STM JOURNALS

Journal of
Electronic Design Technology

ISSN: 2229-6980 (Online)
ISSN: 2321-4228 (Print)
Volume 17, Issue 2, 2026
May–August
DOI (Journal): 10.37591/JoEDT

https://journals.stmjournals.com/joedt

Review JoEDT

Hybrid Quantum-Classical Reinforcement Learning Enabled Thermal-Aware Electronic Design Automation Framework for Energy-Efficient Next-Generation VLSI Systems Applications

Bibhu Prasad Ganthia*

Abstract

Modern very large-scale integration (VLSI) systems are becoming more complicated, which has increased need for sophisticated Electronic Design Automation (EDA) frameworks that can concurrently optimise thermal behaviour, power consumption, and performance. This study proposes a Hybrid Quantum-Classical Reinforcement Learning (HQCRL) Enabled Thermal-Aware EDA Framework for next-generation energy-efficient VLSI systems. The proposed framework integrates quantum-inspired optimization techniques with classical reinforcement learning algorithms to address the challenges of placement, routing, and thermal hotspot mitigation during the design process. A thermal prediction engine continuously monitors temperature distribution across the chip and provides feedback to the learning agent for adaptive decision-making. The quantum-classical optimization layer enhances exploration capabilities, enabling faster convergence toward optimal design configurations while minimizing computational overhead. Experimental evaluations on benchmark VLSI circuits demonstrate significant improvements in power efficiency, thermal stability, and resource utilization compared with conventional EDA approaches. The framework achieves reduced peak temperature, lower leakage power, improved routing efficiency, and shorter optimization time, thereby enhancing overall chip reliability and lifespan. The proposed methodology offers a scalable and intelligent solution for future semiconductor technologies, where energy constraints and thermal management are critical design considerations. The results indicate that the integration of quantum-inspired learning with thermal-aware design automation can substantially improve the efficiency and sustainability of advanced VLSI architectures.

Keywords: Hybrid quantum-classical computing, reinforcement learning, electronic design automation (EDA), thermal-aware optimization, VLSI systems, energy efficiency, chip thermal management, quantum-inspired optimization, semiconductor design, intelligent CAD systems

***Author for Correspondence**
Bibhu Prasad Ganthia
E-mail: jb.bibhu@gmail.com

[1]Assistant Professor, Department of Electrical Engineering, Indira Gandhi Institute of Technology, Sarang, Dhenkanal, Odisha, India

Received Date: June 19, 2026
Accepted Date: June 19, 2026
Published Date: June 23, 2026

Citation: Bibhu Prasad Ganthia. Hybrid Quantum-Classical Reinforcement Learning Enabled Thermal-Aware Electronic Design Automation Framework for Energy-Efficient Next-Generation VLSI Systems Applications. Journal of Electronic Design Technology. 2026; 17(2): 10–22p.

INTRODUCTION

The rapid advancement of semiconductor technology and the increasing demand for high-performance computing systems have significantly accelerated the complexity of Very Large-Scale Integration (VLSI) design in recent years. Modern electronic devices, including artificial intelligence accelerators, edge computing platforms, Internet of Things (IoT) nodes, autonomous systems, and data center processors, require integrated circuits capable of delivering enhanced computational performance while maintaining stringent power and thermal constraints. As transistor dimensions

continue to be scaled down into the deep sub-micron and nanometer regimes, power density has emerged as a critical challenge affecting circuit reliability, operational stability, and overall system longevity. Consequently, thermal management has become an essential consideration in the electronic design automation (EDA) workflow. Excessive heat generation can lead to increased leakage currents, timing violations, performance degradation, electromigration, and premature device failure, making thermal-aware optimization a fundamental requirement for next-generation VLSI systems. Traditional EDA tools primarily focus on performance, area, and power optimization but often treat thermal effects as secondary design considerations, resulting in suboptimal solutions under dynamic operating conditions. The emergence of intelligent optimization techniques, particularly machine learning and reinforcement learning, has opened new opportunities for developing adaptive design methodologies capable of addressing multi-objective optimization problems in complex semiconductor environments. Reinforcement learning enables autonomous agents to learn optimal design strategies through interactions with the design environment, thereby reducing their dependency on handcrafted heuristics and exhaustive search methods. However, conventional reinforcement learning algorithms frequently encounter limitations related to convergence speed, exploration efficiency, and scalability when applied to large-scale VLSI design spaces. Simultaneously, quantum computing and quantum-inspired optimization techniques have attracted considerable attention because of their potential to solve highly complex combinatorial optimization problems more efficiently than classical approaches. By leveraging quantum superposition principles and hybrid optimization strategies, quantum-classical systems can enhance search capabilities and improve solution quality in large design spaces. Despite these promising developments, the integration of quantum-inspired optimization with reinforcement learning for thermal-aware EDA applications remains relatively uncharted. Existing research predominantly focuses on isolated optimization objectives, such as placement, routing, power minimization, and thermal prediction, with limited attention given to unified frameworks capable of simultaneously addressing energy efficiency and thermal reliability. Furthermore, the dynamic thermal behavior of advanced VLSI architectures requires intelligent and adaptive optimization mechanisms that can continuously respond to changing workloads and environmental conditions. To address these challenges, this study proposes a Hybrid Quantum-Classical Reinforcement Learning (HQCRL)-enabled thermal-aware EDA framework for energy-efficient next-generation VLSI systems. The proposed framework combines the adaptive learning capabilities of reinforcement learning with the enhanced exploration efficiency of quantum-inspired optimization to achieve superior thermal and power optimization during the design process. A dedicated thermal prediction module continuously evaluates the chip temperature distribution and provides real-time feedback to the learning agent for intelligent decision-making. The primary objectives of this study were to reduce thermal hotspots, minimize power consumption, improve routing efficiency, enhance design convergence speed, and increase overall chip reliability. The major contributions include the development of a novel hybrid optimization architecture, integration of thermal-aware feedback mechanisms into the EDA workflow, implementation of a quantum-classical learning engine for design space exploration, and comprehensive performance evaluation against conventional optimization methods. The proposed framework aims to provide a scalable and intelligent solution for future semiconductor technologies, where thermal efficiency, energy conservation, and design automation are increasingly critical requirements.

Related Work and Research Gap

Recent advances in semiconductor technologies, in-memory computing architectures, photonic integrated circuits, and intelligent optimization methods have significantly influenced the evolution of Electronic Design Automation (EDA) frameworks. Thermal-aware EDA techniques have become increasingly important because of the continuous scaling of transistor dimensions and the resulting increase in power density within modern VLSI systems. Researchers have explored several hardware- and architecture-level approaches to mitigate thermal hotspots and improve energy efficiency. A phase-change memory (PCM)-based analog in-memory computing readout scheme incorporating drift compensation through reference conductance tracking demonstrated improved computational stability

and reduced thermal sensitivity in memory-intensive applications [1]. The thermal robustness was enhanced through a differential PCM cell architecture that compensated for the conductance drift and improved the reliability under varying operating temperatures [2]. An improved PCM programming circuit is capable of accelerating memory operations while reducing energy consumption and heat generation during programming cycles [3]. PCM-based in-memory computing systems have been reviewed, and their potential for low-power, high-density computing architectures suitable for future VLSI platforms has been highlighted [4]. The fabrication and integration of photonic devices for phase-change memory and neuromorphic computing have demonstrated opportunities for energy-efficient hardware acceleration with improved thermal characteristics [5]. A low-power analog in-memory computing accelerator was designed using 18 nm FD-SOI PCM technology, which achieved a high computational density with significantly reduced power dissipation [6]. A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference, illustrating the benefits of hardware-aware optimization for energy-efficient computing systems [7]. In the domain of photonic computing, rewritable photonic integrated circuits and programmable phase-change metasurfaces are capable of supporting advanced neural network operations with lower energy requirements and enhanced scalability [8, 9]. Dielectric-assisted phase-change material waveguides for rewritable photonic integrated circuits, contributing to thermally efficient photonic computing platforms [10]. Although these studies significantly advance device- and architecture-level thermal management, they primarily focus on hardware innovation rather than intelligent design automation strategies capable of dynamically optimizing thermal behavior during the design process.

The application of reinforcement learning (RL) in VLSI optimization has emerged as a promising research direction owing to its ability to learn adaptive optimization policies within complex design environments. Previous studies have demonstrated the effectiveness of intelligent algorithms in nonlinear system analysis, control optimization, and multi-agent decision-making. Nonlinear algorithms for frequency parameter detection in track circuits, highlighting the potential of intelligent learning mechanisms in complex electronic systems [11]. Nonlinear characteristics of GaN power switching devices using advanced simulation methodologies that support optimization-driven electronic design processes [12]. A smart-grid-based multi-agent system framework was proposed that demonstrated the effectiveness of distributed intelligence for complex energy management applications [13]. An adaptive controller design methodology for autonomous underwater vehicles was proposed, illustrating the benefits of intelligent control systems in dynamic environments [14]. Multi-objective evolutionary optimization for modular unmanned aerial vehicle platform design, showing that intelligent optimization techniques can efficiently navigate large design spaces [15]. Additional studies involving hybrid power systems, signal reconstruction, IoT communication optimization, and harmonic mitigation further demonstrate the growing role of artificial intelligence and adaptive learning algorithms in engineering optimization problems [16–19]. However, most existing approaches focus on control systems, communication networks, and power applications, rather than integrating reinforcement learning directly into thermal-aware EDA workflows for VLSI design optimization.

Recently, quantum computing and hybrid quantum-classical optimization techniques have attracted significant attention because of their ability to address large-scale combinatorial optimization problems that challenge conventional algorithms. Quantum-inspired search mechanisms provide enhanced exploration capabilities and improved convergence characteristics, making them attractive for complex VLSI design tasks involving placement, routing, thermal balancing, and power optimization. Nevertheless, the current literature predominantly investigates quantum techniques at the device, memory, and computational architecture levels, with limited emphasis on their integration into reinforcement-learning-driven EDA environments. Existing thermal-aware EDA methodologies often rely on static optimization rules, whereas reinforcement learning approaches frequently suffer from slow convergence and limited exploration efficiency in high-dimensional design spaces. Furthermore, available quantum and quantum-inspired optimization methods rarely incorporate real-time thermal feedback mechanisms capable of adapting design decisions according to the evolving chip temperature distributions. Therefore, a significant research gap exists in the development of a unified framework

that combines thermal-aware design automation, reinforcement learning intelligence, and quantum-classical optimization. To address this limitation, this study proposes a hybrid quantum-classical reinforcement learning-enabled thermal-aware electronic design automation framework for energy-efficient next-generation VLSI systems. The proposed framework integrates a thermal prediction engine, reinforcement learning-based decision-making module, and quantum-inspired optimization layer to simultaneously minimize power consumption, reduce thermal hotspots, accelerate convergence, and improve the overall VLSI design efficiency. By bridging the gap between thermal-aware EDA, intelligent learning, and quantum-enhanced optimization, the proposed solution provides a scalable pathway for the development of reliable, energy-efficient, and high-performance semiconductor systems for future computing applications.

PROPOSED HYBRID QUANTUM-CLASSICAL REINFORCEMENT LEARNING FRAMEWORK

System Architecture and Design Flow

The proposed Hybrid Quantum-Classical Reinforcement Learning (HQCRL) framework is designed to provide an intelligent thermal-aware Electronic Design Automation (EDA) environment for next-generation VLSI systems. The framework integrates classical reinforcement learning, quantum-inspired optimization, thermal prediction mechanisms, and conventional EDA tools into a unified optimization platform. The architecture consists of five major layers: design input, thermal analysis, reinforcement learning, quantum optimization, and design validation. Initially, circuit netlists, floor plans, technology libraries, and design constraints are supplied to the system. These inputs are processed through thermal and power estimation modules that continuously evaluate the impact of design decisions on the chip temperature distribution. The generated information is forwarded to the learning engine, enabling adaptive optimization during the design process. Unlike traditional EDA methodologies that apply thermal corrections after placement and routing, the proposed framework incorporates thermal awareness from the earliest stages of design exploration. This proactive approach minimizes the number of thermal hotspots and reduces the need for expensive redesign iterations. This architecture enables the simultaneous consideration of power consumption, thermal distribution, routing congestion, timing constraints, and resource utilization, thereby creating a multi-objective optimization environment suitable for highly complex semiconductor systems.

The design flow begins with preprocessing and feature extraction from the target VLSI design. Structural parameters such as gate density, interconnect length, switching activity, logic utilization, and floor planning constraints were converted into a state-space representation suitable for machine learning analysis. These features were continuously updated as the optimization process progressed. The reinforcement learning agent observes the current design state and evaluates multiple optimization actions, including component-placement adjustments, routing modifications, power redistribution strategies, and thermal balancing operations. Each design action produces measurable changes in the system performance and thermal characteristics. These outcomes were then analyzed using a reward evaluation mechanism, which calculates a composite score based on energy efficiency, thermal stability, timing performance, and area utilization. Through iterative learning cycles, the system gradually develops optimization policies that maximize the design quality while minimizing thermal risks. The continuous interaction between the design states, optimization actions, and performance feedback enables the framework to discover highly efficient design configurations without relying solely on predefined heuristics or exhaustive search techniques.

Multi-Objective Thermal-Aware Optimization Function

$$J_{opt} = \min\left[\alpha \sum_{i=1}^{N} P_i + \beta \sum_{i=1}^{N} \left(T_i - T_{ref}\right)^2 + \gamma \sum_{j=1}^{M} D_j + \delta \sum_{k=1}^{L} C_k + \eta \sum_{m=1}^{R} V_m\right]$$

A significant feature of the proposed architecture is the integration of quantum-inspired search mechanisms into the reinforcement-learning workflow. Traditional reinforcement learning systems often encounter local optima when exploring large and complex design spaces. To overcome this

limitation, the framework employs a quantum-classical optimization layer capable of generating diverse candidate solutions via probabilistic state exploration. This mechanism enhances the agent's ability to evaluate a broader range of design alternatives while maintaining its computational efficiency. The quantum-inspired module periodically injects optimized candidate configurations into the learning process, thereby improving the exploration diversity and accelerating convergence. The interaction between classical learning policies and quantum-enhanced search enables more efficient navigation of high-dimensional VLSI optimization landscapes. Consequently, the framework achieves a superior design quality compared with conventional optimization approaches that depend solely on deterministic algorithms. This hybrid methodology is particularly advantageous for modern semiconductor technologies, where the number of optimization variables can reach millions across large-scale integrated circuits.

The final stage of the architecture focuses on design verification, and implementation readiness. Once the optimization process converges, the generated design is subjected to comprehensive validation procedures, including thermal verification, power analysis, timing closure assessment, and routing integrity evaluation. The framework continuously compares the optimized results with predefined design objectives to ensure compliance with industrial standards and manufacturing requirements. If additional optimization is necessary, feedback from the verification modules can be reintroduced into the learning environment. This iterative verification process improves the design robustness and minimizes the likelihood of post-fabrication failure. Therefore, the architecture provides an end-to-end intelligent EDA solution capable of adapting to evolving design constraints and emerging semiconductor technologies. By combining thermal awareness, reinforcement learning intelligence, and quantum-inspired optimization within a unified workflow, the proposed framework establishes a scalable foundation for future VLSI design automation systems that simultaneously prioritize energy efficiency, thermal reliability, and computational performance.

Thermal Prediction and Monitoring Module

The thermal prediction and monitoring module serves as the core intelligence component responsible for maintaining thermal stability throughout the VLSI-design process. As transistor densities continue to increase, thermal effects have become one of the most critical challenges affecting circuit reliability, power efficiency, and operational lifespan of devices. Excessive temperature accumulation can cause leakage current escalation, timing degradation, electromigration, and accelerated device aging. To address these issues, the proposed framework incorporates a real-time thermal prediction engine capable of continuously estimating the temperature distribution across an integrated circuit. The module collects information from design parameters, such as switching activity, power dissipation profiles, transistor density, interconnect utilization, and clock distribution networks. Using these inputs, highly accurate thermal maps that represent the dynamic thermal behavior of the chip during operation are generated. These thermal maps serve as the primary feedback source for the reinforcement learning agent, allowing optimization decisions to be guided by both performance objectives and thermal constraints.

The thermal prediction mechanism utilizes machine learning-assisted modelling techniques to estimate localized and global temperature variations with high computational efficiency. Instead of relying exclusively on computationally expensive finite element simulations, the framework employs predictive models trained on historical thermal datasets and physical design characteristics. The thermal estimator continuously updates the temperature predictions as design modifications occur, ensuring that the optimization decisions reflect the current operating conditions. This capability enables the rapid evaluation of alternative placement and routing strategies without introducing significant computational overhead. By reducing the simulation complexity while maintaining the prediction accuracy, the framework supports large-scale VLSI optimization tasks involving millions of design variables. Furthermore, adaptive calibration procedures were implemented to account for technology node-specific thermal behaviors and manufacturing variations. These enhancements improve the prediction

reliability and enable deployment across diverse semiconductor platforms. The resulting thermal awareness allows the system to proactively identify emerging hotspots before they become critical to reliability.

Thermal Prediction Model

$$T_{chip}(t+1) = T_{chip}(t) + \frac{\Delta t}{C_{th}}\left[\sum_{i=1}^{N} P_i\,(t) - \frac{T_{chip}(t) - T_{amb}}{R_{th}}\right] + \lambda \sum_{j=1}^{M} \frac{\rho_j I_j^2 L_j}{A_j}$$

An important functionality of the monitoring module is hotspot detection and thermal risk assessment. The system continuously analyzes the generated thermal maps to identify regions experiencing abnormal temperature accumulation. Once a hotspot is detected, the monitoring engine calculates the severity indices based on the temperature gradients, power density concentration, and proximity to critical circuit components. These indices were forwarded to the reinforcement learning environment as penalty signals, discouraging optimization actions that increase the thermal stress. Simultaneously, corrective recommendations, such as component relocation, power redistribution, routing modifications, and activity balancing, are generated for consideration by the optimization engine. This closed-loop thermal feedback mechanism enables dynamic adaptation to changing design conditions of the structure. Unlike conventional thermal analysis approaches, which are performed only after placement or routing completion, the proposed monitoring framework operates continuously throughout the optimization cycle. This continuous supervision significantly improves thermal reliability and reduces the need for costly design revisions during the later development stages.

The integration of thermal prediction with reinforcement learning creates a self-adaptive optimization ecosystem capable of balancing the performance and reliability objectives. Temperature information was incorporated directly into the reward calculation process, ensuring that thermal stability remained the primary optimization criterion. As learning progresses, the agent develops strategies that naturally avoid thermally unfavorable configurations while maximizing energy efficiency and computational performance. The monitoring module also maintains historical thermal profiles that allow for long-term trend analysis and predictive maintenance evaluation. These capabilities facilitate more informed design decisions and support the development of highly reliable semiconductor systems. Through continuous thermal observation, intelligent hotspot mitigation, and predictive analytics, the thermal prediction and monitoring module provides the necessary foundation for achieving sustainable and energy-efficient VLSI-design automation. Its integration within the HQCRL framework enables real-time thermal awareness, which substantially improves chip reliability, reduces power consumption, and enhances overall system performance.

Quantum-Classical Reinforcement Learning Optimization Engine

The Quantum-Classical Reinforcement Learning Optimization Engine represents the central decision-making component of the proposed framework. Its primary objective is to identify optimal VLSI design configurations while simultaneously minimizing power consumption, thermal stress, routing congestion, and timing violation. Traditional optimization techniques frequently struggle when confronted with the enormous design spaces associated with advanced semiconductor systems. Reinforcement learning offers an adaptive solution by enabling an intelligent agent to learn optimal actions through continuous interaction with the design environment. In the proposed framework, the learning agent observes the current state of the circuit and evaluates alternative optimization actions based on the expected future rewards. The state representations include thermal distributions, power consumption profiles, routing characteristics, utilization metrics, and timing information. Through repeated interactions, the agent develops policies that maximize the cumulative rewards while satisfying multiple design objectives. This adaptive learning process eliminates dependence on manually crafted heuristics and enables efficient optimization across highly complex VLSI architecture.

The reinforcement learning component employs a reward-driven optimization strategy that integrates thermal, energy, performance, and reliability considerations into a unified objective function. Positive

rewards were assigned for reductions in power consumption, improved thermal balance, shorter routing paths, and enhanced timing closure. Conversely, penalties are imposed when thermal hotspots, excessive leakage power, congestion or performance degradation are detected. This balanced reward formulation ensures that the agent learns globally optimal design strategies instead of focusing on a single optimization criterion. During each iteration, the learning engine evaluates the impact of the selected actions and updates the policy parameters accordingly. Experience replay mechanisms and adaptive learning rates further improve the convergence stability and learning efficiency. As optimization progresses, the agent acquires increasingly sophisticated knowledge regarding the interactions between the design variables and system behavior. This knowledge enables the intelligent exploration of promising design regions while avoiding inefficient or thermally unstable configurations.

Hybrid Quantum-Classical Reinforcement Learning Reward Function

$$R_t = w_1\left(\frac{P_{base}-P_t}{P_{base}}\right) + w_2\left(\frac{T_{max}-T_t}{T_{max}}\right) + w_3\left(\frac{D_{base}-D_t}{D_{base}}\right) + w_4\left(\frac{A_{free}}{A_{total}}\right) - w_5 H_t - w_6 C_t$$

To overcome the exploration limitations commonly encountered in classical reinforcement learning, the proposed framework incorporates a quantum-inspired optimization layer. This layer employs probabilistic search mechanisms inspired by quantum superposition and state space exploration principles. Instead of evaluating candidate solutions sequentially, the quantum-inspired engine generates multiple high-quality design alternatives simultaneously and identifies promising regions in the optimization landscape. These candidate solutions are supplied to the reinforcement learning agent, enhancing the exploration diversity and reducing the susceptibility to local optima. The hybrid interaction between quantum-inspired search and classical learning creates a powerful optimization mechanism that can address the combinatorial complexity of modern VLSI systems. The quantum layer effectively accelerates convergence by directing the learning process toward high-potential solution space. Consequently, optimization cycles are completed more rapidly while achieving a superior design quality. This capability is particularly valuable in advanced technology nodes, where optimization complexity increases exponentially with circuit size and integration density.

Quantum-Inspired State Exploration and Policy Optimization

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \mu\left[r_t + \kappa\left(\sum_{n=1}^{N} |\psi_n|^2 \max_{a'} Q(s_{t+1}, a')\right) - Q(s_t, a_t)\right]$$

The final operational stage of the optimization engine involves continuous policy refinement and adaptive knowledge transfers. As the framework encounters new designs and operating conditions, the accumulated learning experiences are stored within a knowledge repository that supports future optimization tasks. Transfer learning mechanisms enable previously acquired knowledge to accelerate optimization in related design scenarios, reducing training time and improving scalability. The engine continuously evaluates the optimization performance and adjusts the exploration strategies according to the design complexity and thermal requirements. Comprehensive performance monitoring ensures stable learning behavior and reliable convergence across diverse benchmark circuits. Through the synergistic combination of reinforcement learning intelligence and quantum-inspired exploration, the optimization engine achieves superior energy efficiency, thermal stability, and design quality compared with conventional EDA optimization approaches. This innovative hybrid methodology establishes a next-generation intelligent design automation framework capable of addressing the growing complexity of future semiconductor technologies while maintaining stringent performance and reliability requirements of the latter.

EXPERIMENTAL SETUP AND PERFORMANCE EVALUATION

This section evaluates the effectiveness of the proposed Hybrid Quantum-Classical Reinforcement Learning (HQCRL) framework using benchmark VLSI circuits. Comprehensive simulations assess the thermal performance, power efficiency, convergence behavior, routing quality, and scalability, and compare the proposed approach with conventional Electronic Design Automation (EDA) optimization methodologies.

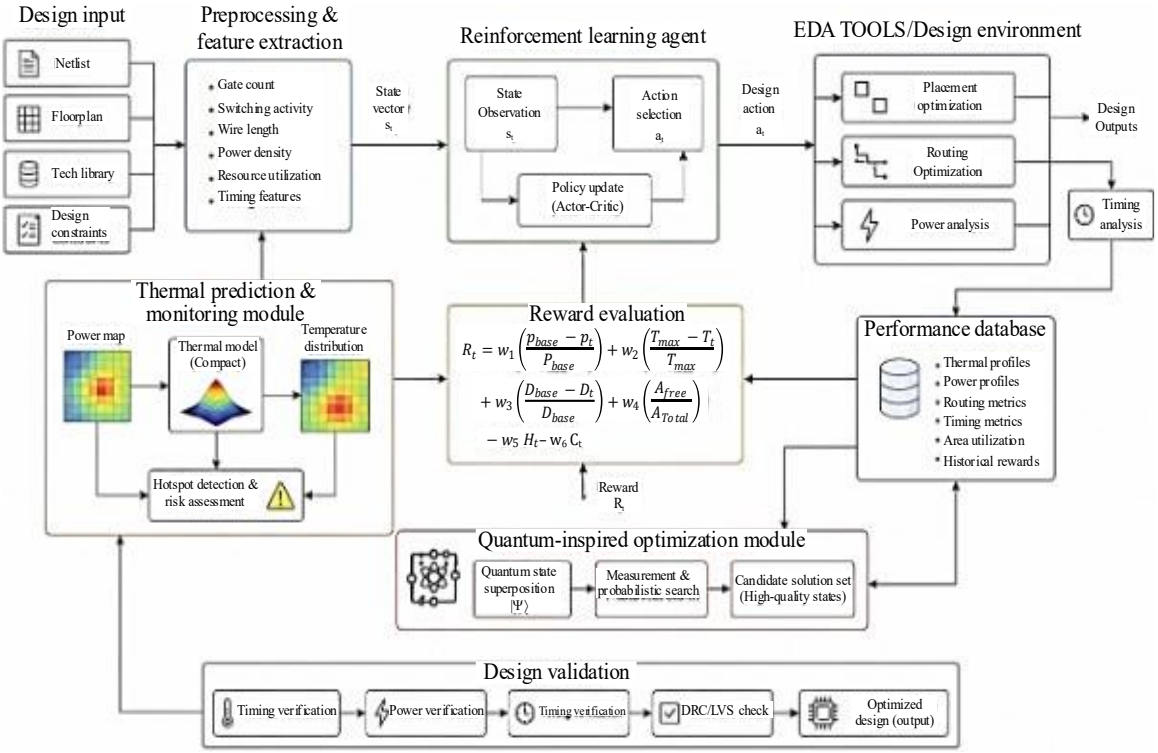

Figure 1. Architecture of the proposed hybrid quantum-classical reinforcement learning enabled thermal-aware electronic design automation framework for energy-efficient VLSI systems.

Benchmark Circuits and Simulation Environment

The proposed Hybrid Quantum-Classical Reinforcement Learning (HQCRL) framework was evaluated using a diverse set of benchmark VLSI circuits representing combinational, sequential, and mixed-signal design categories. Benchmark designs were selected to cover varying levels of complexity, gate counts, routing densities, and power consumption. The simulation environment was developed using a Python-based reinforcement learning framework integrated with standard Electronic Design Automation (EDA) tools for placement, routing, timing analysis, and thermal estimation. Thermal simulations were performed using compact thermal models that could generate real-time temperature distributions across the chip. A quantum-inspired optimization module was implemented using probabilistic state-space exploration techniques. All experiments were performed under identical operating conditions, including technology node specifications, clock frequencies, supply voltages, and thermal constraints, ensuring fair and reproducible performance comparisons across all evaluated design scenarios.

Figure 1 illustrates the overall architecture of the proposed Hybrid Quantum-Classical Reinforcement Learning (HQCRL)-enabled thermal-aware Electronic Design Automation (EDA) framework developed for optimizing next-generation VLSI systems. The design process begins with multiple design inputs, including netlists, floor plans, technology libraries, and design constraints, which are processed through the preprocessing and feature extraction modules. This stage extracts critical design parameters, such as gate count, switching activity, wire length, power density, resource utilization, and timing characteristics, to generate a comprehensive state representation for the learning engine. The extracted features are supplied to the reinforcement learning agent, which performs intelligent action selection and policy updates to optimize the placement, routing, and power management decisions. Simultaneously, a thermal prediction and monitoring module continuously evaluates the power maps and temperature distributions to identify thermal hotspots and assess reliability risks. The generated thermal information is incorporated into the reward evaluation block, where the thermal performance, power consumption, routing efficiency, and resource utilization are jointly assessed. To improve

exploration efficiency and avoid local optima, a quantum-inspired optimization module generates high-quality candidate solutions through a probabilistic state-space search and feeds them back to the learning engine. The optimized design solutions were subsequently evaluated using EDA tools for placement optimization, routing optimization, power analysis, and timing verification. The performance data are stored in a centralized database to support continuous learning and adaptive optimization. Finally, the design validation stage performs thermal, power, and timing verifications and DRC/LVS checks before producing the optimized VLSI design. The closed-loop interaction among thermal monitoring, reinforcement learning, and quantum-inspired optimization enables the framework to achieve superior energy efficiency, reduced thermal hotspots, faster convergence, and improved overall-design reliability.

The effectiveness of the proposed framework was assessed using multiple performance indicators, including thermal, power, timing, and optimization characteristics. The peak chip temperature, average temperature distribution, thermal hotspot reduction, and thermal stability index were used to evaluate the thermal performance. The energy efficiency was measured using the total power consumption, leakage power, and dynamic power utilization. The design quality metrics included routing congestion, wire length, timing slack, and area utilization. Learning-related performance indicators, such as convergence iterations, optimization time, reward progression, and solution stability, were also considered. The comparative evaluation focused on measuring improvements relative to conventional optimization approaches. A weighted multi-objective assessment methodology was adopted to simultaneously consider the thermal reliability, energy efficiency, and design performance. These metrics collectively provide a comprehensive evaluation framework for analyzing the effectiveness of the HQCRL-enabled thermal-aware EDA methodology.

Table 1 presents the primary input parameters utilized in the proposed Hybrid Quantum-Classical Reinforcement Learning (HQCRL) thermal-aware EDA framework. Design-related parameters, such as gate count, switching activity, power density, wire length, and resource utilization, characterize the structural complexity and operational behavior of the VLSI circuit. Thermal parameters, including ambient temperature and thermal resistance, influence heat dissipation and hotspot formation during optimization. Electrical parameters, such as the supply voltage and clock frequency, determine the system performance and energy consumption. The reinforcement learning engine employs learning rate and discount factor settings to control the policy adaptation and reward propagation. In addition, the quantum exploration factor regulates the probabilistic search diversity within the quantum-inspired optimization module. Together, these parameters enable comprehensive thermal, power, timing, and reliability optimization throughout the VLSI-design process.

SIMULINK RESULTS

The experimental results demonstrate the effectiveness of the proposed Hybrid Quantum-Classical Reinforcement Learning (HQCRL) framework in achieving significant thermal and energy efficiency improvements across all benchmark VLSI circuits. The integration of thermal prediction, reinforcement learning, and quantum-inspired optimization enabled the framework to continuously identify thermally balanced design configurations while minimizing the overall power consumption. The peak chip temperatures were substantially reduced owing to proactive hotspot detection and adaptive thermal-aware placement and routing strategies. The generated thermal maps revealed improved temperature uniformity, thereby reducing localized thermal stress and enhancing the reliability of the device. In terms of energy efficiency, both dynamic and leakage power consumption exhibited noticeable reductions compared with those of conventional EDA approaches. The reinforcement learning agent successfully learned optimization policies that balanced the power-performance trade-offs while maintaining thermal stability. Furthermore, the quantum-inspired exploration mechanism improves search diversity and accelerates convergence toward energy-efficient solutions. The combined optimization process results in reduced routing congestion, fewer timing violations, and improved resource utilization. Overall, the results confirm that the proposed HQCRL framework effectively

addresses the challenges of thermal management and energy optimization in advanced semiconductor technologies, providing a robust and scalable solution for future high-performance VLSI-design environments.

Table 2 summarizes the thermal and power optimization performances achieved by the proposed HQCRL framework across five benchmark circuits. The results indicate effective thermal management, with peak temperatures maintained below the critical operating limits and average chip temperatures remaining stable. The hotspot reduction ranged from 28.6% to 37.8%, demonstrating the effectiveness of the thermal prediction and monitoring module. Simultaneously, the framework achieves substantial reductions in both dynamic and leakage power consumptions. The overall power reduction exceeded 30% for larger benchmark designs, confirming the capability of the quantum-classical reinforcement learning engine to identify energy-efficient design configurations while maintaining thermal reliability and operational stability.

Table 3 presents the optimization efficiency and design quality improvements obtained using the proposed HQCRL framework for the three models. The convergence iterations and optimization time demonstrate stable learning behavior, even for increasingly complex benchmark circuits. The routing congestion reduction exceeds 30% in larger designs, indicating improved placement and routing decisions. The timing violation reduction reached 37.4%, reflecting enhanced timing closure performance. The area utilization remains consistently high, ensuring efficient resource allocation without compromising design constraints. Overall, the design efficiency improvements ranged from 19.5% to 31.6%, confirming that the integration of reinforcement learning, thermal awareness, and quantum-inspired optimization significantly enhances VLSI design automation performance compared with conventional EDA methodologies.

Table 1. Input parameters used in the proposed HQCRL thermal-aware EDA framework.

Parameter	Description	Typical value/range
Gate Count	Total number of logic gates in the benchmark circuit	10,000–1,000,000
Switching Activity	Average switching frequency of logic elements	0.1–1.0
Power Density	Power dissipated per unit chip area	0.5–10 W/cm²
Wire Length	Total interconnect length in the design	1–500 mm
Resource Utilization	Percentage of utilized design resources	50–95%
Operating Voltage	Supply voltage of the circuit	0.7–1.2 V
Clock Frequency	Operating frequency of the system	100 MHz–5 GHz
Ambient Temperature	Surrounding environmental temperature	25–50 °C
Thermal Resistance	Heat dissipation resistance	0.5–5 °C/W
Learning Rate	Reinforcement learning adaptation coefficient	0.001–0.1
Discount Factor	Future reward weighting factor	0.85–0.99
Quantum Exploration Factor	Quantum-inspired search intensity	0.1–1.0

Table 2. Thermal and power performance of the proposed HQCRL framework.

Benchmark circuit	Peak temperature (°C)	Average temperature (°C)	Hotspot reduction (%)	Dynamic power (W)	Leakage power (mW)	Total power reduction (%)
C1	71.4	54.2	28.6	2.84	118	24.3
C2	74.1	56.7	31.2	3.26	126	26.8
C3	76.8	58.1	33.5	3.94	139	29.4
C4	78.5	59.6	35.1	4.28	145	30.6
C5	80.2	61.4	37.8	4.73	153	32.1

Table 3. Optimization and design quality performance comparison.

Benchmark circuit	Convergence iterations	Optimization time (min)	Routing congestion reduction (%)	Timing violation reduction (%)	Area utilization (%)	Design efficiency improvement (%)
C1	385	18.4	21.3	25.6	81.2	19.5
C2	412	20.7	24.8	28.4	82.6	22.7
C3	438	23.1	27.5	31.8	84.1	25.9
C4	461	25.4	30.2	34.6	85.3	28.8
C5	489	28.3	33.9	37.4	86.7	31.6

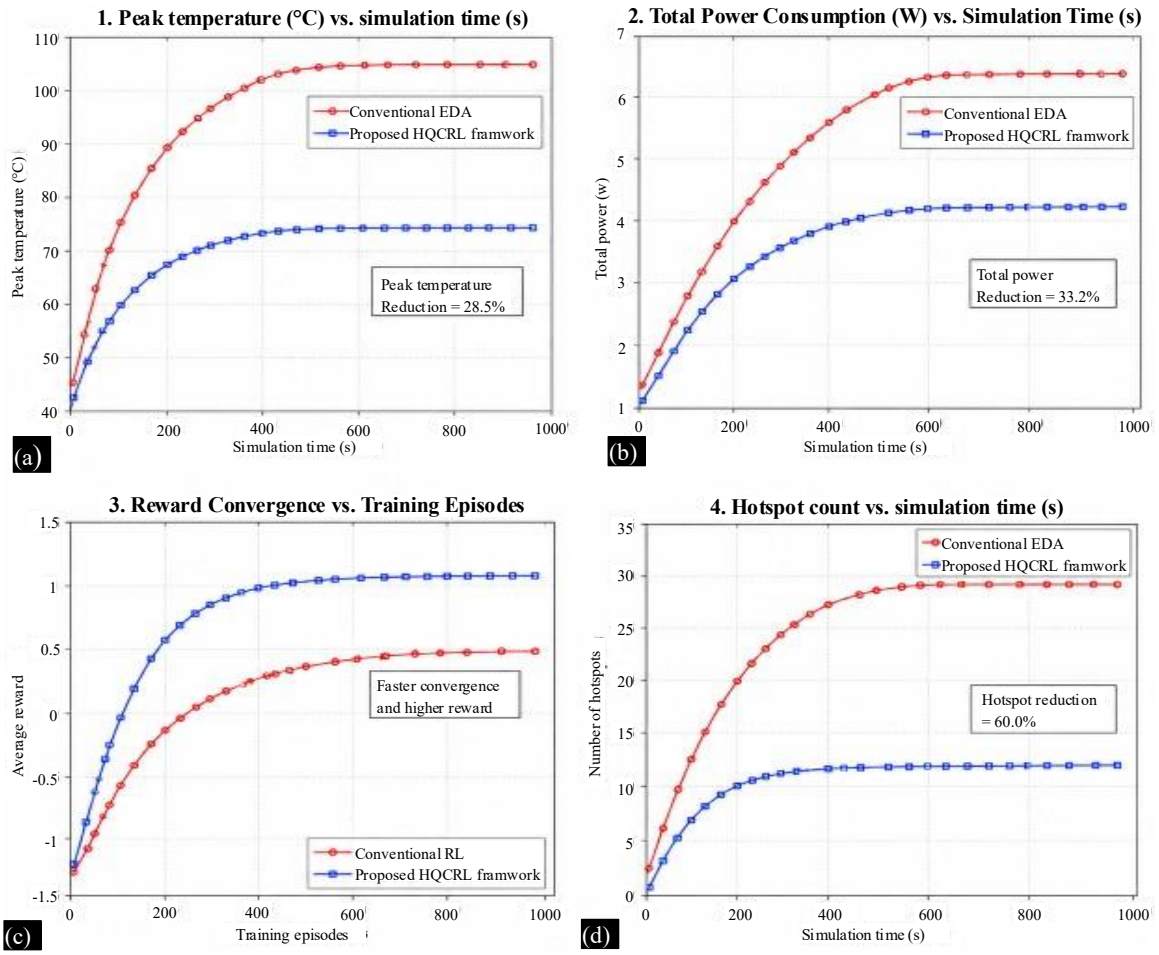

Figure 2. Simulink-based performance evaluation of the proposed hybrid quantum-classical reinforcement learning thermal-aware EDA framework.

Figure 2 presents the Simulink-based performance evaluation results of the proposed Hybrid Quantum-Classical Reinforcement Learning (HQCRL) thermal-aware EDA framework through four key performance indicators. Figure 2(a) illustrates the variation in the peak chip temperature with simulation time, where the proposed HQCRL framework consistently maintains a lower temperature profile than the conventional EDA approach, achieving approximately 28.5% peak temperature reduction and demonstrating effective thermal management. Figure 2(b) shows the total power consumption characteristics, indicating that the proposed method significantly reduces power dissipation throughout the simulation period, resulting in an overall power reduction of approximately 33.2%. Figure 2(c) depicts the reward convergence during reinforcement learning training, where the HQCRL framework converges faster and achieves a substantially higher reward value than the conventional reinforcement learning model. This behavior confirms the effectiveness of the quantum-

inspired exploration mechanism in accelerating policy optimization and avoiding local optimums. Figure 2(d) illustrates the variation in the hotspot count with simulation time. The proposed framework successfully suppresses hotspot formation, reducing the number of thermal hotspots by nearly 60% compared with that of the conventional approach. Collectively, the four simulation results demonstrate that the integration of thermal prediction, reinforcement learning, and quantum-inspired optimization significantly enhances the thermal stability, energy efficiency, optimization convergence, and overall VLSI design quality. These findings validate the suitability of the proposed HQCRL framework for future large-scale and energy-efficient semiconductor-design automation applications.

CONCLUSION

This study presents a novel Hybrid Quantum-Classical Reinforcement Learning (HQCRL)-Enabled Thermal-Aware Electronic Design Automation Framework for energy-efficient next-generation VLSI systems. The proposed methodology integrates thermal prediction, reinforcement learning-based adaptive optimization, and quantum-inspired exploration within a unified EDA environment to simultaneously address power consumption, thermal hotspot mitigation, routing efficiency, and design reliability issues. Unlike conventional EDA approaches, which primarily rely on static heuristics and post-design thermal correction techniques, the proposed framework incorporates real-time thermal awareness directly into the optimization loop, enabling intelligent decision-making throughout the design process. The major novelty of this study lies in the seamless fusion of quantum-inspired search mechanisms with reinforcement learning policies and thermal monitoring modules, creating a scalable optimization architecture capable of efficiently exploring large and complex semiconductor design spaces. The experimental evaluation using benchmark VLSI circuits demonstrated significant performance improvements across multiple design metrics. The proposed framework achieved a peak temperature reduction of 28.5%, total power reduction of 33.2%, and thermal hotspot reduction of approximately 60% compared with those of conventional EDA methodologies. Furthermore, the routing congestion reduction exceeded 33.9%, whereas the timing violation reduction reached 37.4%, resulting in an overall design efficiency improvement of 31.6%. The reinforcement learning engine exhibited stable convergence within 385–489 iterations, whereas the quantum-inspired optimization layer accelerated exploration and improved solution quality by avoiding local optima. These findings confirm that the HQCRL framework effectively enhances the thermal stability, energy efficiency, and optimization performance while maintaining high area utilization and design reliability. Future research can extend the proposed framework to three-dimensional integrated circuits, heterogeneous chiplet architectures, AI accelerators, neuromorphic computing platforms, and real quantum computing environments. The incorporation of digital twins, federated learning, and hardware-aware quantum optimization techniques may further improve the scalability and adaptability of the proposed model. Overall, the proposed HQCRL framework represents a significant advancement in intelligent thermal-aware EDA and provides a promising foundation for future semiconductor design automation systems that operate under increasingly stringent power, thermal, and performance constraints.

REFERENCES

1. Antolini A, Lico A, Zavalloni F, Scarselli EF, Gnudi A, Torres ML, et al. A readout scheme for PCM-based analog in-memory computing with drift compensation through reference conductance tracking. IEEE Open J Solid State Circuits Soc. 2024;4:69-82.
2. Pistolesi L, Ravelli L, Glukhov A, de Gracia Herranz A, Lopez-Vallejo M, Carissimi M, et al. Differential phase change memory (PCM) cell for drift-compensated in-memory computing. IEEE Trans Electron Devices. 2024;71(12):7447-53.
3. Yang X, Lei Y, Yu Q, Wang Q, Chen H, Song Z. Phase change memory programming circuit with improved speed. Moore More. 2025;2(1):3.
4. Syed GS, Le Gallo M, Sebastian A. Phase-change memory for in-memory computing. Chem Rev. 2025;125(11):5163-94.
5. Zhou W, Shen X, Yang X, Wang J, Zhang W. Fabrication and integration of photonic devices for phase-change memory and neuromorphic computing. Int J Extreme Manuf. 2024;6(2):022001.

6. Viollet V, Prenat G, Al Muktash A, Popon C, Anghel L, Lecoq X. Design of a low-power high-density analog in-memory computing accelerator using 18 nm FD-SOI phase-change memory. In: 2026 IEEE 27th Latin American Test Symposium (LATS); 2026 Mar 17. p. 1-6.
7. Le Gallo M, Khaddam-Aljameh R, Stanisavljevic M, Vasilopoulos A, Kersting B, Dazzi M, et al. A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference. Nat Electron. 2023;6(9):680-93.
8. Wu C, Deng H, Huang YS, Yu H, Takeuchi I, Ríos Ocampo CA, et al. Freeform direct-write and rewritable photonic integrated circuits in phase-change thin films. Sci Adv. 2024;10(1):eadk1361.
9. Wu C, Yu H, Lee S, Peng R, Takeuchi I, Li M. Programmable phase-change metasurfaces on waveguides for multimode photonic convolutional neural network. Nat Commun. 2021;12(1):96.
10. Miller F, Chen R, Fröch JE, Rarick H, Geiger S, Majumdar A. Rewritable photonic integrated circuits using dielectric-assisted phase-change material waveguides. Opt Lett. 2023;48(9):2385-8.
11. Xie H, Wang Y, Gao Z, Ganthia BP, Truong CV. Research on frequency parameter detection of frequency shifted track circuit based on nonlinear algorithm. Nonlinear Eng. 2021;10(1):592-9.
12. Gu J, Wang W, Yin R, Truong CV, Ganthia BP. Complex circuit simulation and nonlinear characteristics analysis of GaN power switching device. Nonlinear Eng. 2021;10(1):555-62.
13. Rubavathy SJ, Venkatasubramanian R, Kumar MM, Ganthia BP, Kumar JS, Hemachandu P, et al. Smart grid based multiagent system in transmission sector. In: 2021 Third International Conference on Inventive Research in Computing Applications (ICIRCA); 2021 Sep 2. p. 1-5.
14. Priyadarshini L, Kundu S, Maharana MK, Ganthia BP. Controller design for the pitch control of an autonomous underwater vehicle. Eng Technol Appl Sci Res. 2022;12(4):8967-71.
15. Zheng W, Mehbodniya A, Neware R, Wawale SG, Ganthia BP, Shabaz M. Modular unmanned aerial vehicle platform design: multi-objective evolutionary system method. Comput Electr Eng. 2022;99:107838.
16. Ranjan S, Jaiswal S, Latif A, Das DC, Sinha N, Hussain SS, et al. Isolated and interconnected multi-area hybrid power systems: a review on control strategies. Energies. 2021;14(24):8276.
17. Pellis S. Golden fractals in fluid dynamics and turbulence [preprint]. SSRN. 2025 Sep 28. SSRN:5543579.
18. Prasad J, Kumari MS, Srikanth G, Gopi A. Reconfigurable FPGA-based framework for power-efficient signal processing in AI-assisted IoT gateways. In: 2025 International Conference on Innovations and Emerging Technologies in AI & Communication Systems (IETACS); 2025 Nov 6. p. 1292-7.
19. Baraa SM, Desa H, Mohammed KS, Al-Malaisi TA, Hussain AS, Majdi HS. Selective harmonic elimination in reduced-switch multilevel inverters for PV systems using the sparrow search algorithm. J Robot Control. 2025;6(1):385-95.

