International Journal of Engineering & Technology Sciences (ISSN: 2289-4152)

Volume 2025, 46 pages
Article ID: IJETS-2508132112962
HTPub, an imprint of Spot Publisher LLC

Review Article

Secure and Efficient Signal Processing on FPGA: A Comprehensive Review of Cryptographic, Post-Quantum, and AI-Enhanced DSP Implementations for Embedded Systems

Saman Sadeghi [a*]

Department of Electrical Engineering-Electronics, Mazandaran University of Science and Technology, Babol, Iran

Corresponding authors: Saman Sadeghi, Email: Samansdi30@gmail.com

Date Received: 05-04-2025, Date revised: 08-09-2025, Date accepted: 08-12-2025;

Abstract

The increasing demand for secure and efficient digital signal processing (DSP) in embedded systems has intensified the need for hardware platforms capable of delivering real-time performance without compromising cryptographic robustness. Field-Programmable Gate Arrays (FPGAs) have emerged as a versatile solution, offering fine-grained parallelism, low-latency execution, and reconfigurable logic for integrating both signal processing and cryptographic functions. This review presents a comprehensive synthesis of recent advancements in FPGA-based secure DSP architectures, encompassing classical primitives such as AES and ECC, as well as post-quantum cryptographic algorithms like Kyber, Dilithium, and NTRU. It explores design methodologies ranging from Register Transfer Level (RTL) to High-Level Synthesis (HLS), evaluating trade-offs in power, area, and latency, and detailing hardware-level countermeasures against side-channel attacks. Practical applications are surveyed across domains including healthcare, defense, secure communications, and multimedia processing, with comparative benchmarking on major Xilinx and Intel FPGA platforms. The review identifies key challenges, such as resource constraints, cross-platform portability, and the real-time implementation of cryptographic workloads, and highlights the integration of AI-based threat detection using models like convolutional and graph neural networks. A classification of machine learning applications in secure DSP is provided to contextualize current research directions. Finally, emerging trends are discussed, including post-quantum secure DSP systems, FPGA–AI co-acceleration, secure and reconfigurable hardware architectures, processing-in-memory (PIM), and heterogeneous platforms combining RISC-V SoCs with FPGA fabrics. By bridging cryptographic assurance with signal integrity, this work offers a holistic overview of the secure embedded computing landscape and a roadmap for future innovation at the intersection of signal processing, reconfigurable computing, and hardware-level security.

Keywords: FPGA, Digital Signal Processing (DSP), Cryptographic Primitives, Embedded Systems, Post-Quantum Security (PQC), Side-Channel Attacks, High-Level Synthesis (HLS), AES, RISC-V, Real-Time Processing, AI-Based Intrusion Detection, Reconfigurable Architectures, Lightweight Cryptography, Signal Authentication, FPGA-as-a-Service, Processing-in-Memory (PIM)

Introduction

Efficient and secure signal processing is rapidly becoming a core necessity in today's embedded systems, where data flows are continuous, real-time, and often exposed to adversarial conditions. In an increasingly interconnected digital world, embedded systems underpin a wide range of critical applications from mobile communications and multimedia processing to industrial automation and

mission-critical defense operations. As these systems grow more data-intensive and security-sensitive, the need for Digital Signal Processing (DSP) that delivers both high performance and robust protection becomes more pressing. Modern DSP solutions must not only meet real-time performance requirements, but also safeguard data confidentiality, authenticity, and integrity especially in untrusted or resource-constrained environments. (Motahhir & Maleh, 2023; Kajol & Yu, 2025; Narimani et al., 2025)

Cryptographic signal processing addresses this challenge by embedding conventional DSP algorithms with secure primitives such as Advanced Encryption Standard (AES), Elliptic Curve Cryptography (ECC), secure hash functions, and digital watermarking. These integrations are particularly relevant to embedded systems, which are often deployed in environments where trust cannot be assumed, and where solutions must be lightweight, power-efficient, and low-latency. The ability to protect signal streams, whether audio, video, sensor data, or control signals, from interception, tampering, or reverse engineering is essential to building the next generation of secure embedded platforms (Hosseini & Pilaram, 2024; Alnaseri et al., 2025). Among current hardware options, Field-Programmable Gate Arrays (FPGAs) have become a particularly attractive platform for implementing secure DSP. Unlike general-purpose processors, FPGAs offer massive parallelism, hardware-level reconfigurability, and customizable acceleration, giving designers the flexibility to optimize trade-offs among throughput, power efficiency, resource usage, and security robustness. Furthermore, recent advancements in High-Level Synthesis (HLS), reusable IP cores, and open-source toolchains have significantly lowered the barriers to developing cryptographic DSP architectures on FPGAs. (Kari et al., 2025; Gladis et al., 2025)

This review aims to bridge the gap between the theoretical foundations of secure signal processing and its practical implementation on reconfigurable hardware platforms such as FPGAs. Focusing on cryptographic DSP solutions in embedded systems, it surveys how these designs can simultaneously achieve high performance and robust security, even under stringent real-time and resource-constrained conditions. The paper classifies and analyzes a broad range of architectures and algorithms that integrate cryptographic primitives from traditional methods like AES and ECC to emerging post-quantum candidates such as Kyber and Dilithium directly within DSP pipelines. Beyond architectural considerations, it examines trade-offs in throughput, area efficiency, power consumption, and scalability across FPGA platforms from leading vendors like Xilinx and Intel, while addressing practical challenges related to integration and deployment. Furthermore, the review explores forward-looking research directions including AI-driven threat detection to bolster resilience against sophisticated attacks, the adoption of post-quantum cryptography to future-proof security, and the rise of FPGA-as-a-Service (FaaS) as a scalable approach for cloud and edge environments. Drawing on insights from both academic studies and real-world applications, this paper offers a comprehensive and practical perspective on designing secure, adaptive, and high-performance DSP systems, ultimately supporting researchers, engineers, and system architects in developing trustworthy, efficient, and future-ready signal processing solutions on reconfigurable hardware platforms.

Table 1 provides a comparative overview of FPGA, CPU, and ASIC platforms in the context of secure DSP implementation, highlighting FPGA's strengths in reconfigurability, parallelism, and cost-effective performance scaling.

Table 1. Comparison of Hardware Platforms for Secure DSP Implementation

Feature / Metric	FPGA	CPU	ASIC
Reconfigurability	High – Fully programmable hardware fabric allowing design updates and flexibility	None – Fixed architecture; updates only via software	None – Fixed hardware after fabrication
Parallelism	Massive parallelism through custom logic and pipelining	Limited to instruction-level parallelism	Very high parallelism but fixed at design time
Development Cost	Moderate – Moderate cost and development time	Low – Software-centric development	Very high – Expensive fabrication and design
Power Efficiency	High – Custom hardware can be optimized for low power	Moderate to low – General-purpose architecture	Very high – Optimized hardware, but inflexible
Latency	Low – Hardware acceleration enables real-time processing	Moderate to high – Software overhead	Very low – Dedicated hardware logic
Security Features	Customizable – Hardware-based cryptographic modules and side-channel protections	Software-based security only	Fixed hardware security features
Design Flexibility	High – Supports rapid prototyping and iterative design	High – Software can be updated easily	None – Hardware fixed post-fabrication
Suitability for Real-Time DSP	Excellent – High throughput and low latency	Limited – Depends on processor speed	Excellent – High performance but less flexible

Background and Technical Foundations

Digital Signal Processing (DSP) lies at the core of modern embedded systems, converting raw analog signals into precise digital data in real time. In applications with tight constraints on power, memory, and latency such as wearable health monitors, industrial automation, and automotive control units DSP architectures must be carefully designed to balance energy efficiency with consistent, predictable performance. Hardware implementations, particularly those leveraging FPGAs, offer clear advantages in these resource-limited environments. Comparative studies on FIR filter workloads reveal that pipelined FPGA designs can run up to 27 times faster than CPUs and twice as fast as GPUs (Arucu & Iliev, 2025). In wearable health monitoring, energy-aware DSP architectures combining low-leakage FPGA families with algorithm-level duty cycling have significantly reduced overall energy consumption, enabling multi-day ECG streaming powered by coin-cell batteries. (Khan & Da Silva, 2024)

Predictable performance is critical in real-time applications. For example, an FPGA-based frequency-hopping pseudorandom bit generator using chaotic maps delivers high-quality randomness and reliability, achieving 2 Gbps throughput on a Xilinx ZC702 while passing rigorous NIST and Diehard tests (Ayoub et al., 2024). Similarly, García-Requejo et al. (2025) developed a real-time, device-free indoor localization system based on ultrasound sonar with one receiver and four emitters. Built on an FPGA-based SoC, it processes signals in parallel and achieves localization errors below 13 cm for static points and under 10 cm for moving targets in most cases, broadcasting results via Wi-Fi.

In industrial settings, distributed digital manufacturing increasingly relies on FPGA-based sensor front-ends to enable microsecond-level deterministic control, even under harsh electromagnetic

interference (EMI), while performing feature extraction directly at the edge to reduce upstream network traffic (Khan et al., 2024). In human–robot interaction scenarios, adaptive FPGA-based accelerators integrate PIR and ultrasonic sensor data for posture classification, enabling implicit-communication SLAM navigation; a Verilog-based ZedBoard implementation has been successfully validated in hospital environments (Sravanthi et al., 2024).

Despite these advances, FPGA-as-a-Service deployments still face challenges such as data transfer bottlenecks, clock-rate limitations, and multi-card delays. However, techniques including top-level scaling, deep pipelining, parallelization, and parameterization have demonstrated significant improvements in throughput and response times, as confirmed by a queueing-theoretic model (Perepelitsyn & Kulanov, 2025). Meanwhile, security concerns in embedded DSP pipelines have grown, with low-cost FPGAs serving as accessible platforms for prototyping RISC-V and heterogeneous security architectures that combine lightweight cryptography and physically unclonable functions (PUFs) to protect DSP binaries in IoT applications (Stoyanov et al., 2025). Together, these developments indicate that secure, deterministic, and energy-efficient DSP on FPGAs is not merely a technical advantage but a foundational enabler for next-generation embedded intelligence.

Cryptographic Techniques in DSP (AES, ECC, Lightweight Crypto, Watermarking)

Embedding cryptographic functionality directly into DSP pipelines has become increasingly critical to ensure data integrity, confidentiality, and authenticity, particularly in embedded systems that process sensitive or real-time information. Among established standards, the Advanced Encryption Standard (AES) remains the dominant choice due to its proven security, high throughput, and hardware-efficient block cipher structure. According to Nguyen et al. (2025), the AES-RV architecture, which integrates a RISC-V AES instruction extension on the Xilinx ZCU102 platform, achieves up to 256 times higher processing speed and more than 450 times greater energy efficiency, as illustrated in Figure 1. These results highlight its strong potential for real-time DSP acceleration in embedded environments (Nguyen et al., 2025; Mazouz et al., 2025). Complementary research by Nassim and Zakary (2025) on high-throughput AES hardware co-processors, particularly optimized for the S-box stage, has demonstrated significant latency reductions and improved resource utilization on FPGA platforms.

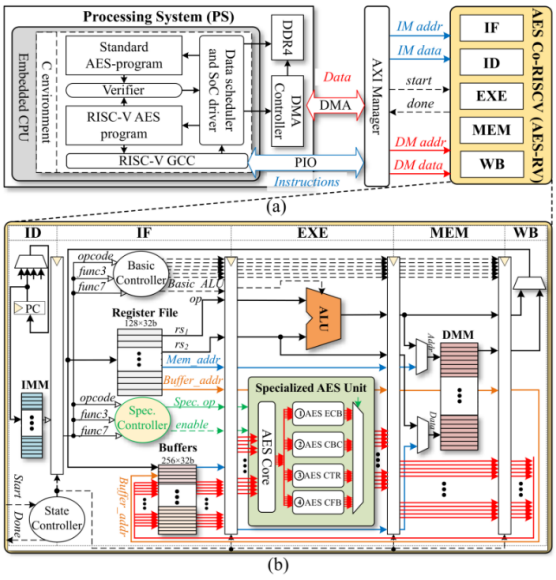

Figure 1. (a) System architecture overview of AES-RV at the SoC. (b) AESRV architecture. (Nguyen et al., 2025)

By comparison, Elliptic Curve Cryptography (ECC) offers asymmetric encryption with significantly smaller key sizes, making it particularly well-suited for resource-constrained embedded devices. Recent FPGA-based dual-layer authentication protocols that combine AES and ECC have demonstrated low-latency, hardware-accelerated security within real UAV communication pipelines (Roy et al., 2024). Similarly, ECC implementations on RISC-V cores have enabled low-power signature generation, supporting authenticated DSP data streams (Preethi et al., 2024). Beyond these established standards, lightweight cryptography algorithms, such as LED and chaotic map-based designs, are being explored for ultra-constrained systems. For example, integrating LED lightweight encryption with least significant bit (LSB) watermarking in medical image processing pipelines has achieved over 86 dB peak signal-to-noise ratio (PSNR), approximately 449 Mbps throughput, and strong watermark imperceptibility (Elhamzi, 2024). Another study presents a real-time speech and video crypto-watermarking prototype using chaos-driven symmetric encryption, where encryption and watermark embedding are fused within FPGA hardware to protect media pipelines with negligible processing overhead (Azzaz et al., 2023).

Additionally, watermarking and fingerprinting techniques embedded directly into DSP architectures provide a robust yet non-intrusive layer of protection for signal provenance and intellectual property (IP) ownership. Hardware watermarking through side-channel signatures, such as power consumption or electromagnetic (EM) emissions, enables compact IP-level authentication without altering DSP functionality (Wikipedia, 2025). FPGA-based real-time watermarking systems employing chaotic encryption have been developed to embed imperceptible identifiers into audio and video streams, which are particularly valuable for content authentication in surveillance or protected media workflows (Azzaz et al., 2023). Collectively, these integrated cryptographic and watermarking methods deliver embedded, hardware-level protection with minimal performance overhead, safeguarding signal authenticity and IP rights without compromising DSP throughput or fidelity.

FPGA Architecture and Design Flow (RTL, HLS, Vivado, Quartus, etc.): FPGAs offer architectural advantages such as inherent parallelism, reconfigurability, and fine-grained control, making them highly suitable for high-performance DSP applications. Core resources, including Look-Up Tables (LUTs), DSP slices, and Block RAMs (BRAMs), enable efficient implementation of custom datapaths tailored for time-critical signal processing tasks. Recent studies highlight these capabilities across diverse use cases: Khan and Da Silva (2024) employed FPGAs in wearable medical devices to facilitate efficient multimodal sensor data processing; Kim et al. (2024) demonstrated FPGA-accelerated convolutional neural networks (CNNs) for real-time satellite network routing, achieving a 3.1 times speedup; Syed et al. (2023) optimized LUT usage for FIR filter multipliers in audio DSP; and Hao et al. (2024) surveyed FPGA-based FIR optimizations focusing on linear phase response and system stability. Together, these efforts underscore the viability of FPGAs in modern DSP systems, though integrating robust security remains a key challenge warranting further research.

High-Level Synthesis (HLS) has emerged as a promising alternative to traditional RTL design, offering accelerated development cycles by translating algorithmic descriptions written in C/C++ or SystemC into hardware logic. While conventional RTL development in VHDL or Verilog allows for meticulous timing and resource control, it is often time-consuming and complex. HLS, by contrast, lowers the entry barrier for DSP developers, albeit sometimes at the expense of optimal quality of results (QoR) in performance-critical scenarios. The HLS design flow, as shown in Figures 2 and 3 from Lahti and Hämäläinen (2025), involves several stages with key steps automated by the HLS tool. They also observed that, despite the higher level of abstraction, HLS designs typically underperform compared to hand-crafted RTL implementations. Similarly, Duarte et al. (2018) reported that HLS facilitates neural network inference design, though with some performance trade-offs. Forelli et al. (2024) demonstrated

faster design times for FPGA-based machine learning accelerators using HLS, although manual fine-tuning was still necessary. In parallel, Curzel et al. (2023) explored MLIR-based loop optimizations to enhance HLS productivity in compute-intensive DSP applications. Collectively, these findings suggest that while HLS significantly streamlines the FPGA-based DSP design process, traditional RTL remains the gold standard when maximum optimization and tight timing closure are required.

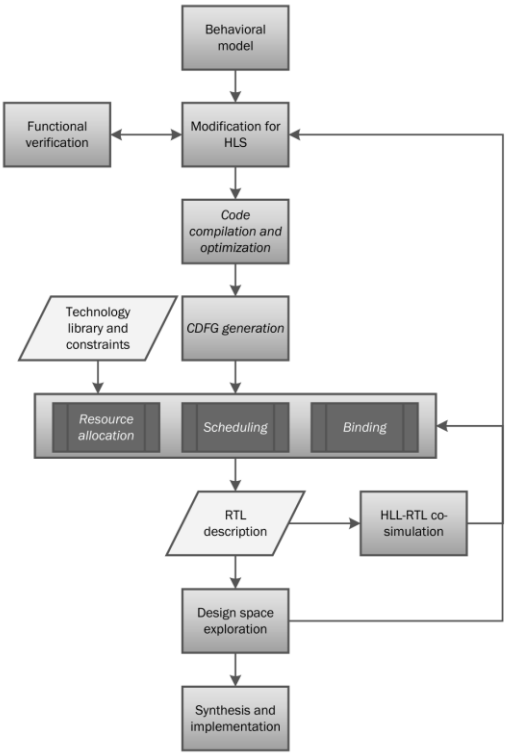

Figure 2. HLS design flow. Steps performed by the HLS tool are in italics. (Lahti and Hämäläinen, 2025)

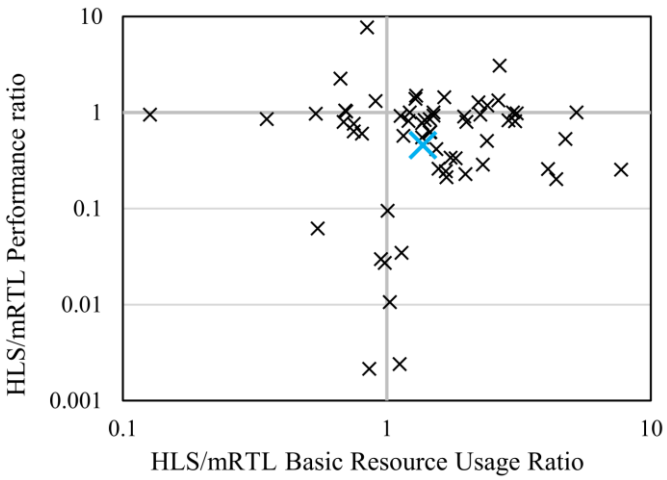

Figure 3. Relative HLS to mRTL performance to basic resource usage by application. (Lahti and Hämäläinen, 2025)

Security concerns have become increasingly critical in FPGA toolchains, especially for cryptographic DSP applications in high-stakes fields such as defense, aerospace, and finance. Modern toolchains now incorporate robust hardware security features, including bitstream encryption,

authentication mechanisms, and secure boot capabilities, to protect FPGA configuration integrity against tampering and unauthorized access. Recent research has driven advances across multiple aspects of FPGA toolchain technology. For instance, Xiao et al. (2024) introduced synthesis techniques based on cycle-deterministic high-level design languages that improve design efficiency and reinforce security guarantees within standard workflows like Vivado and Quartus Prime. Complementing this, Abbaszadeh and How (2024) developed advanced routing algorithms to tackle challenges posed by densely packed, performance-critical DSP designs, essential for reliable timing closure in real-time systems. Park and DeHon (2024) proposed runtime feedback frameworks that iteratively refine FPGA designs, enhancing latency optimization and resource utilization during design space exploration, addressing the growing demand for adaptive, secure DSP architectures. Finally, Pouchet et al. (2024) presented formal verification techniques tailored to high-level synthesis (HLS), providing rigorous correctness assurances that elevate the reliability of cryptographic DSP implementations. Together, these developments demonstrate how FPGA toolchains are evolving to balance demanding performance requirements with stringent security needs, paving the way for the next generation of resilient and scalable embedded systems.

A key architectural consideration in secure FPGA-based DSP design lies in the contrast between SRAM-based and Flash-based FPGA devices. SRAM-based FPGAs, widely offered by vendors like Xilinx and Intel, require external configuration at every power cycle, introducing boot-time latency and posing security risks if bitstreams are not properly encrypted and authenticated. In contrast, Flash-based FPGAs from providers such as Microchip or Lattice store configurations internally, offering instant-on performance and reducing vulnerability at startup. While SRAM-based devices typically deliver higher logic density and scalability, Flash-based counterparts inherently protect against bitstream interception and modification, making them particularly attractive for security-sensitive and power-constrained embedded applications. Xu and Zhang (2020) discussed security risks tied to SRAM-based FPGA bitstreams and highlighted the dual role of machine learning as both a security threat and safeguard when FPGAs are used standalone or within shared systems. Similarly, Proulx et al. (2023) reviewed FPGA cybersecurity challenges, emphasizing special considerations for SoC FPGAs that combine programmable logic with integrated processors, and called for robust, threat-driven security strategies despite existing protective features (Liu et al., 2024; Jung & Choi, 2019). Table 2 summarizes the core distinctions between SRAM-based and Flash-based architectures based on recent literature and discussion.

Table 2. Key Differences Between SRAM-Based and Flash-Based FPGAs

Aspect	SRAM-Based FPGAs	Flash-Based FPGAs
Configuration	Requires external bitstream loading at every power-up	Configuration stored internally; supports instant-on capability
Security	Vulnerable to bitstream interception if not properly encrypted	Inherently more secure against tampering and reverse engineering
Performance	Higher logic density and scalability for complex applications	Lower density; suitable for specific, targeted applications
Power Consumption	Higher static power; lower dynamic power in high-performance use	Lower static power; ideal for low-power, energy-sensitive designs
Application Suitability	Ideal for high-performance DSP, telecommunications, and aerospace	Well-suited for security-critical, low-power embedded systems

Why FPGA for Secure DSP? FPGAs offer a uniquely versatile and compelling solution for secure DSP implementations, especially in scenarios where deterministic execution, scalable throughput, and cryptographic agility are non-negotiable. Unlike general-purpose microcontrollers and CPUs, FPGAs bypass software overhead by directly mapping custom logic into hardware, which minimizes latency and significantly shrinks the attack surface. This hardware-centric approach inherently improves system predictability and strengthens overall security, both vital for embedded DSP systems managing sensitive or real-time data streams. Compared to Application-Specific Integrated Circuits (ASICs), FPGAs bring a decisive edge: post-deployment reconfigurability. This feature allows rapid updates to cryptographic algorithms, security protocols, and policies to address emerging threats or evolving standards without redesigning hardware. Such adaptability is crucial in dynamic security landscapes where long-term flexibility is as important as initial robustness (Sadeghi, 2024).

Moreover, FPGAs enable the development of highly specialized, side-channel-resistant architectures and allow secure key storage to be integrated directly into hardware. These capabilities enhance resilience against invasive attacks, like physical probing, as well as passive threats, such as power analysis and electromagnetic leakage. Consequently, FPGAs have become the platform of choice for deploying real-time, high-trust secure DSP systems in high-stakes domains including defense electronics, satellite communications, aerospace, and critical national infrastructure. For instance, Bommana et al. (2025) showcased an innovative framework combining Deep Learning with Dynamic Partial Reconfiguration (DPR) to counter power side-channel attacks in real time. This adaptive strategy exploits the FPGA's inherent flexibility and low-latency response, making it particularly suitable for DSP applications that require immediate defense against hardware-level threats.

Ultimately, selecting the right hardware platform is pivotal in designing secure DSP systems for embedded or mission-critical applications. FPGAs deliver a rare blend of ultra-low latency, moderate-to-high energy efficiency, and unmatched post-deployment adaptability, offering significant advantages over conventional CPUs, GPUs, or even advanced ASICs (Qasaimeh et al., 2019). Table 3 summarizes these core distinctions across platforms, underscoring why FPGAs remain at the forefront of modern secure DSP architectures.

Table 3. Comparison of FPGA with CPU, GPU, and ASIC for Secure DSP

Platform	Latency	Energy Efficiency	Flexibility	Remarks
FPGA	Very low due to hardware-level parallelism	Moderate to high depending on custom logic	High, supports post-deployment reconfiguration	Well-suited for real-time secure DSP with evolving cryptographic requirements
CPU	High because of sequential execution and OS overhead	Low due to general-purpose architecture	Moderate via software programmability	Easy to develop on, but not ideal for secure or time-critical DSP tasks
GPU	Moderate; parallel but less deterministic	High power usage, especially under full load	Good for flexible software-based tasks	Strong in throughput, but not optimal for embedded low-latency systems
ASIC	Very low due to full hardware optimization	Very high efficiency with minimal overhead	No flexibility after fabrication	Extremely efficient, but cannot adapt to new threats or standards

Post-Quantum Cryptography (PQC) and Its Relevance to DSP Systems

As quantum computing rapidly advances, it poses a significant threat to the mathematical foundations of traditional cryptographic algorithms, putting widely used public-key schemes such as RSA and ECC at risk. In response, integrating post-quantum cryptographic (PQC) primitives directly into digital signal processing (DSP) systems has become both essential and urgent. This need is especially critical for embedded and DSP-enabled systems deployed in mission-critical applications with long operational lifecycles, where hardware-level future-proofing is key to ensuring enduring security and system resilience. Designing DSP architectures that incorporate PQC not only prepares for the emerging quantum era but also guarantees the long-term reliability and trustworthiness of secure signal processing platforms. (Li et al., 2023; Malik & Islam, 2025; Allgyer et al., 2024; Aydeger et al., 2024)

Among post-quantum cryptographic (PQC) algorithms, lattice-based cryptography stands out due to its robust security guarantees and suitability for efficient hardware implementation. The proposed PQ-ALU architecture, illustrated in Figure 4 (Stelzer et al., 2025), employs modular arithmetic with two subtractors, one adder, and a multiplier, each equipped with dedicated modular reduction logic. Modular reduction for the adders and subtractors is performed via conditional additions and subtractions to optimize computational efficiency. Leading PQC candidates include Kyber (a Key Encapsulation Mechanism, KEM), Dilithium (a digital signature scheme), and NTRU (a public-key encryption algorithm), all based on hard mathematical problems such as Learning with Errors (LWE) and Ring-LWE, which provide resistance against quantum attacks. Kyber and Dilithium, both part of the CRYSTALS suite, have been optimized for practical deployment, with Kyber offering fast key encapsulation and Dilithium providing efficient digital signatures (Nguyen et al., 2024; Cheng et al., 2025). NTRU, a mature lattice-based cryptosystem, offers faster private-key operations compared to RSA, making it suitable for resource-constrained embedded systems (Bote & Diaz-Vargas, 2025). Although these algorithms introduce higher computational complexity than classical cryptosystems, their highly regular structure, intrinsic parallelism, and predictable control flow make them well-suited for hardware acceleration on FPGAs. Furthermore, techniques like the Number Theoretic Transform (NTT), depicted in Figures 4, 5 and Table 4 (Stelzer et al., 2025), significantly reduce latency in polynomial multiplication, a core operation in these algorithms, thus enhancing their applicability in DSP-focused FPGA architectures.

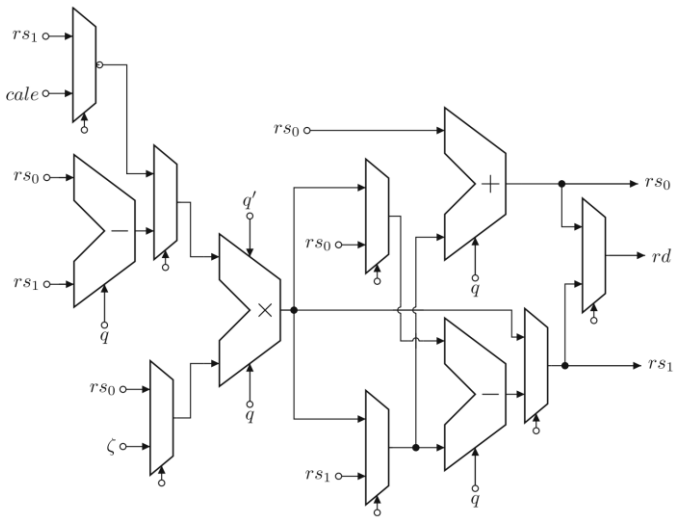

Figure 4. Architecture of PQ-ALU. (Stelzer et al., 2025)

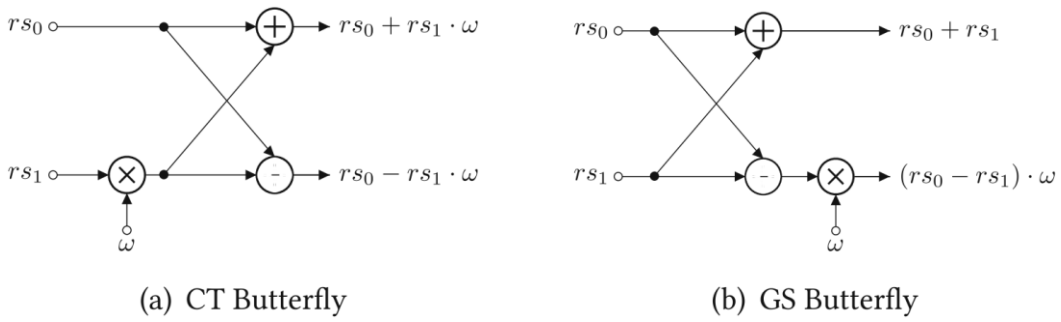

Figure 5. NTT butterfly operations. (Stelzer et al., 2025)

Table 4. Overview of NTT-based polynomial multiplication for Dilithium, Kyber and Falcon. (Stelzer et al., 2025)

Scheme	n	Stages	q	$\log_2 q$	PWM
DILITHIUM	256	8	8380417	23	✓
KYBER	256	7	3329	12	
FALCON-512	512	9	12289	14	✓
FALCON-1024	1024	10	12289	14	✓

NIST Standardization Efforts: The NIST Post-Quantum Cryptography (PQC) initiative has been instrumental in guiding the global shift from traditional cryptographic algorithms, now vulnerable to quantum attacks, towards robust quantum-resistant alternatives (Li et al., 2023; Allgyer et al., 2024; Aydeger et al., 2024; Liu & Moody, 2024). Following a comprehensive, multi-phase evaluation process, NIST selected Kyber and Dilithium as the primary standards for key encapsulation and digital signatures, respectively. To date, three standards have been finalized: FIPS 203 (ML-KEM, derived from Kyber), FIPS 204 (ML-DSA, based on Dilithium), and FIPS 205 (SLH-DSA, derived from SPHINCS+), with a fourth, FALCON (FN-DSA), expected to be released soon (Allgyer et al., 2024). These standards provide hardware engineers with a firm foundation for integrating quantum-secure cryptographic primitives at the architectural level, particularly within secure DSP environments. FPGAs stand out as ideal platforms for early PQC adoption due to their flexibility in rapid prototyping, real-time benchmarking, and evaluating side-channel attack resistance (Chen et al., 2023). Overall, NIST's PQC standardization ensures these algorithms are not only secure but also interoperable, speeding their deployment in mission-critical applications.

Challenges and Opportunities for FPGA-Based PQC Integration: Despite the urgent need for PQC, embedding these primitives into FPGA-based DSP systems presents significant challenges. Lattice-based algorithms like Kyber and Dilithium demand complex modular arithmetic and high memory bandwidth, which increase computation time and vulnerability to side-channel attacks (Nguyen et al., 2024; Dam et al., 2024; Li et al., 2024; Kim et al., 2024). When PQC cores share FPGA resources with DSP functions, designers face tight constraints on logic utilization, memory access, and performance. Striking a balance among low latency, reconfigurability, and robust security is crucial, especially in real-time applications requiring strong protection against evolving threats. Nonetheless, FPGAs' inherent reconfigurability offers unique advantages: secure firmware updates, bitstream encryption, dynamic cryptographic resource allocation, and design-space exploration for performance optimization through

pipelining, loop unrolling, and efficient memory hierarchies (Nguyen et al., 2024; Dam et al., 2024). These features position FPGAs as the preferred platform to meet the demanding security, efficiency, and flexibility requirements of emerging IoT devices, 5G networks, and defense-grade DSP systems. (Li et al., 2024; Kim et al., 2024)

Secure DSP Architectures on FPGA

AES-Based Signal Encryption and Filtering: The Advanced Encryption Standard (AES) remains one of the most widely used block ciphers to ensure data confidentiality, especially in real-time signal processing applications. Implementing AES directly on FPGAs alongside signal filtering allows secure encryption of streaming data while maintaining the low latency critical for domains such as secure communications, industrial control, and defense systems. Recent studies have expanded on this foundation by developing low-power AES architectures and designs resistant to side-channel attacks, thereby enhancing both the efficiency and robustness of FPGA-based secure DSP systems (Kumar et al., 2021; Siddiqui & Sekhar, 2024; Prakashan et al., 2024; Pariya et al., 2022; Dhanda et al., 2022; Sunil et al., 2020). These advancements highlight AES not only as a fundamental cryptographic building block but also as a practical enabler for secure, high-throughput DSP pipelines deployed in embedded and resource-constrained environments.

Elliptic Curve Cryptography (ECC) provides strong asymmetric encryption with significantly smaller key sizes, making it especially suitable for resource-constrained systems such as IoT devices and embedded DSP platforms. FPGA-based ECC accelerators efficiently perform public-key operations, reducing latency and hardware footprint compared to general-purpose processors. For instance, Javeed et al. (2023) proposed an enhanced Elliptic Curve Scalar Multiplication (ECSM) architecture that achieves a 30% reduction in computation time alongside improved area–time efficiency on Virtex-7 FPGAs, as illustrated in Figure 6. Concurrently, lightweight cryptographic algorithms designed for ultra-low-power devices have been effectively accelerated on FPGAs, balancing security requirements with strict resource constraints. The inherent reconfigurability of FPGAs enables these systems to adapt to emerging threats by updating cryptographic modules without necessitating hardware redesign. Looking ahead, the integration of post-quantum cryptographic algorithms alongside ECC in FPGA-based DSP architectures is expected to ensure long-term security in preparation for quantum computing threats. (Hossain et al., 2025; Kumari et al., 2025; Lin et al., 2023; Wang et al., 2023; Dong et al., 2018)

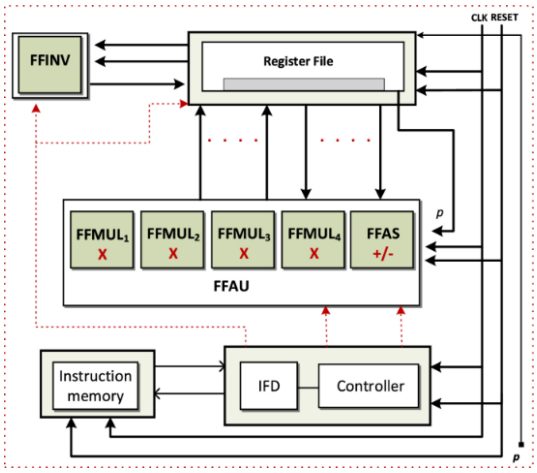

Figure 6. Low-Latency FPGA Accelerator for Twisted Edwards Curve. (Javeed et al., 2023)

Signal Integrity, Authentication, and Hashing: Ensuring signal integrity and verifying data authenticity are critical for protecting digital signal processing (DSP) systems against tampering and unauthorized access. Cryptographic hash functions such as SHA-3 and message authentication codes (HMAC) can be efficiently implemented on FPGAs, leveraging their inherent parallelism to achieve high throughput with minimal latency. These capabilities enable real-time authentication, which is crucial for secure communications and sensor network applications. Ongoing research aims to optimize these implementations for lower power consumption and enhanced resistance to physical and side-channel attacks, thereby strengthening the security foundations of FPGA-based DSP systems. (Holla et al., 2023; Sideris et al., 2024; Sideris & Dasygenis, 2023; Kieu-Do-Nguyen & Hoang, 2022; OpenTitan Project, 2025)

Watermarking and Tamper Detection in Multimedia DSP: Digital watermarking embeds imperceptible and robust signatures into multimedia signals to support tamper detection and authentication. FPGA-based implementations integrate watermarking directly into DSP pipelines, providing real-time protection for audio, video, biometric, and IoT sensor streams without significantly impacting performance. These techniques play a vital role in safeguarding intellectual property, ensuring data provenance, and preventing unauthorized access in connected and security-sensitive systems. Application areas include anti-piracy protection in streaming media (Yasin et al., 2023; Aissaoui et al., 2022; Hussain et al., 2022), secure biometric authentication (Janaki et al., 2024; Altman et al., 2024), and lightweight watermarking for IoT sensor networks balancing resource constraints with data integrity (Mekhfioui et al., 2025). In the medical domain, FPGA-based real-time watermarking for imaging and EEG signals in wearable health devices ensures patient data confidentiality and tamper resistance, meeting growing regulatory requirements for secure medical data transmission and storage. (Gull & Parah, 2023; Mazouz et al., 2025; GAPses, 2024; Abdelaziz et al., 2024)

RISC-V SoC with Integrated FPGA for Secure DSP Pipelines: The advent of heterogeneous System-on-Chip (SoC) platforms combining RISC-V cores with FPGA fabrics enables tightly integrated secure DSP architectures. These platforms facilitate a hybrid task division: control planes execute cryptographic protocols such as TLS and SSH in software, while FPGA fabrics accelerate compute-intensive kernels like AES, SHA-3, and post-quantum cryptographic (PQC) primitives. This design ensures both low-latency signal processing and robust security. Representative platforms include Microchip's PolarFire SoC and Intel's Agilex FPGAs with integrated Nios II soft-core processors, deployed in diverse real-time applications such as medical signal acquisition, secure audio streaming, and authenticated imaging. (Microchip Technology Inc., 2024; Nguyen et al., 2025; Ma et al., 2023; Kieu-Do-Nguyen et al., 2024)

Secure DSP methods present diverse trade-offs involving latency, security strength, hardware resource usage, and implementation complexity on FPGA platforms. A clear understanding of these trade-offs is essential for selecting appropriate cryptographic or protection strategies tailored to application-specific constraints. Table 5 summarizes key cryptographic and signal integrity techniques, highlighting their main characteristics and use cases to guide designers in balancing security, performance, and hardware efficiency.

.

Table 5. Comparison of Secure DSP Techniques and Architectures on FPGA

Technique / Architecture	Security Strength	Latency	Resource Usage	Typical Use Cases	FPGA Suitability
AES-based Filtering	High	Low	Moderate	Secure comms, industrial control	Highly suitable (pipelined)
ECC / Lightweight Crypto	Medium-High	Low	Low	IoT, embedded DSP, wearable devices	Efficient for low-area
SHA-3, HMAC Hashing	High	Low	Moderate	Authentication, signal integrity	Parallelism-friendly
Watermarking (multimedia, medical, biometric)	Medium	Very Low	Low	Video/audio, EEG, IP protection	Customizable + low-power
RISC-V SoC + FPGA Integration	High	Very Low	Configurable	Secure AI-DSP pipelines, post-quantum protocols	Ideal for hybrid designs

The integration of RISC-V System-on-Chips (SoCs) with embedded FPGA fabric on a single die offers a powerful and flexible foundation for secure DSP system design. These heterogeneous platforms enable a synergistic architecture where software-managed cryptographic control can be tightly coupled with hardware-accelerated processing for functions such as encryption, authentication, and post-quantum cryptography. This combined capability makes them particularly well-suited for security-critical applications at the edge, including AI-driven signal analytics, real-time encrypted media streaming, and continuous monitoring in medical wearable devices. Table 6 presents an overview of leading RISC-V SoC and FPGA platforms and their associated secure DSP use cases, illustrating the growing role of these integrated systems in high-assurance embedded computing.

Table 6. Representative RISC-V SoC + FPGA Platforms and Their Applications in Secure DSP

Platform / Vendor	RISC-V Core(s)	FPGA Fabric Type	Key Applications	Distinctive Features
Microchip PolarFire SoC	Quad-core RISC-V	Flash-based FPGA	Secure medical imaging, encrypted industrial DSP	Low power, crypto co-processors, deterministic real-time performance
Quick Logic EOS S3 + RISC-V	Single-core SiFive RISC-V	Embedded FPGA (eFPGA)	Wearable ECG/EEG, biometric authentication	Ultra-low-power sensor fusion, always-on DSP
Intel Agilex + Nios II	Softcore Nios II (RISC)	High-end SRAM-based FPGA	PQC acceleration, encrypted video/audio pipelines	Customizable datapaths, OpenCL/HLS support
SiFive Intelligence SoC + eFPGA	Multi-core RV64 + Vector Extensions	Embedded FPGA fabric	Edge AI + DSP, anomaly detection, PQC frameworks	ML accelerators, secure co-processing engines

Implementation Strategies and Optimization Techniques

Achieving effective implementation of secure DSP algorithms on FPGAs demands thoughtful choices in design methodologies, architectural enhancements, and strong defenses against increasingly sophisticated security threats. In this section, we examine the main strategies and technologies that support high-performance, low-latency, and robust secure signal processing.

RTL vs. High-Level Synthesis (HLS) Approaches in Secure DSP Designs: Register Transfer Level (RTL) design remains the dominant methodology for developing secure DSP modules that demand low latency, tight timing control, and resistance to side-channel attacks. This includes AES encryption cores, ECC accelerators, and cryptographic hashing engines. RTL enables designers to finely tune resource allocation and pipeline stages, which is critical in maintaining both performance and security. Recent synthesis frameworks like PoSyn have demonstrated significant success in minimizing power side-channel leakage in cryptographic implementations such as AES, RSA, and PQC. By optimizing how sensitive RTL blocks are mapped to standard cells, PoSyn reduces mutual information leakage without degrading functionality (Srivastava et al., 2025; Calvo et al., 2024). Complementing this, MaskedHLS introduces a domain-specific high-level synthesis flow for secure designs, providing glitch-resistant masking, balanced execution paths, and improved register allocation. Its performance gains are notable, achieving up to 45% latency reduction and 74% register savings over manual RTL (Sarma et al., 2024).

Despite the advantages of RTL, High-Level Synthesis (HLS) has become increasingly popular for its ability to speed up design iteration and exploration. Tools such as Xilinx Vivado HLS and Intel OpenCL support loop unrolling, pipelining, and resource reuse, enabling fast prototyping of secure datapaths. However, HLS abstractions also introduce potential pitfalls. Studies have revealed that memory-related optimizations, such as those in AES S-boxes synthesized from C code, can leak sensitive data through power side channels if not properly masked or hidden (Zhang et al., 2019; Bhashini et al., 2025). Aggressive memory partitioning and directive-based scheduling further increase the risk of leakage unless mitigated during design space exploration (Koufopoulou et al., 2022; Xiong et al., 2025). To counter these threats, next-generation HLS flows integrate protective mechanisms directly into synthesis. MaskedHLS, for example, automatically applies glitch-resistant masking and switching balance, while PoSyn enforces a leakage-aware cost model to reduce side-channel vulnerability during logic synthesis. (Sarma et al., 2024; Srivastava et al., 2025)

To capitalize on the strengths of both approaches, designers are increasingly adopting hybrid flows. Security-critical components, such as masked AES engines or ECC primitives, are developed in RTL to guarantee robustness and precision, while algorithmic functions, filters, and control units are implemented using HLS to accelerate development. This blend achieves a practical balance between productivity and security. Architectural evaluations have shown that combining HLS-generated logic with RTL-based masking techniques can deliver resilient designs without compromising throughput (Koufopoulou et al., 2022; FPGA '23 Proceedings, 2023). A recent advancement in this domain is HLSPilot, a large language model–driven synthesis framework introduced by Xiong et al. (2025). HLSPilot uses in-context learning and profiling to guide design space exploration on hybrid CPU–FPGA platforms. It provides performance comparable to expert-crafted RTL implementations while greatly improving design productivity. Figure 7 illustrates HLSPilot's synthesis pipeline, from profiling to hardware deployment.

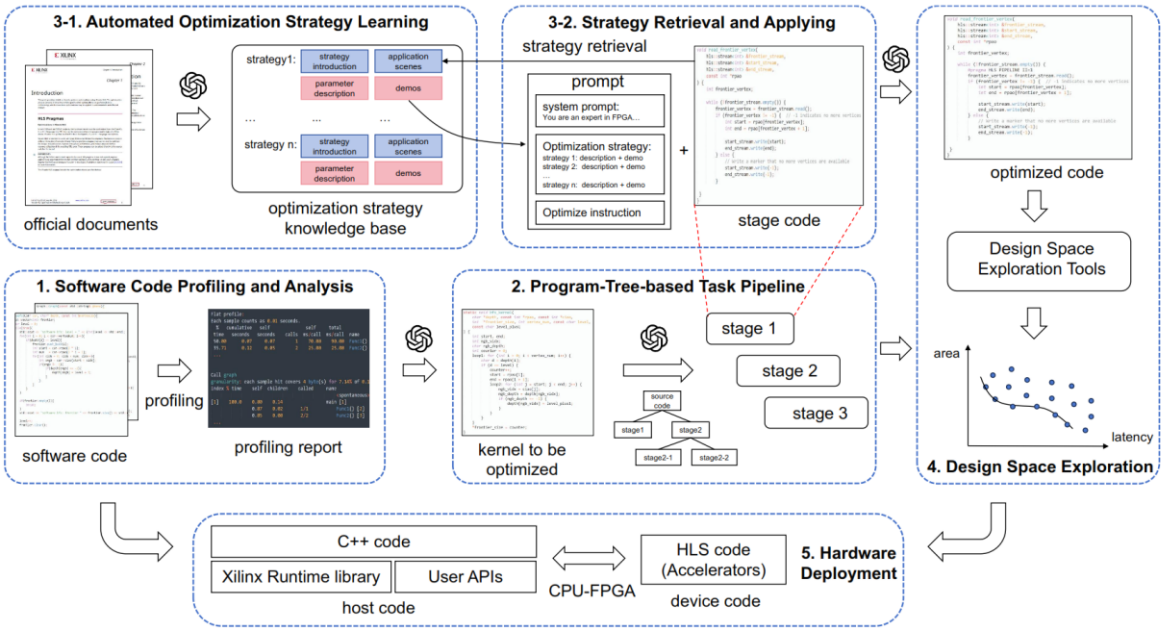

Figure 7. HLSPilot framework. (Xiong et al., 2025)

Pipelined and Parallel Architectures for High Throughput: Pipelining is a highly effective technique for boosting throughput in cryptographic designs by breaking down operations such as AES rounds into sequential stages. This allows new data blocks to enter the pipeline every clock cycle, significantly increasing processing speed while maintaining low latency. For example, the fully pipelined AES-256 implementation WAES 256 achieved an impressive 206 Gbps throughput on a Kintex UltraScale FPGA. This throughput scales up to 426 Gbps with two cores and 742 Gbps using four cores, all without consuming any Block RAM (BRAM) resources (Malal & Tezcan, 2025). Another design featuring a four-stage AES pipeline reached 105.7 Gbps with an efficiency of 31.48 Mbps per slice (Tran et al., 2024). Similarly, low-latency AES architectures demonstrating around 2.34 Gbps throughput with minimal resource usage have also been reported (Sumit et al., 2022). Furthermore, architectures leveraging multi-core parallelism combined with composite field arithmetic offer further throughput enhancements alongside reduced area overhead (Calvo et al., 2024; Malal & Tezcan, 2025). Collectively, these works confirm that aggressive pipelining in FPGA-based AES implementations can achieve very high throughput paired with low latency.

Beyond AES, parallelism is equally crucial for accelerating lattice-based post-quantum cryptography (PQC) algorithms. For instance, KyberMat employs a parallel feed-forward Number Theoretic Transform (NTT) architecture for polynomial multiplication, reducing execution time by 90% and increasing throughput by a factor of 66 (Tan et al., 2023). Similarly, a unified pipelined NTT design supporting both Kyber and Dilithium utilizes configurable radix butterfly units to balance area and timing performance across diverse FPGA platforms (Mandal & Basu Roy, 2023). Additional studies introduce highly pipelined butterfly units for high-speed NTT and inverse NTT operations, achieving a 62% improvement in throughput-per-slice compared to prior designs (Rashid et al., 2025). Other optimizations target forward and inverse NTT for ML-DSA and ML-KEM workloads using parallel FPGA pipelines to boost performance (Taghavi et al., 2025). Moreover, hardware accelerators deployed on platforms like the Alveo U280 demonstrate speedups ranging from 3 times to 9 times over CPU implementations for Kyber and Dilithium via batch processing and parallel execution (Carril et al., 2024). Research on unified PQC accelerators further highlights the importance of compact, high-performance parallel architectures in efficiently implementing lattice-based cryptosystems (Bao et al.,

2025). Altogether, these advances establish FPGA-enabled parallelism as a key enabler for low-latency, high-throughput PQC acceleration.

Resource Sharing and Its Latency Implications: Resource sharing is a common approach in resource-constrained FPGA designs aimed at reducing area and power by allowing multiple functions such as multiply-accumulate (MAC) units or cryptographic cores to reuse the same logic blocks. However, this approach inherently serializes execution, increasing latency due to task queuing. For example, a wearable FPGA-based qMLP accelerator used time-multiplexed reuse of MAC blocks to reduce resource use and power consumption but experienced added delays when multiple tasks contended for hardware access (Khan & Da Silva, 2024). Optimization frameworks employing clock gating and task scheduling similarly indicate that while resource reuse significantly lowers area and power, it must be managed carefully to avoid bottlenecks that degrade performance (Harvie, 2023). Comprehensive reviews of embedded IoT systems also point out that resource sharing benefits power-constrained designs but can negatively affect timing-sensitive applications (Barge & Gerardine, 2024).

Power Reduction Methods: Clock gating and power gating are critical techniques used to reduce power consumption by selectively disabling clock signals or power supply to idle modules, thereby cutting dynamic and static power dissipation respectively. Studies in FPGA and hardware accelerators confirm that clock gating effectively lowers switching activity and dynamic power, while power gating minimizes leakage current in inactive components. Reviews focusing on low-power mobile and embedded FPGA designs highlight that combining these gating techniques achieves notable power savings with minimal area overhead (Vaithianathan et al., 2024). Furthermore, realistic thermal- and power-aware FPGA systems integrate gating into system-level power management strategies to ensure reliability amid varying workloads (Chowdhury & Schafer, 2021).

Another approach involves adopting lightweight cryptographic algorithms or approximate arithmetic units to reduce complexity, hardware footprint, and power, potentially at the expense of security margins or accuracy. Studies on IoT-targeted FPGA implementations recommend tailored lightweight AES variants that significantly cut power and area requirements while maintaining adequate throughput and latency (Barge & Gerardine, 2024). Additionally, frameworks that optimize the reuse of DSP slices, BRAM, and multipliers demonstrate appreciable gains in area efficiency and energy per bit processed (Grycel & Walls, 2019; Tibaldi & Pilato, 2023). Industry analyses also emphasize that optimizing block RAM and multiplier unit mapping can deliver high throughput with low resource usage in secure FPGA pipelines (Gu et al., 2025).

To dynamically balance performance and power in FPGA-based secure DSP systems, techniques such as dynamic voltage and frequency scaling (DVFS) and adaptive clocking prove effective. Recent research shows DVFS can substantially reduce both dynamic and static power by adjusting supply voltage and clock frequency according to workload demands. For example, one study reports energy savings up to 47.74% in ultra-low-power embedded systems using DVFS (Zidar et al., 2024). Vendor-supported FPGA DVFS flows, especially for IoT and edge devices, incorporate device-level gating, clock domain management, and adaptive voltage scaling to optimize throughput-per-watt trade-offs (Khan & Da Silva, 2024).

IP Cores and FPGA Toolchains for Secure DSP Systems: Modern FPGA development flows such as Xilinx Vivado HLS and Intel Quartus Prime provide integrated toolchains covering synthesis, simulation, timing analysis, and debugging, all with streamlined support for cryptographic and DSP IP cores. For instance, the AES-RV architecture embeds custom AES instructions directly into a RISC-V SoC on a Xilinx FPGA, demonstrating how advanced toolchain integration accelerates development and delivers low-latency, pipelined cryptographic throughput (Nguyen et al., 2025). Likewise, the SPiME parallel AES architecture leverages an extensive Verilog and RTL-based toolchain to instantiate

thousands of lightweight AES cores, achieving predictable latency alongside area-efficient resource utilization (Karakchi et al., 2025). Additionally, research into accelerating zero-knowledge proofs on Intel FPGAs highlights how deep integration of IP libraries within the OneAPI toolchain can securely scale modular arithmetic across FPGA platforms (Butt et al., 2024).

Secure DSP systems greatly benefit from robust, NIST-certified cryptographic IP cores that implement AES, ECC, SHA-3, and post-quantum cryptography (PQC) algorithms, many equipped with built-in side-channel countermeasures. For example, Xiphera's Crypto Module IP portfolio includes AES, SHA-3, TRNG, ECC (NIST curves), Curve25519, Kyber, and Dilithium, all designed using constant-time logic to defend against Differential Power Analysis (DPA) attacks (Xiphera, 2024). A dual-layer AES and ECC authentication framework for UAVs, built with hardened crypto cores on FPGA, has demonstrated strong resilience against timing and power analysis attacks (Roy et al., 2024). Similarly, the SPiME architecture shows how large arrays of hardened AES cores can be efficiently deployed with minimal area overhead while maintaining strong security guarantees (Karakchi et al., 2025). Moreover, specialized low-latency AES cores have been successfully integrated within secure RISC-V platforms supporting energy-harvesting DSP, showcasing how modern toolchains enable seamless use of hardened IP cores (Nguyen et al., 2025).

In parallel, open-source hardware initiatives like OpenTitan contribute transparent and auditable roots of trust (RoT) and cryptographic IP to FPGA-based designs, fostering openness and community-driven validation. Reviews of OpenTitan's AES, HMAC, and big-number accelerator (OTBN) cores report speedups of roughly 4.3 to 12.5 times over software-only implementations, demonstrating the tangible benefits of IP reuse paired with formal verification (Parisi et al., 2024). Figure 8 illustrates how OpenTitan is architecturally integrated as a secure subsystem. Although OpenTitan has yet to be fully realized on FPGA silicon, ongoing efforts are bringing its secure boot controllers, TRNGs, and ECC modules into FPGA-based RISC-V ecosystems. At the same time, Xiphera complements these trends by adopting open-source verification practices and developing secure IP cores aimed at future-proof compliance and alignment with open standards (Xiphera, 2024). Collectively, these advances in cryptographic IP cores and FPGA development toolchains enable high-performance, standards-compliant, and secure DSP systems, reducing design time and risk while providing a solid architectural foundation for next-generation real-time embedded applications that demand robust cryptographic assurances.

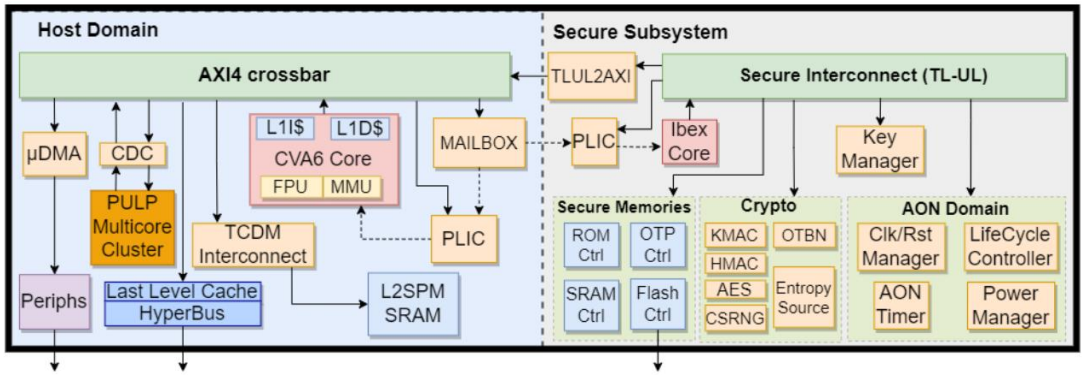

Figure 8. Diagram depicting the architectural integration of OpenTitan as a secure subsystem. (Parisi et al., 2024)

Hardware-Based Countermeasures to Prevent Side-Channel Attacks: As physical side-channel attacks such as power analysis and electromagnetic (EM) leakage become increasingly sophisticated, robust hardware-level protections are vital for safeguarding sensitive data within secure DSP pipelines. These attacks exploit subtle correlations between device power consumption or EM emissions and secret values like cryptographic keys, posing significant threats to system confidentiality and integrity. Without effective countermeasures, embedded cryptographic modules remain highly vulnerable.

At the hardware level, primary defenses include masking, dual-rail logic, clock and instruction randomization, and threshold implementations, all aimed at mitigating power and EM side-channel leakage in secure DSP architectures. For example, the PoSyn synthesis framework introduces power side-channel-aware cell mapping, optimizing standard-cell allocation to achieve up to 3.79 times area efficiency and more than 70% reduction in Differential Power Analysis (DPA) and Correlation Power Analysis (CPA) success rates for AES, Saber, and Kyber designs (Srivastava et al., 2025). Similarly, threshold implementations optimized for FPGA pipelines enable masked AES and Kyber cores to achieve mathematically guaranteed first-order DPA resistance without degrading throughput (Ji & Dubrova, 2025). These strategies increasingly target DSP pipelines to balance high security with performance on FPGA platforms.

Beyond static protections, runtime or dynamic countermeasures like partial reconfiguration enhance resilience by switching between functionally identical but structurally distinct hardware variants. Ahmadi et al. (2023) demonstrated the automatic generation of AES and Kyber hardware variants that increase the minimum traces-to-disclosure (MTD) by over 1000 times, with only around 14% area overhead and no performance loss on Xilinx and AMD FPGAs. Additionally, combining deep learning-based threat detection with partial reconfiguration enables real-time adaptation of hardware configurations, dynamically countering detected attack patterns (Bommana et al., 2025). These adaptive methods promise strong resilience in secure DSP pipelines without sacrificing throughput.

Prior to deployment, rigorous statistical testing such as Test Vector Leakage Assessment (TVLA) and experimental tools like ChipWhisperer are critical to identify potential leakages early. Notably, Ji and Dubrova (2025) presented the first successful side-channel attack on a masked FPGA implementation of CRYSTALS-Kyber, achieving full key recovery via deep learning-aided CPA, underscoring the need for early leakage detection and robust defenses. Likewise, Bhashini et al. (2025) revealed how synthesis choices in Vivado significantly affect AES leakage profiles and side-channel resistance, highlighting the importance of toolchain-aware evaluation. Modern synthesis flows and secure hardware design pipelines are increasingly integrating formal leakage verification, information-theoretic leakage minimization, and secure RTL-to-netlist mappings. Srivastava et al. (2025) and Grosso & Lara Nino (2025) have developed methods that not only minimize leakage but also embed TVLA-like analysis directly into synthesis workflows, mathematically bounding mutual information leakage in modules like AES and Kyber. Complementary toolchain-based assessments further evaluate side-channel resistance in masked implementations generated by High-Level Synthesis (HLS) and RTL, helping to identify vulnerabilities introduced by optimization passes (Koufopoulou et al., 2022). Collectively, these methods form a comprehensive verification framework, ensuring secure FPGA-based DSP pipelines meet stringent industrial security standards.

.

Applications and Industrial Use Cases

The convergence of secure DSP architectures with FPGA platforms is profoundly transforming a range of high-impact domains. Thanks to their inherent parallelism, reconfigurability, and built-in

hardware-level security features, FPGAs enable real-time, low-latency, and energy-efficient processing of sensitive signals under demanding operational constraints. This section explores the main application areas, highlighting recent advancements, practical challenges, and key trends shaping the future of secure DSP systems.

Medical Devices and Wearable Signal Security: Wearable medical devices such as ECG monitors and continuous glucose sensors are increasingly leveraging FPGAs to embed real-time encryption, authentication, and watermarking directly within the signal processing pipeline. This integration helps ensure that raw physiological data remains protected in transit, while maintaining ultra-low latency critical for medical applications. For instance, Sengupta and Chaurasia (2025) describe secure FPGA hardware IP that inserts AES-256 encrypted signatures into pacemaker filter and pulse detection modules, effectively safeguarding against tampering and intellectual property theft, as shown in Figure 9. Similarly, Vaithianathan et al. (2024) demonstrate a Spartan 7 FPGA-based smart health monitoring system capable of performing feature extraction on ECG, PPG, and motion sensor data with exceptional power efficiency and minimal latency, achieved through clock gating. The use of chaotic encryption in real-time ECG monitoring further enhances security, while still supporting on-device deep learning diagnostics without introducing latency penalties (Yuksel & Metin, 2024). Khan and Da Silva (2024) also highlight ultra-low-power FPGA platforms for physiological monitoring that incorporate partial reconfiguration, dynamic voltage scaling, and compliance with privacy regulations such as HIPAA and GDPR.

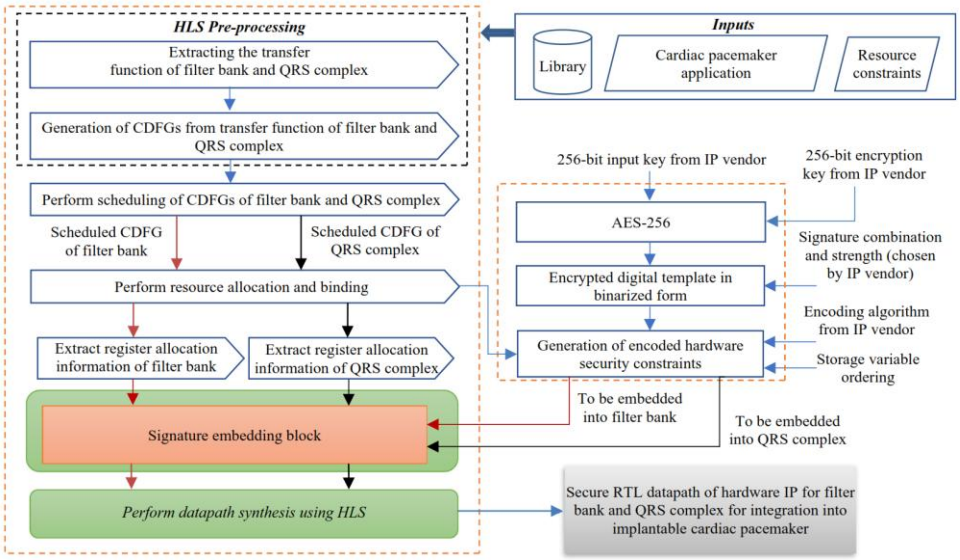

Figure 9. High level design flow for generating secure hardware IPs for implantable cardiac pacemaker. (Sengupta & Chaurasia, 2025)

Power efficiency is critical for wearable and implantable devices operating under limited battery capacity. Vaithianathan et al. (2024) investigate FPGA power optimization strategies such as power gating and dynamic voltage scaling to effectively balance energy consumption with throughput demands. Complementing this, Khan and Da Silva (2024) survey additional low-power techniques including compressed learning and novel hardware like flash-based FPGAs. Benchmarking post-quantum cryptographic (PQC) algorithms such as Kyber and NTRU on FPGA platforms highlights inherent trade-offs between throughput and power efficiency in secure embedded systems (Cano Aguilera et al., 2024). For example, Baidya et al. (2025) propose an optimized Number Theoretic Transform (NTT) sampler for Kyber on Artix 7 FPGAs, achieving approximately 33% improvements in latency, energy efficiency, and area, thus suiting resource-constrained medical edge applications.

Ensuring the long-term security of wearable medical devices requires area-efficient FPGA implementations of lattice-based PQC schemes. Stelzer et al. (2025) integrate Dilithium and Falcon signature schemes into the OpenTitan root-of-trust, achieving minimal area overhead and verification latencies under 10 ms. Similarly, Rudraksh, a compact lattice-based PQC key encapsulation mechanism (KEM) optimized for constrained devices, demonstrates nearly threefold improvements in area and operating frequency compared to Kyber on FPGA hardware (Kundu et al., 2025). Kieu-Do-Nguyen et al. (2024) introduce a small, high-frequency (approximately 417 MHz) FPGA NTT engine for CRYSTALS-Kyber that balances low LUT usage with state-of-the-art area-time performance. Meanwhile, Dang et al. (2022) provide a comparative analysis of Kyber, NTRU-HPS, and Saber across FPGA platforms, illustrating practical performance and resource trade-offs for real-time secure workloads.

Advances in edge AI and federated learning (FL), combined with secure FPGA architectures, are transforming real-time patient monitoring. Yan et al. (2024) present FedEYE, a scalable federated learning system enabling privacy-preserving model training across heterogeneous endpoints including FPGA-based devices. Kaur et al. (2023) review lightweight cryptography and secure Internet of Medical Things (IoMT) designs, emphasizing the incorporation of symmetric algorithms such as ASCON alongside side-channel attack countermeasures in wearable FPGA firmware. Aminifar et al. (2024) propose a privacy-preserving edge FL system for seizure detection on resource-constrained mobile and wearable platforms, analyzing trade-offs between computation, energy consumption, and privacy on FPGA endpoints. Similarly, an IEEE BigData (2024) study examines federated learning in medical applications, focusing on performance, fairness, and privacy in small dataset scenarios, offering valuable design insights for secure FPGA edge systems (BigData, 2024; Kim et al., 2025).

Secure Communication Systems and Encrypted Audio: Robust encryption is critical for applications such as military command networks, emergency coordination, and secure conferencing. FPGA-accelerated DSP systems provide the low-latency, high-throughput performance required to implement cryptographic primitives like AES-GCM, SHA-3, and emerging post-quantum cryptography (PQC) algorithms such as CRYSTALS-Kyber and Dilithium. For example, Nguyen et al. (2024) demonstrate an efficient hardware accelerator for NTT and INTT operations on CRYSTALS-Kyber using an Artix-7 FPGA running at 262 MHz with about 1,405 LUTs. The dynamic reconfigurability of FPGAs allows real-time switching of encryption protocols and key sizes, enhancing resilience to evolving threats and zero-day vulnerabilities. Papalamprou et al. (2025) propose a hybrid hardware-software remote attestation architecture combining PQC and blockchain to secure FPGA configuration and provide verifiable logging. FPGA designs also support multi-protocol cryptographic suites, ensuring compatibility with legacy standards like AES and SHA-3 as well as upcoming quantum-resistant protocols. Surveys on secure Internet of Medical Things (IoMT) architectures emphasize embedding lightweight symmetric ciphers like ASCON alongside side-channel countermeasures within wearable FPGA firmware (Kaur et al., 2023; Mirigaldi et al., 2025).

Hardware-software co-design strategies improve flexibility by running encryption control and key management on embedded processors while offloading computationally intensive encryption to dedicated FPGA logic. The BYOTee framework, introduced by Armanuzzaman et al. (2022), Perkins et al. (2024), and Zou et al. (2025), enables creation of tamper-resistant enclaves on FPGA platforms for sensitive workload execution. Additionally, integrating secure key storage, hardware roots of trust, and side-channel mitigations directly into FPGA fabric further enhances security. However, reviews by Moraitis et al. (2023) and Mishra et al. (2025) highlight risks associated with bitstream manipulation, underscoring the need for robust integrity checks and trust anchor mechanisms to protect FPGA deployments.

Military, Tactical, and Defense Signal Processing: Military and tactical applications demand secure DSP platforms that not only incorporate strong cryptography but also withstand physical probing, environmental stress, and deliver strict real-time performance guarantees. FPGAs are particularly well-suited to these needs by enabling the deployment of custom cryptographic cores, hardware obfuscation strategies, and precise timing control that collectively strengthen resistance to side-channel and fault-injection attacks. Academic initiatives exploring Trusted Execution Environment (TEE)-enabled FPGA platforms such as Keystone focus on secure boot mechanisms, runtime attestation processes, and side-channel mitigation techniques, all tailored to mission-critical embedded systems (Perkins et al., 2025). At the hardware level, security mechanisms like secure bootloaders, runtime attestation modules, and anti-tamper logic are essential to protect cryptographic keys and prevent unauthorized firmware modification. In addition, recent hybrid attestation designs incorporate post-quantum secure configurations to enhance resilience in defense FPGA deployments (Papalamprou et al., 2025). The emergence of heterogeneous architectures combining FPGA fabrics with embedded CPUs and AI accelerators further enables adaptive, autonomous security systems by facilitating simultaneous cryptographic processing and AI-driven analytics in real time (Wang et al., 2025; Silvano et al., 2025).

In preparation for the advent of quantum threats, defense agencies are proactively investing in FPGA-based implementations of post-quantum cryptography (PQC). Extensive research efforts are dedicated to optimizing PQC algorithms to ensure both high performance and robustness, tailored for the stringent requirements of defense environments (Li et al., 2022; Stelzer et al., 2025). Notably, hardware implementations of CRYSTALS-Kyber on FPGAs have demonstrated remarkable high-throughput capabilities, positioning them as compelling candidates for securing communication networks within critical defense and infrastructure systems (Cano Aguilera et al., 2023). For instance, Dang et al. (2022) present the top-level datapath of a CRYSTALS-Kyber hardware accelerator, as illustrated in Figure 10, which details its modular components and optimized data flows designed to efficiently support key generation, encapsulation, and decapsulation on FPGA platforms.

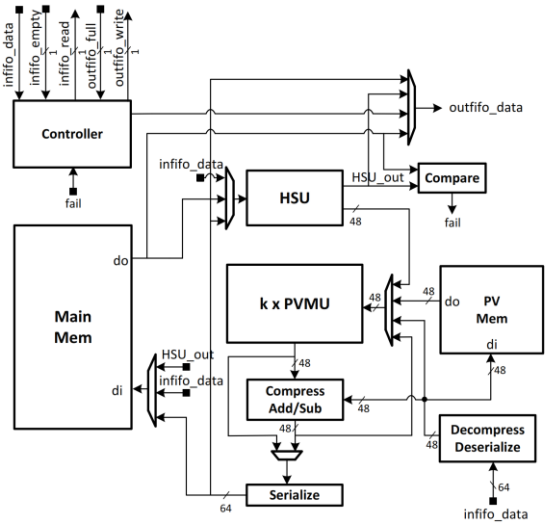

Figure 10. Block diagram of the Kyber top-level datapath. All bus widths are 64 bits unless specified otherwise. (Dang et al., 2022)

Secure Video Processing in Surveillance and Drones: High-bandwidth drone surveillance systems and expansive video networks increasingly rely on end-to-end secure video pipelines. FPGA-based secure DSP architectures enable this by embedding hardware-accelerated encryption, watermarking,

and tamper detection directly into real-time video processing chains, effectively protecting against interception, manipulation, and intellectual property theft. Implementations include innovative techniques such as learned image compression combining quantization-aware watermarking with public-key encryption for low-latency drone video streams (Mazouz et al., 2025), robust watermarking based on DCT and Arnold transforms on Artix 7 FPGAs (Patil et al., 2024), and broader surveys of video security methods including watermarking and tamper detection in surveillance (Asghar et al., 2024). Fusion architectures integrating FPGA and MPSoC hardware have also enabled secure video streaming using SM4 encryption with isolated key storage, especially for IoT surveillance deployments (Liu et al., 2025). Trusted multi-camera outdoor tracking systems powered by FPGAs further demonstrate enhanced tamper resistance in surveillance networks (Al-Shamma & Fadhel, 2024).

Considering the strict constraints of UAVs, such as limited power, weight, and thermal capacity, FPGA solutions optimize secure video analytics by balancing throughput, latency, and resource use. For example, FPGA-accelerated segmentation models achieve 33.63 fps for wildfire drone video analysis on Ultra96-v2 boards (Guarduño-Martinez et al., 2023). Other approaches introduce privacy-preserving live video analytics via split-model inference and obfuscation layers to minimize data leakage risks (Nagasubramaniam et al., 2024). Additionally, adaptive resource scheduling and resilient FPGA mesh network designs improve communication robustness for drone swarms, ensuring tamper-resistant encrypted video streaming (Roy et al., 2024; Gilani et al., 2024). The integration of AI accelerators within FPGA platforms also facilitates encrypted-domain analytics such as real-time anomaly detection and object classification without compromising confidentiality. Examples include lightweight 3D object detection pipelines processing homomorphically encrypted LiDAR and drone video data (Lis et al., 2025), as well as noise-injected feature selection methods that protect privacy in split-model drone video analytics (Nagasubramaniam et al., 2024). Moreover, FPGA-accelerated mesh crypto-communication protocols secure multi-node drone swarm video networks, delivering high-assurance, low-latency encrypted streaming (TrustCom, 2024).

Comparative Analysis of DSP Algorithms, Cryptographic Primitives, and FPGA Platforms

AES Accelerators: Recent FPGA AES implementations target multi-gigabit throughput, ultra-low latency, and efficient resource use within secure DSP systems. Nguyen et al. (2025) introduced AES-RV, an AES accelerator embedded in a RISC-V core on the Xilinx ZCU102, achieving up to 255.97 times speedup and 453.4 times energy efficiency gains, with detailed LUT and slice metrics. Calvo et al. (2024) proposed a dynamic S-Box AES design on Xilinx XC7Z020 via high-level synthesis, delivering pipelined performance with strong side-channel resistance and low area overhead. Other notable works include Prakashan et al.'s configurable AES architecture on Virtex UltraScale (175 MHz operation), Siddiqui and Sekhar's resource-optimized AES on Virtex-7 using minimal LUTs and flip-flops, and Rao and Rao's customized FPGA-friendly AES algorithm targeting Spartan, Virtex, and Kintex platforms that balance latency, power, and area.

ECC Accelerators: FPGA ECC cores are advancing to meet authentication demands in real-time DSP. Nguyen et al. (2025) presented a constant-time ECC processor for Secp256k1 on Kintex-7 operating at 202.3 MHz with 16.2k slices and 54 DSPs. Banerjee and Banerjee (2024) developed an ASIC accelerator supporting Curve25519 and Curve448 with scalar multiplication latencies of 10.36 μs and 54.01 μs, integrating robust side-channel defenses. Hossain et al. (2025) implemented Edwards25519 point multiplication on Virtex-5 with 1.4 ms latency and strong side-channel resistance. Kumari et al. (2025) deployed ECC over GF(2m) on Virtex-7 using hybrid Karatsuba multipliers at 213 MHz, suitable

for IoT DSP authentication. Javeed (2023) designed an ECC scalar core optimizing area-latency trade-offs using Montgomery ladder and Jacobian coordinates across FPGA families.

Kyber and NTT Integration: Efficient Number Theoretic Transform (NTT) cores are critical for embedding CRYSTALS-Kyber in secure DSP pipelines. Rashid et al. (2025) introduced a unified FNTT/INTT butterfly unit with pipelining that optimizes latency and throughput-to-area ratio while minimizing multiplier usage. Saoudi et al. (2024) developed a low-latency NTT tailored for Kyber with strict throughput and area goals. Cheng et al. (2025) proposed an MDC-NTT architecture reducing latency by 38.9% and improving hardware efficiency up to 1.9 times for Kyber encapsulation, enabling practical PQC in real-time contexts. Ahmadi et al. (2024) integrated algorithm-level fault detection into Kyber NTT implementations on Zynq UltraScale+ and Artix-7, achieving near 100% error coverage with minimal area and latency penalties. Sonbul et al. (2025) contributed a shift-add modular reduction and unified butterfly design on Virtex FPGAs to enhance frequency and throughput for secure DSP and PQC workloads, as illustrated in Figures 11 and 12.

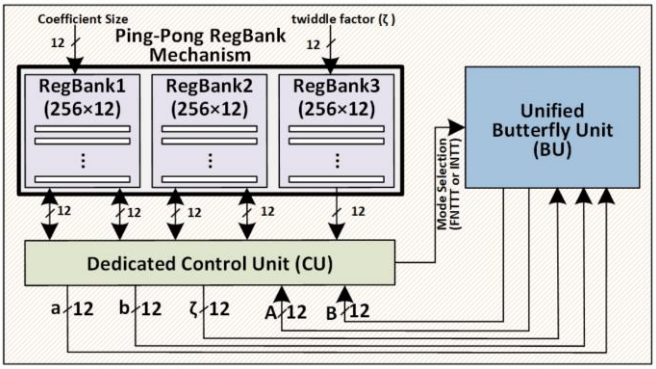

Figure 11. Block Diagram of the Proposed NTT Accelerator Architecture Featuring RegBanks, Unified Butterfly Unit, and Dedicated Control Unit. (Sonbul et al., 2025)

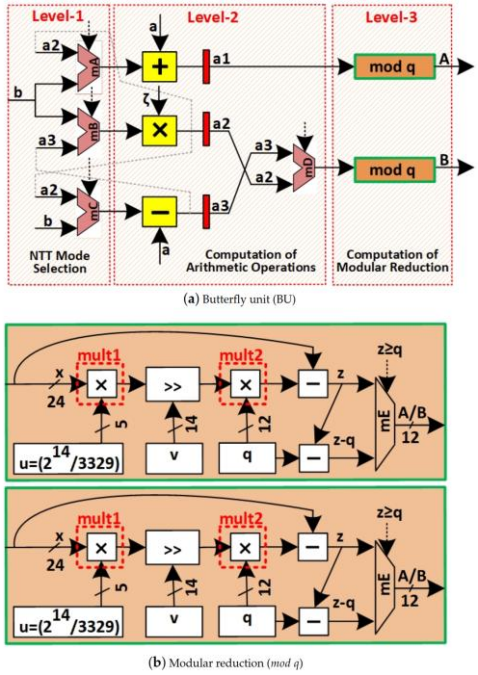

Figure 12. Proposed unified butterfly unit architecture of CT and GS configurations for forward and inverse NTT computation. (Sonbul et al. 2025)

Figure 12a illustrates a design structured across three functional levels: Level-1, Level-2, and Level-3. At Level-1, the control logic determines the computation path based on the selected NTT mode. Specifically, multiplexer mB directs the data flow for the forward NTT (FNTT), while multiplexers mA and mC manage the inverse NTT (INTT) execution. Subsequently, Levels 2 and 3 carry out the required arithmetic operations sequentially, guided by the selection signals and outputs generated at Level-1. Figure 12b presents the architectural design implementing the optimized Barrett reduction algorithm.

To evaluate the practical viability of cryptographic algorithms within secure DSP systems, key attributes such as memory footprint, cryptographic security level, and FPGA deployment suitability are compared. Additionally, the selection of FPGA platforms, from high-performance SoCs to low-power embedded devices, directly affects system performance, power efficiency, and flexibility. Table 7 summarizes widely-used cryptographic schemes, including post-quantum candidates, alongside their compatibility with leading FPGA platforms, offering system designers a comprehensive overview of trade-offs and deployment strategies for embedded secure signal processing applications.

.

Table 7. Comparative Table of DSP Algorithms, Cryptographic Primitives, and FPGA Platforms

Algorithm / Platform	Memory Footprint	Security Level	FPGA Suitability	Typical Applications	Notes / Enhancements
AES (Symmetric)	Low	High (Classical)	Excellent – pipelining & parallelism	Encrypted real-time DSP, low-latency processing	Widely supported in IP cores; power-efficient
ECC (Asymmetric)	Moderate	High (Public-key)	Good – needs optimized modular arithmetic	Key exchange, authentication in embedded systems	Uses projective coordinates, Montgomery multiplication
Kyber (PQC – Lattice)	High	Very High (Post-quantum)	Challenging – demands advanced architectures	Future-proof signal security, government-grade systems	Polynomial ops, memory-banking needed
Dilithium (PQC – Signature)	High	Very High	Moderate – signature schemes on FPGA	Firmware signing, secure boot for DSP hardware	Typically, more latency than encryption schemes
Xilinx Zynq UltraScale+	N/A	N/A	Excellent – HW/SW co-design with ARM cores	Hybrid DSP-crypto workloads	High performance with rich DSP resources
Intel Agilex	N/A	N/A	Strong – includes AI acceleration	AI + DSP + crypto integration	OpenCL/HLS support for flexibility

Metrics: Throughput, Latency, Resource Usage, Power, and Scalability

Throughput and Latency: FPGA architectures inherently enable deep pipelining and loop unrolling, which allow AES implementations to achieve high throughput with low latency. For instance, recent designs such as a 4-stage pipelined AES architecture deliver impressive throughput of 105.7 Gbps and an efficiency of 31.48 Mbps per slice, outperforming many existing implementations in both throughput and latency (Nam et al., 2024). In the domain of elliptic-curve cryptography (ECC), accelerators utilizing Karatsuba-optimized multipliers report key exchange latencies in the range of

hundreds of microseconds. Banerjee and Banerjee (2024) introduced a unified ASIC accelerator for Curve25519 and Curve448, achieving scalar multiplication latencies as low as 10.38 μs and 54.01 μs, respectively, while incorporating energy-efficient side-channel countermeasures.

Resource Utilization: The resource consumption of cryptographic cores on FPGAs varies based on the algorithmic complexity. ECC implementations tend to be compact, making them suitable even for IoT-scale devices. For example, Kumari et al. (2025) presented an ECC design over GF(2^{163}) requiring approximately 14,195 LUTs on a Virtex-7 FPGA, achieved through efficient Karatsuba multipliers. On the other hand, post-quantum cryptography (PQC) accelerators generally demand more Block RAM (BRAM) to manage polynomial vectors. However, optimized designs such as a compact Kyber NTT core utilize only 541 LUTs, four 18 Kb BRAMs, and a limited number of DSP blocks by leveraging multi-port memory and banked pipeline architectures to mitigate resource bottlenecks (Kieu-Do-Nguyen et al., 2024), as detailed in Figure 13 and Table 8. High-throughput AES implementations may consume tens of thousands of LUTs but maintain favorable throughput-per-area efficiency owing to their deeply pipelined architectures (Nam et al., 2024).

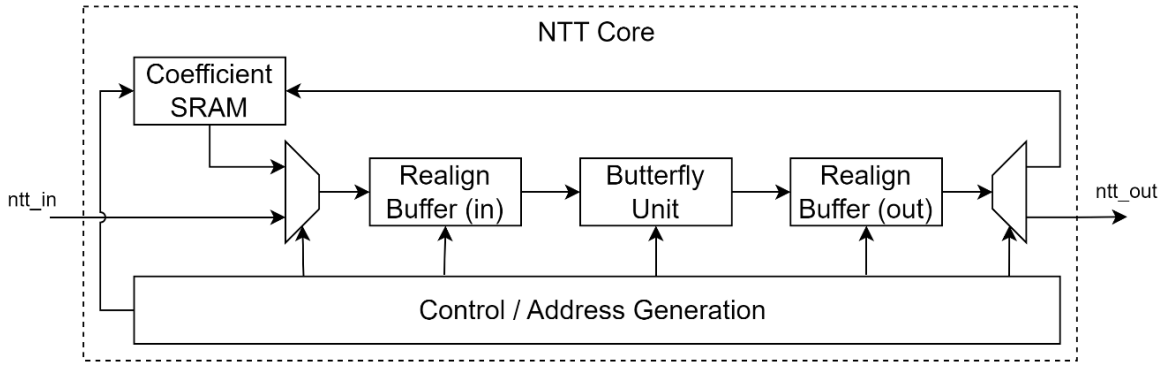

Figure 13. Control/data flow of the proposed NTT core. (Kieu-Do-Nguyen et al., 2024)

Table 8. Result comparison with other lightweight NTT-core methods. (Kieu-Do-Nguyen et al., 2024)

	Mode	LUTs	FFs	DSPs	BRAMs (18 Kb)	Freq (MHz)	Cycles	Time (μs)	Platform	LUT ATP	FF ATP
This (α) ×2	◐	429	538	0	4	446	459	1.03	Artix-7	441.87	554.14
This (β) ×1 (quarter sq.)	◐	379	414	0	1	435	910	2.05	Artix-7	792.11	865.26
[23] (FNTT)	◐	9187	9328	0	0	100	1410	14.10	Virtex-7	129,536.7	131,524.8
This (γ) ×2	●	541	680	0	4	417	461	1.10	Artix-7	595.10	748.00
[27]	●	810	717	4	2	222	324	1.46	Artix-7	1182.60	1046.82
[21]	●	609	640	2	4	257	490	1.91	Artix-7	1163.19	1222.40
[22]	●	1154	1031	2	0	300	456	1.52	Zynq US+	1754.08	1567.12
[26] ×1	●	360	145	3	2	115	940	8.17	Artix-7	2941.20	1184.65
[26] ×2	●	737	290	6	4	115	474	4.12	Artix-7	3036.44	1194.80
[28] ×1	●	948	325	1	5	190	904	4.76	Artix-7	4512.48	1547.00
[28] ×4	●	2543	792	4	18	182	232/233	1.27	Artix-7	3229.61	1005.84
[28] ×16	●	9508	2684	16	70	172	69/71	0.40	Artix-7	3803.20	1073.60
[29]	●	1737	1167	2	3	161	512	3.18	Artix-7	5523.66	3711.06
[23] (UNTT)	●	9298	9402	0	0	20	1410	70.50	Virtex-7	655,509.0	662,841.0
[24]	●	4969	1616	9	35	N/A	126/127	N/A	Artix-7	N/A	N/A
[30]	●	1243	562	11	7	118	933	7.91	Artix-7	9832.13	4445.42

◐: Only NTT, ●: both NTT and INTT

Power Consumption: Power efficiency is a critical factor in embedded cryptographic DSP systems, especially for battery-constrained applications. For example, Baidya et al. (2025) introduced the Modified SampleNTT technique, which reduces the number of SHAKE-128 squeezes from three to two, achieving a 99.16% polynomial success rate while delivering a 33.14% reduction in energy consumption, a 33.32% decrease in latency, and a 0.28% reduction in slice count on Artix-7 FPGAs. Surveys of dynamic voltage and frequency scaling (DVFS) and other power-aware methods for FPGA and ASIC platforms report power savings up to 40% through techniques such as fine-grained power gating and architectural optimizations tailored for cryptographic modules (Parikh, 2025). Additionally, Tran et al. (2024) proposed a 4-stage pipelined AES architecture achieving 105.7 Gbps throughput and 31.48 Mbps per slice efficiency on FPGA, outperforming most existing designs in both throughput and latency.

Scalability and Flexibility: To future-proof secure DSP systems, dynamic scalability is essential for adapting to evolving cryptographic standards and increasing key sizes. Designs leveraging partial and dynamic reconfiguration enable runtime switching between AES and post-quantum cryptography (PQC) engines without requiring full system reboots. Bommana et al. (2025) combined dynamic partial reconfiguration with deep learning to detect and mitigate side-channel attacks in real time during continuous operation. Similarly, reconfigurable bitstream approaches such as SPREAD provide hardware flexibility for adjusting security policies with minimal overhead (Bow et al., 2020). Open-source toolchains designed for side-channel security research also offer modular SoC-based FPGA platforms featuring RISC-V cores and dynamic counters, supporting adaptability to emerging cryptographic requirements (Zoni et al., 2025).

Security and Side-Channel Resistance: Robust hardware countermeasures remain vital for defending against physical side-channel attacks. Techniques including masking, hiding, threshold implementations, and dual-rail logic are increasingly adopted to mitigate leakage. Bommana et al. (2025) introduced a deep learning-driven dynamic partial reconfiguration method that obfuscates side-channel signatures at runtime, substantially increasing attack difficulty. The SPREAD framework employs delay-domain diversity via reconfigurable S-box paths to mitigate Differential Power Analysis (DPA) and Correlation Power Analysis (CPA) attacks with minimal overhead (Bow et al., 2020). Furthermore, Zoni et al. (2025) presented an open-source hardware–software framework for side-channel attack (SCA) research on IoT-class FPGA SoCs, integrating a RISC-V core, debug infrastructure, and built-in support for evaluating attacks and deploying countermeasures.

Cross-Platform Challenges: Despite technological progress, benchmarking cryptographic implementations across different FPGA vendors and platforms remains challenging due to variations in toolchains, IP core availability, and architectural differences. The scarcity of open-source PQC FPGA cores complicates standardization and comprehensive evaluation. For instance, Akçay and Yalçın (2025) utilized the PQClean software as a reference framework to ensure open-source portability in their lightweight ASIP design targeting lattice-based PQC on FPGAs.

Emerging Trends: The convergence of post-quantum cryptography, AI-enhanced security protocols, and adaptive DSP workloads on FPGA platforms is shaping a rapidly evolving research landscape. Architectures capable of dynamically reconfiguring cryptographic kernels in response to changing threat models and workloads are under active development. Concurrently, machine learning-based anomaly detection modules embedded within FPGA fabric enable real-time threat monitoring and mitigation. For example, the fSEAD streaming anomaly detector employs dynamic partial reconfiguration to adapt at runtime to workload variations and potential attacks (Lou et al., 2023). Research also explores trade-offs between security and power efficiency using fine-grained voltage and frequency scaling, approximate computing, and hardware obfuscation. These innovations aim to create

resilient, scalable, and highly efficient secure DSP platforms for future embedded and edge computing systems. Parikh and Parikh (2025) proposed a power-aware, machine learning-driven monitoring framework that combines lightweight sensors, anomaly detection, and hierarchical management to enable scalable security awareness on FPGAs with minimal overhead.

FPGA Family Selection: Choosing the appropriate FPGA family is critical for implementing secure DSP pipelines, particularly as PQC and AI-driven threat detection become integral to next-generation embedded systems. FPGA architectures differ in DSP resource density, AI integration capabilities, and suitability for PQC algorithm implementation, all of which impact overall system performance, power consumption, and security posture. Table 9 summarizes key characteristics of prominent FPGA families from Xilinx and Intel, assisting designers in making informed choices for cryptographically secure, real-time DSP applications.

Table 9. Comparative Analysis of FPGA Families for Secure DSP and Post-Quantum Cryptography Integration

FPGA Family	DSP Resource Density	AI Acceleration	PQC Support (Feasibility)	Power Profile	Typical Use Cases
Xilinx 7-series	Medium	No	Limited	Moderate	Legacy DSP with AES
Zynq UltraScale+	High + ARM	Partial	Good	Moderate	Secure Hybrid Systems
Intel Agilex	Very High	Yes	Excellent	High	PQC + AI fusion
Cyclone V	Low	No	Low	Low	Lightweight crypto-DSP
Arria 10	Moderate	No	Moderate	Medium	Mid-tier embedded security

Challenges and Open Research Problems

Secure DSP systems implemented on FPGAs are gaining increasing traction in real-time and embedded applications. However, several critical challenges and open research problems remain unresolved. Addressing these issues effectively is crucial for advancing the field and developing robust, practical, and scalable solutions that fulfill the demands of modern secure signal processing.

Balancing Cryptographic Strength and Real-Time Performance: As FPGA-based secure DSP systems become central to embedded and real-time applications, numerous research challenges arise in balancing cryptographic strength with system performance. High security, especially with post-quantum cryptography (PQC), typically requires significant computational resources, which can increase latency and power consumption. For example, lightweight PQC schemes proposed by Kundu et al. (2025) demonstrate approximately a twofold improvement in the time-area product compared to traditional Kyber implementations, while maintaining NIST Level I security. These advances position such schemes as promising candidates for latency-sensitive DSP environments, as illustrated in Figure 14 and Table 10.

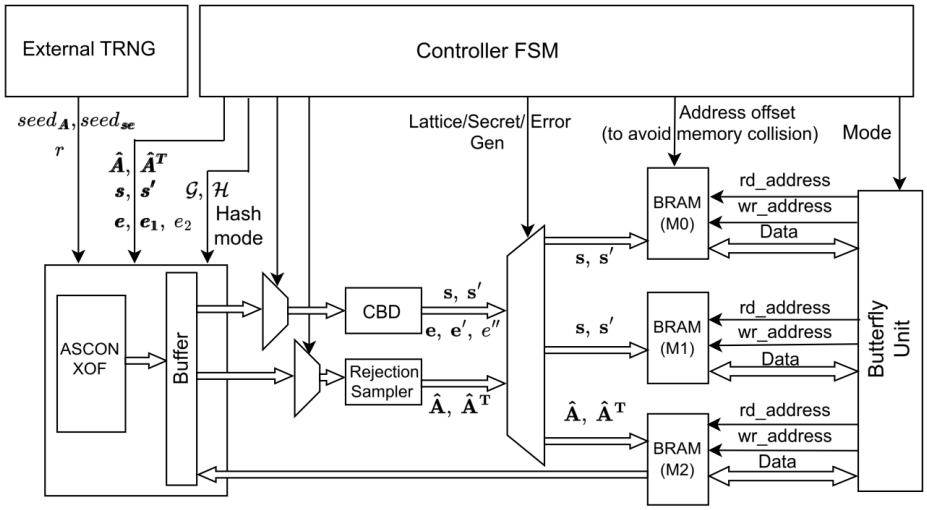

Figure 14. Full system architecture. (Kundu et al., 2025)

Table 10. Comparison of implementation of Rudraksh (KEM-poly64) with the state-of-the-art schemes. Freq. represents frequency, Exec time represents execution time, and k denotes 1000x. KG, Enc, and Dec represent key-generation, encapsulation, and decapsulation, respectively. All the KEMs, except the * and † ones, belong to the NIST-level-1 category. (Kundu et al., 2025)

Scheme	Platform	client/ server	Area LUT/FF/Slice/BRAM/DSP	ENS**	Freq. (MHz)	Exec time(μs) KG/Enc/Dec	T×A (ENS×ms)
Rudraksh (KEM-poly64)	**Virtex-7**	client server	**2813/1494/0/1.5/1** **2869/1413/0/1.5/1**	**1098** **1112**	**400** **431**	**54/70/81**	**60/77/90**
	Artix-7	client server	**2776/1487/0/1.5/1** **2839/1413/0/1.5/1**	**1088** **1104**	**387** **367**	**64/73/96**	**71/79/106**
Kyber[HLLM24]	Kintex-7	client server	4777/2661/1395/2.5/0 4993/2765/1452/2.5/0	3080 3191	244	278/416/552	887/1281/1761
Kyber[ZLZ+22]	Artix-7		8966/9173/3186/10.5/6	8086	204	11.5/17.3/23.5	93/140/190
Kyber[XL21]	Artix-7	client server	6785/3981/1899/3/2 7412/4644/2126/3/2	4384 4767	161	23.4/30.5/41.3	112/134/197
Kyber[DFA+20]	Virtex-7		13745/11107/4590/14/8	11571	245	8.8/12.2/17.9	102/141/207
	Artix-7		11864/10348/3989/15/8	10695	210	-/14.3/20.9	-/153/224
Kyber[BUC19]	Artix-7		14975/2539/4173/14/11	11761	25	2980/5268/5692	35k/62k/67k
Frodo[HOKG18]	Artix-7	client server	6745/3528/1855/1/11 7220/3549/1992/1/16	4838 5593	167 162	20k/20k/21k	110k/96k/116k
NewHope[ZYC+20]	Artix-7		6780/4026/-/7/2	3267	200	21/33/12.5	69/108/41
LightSaber[RB20]	Zynq Ultra-scale+		23686/9805/0/2/0	6314	150	18.4/26.9/33.6	116/170/212
Espada†[KNK+24]			18741/18823/-/14/48	12229	250	92.2/154.3/219.3	1128/1887/2682
Sable†[KNK+24]			17092/11280/-/2/0	4665	250	18.9/23.6/29.0	88/110/135
Florete†[KNK+24]			28281/16029/-/2/140	21462	250	28.3/56.7/84.4	607/1217/1811
NTRU-HRSS701 [DMG23]	Zynq Ultra-scale+	KG Enc Dec	49001/39957/9357/2.5/45 31494/25120/6652/2.5/0 37702/34441/8032/2.5/45	26598 15016 22448	300	172.7/7.4/29.4	4593/111/660
NTRU-HPS677 [DMG23]		KG Enc Dec	41047/39037/7968/6/45 26325/17568/4638/5/0 29935/19511/5217/2.5/45	23906 12200 17691	250 300	192.7/14.7/25.1	4607/179/444
NTRUEncrypt* [KY09]	Virtex-E		27292/5160/14352/-/-	21175	62	-/1.54/1.41	-/33/30
InvRBLWE* [EBSMB19]	Virtex-7	client server	5000/5000/1292/0/0	2542	443 455	0.95/1.97/0.95	2.4/5/2.4
RLWE*[RVM+14]	Virtex-6		1536/953/-/1.5/1	778	278	-/47.9/21	-/37/16
RLWE*[PG13]	Virtex-6		5595/4760/1887/7/1	4757	251	57.9/54.9/35.4	71/67/43

* This scheme is PKE not KEM and only provides CPA security. All other schemes are CCA secure.
† These schemes provides NIST-level-3 security.
** ENS [NKLO23] = Slice + DSP×100 + BRAM×196 + LUT/4

Portability Across FPGA Families and Generations: Achieving portability across different FPGA families and generations remains a significant challenge due to inherent architectural differences and variations in vendor tooling. Current research efforts focus on developing standardized, vendor-agnostic design methodologies and modular, parameterized IP cores to mitigate these issues. For instance, Shrivastava et al. (2025) propose a unified FFT and NTT accelerator that supports both DSP and post-quantum cryptography workloads with minimal resource overhead, approximately a 62% increase in LUT usage but no additional DSP blocks or BRAM. This is achieved by extending conventional

FFT engines to accommodate NTT operations tailored for ML-KEM and ML-DSA applications. Crucially, their design is scalable across both Xilinx and Intel platforms without requiring significant re-engineering, as shown in Figure 15.

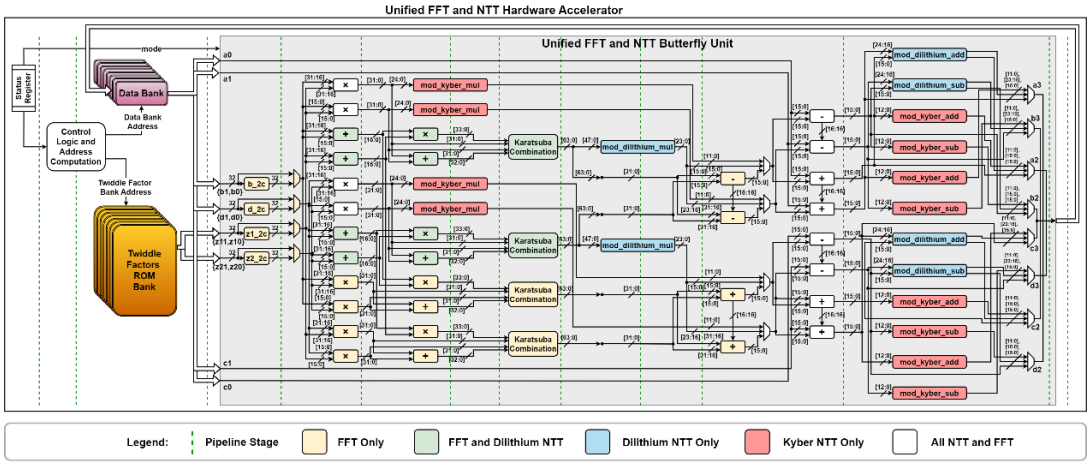

Figure 15. Detailed architecture of the proposed unified butterfly unit and memory organization for FFT and ML-KEM (Kyber) + ML-DSA (Dilithium) NTT. (Shrivastava et al., 2025)

Hardware Resource Constraints on Embedded Platforms: Embedded FPGA platforms face stringent limitations in power, area, and memory, creating significant challenges for the implementation of modern DSP and post-quantum cryptography (PQC) algorithms. To address these constraints, resource-sharing techniques and approximate computing methods offer promising solutions. For example, Baidya et al. (2025) report approximately 33% reductions in both energy consumption and latency for Kyber sampling operations, along with decreased FPGA slice utilization, making their design well-suited for resource-limited DSP applications. Furthermore, approximation-aware synthesis techniques such as QUADOL+ leverage dual-output LUTs to reduce overall LUT usage by up to 18% on contemporary FPGAs, demonstrating an effective approach for integrating DSP and cryptographic acceleration within strict area constraints (Shi et al., 2024).

AI-Based Security: Leveraging Machine Learning for Threat Detection in DSP Systems

AI Models for Anomaly Detection in Signal Streams: Convolutional Neural Networks (CNNs) are particularly effective at capturing subtle spatial and temporal patterns within side-channel leakage or fault-induced anomalies in real-time signal traces. Feng et al. (2025) introduced HACNN-SCA, a hybrid attention-enabled CNN architecture that significantly improves side-channel attack detection on AES power traces, achieving about 75% lower power consumption and faster convergence on the ASCAD dataset compared to conventional methods. Building on this, Bommana et al. (2025) demonstrated the integration of CNN-based detectors within FPGA fabric combined with dynamic partial reconfiguration (DPR) triggered by anomaly detection, enabling seamless real-time mitigation of side-channel leakage without disrupting system operation. Similarly, Parikh and Parikh (2025) implemented a fully on-chip CNN utilizing DSP slices capable of detecting power-transient anomalies in real time independently of CPU resources. Additionally, Li et al. (2024) showcased 1D CNNs applied to images and signal streams for FPGA-based fault detection, achieving sub-millisecond latency and high accuracy with low power overhead.

Autoencoders provide an unsupervised learning approach by capturing compact representations of normal behavior and flagging deviations, making them suitable for identifying zero-day threats in live signal pipelines. Gupta (2025) reported NomAD, an FPGA-based unsupervised anomaly detector combining a Variational Autoencoder and Decision Tree Regression for the ATLAS L1Topo trigger, achieving a 21% increase in unique acceptance on dimuon events while operating at a tunable approximately1.8 kHz trigger rate. Roche et al. (2024) presented a deep decision tree-enhanced autoencoder on the Virtex UltraScale platform, enabling real-time anomaly detection with just 30 ns latency and minimal resource use, suitable for FPGA-based trigger systems. Yang et al. (2024) developed a baseline-optimized autoencoder trained directly on guided wave data for structural health monitoring, showing robust detection even under dynamic environmental conditions with contaminated training sets. Najafi et al. (2024) proposed a hybrid attention autoencoder combining reconstruction and attention mechanisms for time-series anomaly detection, significantly improving accuracy without requiring validation datasets.

Graph Neural Networks (GNNs) have gained attention for modeling complex FPGA interconnects or sensor networks, facilitating detection of coordinated anomalies such as hardware Trojan insertions or fault-induced deviations. Alrahis et al. (2024) modeled FPGA netlists as graphs to identify malicious configurations pre-runtime with 98.24% accuracy by detecting anomalies at the sub-circuit level. Chen et al. (2025) proposed GNN4HT, a two-stage graph neural network that first localizes hardware Trojans (HTs) with a 94.28% true positive rate on the Trust-Hub dataset and then classifies their functionalities via HT information graphs, achieving 80.95% accuracy at the gate level and 62.96% at RTL. Thorat et al. (2024) introduced a two-stage GNN leveraging graph- and node-level classification to detect HTs in large gate-level netlists, reaching up to 98.66% precision and 92.30% recall post-quantization. Ma et al. (2025) proposed a GNN-based HT detection method using harmonic centrality and weight optimization to improve feature representation, with GraphSAGE models achieving an F1 score of 98.59%, surpassing state-of-the-art benchmarks. The Long Short GNN (2025) was designed for high-assurance detection of modified FPGA IP circuits within hardware validation workflows. Lastly, Kose et al. (2024) surveyed GNN applications for anomaly detection, including quantization techniques and FPGA acceleration strategies that enable efficient structured graph modeling in resource-constrained hardware environments.

FPGA Implementations of AI-Based Intrusion Detection Systems: Machine learning (ML) inference engines deployed directly on FPGA fabric enable low-latency, power-efficient threat detection co-located with secure DSP cores. To accommodate complex ML models within the limited resources of FPGAs, hardware-aware neural network design techniques such as quantization, pruning, and neural architecture search are widely used. Frameworks like Xilinx Vitis AI and Intel OpenVINO facilitate optimized deployment by managing dataflow and memory access efficiently. These on-chip accelerators continuously monitor side-channel signals, communication buses, or sensor streams to detect anomalies and trigger proactive countermeasures in real time (Rangsikunpum et al., 2025; Kim et al., 2025; Rak et al., 2024; Ghoi et al., 2025).

Integrating AI Inference with Secure DSP Pipelines: A modular system architecture that tightly integrates AI inference cores, cryptographic accelerators, and DSP units establishes a dynamic, adaptive security environment. Embedding anomaly detection models within the FPGA fabric enables runtime reconfiguration of cryptographic modules, for instance, adjusting key sizes, activating masking layers, or initiating side-channel countermeasures in response to detected threats. This closed feedback loop maintains high system performance during normal operation while escalating protection under attack. Recent works demonstrate power-aware FPGA health monitoring combining lightweight on-chip sensors with ML-based detectors such as One-Class SVMs, isolation forests, and shallow neural

networks coordinated by hierarchical managers (Parikh and Parikh, 2025). Similarly, Lou et al. (2023) introduced a composable streaming ensemble anomaly detection library using partially reconfigurable FPGA regions for dynamic function exchange, while Mao et al. (2025) presented a dynamic Tsetlin machine accelerator capable of rearranging learning logic at runtime without full resynthesis. Nechi et al. (2023) surveyed over 120 FPGA-based deep learning accelerators, benchmarking performance on ResNet-2 and LSTM models. Additionally, Alrahis et al. (2024) proposed a graph neural network (GNN)-based classifier to screen FPGA netlists for malicious cryptographic configurations before deployment, facilitating preemptive reconfiguration or blocking.

AI-Augmented Watermarking for Multimedia and Biometric DSP Applications: In multimedia and biometric DSP contexts, coupling watermark detection with AI inference significantly enhances tamper detection capabilities. AI-based watermark verifiers integrated into DSP pipelines can identify subtle manipulations and activate adaptive protections accordingly. Taj et al. (2024) proposed a zero-watermarking scheme that preserves image integrity while achieving high authentication and tamper detection accuracy in sensitive fields such as medical imaging. Siryeh et al. (2025) developed a DCT-based fragile watermarking method combined with a deep CNN for fingerprint template tamper detection, achieving real-time operation (12–18 ms per template), a detection rate of 98.3%, and strong robustness against compression and noise, as illustrated in Figures 16 and 17 and Table 11. Madhushree et al. (2023) conducted a comprehensive review of watermarking techniques for tamper localization and recovery in medical image authentication, underscoring the need for hybrid AI-augmented schemes. Furthermore, Sengupta and Anshul (2025) introduced a behavioral synthesis-based, cost-effective fingerprint biometric watermarking method optimized for FPGA-accelerated biometric pipelines, enabling efficient watermark detection in loop circuits.

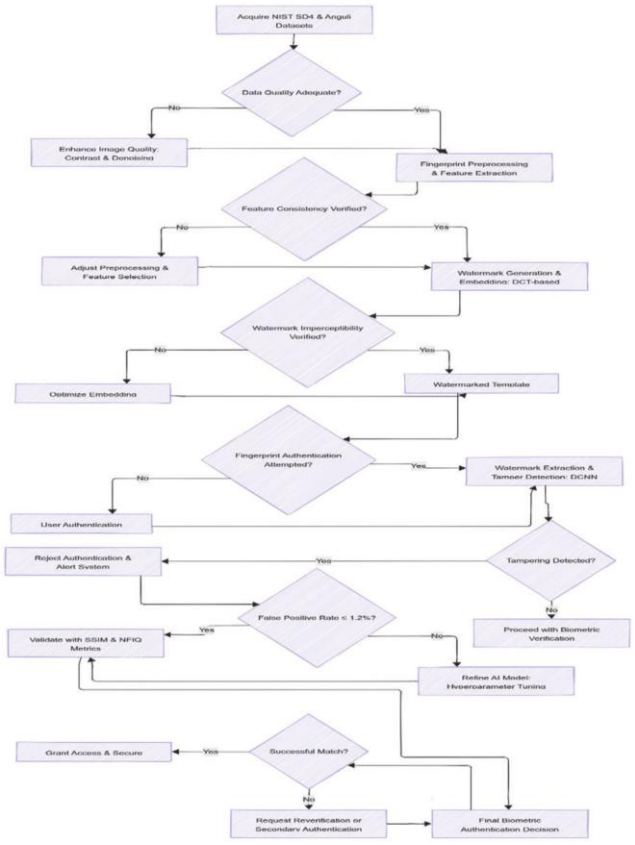

Figure 16. AI-Driven fragile watermarking and tamper detection process, illustrating the end-to-end biometric authentication pipeline, from dataset acquisition to real-time watermark verification and decision-making. (Siryeh et al., 2025)

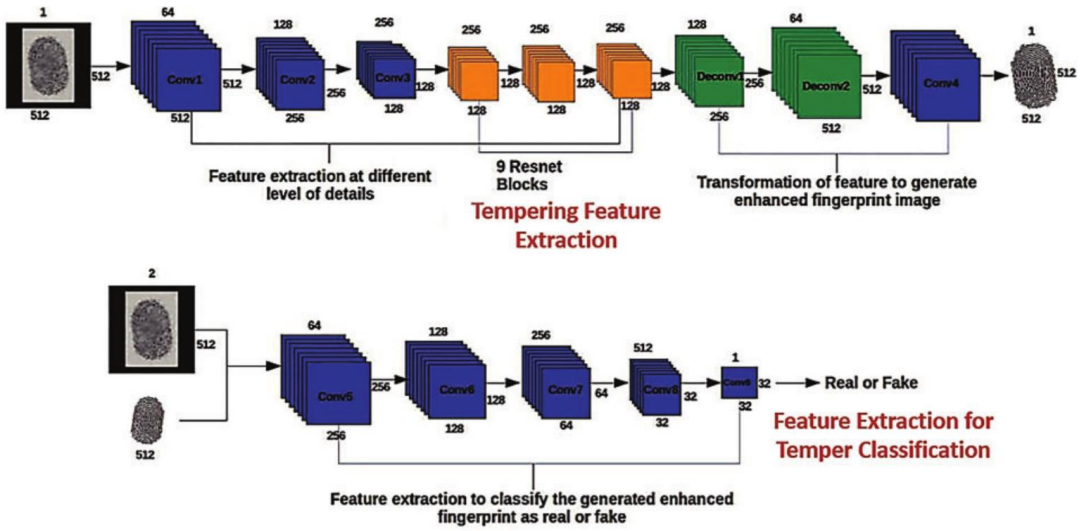

Figure 17. Deep learning-based tamper detection architecture, illustrating feature extraction, transformation, and classification of fingerprint templates to detect tampering. (Siryeh et al., 2025)

Table 11. Performance measures for real-time tamper verification, with important parameters influencing processing speed, feature extraction, and tamper detection accuracy. (Siryeh et al., 2025)

Parameter	Value	Impact on Verification
Input image size	256 × 256 pixels	Ensures uniform fingerprint processing
Processing time per template	12–18 ms	Enables real-time authentication
Feature extraction time	5–8 ms	Efficient extraction of watermark data
False positive rate (%)	≤ 1.2 %	Reduces incorrect tamper flagging
Threshold (τ)	0.9	Correlation score for tamper decision

Machine learning (ML) algorithms are increasingly integrated into secure DSP architectures on FPGAs to enhance threat detection, adaptive security, and data integrity verification. Various ML models are tailored for specific applications such as anomaly detection, side-channel attack mitigation, watermark verification, and signal classification. Evaluating the feasibility of these models for FPGA implementation is crucial, especially regarding resource utilization, latency, and overall architectural impact. Such assessments are vital for designing robust, real-time, and secure DSP systems. Table 12 provides a classification of prominent ML models along with their corresponding application domains within cryptographic signal processing pipelines.

Table 12. Classification of ML Models for Secure DSP Applications on FPGA

Use Case	ML Model Type	Description	FPGA Implementation Considerations	Example Applications
Anomaly Detection	CNN (Convolutional Neural Network)	Detects unusual patterns in signal streams using spatial feature extraction	Moderate resource use; can leverage FPGA parallelism	Intrusion detection in communication signals
Side-Channel Attack (SCA) Mitigation	Autoencoder	Learns normal signal behavior to identify deviations caused by attacks	Lightweight; suitable for on-chip real-time monitoring	Detection of power/EM leakage anomalies
Watermark Detection	Graph Neural Network (GNN)	Models' relationships and structures in multimedia or sensor data	Emerging; may require custom accelerators and optimized data flow	Verifying watermark integrity in multimedia streams
Signal Classification	Recurrent Neural Network (RNN) / LSTM	Processes time-series signal data to classify and predict states	Higher latency; optimized architectures needed for real-time	Biometric authentication, EEG pattern recognition
Adaptive Security	Reinforcement Learning (RL)	Dynamically adjusts cryptographic parameters or filter settings	Complex training; inference can be deployed on FPGA for adaptive control	Real-time tuning of crypto-DSP pipelines

Future Research Directions

The field of secure digital signal processing (DSP) on field-programmable gate arrays (FPGAs) is rapidly evolving, propelled by emerging threats, novel cryptographic paradigms, and advances in heterogeneous computing architectures. This section highlights promising future research directions focused on enhancing the resilience, scalability, and adaptability of FPGA-based DSP systems.

Trends Towards Post-Quantum Secure DSP Systems: With the rise of quantum computing, traditional cryptographic algorithms face obsolescence risks, making the integration of post-quantum cryptography (PQC) into FPGA-based DSP pipelines essential for long-term security. Current research prioritizes optimizing resource-intensive lattice-based schemes such as Kyber and Dilithium to satisfy real-time constraints while minimizing area and power overhead. Novel architectural solutions, including hybrid hardware-software co-design and dynamically reconfigurable cryptographic cores, are actively explored to strike an optimal balance among security, throughput, latency, and energy efficiency in embedded and edge DSP applications.

FPGA + AI Co-Acceleration for Adaptive Signal Security: The synergy of FPGA programmability with AI-driven analytics unlocks adaptive, intelligent security capabilities. Co-accelerating AI inference engines alongside secure DSP cores enables real-time anomaly detection, intrusion prevention, and dynamic cryptographic workload optimization. Research increasingly focuses on embedding machine learning models such as convolutional neural networks (CNNs) and graph neural networks (GNNs) directly within FPGA fabric to deliver low-latency, continuous security adaptation.

Reliable and Reconfigurable Hardware Design: Trustworthy hardware remains fundamental for secure DSP, especially in mission-critical sectors like defense and healthcare. Emerging efforts investigate runtime attestation protocols, anti-tamper boot mechanisms, and secure patching implemented at the FPGA level to ensure system integrity and prevent unauthorized modifications. Reconfigurable architectures facilitate seamless security updates without disrupting service, thus

extending device lifecycles. Additionally, physically unclonable functions (PUFs) and secure key storage embedded in FPGA fabrics are gaining momentum as robust root-of-trust primitives.

FPGA-as-a-Service for Cloud-Edge Cryptographic DSP: Rising demand for scalable cryptographic DSP in cloud and edge environments drives the development of FPGA-as-a-Service (FaaS) frameworks. These platforms enable remote deployment, configuration, and execution of secure DSP workloads on shared FPGA infrastructures with hardware-level isolation and strong security assurances. Key research challenges include secure multi-tenancy, encrypted bitstream management, and ultra-low-latency networking, facilitating cost-effective, privacy-preserving deployments in domains like healthcare analytics and secure communications.

Processing-in-Memory (PIM) and Neuromorphic Approaches: PIM and neuromorphic computing represent emerging paradigms targeting the mitigation of memory bottlenecks and power inefficiencies typical in conventional secure DSP designs. Though at an early stage within FPGA research, recent studies explore integrating computation and memory within FPGA fabrics to accelerate cryptographic primitives and DSP kernels. Neuromorphic-inspired architectures further promise energy-efficient, massively parallel processing suited for AI-enhanced security features. Applying these concepts to FPGA-based secure DSP could yield breakthroughs in low-power, real-time cryptographic processing.

Secure Heterogeneous Architectures: Integrating RISC-V SoCs with FPGA fabric is reshaping secure DSP system design through hardware-software co-design of cryptographic and signal processing functions. Heterogeneous platforms allow control-plane cryptographic protocols to run on RISC-V processors while delegating data-plane intensive workloads to FPGA accelerators. The openness and customizability of the RISC-V ecosystem accelerate security auditing and innovation. Future research will likely focus on modular, scalable platforms that unify AI inference, PQC acceleration, and trusted execution environments, paving the way for versatile, future-proof secure DSP solutions.

Conclusion

This review has presented a comprehensive overview of the current landscape of secure digital signal processing (DSP) implementations on FPGA platforms. Leveraging the intrinsic parallelism and dynamic reconfigurability of FPGAs, cryptographic primitives such as AES, ECC, and emerging post-quantum algorithms can be seamlessly integrated to enable real-time secure processing critical for embedded and edge computing applications. The advent of heterogeneous system-on-chip (SoC) architectures, especially those combining RISC-V cores with FPGA fabrics, opens promising pathways for flexible, low-latency, and scalable secure DSP designs. Comparative analyses of cryptographic algorithms and FPGA families reveal the complex trade-offs designers must consider to optimize performance, resource efficiency, and security robustness.

Industrial deployments of secure DSP on FPGAs span vital domains including medical devices, secure communications, defense, and IoT networks, where guaranteeing data integrity, confidentiality, and real-time responsiveness is paramount. The integration of AI accelerators within FPGA platforms further expands these possibilities by introducing intelligent threat detection and adaptive security mechanisms. From a research perspective, significant opportunities exist to advance the integration of post-quantum cryptography, develop low-power and resource-efficient architectures, and explore runtime-reconfigurable designs. Emerging paradigms such as cloud-edge computing and FPGA-as-a-Service (FaaS) stimulate interdisciplinary innovation at the intersection of hardware architecture, cryptography, and machine learning.

Looking ahead, the fusion of secure DSP pipelines with AI-driven adaptive security, trustworthy hardware design, including runtime attestation and anti-tamper techniques, and scalable post-quantum cryptographic solutions will be pivotal. Overcoming existing challenges such as balancing cryptographic strength with real-time constraints, ensuring portability across heterogeneous FPGA platforms, and managing resource limitations demands novel hardware-software co-design strategies. Close collaboration between academia and industry will be essential to translate research breakthroughs into deployable, resilient systems. The future of secure and efficient DSP on FPGA platforms promises to be dynamic, intelligent, and robust, underpinning the next generation of embedded and edge computing applications.

References

Motahhir, S., & Maleh, Y. (Eds.). (2023). Security engineering for embedded and cyber-physical systems (1st ed.). Boca Raton, FL: CRC Press.

Kajol, M., & Yu, Q. (2025). New security challenges towards in-sensor computing systems. arXiv preprint arXiv:2502.05046.

Narimani, P., Wang, M., Planta, U., & Abbasi, A. (2025). Exploring power side-channel challenges in embedded systems security. arXiv preprint arXiv:2410.11563.

Hosseini, S. M. R., & Pilaram, H. (2024). A comprehensive review of post-quantum cryptography: Challenges and advances. IACR Cryptology ePrint Archive, Paper 2024/1940.

Alnaseri, O., Himeur, Y., Atalla, S., & Mansoor, W. (2025). Complexity of post-quantum cryptography in embedded systems and its optimization strategies. arXiv preprint arXiv:2504.13537.

Karl, P., Antognazza, F., Barenghi, A., Pelosi, G., & Sigl, G. (2025). High-performance FPGA accelerator for the post-quantum signature scheme CROSS. IACR Cryptology ePrint Archive, Paper 2025/1161.

Gladis Kurian, M., & Chen, Y. (2025). Ascon on FPGA: Post-quantum safe authenticated encryption with replay protection for IoT. Electronics, 14(13), 2668.

Arucu, M., & Iliev, T. (2025). Performance evaluation of FPGA, GPU, and CPU in FIR filter implementation for semiconductor-based systems. Journal of Low Power Electronics and Applications, 15(3), 40.

Khan, M. I., & da Silva, B. (2024). Harnessing FPGA technology for energy-efficient wearable medical devices. Electronics, 13(20), 4094.

Ayoub, H. G., Abdulrazaq, Z. A., Fehr, A., & Hasso, S. (2024). Unveiling robust security: Chaotic maps for frequency hopping implementation in FPGA. *Ain Shams Engineering Journal, 15*(11), 103016.

García-Requejo, Á., Hernández, Á., & Pérez-Rubio, M. C. (2025). FPGA-based SoC architecture for real-time device-free localization using a multistatic US sonar. Measurement, 256, 118320.

Khan, L., Isanaka, S. P., & Liou, F. (2024). FPGA-based sensors for distributed digital manufacturing systems: A state-of-the-art review. Sensors, 24(23), 7709.

Sravanthi, M., Gunturi, S. K., Chinnaiah, M. C., Lam, S.-K., Vani, G. D., Basha, M., Janardhan, N., Krishna, D. H., & Dubey, S. (2024). Adaptive FPGA-based accelerators for human–robot interaction in indoor environments. Sensors, 24(21), 6986.

Perepelitsyn, A., & Kulanov, V. (2025). Methods of deployment and evaluation of FPGA as a service under conditions of changing requirements and environments. Technologies, 13(7), 266.

Stoyanov, S., Kakanakov, N., & Marinova, M. (2025). FPGA prototyping of heterogeneous security architecture for educational purposes. Engineering Proceedings, 10, 18.

Nguyen, V. T., Pham, P. H., Le, V. T. D., Pham, H. L., Vu, T. H., & Tran, T. D. (2025). AES-RV: Hardware-efficient RISC-V accelerator with low-latency AES instruction extension for IoT security. arXiv preprint arXiv:2505.11880.

Mazouz, A., Tria, C. D. S., Chaudhuri, S., Fiandrotti, A., Cagnanzzo, M., Mitrea, M., & Tartaglione, E. (2025). Security and real-time FPGA integration for learned image compression. arXiv preprint arXiv:2503.04867.

Nassim, M. K. A., & Zakarya, Z. (2025). FPGA-based implementation of a substitution box cryptographic co-processor for high-performance applications. International Journal of Reconfigurable Embedded Systems, 14(2), 587–596.

Roy, K. S., Sujith, M., Bhanu, B., et al. (2024). FPGA-based dual-layer authentication scheme utilizing AES and ECC for unmanned aerial vehicles. Journal of Wireless Communications and Networking, 2024, 91.

Preethi, P., Ulla, M. M., Yadav, G. P. K., Roy, K. S., Hazarika, R. A., & K. S. K. (2024). Elliptic-curve cryptography implementation on RISC-V processors for Internet of Things applications. Journal of Sensors, 2024, 5116219.

Elhamzi, W. (2024). Enhancing medical image security with FPGA-accelerated LED cryptography and LSB watermarking. Traitement du Signal, 41(1), 85–97.

Azzaz, M. S., Kaibou, R., & Smahi, A. (2023). FPGA implementation using novel co-design approach of real-time speech chaos-based crypto-watermarking prototype. In Proceedings of the Conference on Advances in Electronics, Control and Communication Systems (ICAECCS).

Wikipedia contributors. (2025). Hardware watermarking. In Wikipedia, The Free Encyclopedia. Retrieved July 27, 2025

Kim, H., Park, J., Lee, H., Won, D., & Han, M. (2024). An FPGA-accelerated CNN with parallelized sum pooling for onboard real-time routing in dynamic low-orbit satellite networks. Electronics, 13(12), 2280.

Syed, F., Ali, W., Kumar, A., & Bakhsh, F. I. (2023). Implementation of FIR digital filters on FPGA board for real-time audio processing. In Proceedings of the 2023 International Conference on Technology and Policy in Energy and Electric Power (ICT-PEP) (pp. 233–237).

Hao, Z., Liu, C., & Ouyang, S. (2024). Study on the optimization process and application of FIR digital filter. Applied Computing Engineering, 72(1), 107–113.

Lahti, S., & Hämäläinen, T. D. (2025). High-level synthesis for FPGAs—A hardware engineer's perspective. IEEE Access, 13, 28574–28593.

Duarte, J., Han, S., Harris, P., Jindariani, S., Kreinar, E., Kreis, B., Ngadiuba, J., Pierini, M., Rivera, R., Tran, N., & Wu, Z. (2018). Fast inference of deep neural networks in FPGAs for particle physics. Journal of Instrumentation, 13(07), P07027.

Forelli, R. F., Shi, R., Ogrenci, S., & Agar, J. (2024). A high-level synthesis methodology for dynamic monitoring of FPGA ML accelerators. In Proceedings of the 2024 IEEE 42nd VLSI Test Symposium (VTS) (pp. 1–5).

Curzel, S., Jovic, S., Fiorito, M., Tumeo, A., & Ferrandi, F. (2023). MLIR loop optimizations for high-level synthesis: A case study. In Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT) (pp. 544–545).

Xiao, Y., Luo, Z., Zhou, K., & Liang, Y. (2024). Cement: Streamlining FPGA hardware design with cycle-deterministic eHDL and synthesis. In Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA) (pp. 211–222).

Abbaszadeh, M., & How, D. L. (2024). From topology to realization in FPGA/VPR routing. In Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA) (pp. 85–96).

Park, D., & DeHon, A. (2024). REFINE: Runtime execution feedback for incremental evolution on FPGA designs. In Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA) (pp. 108–118).

Pouchet, L.-N., Tucker, E., Zhang, N., Chen, H., Pal, D., Rodríguez, G., & Zhang, Z. (2024). Formal verification of source-to-source transformations for HLS. In Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA) (pp. 97–107).

Xu, X., & Zhang, J. (2020). Rethinking FPGA security in the new era of artificial intelligence. In Proceedings of the 2020 21st International Symposium on Quality Electronic Design (ISQED) (pp. 46–51).

Proulx, A., Chouinard, J.-Y., Fortier, P., & Miled, A. (2023). A survey on FPGA cybersecurity design strategies. ACM Transactions on Reconfigurable Technology and Systems, 16(2), 20.

Liu, Z., Lu, Z., Huang, L., Yao, Z., Lu, Z., & Zhang, J. (2024). Recent advances on reliability of FPGAs in a radiation environment. Microelectronics Journal, 148, 106176.

Jung, S., & Choi, J. P. (2019). Predicting system failure rates of SRAM-based FPGA on-board processors in space radiation environments. Reliability Engineering & System Safety, 183, 374–386.

Sadeghi, S. (2024). Classifying FPGA technology in digital signal processing: A review. International Journal of Engineering Technology Science, 10.

Bommana, S. R., Veeramachaneni, S., Ershad, S., & Srinivas, M. B. (2025). Mitigating side channel attacks on FPGA through deep learning and dynamic partial reconfiguration. Scientific Reports, 15(1).

Qasaimeh, M., Denolf, K., Lo, J., Vissers, K., Zambreno, J., & Jones, P. H. (2019). Comparing energy efficiency of CPU, GPU and FPGA implementations for vision kernels. In Proceedings of the 2019 IEEE International Conference on Embedded Software and Systems (ICESS) (pp. 1–8).

Li, S., Chen, Y., Chen, L., Liao, J., Kuang, C., Li, K., Liang, W., & Xiong, N. (2023). Post-quantum security: Opportunities and challenges. Sensors, 23(21), 8744.

Malik, A., & Islam, S. M. N. (2025). Quantum computing and cybersecurity: Navigating threats and opportunities. In S. B. Goyal, V. Kumar, S. M. N. Islam, & D. Ghai (Eds.), Quantum computing, cyber security and cryptography (pp. xx–xx). Singapore: Springer.

Allgyer, W., White, T., & Youssef, T. A. (2024). Securing the future: A comprehensive review of post-quantum cryptography and emerging algorithms. In Proceedings of SoutheastCon 2024 (pp. 1282–1287).

Aydeger, A., Zeydan, E., Yadav, A. K., Hemachandra, K. T., & Liyanage, M. (2024). Towards a quantum-resilient future: Strategies for transitioning to post-quantum cryptography. In Proceedings of the 2024 15th International Conference on Network of the Future (NoF) (pp. 195–203).

Stelzer, T., Oberhansl, F., Schupp, J., et al. (2025). Enabling lattice-based post-quantum cryptography on the OpenTitan platform. Journal of Cryptographic Engineering, 15, 11.

Nguyen, T.-H., Kieu-Do-Nguyen, B., Pham, C.-K., & Hoang, T.-T. (2024). High-speed NTT accelerator for CRYSTAL-Kyber and CRYSTAL-Dilithium. IEEE Access, 12, 34918–34930.

Cheng, S., Chen, J., Li, J., Yao, K., Gao, S., Rui, K., & Cui, Y. (2025). Optimized design and implementation of CRYSTALS-KYBER based on MLWE. Security and Privacy.

Bote, M., & Diaz-Vargas, J. (2025). STRU: A variant of NTRU. AIMS Mathematics and Computation, 19(3), 853–871.

Liu, Y. K., & Moody, D. (2024). Post-quantum cryptography and the quantum future of cybersecurity. Physical Review Applied, 21(4), 040501.

Dam, D.-T., Nguyen, T.-H., Tran, T.-H., Kieu-Do-Nguyen, B., Hoang, T.-T., & Pham, C.-K. (2024). An efficient method for accelerating Kyber and Dilithium post-quantum cryptography. Proceedings of the 21st Annual International Conference on Privacy, Security and Trust (PST), 1–5.

Li, A., Li, Z., Tang, J., & Lu, Y. (2024). KDA: Kyber and Dilithium accelerator for CRYSTALS suite of post-quantum cryptography in hybrid multipath delay commutator pipelined architecture. Proceedings of the IEEE Asian Solid-State Circuits Conference (A-SSCC), 1–3.

Kim, Y., Yoon, S., & Seo, S. C. (2024). Vectorized implementation of Kyber and Dilithium on 32-bit Cortex-A series. IEEE Access, 12, 104414–104428.

Kumar, T. M., Reddy, K. S., Rinaldi, S., Parameshachari, B. D., & Arunachalam, K. (2021). A low area high speed FPGA implementation of AES architecture for cryptography application. Electronics, 10(16), 2023.

Siddiqui, A. F., & Sekhar, P. C. (2024). FPGA acceleration of AES algorithm for high-performance cryptographic applications. International Journal of Information Technology and Computer Engineering, 4(4), 1–11.

Prakashan, P., Gupta, H., Sivanandan, N., & K. S. (2024). A configurable AES implementation on FPGA for secure 5G systems. Proceedings of the IEEE Recent Advances in Intelligent Computational Systems (RAICS).

Priya, S. S. S., Karthigaikumar, P., & Teja, N. R. (2022). FPGA implementation of AES algorithm for high-speed applications. Analog Integrated Circuits and Signal Processing, 112, 115–125.

Dhanda, S., Singh, B., Jindal, P., & Panwar, D. (2022). A highly efficient FPGA implementation of AES for high throughput IoT applications. Journal of Discrete Mathematical Sciences and Cryptography, 25(7), 2029–2038.

Sunil, J., H. S. S., B. K. S., & Santhameena, S. (2020). Implementation of AES algorithm on FPGA and on software. Proceedings of the IEEE International Conference for Innovation in Technology (INOCON), 1–4.

Hossain, M. R., Rahman, M. S., Zaman, K. S., Bhuiyan, W. E. F., Kang, C. C., Yew, T. J., & Miraz, M. H. (2025). Low latency FPGA implementation of twisted Edward curve cryptography hardware accelerator over prime field. Scientific Reports, 15, 15097.

Kumari, R., Purohit, G., & Karmakar, A. (2025). Efficient hardware implementation of modular multiplier over GF(2^m) on FPGA. arXiv preprint arXiv:2506.09464.

Javeed, K., El-Moursy, A., & Gregg, D. (2024). E^2CSM: Efficient FPGA implementation of elliptic curve scalar multiplication over generic prime field GF(p). Journal of Supercomputing, 80, 50–74.

Lin, J.-L., Zheng, P.-Y., & Chao, P. C.-P. (2023). A new ECC implemented by FPGA with favorable combined performance of speed and area for lightweight IoT edge devices. Microsystem Technologies, 30(12), 1537–1546.

Wang, D., Lin, Y., Hu, J., Zhang, C., & Zhong, Q. (2023). FPGA implementation for elliptic curve cryptography algorithm and circuit with high efficiency and low delay for IoT applications. Micromachines, 14(5), 1037.

Dong, X., Zhang, L., & Gao, X. (2018). An efficient FPGA implementation of ECC modular inversion over F256. Proceedings of the 2nd International Conference on Cryptography, Security and Privacy (ICCSP), 29–33.

Holla, A., Shetty, A. S., Santhameena, S., & Harshraj, S. (2023). Implementation of a modified SHA-3 hash function on FPGA. Proceedings of the 4th IEEE Global Conference for Advancement in Technology (GCAT).

Sideris, A., Sanida, T., & Dasygenis, M. (2024). Hardware acceleration design of the SHA-3 for high throughput and low area on FPGA. Journal of Cryptographic Engineering, 14, 193–205.

Sideris, A., & Dasygenis, M. (2023). Enhancing the hardware pipelining optimization technique of the SHA-3 via FPGA. Computation, 11(8), 152.

Kieu-Do-Nguyen, B., Hoang, T.-T., Tsukamoto, A., et al. (2022). High-performance multi-function HMAC-SHA2 FPGA implementation. Proceedings of the IEEE Interregional NEWCAS Conference (NEWCAS).

OpenTitan Project. (2025). HMAC IP. OpenTitan Documentation.

Yasin, H. M., Sallow, A. B., & Mahmood, R. Z. (2023). High-speed FPGA-based video watermarking using LSB technique in the spatial domain. International Journal of Intelligent Systems and Applications in Engineering, 12(8s), 644–653.

Aissaoui, N., Azzaz, M. S., & Kaibou, R. (2022). FPGA implementation of a digital watermarking system using DWT transformation. Proceedings of the 2nd International Conference on Advanced Electrical Engineering (ICAEE).

Hussain, S., Sheybani, N., Neekhara, P., Zhang, X., Duarte, J., & Koushanfar, F. (2022). FastStamp: Accelerating neural steganography and digital watermarking of images on FPGAs. Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD).

Janaki, K., Srinivasan, C., & Malini, A. H. (2024). FPGA-enhanced system-on-chip for finger vein-based biometric system using novel DL model. Integration, 102231.

Altman, M. B., Wan, W., Hosseini, A. S., et al. (2024). Machine learning algorithms for FPGA implementation in biomedical engineering applications: A review. Heliyon, 10(6), e26652.

Mekhfioui, M., El Bazi, N., Laayati, O., Satif, A., Bouchouirbat, M., Kissi, C., Boujiha, T., & Chebak, A. (2025). Optimized digital watermarking for robust information security in embedded systems. Information, 16(4), 322.

Gull, S., & Parah, S. A. (2023). Advances in medical image watermarking: A state-of-the-art review. Multimedia Tools and Applications, 1–41.

Mazouz, A., De Sousa Tria, C., Chaudhuri, S., Fiandrotti, A., Cagnanzzo, M., Mitrea, M., & Tartaglione, E. (2025). Security and real-time FPGA integration for learned image compression. arXiv preprint arXiv:2503.04867.

GAPses Project. (2024). Versatile smart glasses for comfortable and fully-dry acquisition and parallel ultra-low-power processing of EEG and EOG. arXiv preprint arXiv:2406.07903.

Abdelaziz, A., Fathi, A., & Fares, A. (2025). Protecting intellectual property of EEG-based neural networks with watermarking. arXiv preprint arXiv:2502.05931.

Microchip Technology Inc. (2024). PolarFire® SoC FPGAs overview & architecture. Datasheet DS00003292B.

Nguyen, V. T., Pham, P. H., Le, V. T. D., Pham, H. L., Vu, T. H., & Tran, T. D. (2025). AES-RV: Hardware-efficient RISC-V accelerator with low-latency AES instruction extension for IoT security. arXiv preprint arXiv:2505.11880.

Ma, K.-M., Le, D.-H., Pham, C.-K., & Hoang, T.-T. (2023). Design of an SoC based on 32-bit RISC-V processor with low-latency lightweight cryptographic cores in FPGA. Future Internet, 15(5), 186.

Kieu-Do-Nguyen, B., Nguyen, K.-D., Dang, T.-K., The Binh, N., Pham-Quoc, C., Tran, N.-T., Pham, C.-K., & Hoang, T.-T. (2024). A trusted execution environment RISC-V system-on-chip compatible with transport layer security 1.3. Electronics, 13, 2508.

Srivastava, A., Miftah, S. S., Kim, H., Pal, D., & Basu, K. (2025). PoSyn: Secure power side-channel aware synthesis. arXiv preprint arXiv:2506.08252.

Calvo, H., David, N., Madani, M., & Bourennane, E.-B. (2024). FPGA implementation of AES-based on optimized dynamic S-box. Proceedings of the 21st International Conference on Security and Cryptography (SECRYPT).

Sarma, N., Thakur, A. S., & Karfa, C. (2024). MaskedHLS: Domain-specific high-level synthesis of masked cryptographic designs. arXiv preprint arXiv:2407.11806.

Zhang, L., Mu, D., Hu, W., et al. (2019). Memory-based high-level synthesis optimizations security exploration on the power side-channel. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, PP (99), 1–1.

Bhashini, H. M., Venkateshwarlu, C., Nagaraju, S., & Puppala, G. B. (2025). Analysis of Vivado implementation strategies regarding side-channel leakage for FPGA-based AES implementations. International Journal of Scientific Research in Engineering and Management, 9(5), 9.

Koufopoulou, A.-A., Xevgeni, K., Papadimitriou, A., et al. (2022). Security and reliability evaluation of countermeasures implemented using high-level synthesis. Proceedings of the IEEE International Symposium on On-Line Testing and Robust System Design (IOLTS).

Xiong, C., Liu, C., Li, H., & Li, X. (2025). HLSPilot: LLM-based high-level synthesis. Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD).

ACM/SIGDA. (2023). Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA '23).

Malal, A., & Tezcan, C. (2025). First fully pipelined high throughput FPGA implementation and GPU optimization of wider variant of AES. Preprint.

Tran, S. N., Dung, L. T., & Long, N. V. (2024). A high throughput, low latency 105 Gbps four-pipeline stage AES. Journal of Science and Technology on Information Security.

Sumit, D., Singh, B., Jindal, P., & Panwar, D. (2022). A highly efficient FPGA implementation of AES for high throughput IoT applications. Journal of Discrete Mathematical Sciences and Cryptography, 25(7), 2029–2038.

Tan, W., Lao, Y., & Parhi, K. K. (2023). KyberMat: Efficient accelerator for matrix-vector polynomial multiplication in CRYSTALS-Kyber scheme via NTT and polyphase decomposition. Proceedings of the IEEE/ACM International Conference on Computer Aided Design (ICCAD).

Mandal, S., & Basu Roy, D. (2023). KiD: A hardware design framework targeting unified NTT multiplication for CRYSTALS-Kyber and CRYSTALS-Dilithium on FPGA. arXiv preprint arXiv:2311.04581.

Rashid, M., Sonbul, O. S., Jamal, S. S., Jaffar, A. Y., & Kakhorov, A. (2025). Hardware design of FNTT and INTT of CRYSTALS-Kyber PQC algorithm. Information, 16(1), 17.

Taghavi, B., Azarderakhsh, R., & Mozaffari Kermani, M. (2025). ParallelNTT: Maximizing performance of forward and inverse NTT on FPGA for ML-DSA and ML-KEM. Proceedings of the Great Lakes Symposium on VLSI (GLSVLSI '25), 372–378.

Carril, X., Kardaris, C., Ribes-González, J., Farràs, O., Hernandez, C., Kostalabros, V., González-Jiménez, J. U., & Moretó, M. (2024). Hardware acceleration for high-volume operations of CRYSTALS-Kyber and CRYSTALS-Dilithium. ACM Transactions on Reconfigurable Technology and Systems, 17(3), Article 41, 1–26.

Bao, T., He, P., Fujimoto, D., Hayashi, Y., & Xie, J. (2025). CHIRP: Compact and high-performance FPGA implementation of unified hardware accelerators for ring-binary-LWE-based PQC. ACM Transactions on Reconfigurable Technology and Systems, 18(2), Article 19, 1–27.

Harvie, L. (2024). Optimizing FPGA designs for high-speed networking applications. Embedded Development Blog.

Barge, S., & Gerardine, M. (2024). Low power techniques for internet of things implementation: A review. Multidisciplinary Reviews, 7(12), 2024306.

Vaithianathan, M., Patil, M., Ng, S. F., & Udkar, S. (2024). Low-power FPGA design techniques for next-generation mobile devices. ESP International Journal of Advancements in Computational Technology, 2(2), 82–93.

Chowdhury, P., & Schafer, B. C. (2021). Leveraging automatic high-level synthesis resource sharing to maximize dynamical voltage overscaling with error control. ACM Transactions on Design Automation of Electronic Systems, 27(2), Article 14, 1–18.

Grycel, J. T., & Walls, R. J. (2019). DRAB-LOCUS: An area-efficient AES architecture for hardware accelerator co-location on FPGAs. arXiv preprint.

Tibaldi, M., & Pilato, C. (2023). A survey of FPGA optimization methods for data center energy efficiency. IEEE Transactions on Sustainable Computing, 8(3), 343–362.

Gu, J., Wang, H., Guo, X., Schulz, M., & Gerndt, M. (2025). VersaSlot: Efficient fine-grained FPGA sharing with big.little slots and live migration in FPGA cluster. arXiv preprint.

Zidar, J., Matic, T., Aleksi, I., & Hocenski, Z. (2024). Dynamic voltage and frequency scaling as a method for reducing energy consumption in ultra-low-power embedded systems. Electronics, 13(5), 826.

Nguyen, V. T., Pham, P. H., Le, V. T. D., Pham, H. L., Vu, T. H., & Tran, T. D. (2025). AES-RV: Hardware-efficient RISC-V accelerator with low-latency AES instruction extension for IoT security. arXiv preprint.

Karakchi, R., Stahle-Smith, R., Chinnasami, N., & Yu, T. (2025). Toward a lightweight, scalable, and parallel secure encryption engine. arXiv preprint.

Butt, S. A., Reynolds, B., Ramamurthy, V., Xiao, X., Chu, P., Sharifian, S., Gribok, S., & Pasca, B. (2024). if-ZKP: Intel FPGA-based acceleration of zero knowledge proofs. arXiv preprint.

Xiphera. (2024). Xiphera's crypto module offers customisable offload and acceleration solutions.

Roy, K. S., Sujith, M., Bhanu, B., et al. (2024). FPGA-based dual-layer authentication scheme utilizing AES and ECC for unmanned aerial vehicles. Journal of Wireless Communications and Networking, 2024(91).

Parisi, E., Musa, A., Ciani, M., Barchi, F., Rossi, D., Bartolini, A., & Acquaviva, A. (2024). Assessing the performance of OpenTitan as cryptographic accelerator in secure open-hardware system-on-chips. In Proceedings of the 21st ACM International Conference on Computing Frontiers (pp. 172–179).

Ji, Y., & Dubrova, E. (2025). A side-channel attack on a masked hardware implementation of CRYSTALS-Kyber. Journal of Cryptographic Engineering, 15(7).

Ahmadi, M. M., Alrahis, L., Sinanoglu, O., & Shafique, M. (2023). FPGA-Patch: Mitigating remote side-channel attacks on FPGAs using dynamic patch generation. arXiv preprint.

Bommana, S. R., Veeramachaneni, S., Ershad, S., et al. (2025). Mitigating side channel attacks on FPGA through deep learning and dynamic partial reconfiguration. Scientific Reports, 15, 13745.

Grosso, V., & Lara Nino, C. A. (2025). Leveraging the resources of modern FPGAs in the design of high-performance masked architectures. In Proceedings of XVIII Reunión Española sobre Criptología y Seguridad de la Información (pp. 211–216).

Sengupta, A., & Chaurasia, R. (2025). Secure implantable cardiac pacemaker for medical consumer electronics. npj Biomedical Innovation, 2, 5.

Vaithianathan, M., Udkar, S., Reddy, M., Rajasekaran, S., et al. (2024). FPGA-based smart health monitoring systems for wearable devices.

Yuksel, B. B., & Metin, A. Y. (2024). Advancing biomedical signal security: Real-time ECG monitoring with chaotic encryption. arXiv preprint.

Vaithianathan, M., Patil, M., Ng, S. F., & Udkar, S. (2024). Energy-efficient FPGA design for wearable and implantable devices. International Journal of Advanced Science and Technology, 2(2), 1–10.

Cano Aguilera, A., Rubio Garcia, C., Lawo, D., et al. (2024). In-line rate encrypted links using pre-shared post-quantum keys and DPUs. Scientific Reports, 14(21227), 1–10.

Baidya, P., Paul, R., Srivastava, V., & Debnath, S. K. (2025). Energy-efficient NTT sampler for Kyber benchmarked on FPGA. arXiv preprint.

Stelzer, T., Oberhansl, F., Schupp, J., et al. (2025). Extended version: enabling lattice-based post-quantum cryptography on the OpenTitan platform. Journal of Cryptographic Engineering, 15, 11.

Kundu, S., Ghosh, A., Karmakar, A., Sen, S., & Verbauwhede, I. (2025). Rudraksh: A compact and lightweight post-quantum key-encapsulation mechanism. arXiv preprint.

Kieu-Do-Nguyen, B., The Binh, N., Pham-Quoc, C., Nghi, H. P., Tran, N.-T., Hoang, T.-T., & Pham, C.-K. (2024). Compact and low-latency FPGA-based number theoretic transform architecture for CRYSTALS Kyber postquantum cryptography scheme. Information, 15(7), 400.

Dang, V. B., Mohajerani, K., & Gaj, K. (2022). High-speed hardware architectures and FPGA benchmarking of CRYSTALS-Kyber, NTRU, and Saber. IEEE Transactions on Computers, 1–14.

Yan, B., Cao, D., Jiang, X., Chen, Y., Dai, W., Dong, F., Huang, W., Zhang, T., Gao, C., Chen, Q., Yan, Z., & Wang, Z. (2024). FedEYE: A scalable and flexible end-to-end federated learning platform for ophthalmology. Patterns, 5(2), 100928.

Kaur, J., Canto, A. C., Mozaffari Kermani, M., & Azarderakhsh, R. (2023). A comprehensive survey on the implementations, attacks, and countermeasures of the current NIST lightweight cryptography standard. arXiv preprint.

Aminifar, A., Shokri, M., & Aminifar, A. (2024). Privacy-preserving edge federated learning for intelligent mobile-health systems. Future Generation Computer Systems, 161, 625–637.

2024 IEEE International Conference on Big Data (BigData 2024). (2024). Washington, DC, USA, Dec. 15–18.

Kim, D., Oh, K., Lee, Y., & Woo, H. (2025). Overview of fair federated learning for fairness and privacy preservation. Expert Systems with Applications, 293, 128568.

Nguyen, T., Anh, H., Nguyen, H., & Kiet, T. (2024). Efficient number theoretic transform accelerator for CRYSTALS-Kyber. Indonesian Journal of Electrical Engineering and Computer Science, 33(2), 795–803.

Papalamprou, I., Fotos, N., Chatzivasileiadis, N., Angelogianni, A., Masouros, D., & Soudris, D. (2025). Post-quantum and blockchain-based attestation for trusted FPGAs in B5G networks. In Proceedings of the 2025 IEEE International Symposium on Circuits and Systems (pp. 1–5).

Al-Dabbagh, R., Alkhatib, M., & Albalawi, T. (2025). Efficient post-quantum cryptography algorithms for auto-enrollment in public key infrastructure. Electronics, 14, 1980.

Mirigaldi, M., Piscopo, V., Martina, M., & Masera, G. (2025). The quest for efficient ASCON implementations: A comprehensive review of implementation strategies and challenges. Chips, 4, 15.

Armanuzzaman, M., Sadeghi, A.-R., & Zhao, Z. (2022). Building your own trusted execution environments using FPGA. arXiv preprint.

Perkins, G., Macht, B., Ritzdorf, L., Running Crane, T., LaMeres, B., Izurieta, C., & Reinhold, A. M. (2024). SoK: Trusted execution in SoC-FPGAs. arXiv preprint.

Zou, Y., Li, Y., Wang, S., Su, L., Gu, Z., Lu, Y., Guan, Y., Niu, D., Gao, M., Xie, Y., & Li, F. (2025). Salus: A practical trusted execution environment for CPU-FPGA heterogeneous cloud platforms. In Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (Vol. 4, pp. 252–266).

Moraitis, M. (2023). FPGA bitstream modification: Attacks and countermeasures. IEEE Access.

Mishra, J., & Sahay, S. K. (2025). Modern hardware security: A review of attacks and countermeasures. arXiv preprint.

Perkins, G., Macht, B., Ritzdorf, L., Running Crane, T., LaMeres, B., Izurieta, C., & Reinhold, A. M. (2025). SoK: Trusted execution in SoC-FPGAs. arXiv preprint.

Wang, Z., Zhang, J., Huang, H., Li, Y., Zhu, X., Sun, M., Yang, Z., Ma, D., Tang, H., Pan, G., Wu, F., He, B., & Alonso, G. (2025). FpgaHub: FPGA-centric hyper-heterogeneous computing platform for big data analytics. arXiv preprint.

Silvano, C., Ielmini, D., Ferrandi, F., Fiorin, L., Curzel, S., Benini, L., Conti, F., Garofalo, A., Zambelli, C., Calore, E., et al. (2025). A survey on deep learning hardware accelerators for heterogeneous HPC platforms. ACM Computing Surveys, 57(11), Article 286.

Li, H., Tang, Y., Que, Z., & Zhang, J. (2022). FPGA accelerated post-quantum cryptography. IEEE Transactions on Nanotechnology, 21, 685–691.

Ni, Z., Khalid, A., & O'Neill, M. (2022). High performance FPGA-based post quantum cryptography implementations. In Proceedings of the 32nd International Conference on Field-Programmable Logic and Applications (pp. 456–457).

Andrzejczak, M. (2019). The low-area FPGA design for the post-quantum cryptography proposal Round5. In Federated Conference on Computer Science and Information Systems (pp. 213–219).

Stelzer, T., Oberhansl, F., Schupp, J., et al. (2025). Extended version: enabling lattice-based post-quantum cryptography on the OpenTitan platform. Journal of Cryptographic Engineering, 15, 11.

Cano Aguilera, A., Rubio Garcia, C., Lawo, D., et al. (2024). In-line rate encrypted links using pre-shared post-quantum keys and DPUs. Scientific Reports, 14, 21227.

Patil, A. N., Ingale, V. V., Dalal, F. A., & Agarwal, V. (2024). Implementation of robust and secure watermarking algorithm on FPGA using DCT. WSEAS Transactions on Signal Processing, 20, 54–59.

Asghar, A., Shifa, A., & Asghar, M. (2024). Survey on video security: Examining threats, challenges, and future trends. Computers, Materials & Continua, 80(3), 3591–3635.

Liu, Z., Wang, Z., Tu, S., Wang, H., Fan, J., & Ren, C. (2025). Real-time secure video streaming system based on FPGA and CUDA technology. In Proceedings of the 14th International Conference on Communication and Network Security (ICCNS ’24) (pp. 146–152).

Al-Shamma, O., & Fadhel, M. A. (2024). Trusted outdoor multi-camera tracking system powered by FPGA. Journal of Engineering Research, In Press.

Guarduño-Martinez, E., Ciprian-Sanchez, J., Valente, G., Garcia, V., Rodriguez-Hernandez, G., Palacios-Rosas, A., Rossi-Tisson, L., & Ochoa-Ruiz, G. (2023). An FPGA smart camera implementation of segmentation models for drone wildfire imagery. In Proceedings of the 22nd Mexican International Conference on Artificial Intelligence (MICAI).

Nagasubramaniam, P., Wu, C., Sun, Y., Karamchandani, N., Zhu, S., & He, Y. (2024). Privacy-preserving live video analytics for drones via edge computing. Applied Sciences, 14(22), 10254.

Han, Y., Yu, J., Zuo, P., Wei, Z., Jin, X., Dou, K., Guo, C., & Xu, H. (2023). Security resource scheduling algorithm for intelligent UAV communication network based on optimized subgraph isomorphism and link partition. Electronics, 12(9), 2096.

Gilani, S. Y., Anjum, A., Khan, A., Khan, A., Syed, M. H., Moqurrab, S. A., & Srivastava, G. (2024). A robust internet of drone’s security surveillance communication network based on IOTA. Internet of Things, 1–21.

Lis, K., Kryjak, T., & Gorgoń, M. (2025). LiFT: Lightweight, FPGA-tailored 3D object detection based on LiDAR data. In J. Lorandel & A. Kamaleldin (Eds.), Design and Architecture for Signal and Image Processing (DASIP 2025) (Lecture Notes in Computer Science, Vol. 15569).

IEEE. (2024). 2024 IEEE 23rd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom 2024), Sanya, China, 17–21 December 2024. IEEE.

Calvo, H., David, N., Madani, M., & Bourennane, E.-B. (2024). FPGA implementation of AES-based on optimized dynamic s-box. In Proceedings of the 21st International Conference on Security and Cryptography.

Siddiqui, A. F., & Sekhar, P. C. (2024). FPGA acceleration of AES algorithm for high-performance cryptographic applications. International Journal of Information Technology and Computer Engineering, 4(44), 1–11.

Rao, G. M., & Rao, K. D. (2025). An optimised AES algorithm and its FPGA implementation for secure information. International Journal of Engineering Systems Modelling and Simulation.

Nguyen, T., Nguyen, K., Nguyen, V., Nguyen, H., & Tran, L. (2025). Next-generation ECC processor on FPGA: Leveraging Koblitz curves for enhanced performance. Ain Shams Engineering Journal, 16(9), 103495.

Banerjee, A., & Banerjee, U. (2024). A high-performance Curve25519 and Curve448 unified elliptic curve cryptography accelerator. In Proceedings of the 2024 IEEE High Performance Extreme Computing Conference (HPEC) (pp. 1–7).

Kumari, R., Rout, T., Saini, B., Pandey, J. G., & Karmakar, A. (2025). An efficient hardware implementation of elliptic curve point multiplication on FPGA. In Lecture Notes in Electrical Engineering (Vol. 1210, pp. 223–236).

Javeed, K. (2023). FPGA implementation of area-time aware ECC scalar multiplication core. In Proceedings of the 2023 30th IEEE International Conference on Electronics, Circuits and Systems (ICECS).

Saoudi, M., Kermiche, A., Benhaddad, O. H., Guetmi, N., & Allailou, B. (2024). Low latency FPGA implementation of NTT for Kyber. Microprocessors and Microsystems, 107, 105059.

Ahmadi, K., Aghapour, S., Kermani, M. M., & Azarderakhsh, R. (2024). Efficient algorithm level error detection for number-theoretic transform used for Kyber assessed on FPGAs and ARM. arXiv preprint.

Sonbul, O. S., Rashid, M., & Jaffar, A. Y. (2025). Accelerating CRYSTALS-Kyber: High-speed NTT design with optimized pipelining and modular reduction. Electronics, 14, 2122.

Nam, T. S., Dung, L. T., & Long, N. V. (2024). A high throughput, low latency 105 Gbps four-pipeline stage AES. Journal of Science and Technology on Information Security.

Parikh, R. (2025). Low-power techniques for FPGA and ASIC design: A comprehensive survey. Preprints.

Bommana, S. R., Veeramachaneni, S., Ershad, S., & Srinivas, M. B. (2025). Mitigating side channel attacks on FPGA through deep learning and dynamic partial reconfiguration. Scientific Reports, 15(1), 13745.

Bow, I., Bete, N., Saqib, F., Che, W., Patel, C., Robucci, R., Chan, C., & Plusquellic, J. (2020). Side-channel power resistance for encryption algorithms using implementation diversity. Cryptography, 4, 13.

Zoni, D., Galimberti, A., & Galli, D. (2025). An FPGA-based open-source hardware-software framework for side-channel security research. IEEE Transactions on Computers, 74(6), 2087–2100.

Akçay, L., & Yalçın, B. Ö. (2025). Lightweight ASIP design for lattice-based post-quantum cryptography algorithms. Arabian Journal for Science and Engineering, 50, 835–849.

mupq. (2025). pqm4 [Source code].

Lou, B., Boland, D., & Leong, P. H. W. (2023). fSEAD: a composable FPGA-based streaming ensemble anomaly detection library. ACM Transactions on Reconfigurable Technology and Systems (TRETS), 16(3), Article 42.

Parikh, R., & Parikh, K. (2025). A scalable and power-aware security health monitor for FPGAs using lightweight sensors and ML-based inference. Preprints.

Kundu, S., Ghosh, A., Karmakar, A., Sen, S., & Verbauwhede, I. (2025). Rudraksh: A compact and lightweight post-quantum key-encapsulation mechanism. IACR Transactions on Cryptographic Hardware and Embedded Systems, (2), 647–680.

Shi, J., Wang, X., Meng, C., & Qian, W. (2024). QUADOL: A quality-driven approximate logic synthesis method exploiting dual-output LUTs for modern FPGAs. arXiv preprint.

Feng, T., Gao, H., Li, X., & Liu, C. (2025). Side channel attacks on convolutional neural networks based on the hybrid attention mechanism. SN Applied Sciences, 7(5), 390.

Shrivastava, R., Ratnala, C. P., Puli, D. M., & Banerjee, U. (2025). A unified hardware accelerator for fast Fourier transforms and number theoretic transform. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 1–5.

Li, J., Fu, Y., Yan, D., Ma, S. L., & Sham, C.-W. (2024). An edge AI system based on FPGA platform for railway fault detection. In Proceedings of IEEE 13th Global Conference on Consumer Electronics (GCCE).

Gupta, R. (2025). NomAD: Real-time unsupervised anomaly detection at the ATLAS Level-1 Trigger. Paper presented at the Phenomenology Symposium (PHENO).

Roche, S. T., Bayer, Q., Carlson, B. T., et al. (2024). Nanosecond anomaly detection with decision trees and real-time application to exotic Higgs decays. Nature Communications, 15, 3527.

Yang, K., Liu, T., Yang, Z., & Zhou, Y. (2025). Baseline optimized autoencoder-based unsupervised anomaly detection in uncontrolled dynamic structural health monitoring. Structural Health Monitoring.

Najafi, S. A., Asemani, M. H., & Setoodeh, P. (2024). Attention and autoencoder hybrid model for unsupervised online anomaly detection. arXiv preprint.

Alrahis, L., Nassar, H., Krautter, J., Gnad, D., Bauer, L., Henkel, J., & Tahoori, M. (2024). MaliGNNoma: GNN-based malicious circuit classifier for secure cloud FPGAs. In Proceedings of IEEE International Symposium on Hardware Oriented Security and Trust (HOST).

Chen, L., Dong, C., Wu, Q., Liu, X., Guo, X., Chen, Z., Zhang, H., & Yang, Y. (2025). GNN4HT: A two-stage GNN-based approach for hardware Trojan multifunctional classification. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 44(1), 172–185.

Thorat, K., Hasan, A., Ding, C., & Shi, Z. (2025). TROJAN-GUARD: Hardware Trojans detection using GNN in RTL designs. arXiv preprint.

Ma, P., Li, J., Liu, H., Shi, J., Zhang, S., Pan, W., & Hao, Y. (2025). Hardware Trojan detection methods for gate-level netlists based on graph neural networks. IEEE Transactions on Computers, 74, 1470–1481.

"Long Short GNN," in Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS), 2025

Kose, H. T., Nunez-Yanez, J., Piechocki, R., & Pope, J. (2024). A survey of computationally efficient graph neural networks for reconfigurable systems. Information, 15(7), 377.

Rangsikunpum, A., Amiri, S., & Ost, L. (2025). Ultra-lightweight and highly efficient pruned binarised neural networks for intrusion detection in in-vehicle networks. Electronics, 14(9), 1710.

Kim, D., Im, H., & Lee, S. (2025). Adaptive autoencoder-based intrusion detection system with single threshold for CAN networks. Sensors, 25, 4174.

Rak, T., & Rzonca, D. (2024). Security and privacy in networks and multimedia. Electronics, 13(15), 2887.

Choi, M., Lee, M., Im, H., Lee, J., & Lee, S. (2025). Shallow learning-based intrusion detection system for in-vehicle network: ASIC implementation. Electronics, 14(4), 683.

Mao, G., Rahman, T., Maheshwari, S., Pattison, B., Shao, Z., Shafik, R., & Yakovlev, A. (2025). Dynamic Tsetlin machine accelerators for on-chip training at the edge using FPGAs. arXiv preprint.

Nechi, A., Groth, L., Mulhem, S., & Merchant, F. (2023). FPGA-based deep learning inference accelerators: Where are we standing? ACM Transactions on Reconfigurable Technology and Systems, 16(4), Article 3613963.

Taj, R., Tao, F., Kanwal, S., et al. (2024). A reversible-zero watermarking scheme for medical images. Scientific Reports, 14, 17320.

Siryeh, F. A., Alrammahi, H., & İbrahim, A. A. (2025). Tamper detection in multimodal biometric templates using fragile watermarking and artificial intelligence. Computers, Materials & Continua.

Madhushree, B., Basanth Kumar, H. B., & Chennamma, H. R. (2023). An exhaustive review of authentication, tamper detection with localization and recovery techniques for medical images. Multimedia Tools and Applications, 83(13), 1–43.

Sengupta, A., & Anshul, A. (2025). Exploring low overhead fingerprint biometric watermark for loop pipelined hardware IPs during behavioral synthesis. Journal of Information Security and Applications, 90, 104041.