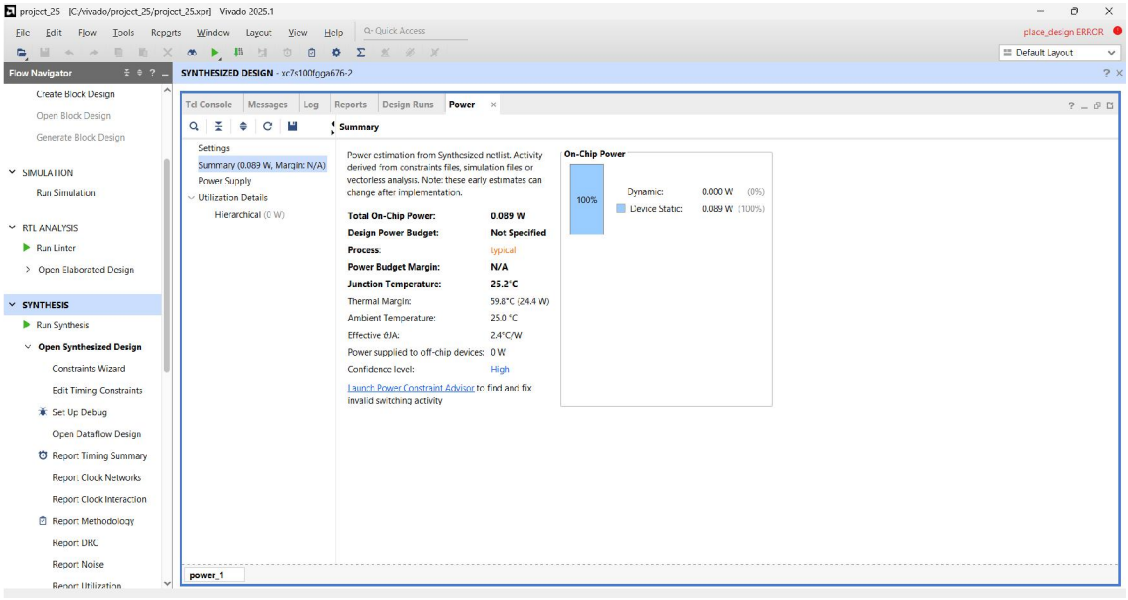

Fig. 12. Vivado post-synthesis power estimation report; reported total on-chip power is 0.089 W.

Metric / observation	Result	Evidence
Simulation environment	Xilinx Vivado/XSim	Project documentation
Proof-of-concept parameters	n=8, q=256	Project documentation
Final visible cipher coefficient	0x92 = 146	Behavioral waveform and console
Final decoded message bit	1	XSim console
RTL elaboration	Successful top-level connectivity	Elaborated schematic
Synthesis	Accepted and mapped to selected FPGA	Synthesized device view
Reported on-chip power estimate	0.089 W	Post-synthesis power report

E. Complexity and Latency Interpretation

The arithmetic advantage of KINA is visible before device-specific synthesis. Schoolbook multiplication requires O(n^2) coefficient interactions. The one-level Karatsuba structure replaces the four half-size products of a direct split by three products, yielding an approximate coefficient-multiplication complexity of 3n^2/4. The extra cost is addition, subtraction, coefficient routing, and middle-term width growth. For binary RBLWE operands, these additions are relatively inexpensive compared with a large set of general multipliers, and T_L/T_H point-wise products are particularly simple.

DOI: 10.48175/IJARSCT-38401

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

The basic serial/column architecture uses n/2 accumulation cycles. Higher-speed KINA reduces this to n/(2u) cycles by processing u columns in parallel. This scaling law explains the published FPGA behavior reported in the reference study: increasing u reduces latency and improves throughput, but LUT/FF/slice use rises because the datapath is replicated [1]. The source project report specifically notes that u=2 provides an attractive throughput/area-delay balance relative to selected prior RBLWE implementations. Because the present n=8 RTL report does not contain a full numerical area/frequency comparison, the paper retains that observation as a literature-backed comparison rather than presenting it as a newly measured result.

Architecture	Multiplication work	Accumulation latency	Main trade-off
Schoolbook RBLWE multiplier	O(n^2)	Depends on serial/parallel organization	Simple structure; latency rises rapidly with n.
Basic KINA (u=1)	Approximately O(3n^2/4)	n/2 accumulation cycles	Three parallel half-size sub-products.
Higher-speed KINA	Approximately O(3n^2/4) arithmetic, more hardware parallelism	n/(2u) accumulation cycles	Lower latency with additional logic and routing.

F. Comparison With Prior Hardware Directions

The most important distinction between KINA and earlier compact schoolbook designs is not that KINA eliminates all quadratic work; instead, it reorganizes the product so that three half-size computations replace four, then exposes a controlled parallelism parameter. Compared with aggressive schoolbook parallelization, this reduces the multiplication structure before hardware replication. Compared with an NTT design, it avoids transform-specific modulus and twiddle-factor structures that are inconvenient for the selected RBLWE parameters. Compared with recursively applying Karatsuba, it limits coefficient growth and keeps the middle term manageable.

The reference KINA implementation reports results across AMD-Xilinx and Intel FPGA families and compares its area-delay and throughput behavior with prior RBLWE accelerators [1]. The project report summarizes the main trend: u=2 achieves higher throughput and lower area-delay product than selected state-of-the-art RBLWE-ENC implementations, while larger u values continue to trade more resources for lower latency. This published evidence supports the scalability argument of the Verilog proof of concept, but practical n=256/n=512 synthesis of the present RTL remains necessary before claiming equivalent performance.

IX. SECURITY AND IMPLEMENTATION CONSIDERATIONS

A. Constant-Cycle Operation

The FSM follows a fixed sequence whose number of cycles does not depend on secret coefficient values. This gives the datapath constant-cycle behavior and removes a direct class of data-dependent timing variations. Such behavior is desirable in cryptographic hardware because timing differences can reveal information about secret-dependent branches or loops.

Constant time should not be interpreted as complete side-channel security. Power consumption, electromagnetic emanation, glitches, and fault behavior can still depend on processed data. Prior work has specifically considered power side-channel countermeasures for binary ring-LWE hardware [26]. A deployment-oriented KINA implementation should therefore consider masking, hiding, balanced logic, randomization, fault detection, redundancy, and careful physical design.

B. Scope of Cryptographic Validation

The n=8 proof of concept validates RTL behavior rather than cryptographic security. Secure parameter selection, noise distribution, failure probability, key/ciphertext formatting, randomness quality, and resistance to cryptanalysis require a

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

complete scheme-level analysis. The n=256 and n=512 values referenced by the source KINA work are appropriate targets for architectural scaling studies, but any production system must follow the security analysis and parameter definitions of the exact RBLWE construction being implemented.

X. APPLICATIONS, LIMITATIONS, AND FUTURE WORK

A. Application Areas

The accelerator is most relevant where post-quantum polynomial arithmetic must be repeated under constrained latency or power budgets. Potential application domains include IoT gateways, industrial sensor networks, smart-grid communication, embedded security processors, secure edge servers, vehicle communication, and academic FPGA/VLSI research. The serial-in/serial-out interface also makes the core suitable as a coprocessor attached to a wider cryptographic controller rather than as a stand-alone system.

B. Present Limitations

The current implementation has several explicit limitations. First, n=8 is a verification parameter and does not establish practical cryptographic security. Second, the available textual report does not provide measured LUT count, flip-flop count, maximum frequency, post-route timing, or board-level power, so these values are intentionally not invented. Third, the 0.089 W power value is a post-synthesis estimate and should be refreshed after implementation with realistic switching activity. Fourth, the work focuses on the dominant multiplication/recombination accelerator and does not document a complete production RBLWE key-generation, sampling, encoding, random-number, and protocol stack. Fifth, no experimental power/EM side-channel or fault-injection assessment has yet been performed.

C. Future Work

Future work should scale the RTL to n=256 and n=512, synthesize the same code for multiple u values, and report complete FPGA metrics including LUTs, FFs, slices/ALMs, Fmax, cycle latency, delay, area-delay product, throughput, dynamic power, and energy per operation. Hardware validation on Artix-7, Kintex-7, Virtex-7, Zynq, Cyclone, or Stratix-class platforms would allow the post-synthesis estimates to be replaced by measured implementation data.

Area and power can be further optimized through resource sharing, clock gating, accumulator-width analysis, register balancing, memory-efficient shift registers, and selective pipelining. The architecture can also be extended into a complete RBLWE cryptoprocessor that integrates key generation, encryption, decryption, message encoding, threshold decoding, random binary polynomial generation, and error generation. A subsequent ASIC study could evaluate performance per watt and silicon area.

Security-oriented extensions should incorporate masking or hiding for power analysis, redundancy or consistency checks for fault injection, and a structured side-channel evaluation methodology. Finally, the arithmetic ideas can be compared with standardized lattice-based hardware such as ML-KEM/ML-DSA accelerators, while carefully recognizing that their rings, moduli, transforms, and security requirements differ from RBLWE.

XI. CONCLUSION

This paper presented a Verilog-based FPGA realization of the KINA concept for Ring-Binary-LWE post-quantum cryptography. The architecture targets the dominant polynomial multiplication T=B·D mod (x^n+1), for which conventional schoolbook hardware has O(n^2) coefficient complexity. A one-level Karatsuba decomposition generates T_L, T_H, and T_M, after which a linear-combination block reconstructs the negacyclic product. Circular shift registers support column-wise accumulation, binary coefficients simplify the T_L and T_H multipliers, and a compact 0/d/2d selection network handles the two-bit middle term.

The RTL architecture was organized into reusable modules for input processing, circular shifting, multiply-accumulate operations, linear recombination, decoding, control, and top-level integration. An educational n=8, q=256 configuration was used for verification. The documented XSim run completed successfully with a final visible cipher coefficient of

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

146 (0x92) and a decoded message bit of 1. RTL elaboration confirmed module connectivity, synthesis mapped the design to the selected FPGA target, and the available post-synthesis report estimated total on-chip power at 0.089 W.

The study also showed how KINA scales through the parallelism factor u: multiplication accumulation latency falls from n/2 to n/(2u) cycles as more columns are processed concurrently, while hardware resource cost rises. The present implementation should therefore be viewed as a verified RTL foundation rather than a final security-grade cryptoprocessor. Scaling the design to n=256/n=512, reporting complete post-route FPGA metrics, integrating the full RBLWE protocol, and adding side-channel countermeasures are the next steps toward a practical lightweight post-quantum accelerator.

REFERENCES

[1] P. He, Y. Tu, J. Xie, and H. S. Jacinto, "KINA: Karatsuba Initiated Novel Accelerator for Ring-Binary-LWE (RBLWE)-Based Post-Quantum Cryptography," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 31, no. 10, pp. 1551–1564, Oct. 2023.

[2] D. J. Bernstein, "Introduction to post-quantum cryptography," in *Post-Quantum Cryptography*. Springer, 2009, pp. 1–14.

[3] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proc. Annual Symposium on Foundations of Computer Science*, 1994, pp. 124–134.

[4] D. J. Bernstein and T. Lange, "Post-quantum cryptography," *Nature*, vol. 549, no. 7671, pp. 188–194, 2017.

[5] D. Micciancio and O. Regev, "Lattice-based cryptography," in *Post-Quantum Cryptography*. Springer, 2009, pp. 147–191.

[6] M. D. Javeed, B. Hari Krishna, and M. V. S. R. Kishore, "Low Power Viterbi Decoder Design for TCM Decoders Using T-Algorithm," *ciiT International Journal of Programmable Device Circuits and Systems*, vol. 4, no. 15, pp. 764–768, 2012.

[7] National Institute of Standards and Technology, "Module-Lattice-Based Key-Encapsulation Mechanism Standard," *FIPS 203*, Aug. 2024.

[8] National Institute of Standards and Technology, "Module-Lattice-Based Digital Signature Standard," *FIPS 204*, Aug. 2024.

[9] National Institute of Standards and Technology, "Stateless Hash-Based Digital Signature Standard," *FIPS 205*, Aug. 2024.

[10] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *Journal of the ACM*, vol. 56, no. 6, pp. 1–40, 2009.

[11] V. Lyubashevsky, C. Peikert, and O. Regev, "On ideal lattices and learning with errors over rings," in *EUROCRYPT*, 2010, pp. 1–23.

[12] D. D. Chen et al., "High-speed polynomial multiplication architecture for ring-LWE and SHE cryptosystems," *IEEE Transactions on Circuits and Systems I*, vol. 62, no. 1, pp. 157–166, 2015.

[13] D. Liu, C. Zhang, H. Lin, Y. Chen, and M. Zhang, "A resource-efficient and side-channel secure hardware implementation of ring-LWE cryptographic processor," *IEEE Transactions on Circuits and Systems I*, vol. 66, no. 4, pp. 1474–1483, 2019.

[14] T. Pöppelmann and T. Güneysu, "Towards practical lattice-based public-key encryption on reconfigurable hardware," in *Selected Areas in Cryptography*, 2013, pp. 68–85.

[15] Mohammad Javeed and B. Swapna, "High Speed Convolution Encoding and Viterbi Decoding Using Dynamic Shift Register," *International Journal of Electronic and Electrical Engineering*, vol. 7, no. 5, pp. 483–490, 2014.

[16] A. Aysu, C. Patterson, and P. Schaumont, "Low-cost and area-efficient FPGA implementations of lattice-based cryptography," in *Proc. IEEE HOST*, 2013, pp. 81–86.

[17] S. Roy et al., "Compact ring-LWE cryptoprocessor," in *CHES*, 2014, pp. 371–391.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 2, September 2026 Impact Factor: 8.2

[18] J. Howe et al., "Standard lattice-based key encapsulation on embedded devices," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 372–393, 2018.
[19] S. Bian, M. Hiromoto, and T. Sato, "Filianore: Better multiplier architectures for LWE-based post-quantum key exchange," in *Proc. DAC*, 2019, pp. 1–6.
[20] T. Fritzmann and J. Sepúlveda, "Efficient and flexible low-power NTT for lattice-based cryptography," in *Proc. IEEE HOST*, 2019, pp. 141–150.
[21] D. Micciancio and C. Peikert, "Hardness of SIS and LWE with small parameters," in *CRYPTO*, 2013, pp. 21–39.
[22] J. Buchmann, F. Göpfert, R. Player, and T. Wunderer, "On the hardness of LWE with binary error," in *AFRICACRYPT*, 2016, pp. 24–43.
[23] D. Micciancio, "On the hardness of learning with errors with binary secrets," *Theory of Computing*, vol. 14, no. 1, pp. 1–17, 2018.
[24] A. Aysu, M. Orshansky, and M. Tiwari, "Binary ring-LWE hardware with power side-channel countermeasures," in *DATE*, 2018, pp. 1253–1258.
[25] S. Ebrahimi, S. Bayat-Sarmadi, and H. Mosanaei-Boorani, "Post-quantum cryptoprocessors optimized for edge and resource-constrained devices in IoT," *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5500–5507, 2019.
[26] D. Srinivas Reddy, G. Satyanarayana, Javeed M. D., and K. Mamatha, "Hybrid Machine Learning and Deep Learning-Based Intrusion Detection System for IoT Network Security Using Python," *International Journal of Computer Information Systems and Industrial Management Applications*, vol. 18, no. 6s, pp. 1133–1141, 2026, doi: 10.70917/ijcisim-2026-3022.
[27] P. He, U. Guin, and J. Xie, "Novel low-complexity polynomial multiplication over hybrid fields for efficient implementation of binary ring-LWE post-quantum cryptography," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 11, no. 2, pp. 383–394, 2021.
[28] J. Xie, P. He, and W. Wen, "Efficient implementation of finite field arithmetic for binary ring-LWE post-quantum cryptography through a novel lookup-table-like method," in *Proc. DAC*, 2021, pp. 1279–1284.
[29] K. Shahbazi and S.-B. Ko, "Area and power efficient post-quantum cryptosystem for IoT resource-constrained devices," *Microprocessors and Microsystems*, vol. 84, 2021.
[30] J. Xie, P. He, X. Wang, and J. L. Imaña, "Efficient hardware implementation of finite field arithmetic for binary ring-LWE based post-quantum cryptography," *IEEE Transactions on Emerging Topics in Computing*, vol. 10, no. 2, pp. 1222–1228, 2022.
[31] B. J. Lucas et al., "Lightweight hardware implementation of binary ring-LWE PQC accelerator," *IEEE Computer Architecture Letters*, vol. 21, no. 1, pp. 17–20, 2022.
[32] J. L. Imaña, P. He, T. Bao, Y. Tu, and J. Xie, "Efficient hardware arithmetic for inverted binary ring-LWE based post-quantum cryptography," *IEEE Transactions on Circuits and Systems I*, vol. 69, no. 8, pp. 3297–3307, 2022.
[33] D. Xu et al., "A more accurate and robust binary ring-LWE decryption scheme and its hardware implementation for IoT devices," *IEEE Transactions on VLSI Systems*, vol. 30, no. 8, pp. 1007–1019, 2022.
[34] V. B. Raju, J. MD, K. Keshamoni, and D. S. Reddy, "AI-Driven Harmful Algal Bloom Prediction in Aquatic Ecosystems Using Machine Learning and Remote Sensing Data," *International Journal of Aquatic Research and Environmental Studies*, vol. 6, no. S2, pp. 860–867, 2026, doi: 10.70102/q4xs0978.
[35] P. He, T. Bao, J. Xie, and M. Amin, "FPGA implementation of compact hardware accelerators for ring-binary-LWE based post-quantum cryptography," *ACM Transactions on Reconfigurable Technology and Systems*, vol. 16, no. 3, pp. 1–23, 2023.
[36] J. M. Pollard, "The fast Fourier transform in a finite field," *Mathematics of Computation*, vol. 25, no. 114, pp. 365–374, 1971.
[37] A. Karatsuba and Y. Ofman, "Multiplication of multidigit numbers on automata," *Soviet Physics Doklady*, vol. 7, no. 7, pp. 595–596, 1963.