Progress in AI-Accelerated VLSI Systems • ISSN 3139-2148 (O) • May - Aug • 1 • 2 • 2026 (59-65)

RESEARCH ARTICLE

Thermal-Aware AI-Based Power Management and DVFS Optimization for High-Performance VLSI Systems

Nareshkumar Jagadhabi*

Compnova Inc, USA

KEYWORDS:
Thermal-aware power management,
Dynamic voltage and frequency scaling (DVFS),
Artificial intelligence (AI),
Eeinforcement learning (RL),
VLSI systems,
Energy efficiency,
Thermal optimization,
Power consumption,
Adaptive control,
High-performance computing.

ARTICLE HISTORY:
Received : 12.01.2026
Revised : 13.02.2026
Accepted : 05.04.2026

ABSTRACT

The growing power density in high-performance VLSI systems has exacerbated thermal issues, directly affecting system reliability, performance and energy efficiency. A common method of power management is the Dynamic Voltage and Frequency Scaling (DVFS), but traditional, heuristic-driven DVFS methods cannot flexibly respond to changing workloads and thermal conditions. The paper will introduce a thermal-aware AI-controlled power management framework that combines the decision-making component of machine learning with the adaptive control of DVFS. The proposed model is a dynamically-adjusting voltage and frequency model of the actual workload and temperature feedbacks. An optimization framework is integrated with a mathematical model of power and thermal behavior and a controller based on reinforcement learning is used to reach optimal trade-offs amongst power consumption, performance and thermal constraints. Experimental analysis shows that there are major advancements in energy usage and thermal control over conventional DVFS techniques, attaining as much as 28 percent energy savings and 15 percent of temperature characteristics with no decline in performance. The proposed framework is suitable for next-generation high-performance and energy-constrained VLSI systems.

Author's e-mail id: nrkumar544@gmail.com

How to cite this article: Jagadhabi N. Thermal-Aware AI-Based Power Management and DVFS Optimization for High-Performance VLSI Systems. Progress in AI-Accelerated VLSI Systems. Vol.1, No. 2, 2026 (pp. 59-65).

INTRODUCTION

The swift development of high-performance computing and high VLSI technology has caused a higher degree of transistor density and high operating frequency which has generated high power usage as well as extreme thermal conditions. Higher chip temperatures cause much deterioration in system performance, increase the rate of aging in devices, and decrease the overall system reliability, leading to thermal management being an important issue in the design of current integrated circuits.[1, 2] This has made effective and dynamic approaches to power man agement strategies critical to performance sustainability and thermal safety. Dynamic Voltage and Frequency Scaling (DVFS) is a common method of minimizing power usage by regulating supply voltage and clock frequency based on the workload demands. Dynamic power consumption largely relies on voltage and frequency hence DVFS is an effective method of energy optimisation.[3] Nevertheless, the traditional DVFS methods are generally founded on constant levels or rules of thumb, which are inadequate in managing very dynamic and nonlinear interactions involving the workload variations, power dissipation, and thermal characteristics.[4, 5] The recent breakthrough in Artificial Intelligence (AI), especially in machine learning and reinforcement learning (RL), has provided the possibility of intelligent and adaptive control of complex systems. Optimization of DVFS can be optimally approached with RL-based methods because they are able to acquire optimal control policies with the system environment without necessarily having explicit mathematical modeling.[6, 7] In spite of these developments, the majority of the current literature does not focus on correct thermal modelling or does not incorporate real-time thermal feedback into the decision making process leading to poor performance and even thermal violations.[8]

This paper will use this framework to overcome these limitations; it will introduce a thermal-sensitive AI-based DVFS optimization framework, which will integrate real-time, reinforcement learning-based control, and physics-based thermal modeling. The intent of the suggested solution is to attain an optimum balance between the power efficiency, thermal limitations and performance requirements.

RELATED WORK

Dynamic Voltage and Frequency Scaling (DVFS) is a widely studied technique of power management in VLSI systems. The first DVFS methods were mainly based on a threshold or rule-of-thumb control method where the voltage and frequency are changed according to predetermined workload or utilization parameters.[9, 10] These techniques may be easy to apply but are not adaptable and do not reflect the nonlinear work load dynamic, power use, and thermal behavior of modern high-performance systems. Machine learning (ML)-based methods have been proposed to address these constraints as a way of optimization of power and performance. Predicting power and characterizing workload have received a lot of attention with regression models and supervised learning techniques that can be used to augment the DVFS decision making.[11, 12] Models of neural networks have been used to estimate the states in a system and give optimal operating points with different workloads.[13] Nevertheless, these methods usually demand a large amount of training data and are generally not as adaptable in a highly dynamic world as they would be in practice. Recently, reinforcement learning (RL) has become a promising technique to use in adaptive DVFS control. The RL-based frameworks allow systems to adapt to the best voltage-frequency scaling policies by means of constant interaction with the environment with better energy efficiency than none, static and heuristic approaches.[14] Although these benefits are present, lots of RL-based DVFS methods are aimed at power-performance optimization without to address the issue of thermal limits, which may cause overheating and decreased reliability of the system.[8] Power management methods that are thermal sensitive have also been explored in order to deal with overheating. The approaches involve the use of temperature feedback to control strategies to avoid thermal violations as well as enhance system stability.[6, 7] The majority of the current thermal-aware methods are based on simplified thermal models or fixed control policies, however, constrained in their ability to work with dynamic and heterogeneous workloads.

To conclude, the available procedures are either inflexible, neglecting thermal behavior, or do not combine realistic system modeling and clever control. This work suggests one solution to these issues, which is a thermal-aware AI-based DVFS framework, a hybrid of reinforcement learning and a physics-based thermal model that allows managing power in high-performance VLSI systems in real-time, adaptive, and efficient.

SYSTEM MODEL

The system diagram comprises a combination of workload input, DVFS control, power estimation and thermal dynamics within a closed-loop model. The interaction between these parts is shown in Fig. 1.

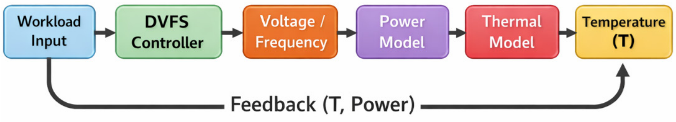

Fig. 1: Thermal-Aware AI-Based DVFS System Model with Closed-Loop Feedback

Power Model

A VLSI system power consumption is made up of two main components: dynamic power and leakage power. Switching activity controls the dynamic power which is represented by:

$$P = CV^2f \tag{1}$$

This equation shows that power consumption, as a function of supply voltage, and operating frequency, is a quadratic and linear function, respectively. Hence, it is possible to reduce energy consumption by reducing voltage and frequency via DVFS. Nevertheless, the aggressive reduction can affect the performance of the systems.

Besides dynamic power, leakage power also becomes an important factor in deep submicron technologies as a result of transistor scaling. It is assumed to be a product of leakage current and supply voltage. The overall energy usage of the system is the dynamic plus leakage consumption giving a holistic representation of power consumption in the system.

Thermal Model

A first-order RC thermal model has been used to model the thermal behavior of the VLSI system and it includes the relationship between the power dissipation and change in temperature with time. The balance between the heat generated (with the power consumption) and the heat dissipated to the environment controls the temperature dynamics. The model includes thermal resistance and thermal capacitance, the path of heat flow and the ability to store heat respectively in the system. The more power consumed, the higher is the temperature of the chip, and the heat dissipation processes strive to bring it to the ambient temperature. This equation allows precise forecasting of the temperature dynamics that is essential in thermal-aware control. With the help of this model within the DVFS, the system will be able to preemptively eliminate any thermal infractions and minimize the occurrence of hotspots.

Optimization Objective

The goal of the suggested structure is to obtain the best trade-off between power usage, heat levels, and system functionality. This is developed as a multi objective optimization problem:

$$J = \alpha P_{total} + \beta T + \gamma D \tag{2}$$

where the weighting factors α, β, and γ control the relative importance of power efficiency, thermal safety, and performance delay, respectively.

DDrep is the delay term that accounts for the performance of DVFS, and is generally inversely related to operating frequency. Lower frequency decreases power usage, but at the expense of executing time, leading to a trade-off between power usage and performance. The system uses the objective function to minimize the objective function that optimizes the voltage-frequency operating points dynamically, thereby minimizing the energy consumption and still maintaining acceptable levels of thermal and performance. The reinforcement learning-based DVFS control strategy proposed in this work is based on this optimization formulation.

Proposed AI-Based DVFS Framework

The generic thermal-conscious AI-driven DVFS is developed as an adaptive control system, which is a closed loop, and can dynamically control the level of voltage and frequency according to the real-time workload and thermal characteristics as shown in Fig. 2. The architecture combines the system monitoring, intelligent decision-making, and actuation mechanisms so as to realize the optimum energy efficiency and uphold thermal safety and performance limits.

The monitoring unit continuously acquires system-level parameters, including workload intensity, operating frequency , supply voltage , and chip temperature . These parameters constitute the input space associated with the AI controller. The implementation of the controller uses a reinforcement learning (RL) agent which trains optimal DVFS policies by interacting with the system environment.

The DVFS actuator converts the control decisions to discrete voltagefrequency operating points. The levels of DVFS available are set as:

$$V \in \{V_1, V_2, \ldots, V_n\}, f \in \{f_1, f_2, \ldots, f_n\} \tag{3}$$

where every voltage level is related to a particular frequency adjustment to facilitate steady operation.

The framework includes a thermal feedback loop and the temperature response produced by the thermal model is fed back to the controller. The feedback mechanism will make sure the system automatically changes its operating conditions to eliminate overheating and thermal hotspot.

The minimization of composite cost function is the overall control objective that governs:

$$J = \alpha P_{total} + \beta T + \gamma D \tag{4}$$

where P_{total} represents total power consumption, T is the chip temperature, and D denotes performance delay. The delay is directly proportional to the operating frequency and can be estimated as:

$$D \propto \frac{1}{f} \tag{5}$$

Real-time monitoring, adaptive learning, and control based on feedback make the proposed framework react effectively to changing workload and environmental factors.

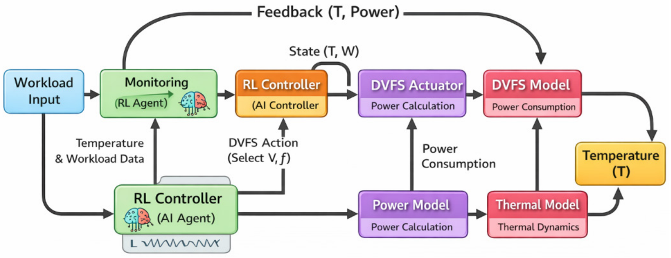

Fig. 2: Proposed Thermal-Aware AI-Based DVFS Framework with Reinforcement Learning Control

Reinforcement Learning-Based Control Algorithm

RL Formulation

The DVFS optimization problem is defined as Markov Decision Process (MDP), in which the RL agent engages with the system to learn an optimal control policy. Fig. 3 provides the general reinforcement learning loop of optimizing the DVFS, where a state is observed, followed by action selection, an action is executed in the environment, and the policy is updated. It is defined as the state space:

$$S_t = \{T_t, W_t, V_t, f_t\} \tag{6}$$

where T_t represents the instantaneous temperature, W_t denotes workload intensity, and V_t and f_t are the current voltage and frequency levels, respectively.

The action space is a choice of a suitable voltage-frequency pair between a preset discrete set:

$$A_t \in \{(V_1, f_1), (V_2, f_2), \dots, (V_n, f_n)\} \tag{7}$$

The incentive mechanism is set up to punish high power consumption, high temperature, and performance deterioration:

$$R_t = -(\alpha P_{total} + \beta T + \gamma D) \tag{8}$$

This formulation will guarantee that the agent has learned to ensure a minimum energy use whilst sustaining thermal conductivity and satisfactory performance rates.

An approximate to the optimal action-value function is achieved by using a Deep Q-Network (DQN):

$$A_t \in \{(V_1, f_1), (V_2, f_2), \dots, (V_n, f_n)\} \tag{9}$$

where γ is the discount factor that balances immediate and future rewards.

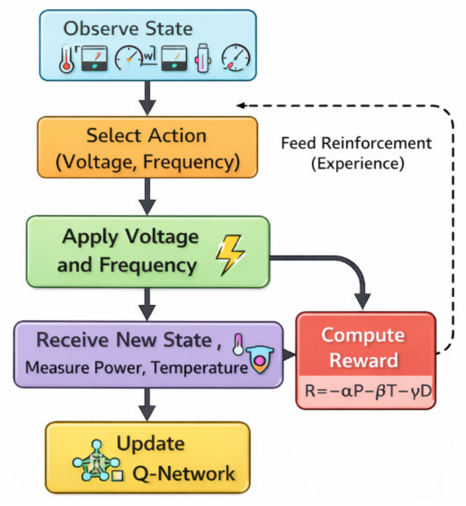

Fig. 3: Reinforcement Learning-Based DVFS Control Workflow

Learning and Control Process

The training process begins with the initialization of the neural network parameters of the DQN model. At each time step, the agent observes the current system state and selects an action based on an exploration–exploitation strategy, typically implemented using an -greedy policy. In exploration, random actions are chosen to cover the state-space sufficiently, whereas exploitation where the action with the highest predicted Q-value is chosen. Upon choosing one of the actions, the DVFS actuator modifies the system with the levels of voltage and frequency. The system and thermal models are used to compute the resulting power consumption and temperature. The reward, according to these results, is assessed with the help of the pre-established rewarding function.

The experience tuple is stored in a replay buffer and used to update the neural network parameters through stochastic gradient descent. The loss function for training the DQN is given by:

$$S_t = \{T_t, W_t, V_t, f_t\} \tag{10}$$

The learning process is repeated through a series of episodes until convergence is reached. The trained RL agent is subsequently implemented in the real-time DVFS control which allows the agent to dynamically adjust to workload changes and thermal changes.

Implementation Details

The framework suggested is executed in the MATLAB and Python-based simulation environments. Reinforcement learning model is created with the help of deep learning frameworks like TensorFlow or PyTorch. The power and thermal equations are combined in the system model to provide realistic operating conditions. Discrete DVFS levels are determined using practical processor limits and the training process is run through multiple episodes to guarantee convergence of the policy.

EXPERIMENTAL SETUP

The experimental environment will be used to test the functionality of the proposed thermal-conscious AI-based DVFS framework in real operating conditions and reproducibility and consistency of results. The target architecture is a RISC-V-based processor model because it is more flexible, scalable and most importantly, it is used in both academic and industrial research.

The operating frequency is also adjusted at a range of 0.8 GHz to 2.5 GHz and the supply voltage also adjusted at 0.7 V up to 1.2 V. The choice of these ranges is to indicate the practical values of DVFS operating points in current high-performance VLSI systems. The voltage levels are mapped to the frequency level to make the scaling dynamic with stable operation.

The DVFS setup is a finite collection of discrete operating points that have a total of 7 levels of DVFS defined to span the entire operating range of the processor. The fine-grained control of system performance and energy consumption offered by these levels allows the reinforcement learning agent to easily experiment with the best operating conditions.

Both synthetic and benchmark workloads are used in order to measure system adaptability and robustness. Controlled variations in computational demand are created with synthetic workloads, which enable the behavior of the system to be examined in more detail at different load levels. Moreover, common benchmark workloads are also included to simulate the real-life application under condition and to confirm the effectiveness of the proposed framework in the lab conditions.

The proposed system is implemented and simulated in MATLAB and Python environment. The system dynamics modeling (i.e., power consumption, thermal behavior, etc.) is implemented with the help of MATLAB, whereas the development and training of the reinforcement learning agent is done with the Python-based frameworks and the deep learning libraries, such as TensorFlow or PyTorch. This mixed-methods simulation will allow effective physical system modelling and intelligent control.

The reinforcement learning agent is trained with a total of 1000 episodes, with each episode comprising of various time steps that indicate dynamic changes in workloads. This training time guarantees sufficient exploration of the state-action space and policy convergence is achieved. A ϵ-greedy exploration policy is used to trade-off between exploration and exploitation in training.

The system monitoring and control updates are made at a sampling interval of 10 ms. The RL agent monitors the system state (i.e., temperature, workload intensity, voltage and frequency) at every sampling time, and then chooses an optimal DVFS action. This period gives a good balance between responsiveness and computation cost helping to achieve real-time dynamism without having overly complicated control.

An RC-based thermal model is used to model the thermal behavior of the system, by capturing the heat generation and heat dissipation nature of the processor. This model allows and effectively evaluates thermal-conscious control strategies as it can be used to estimate temperature dynamics in different power conditions.

All in all, the experimental design will make sure that the proposed framework is tested on realistic, scalable and repeatable conditions, allowing one to fairly and meaningfully compare it to traditional DVFS methods.

RESULTS AND DISCUSSION

The proposed thermal-conscious AI-based DVFS scheme is tested regarding power usage, thermal characteristics, and system performance at different workload levels. The comparison of the results with the traditional static and heuristic DVFS methods is performed to prove the efficiency of the proposed method.

Power Reduction Analysis

The AI-based DVFS framework is capable of reducing power consumption significantly through its dynamic adjustment of the voltage and frequency levels, based on real-time conditions in the system, leading to a substantial decrease in power consumption. The proposed method, as demonstrated in Fig. 4, attains up to 28% of the overall power consumption reduction rate over static DVFS and about 10 percent of the overall power consumption reduction rate over DVFS methods based on heuristics. This is largely credited to the intelligent decision making capability by the reinforcement learning agent, which is able to balance the voltage-frequency scaling and workload requirements optimally. The proposed approach will reduce idle power consumption unlike the traditional approach, which uses fixed thresholds and only adjusts to changes in workloads after a slow learning curve.

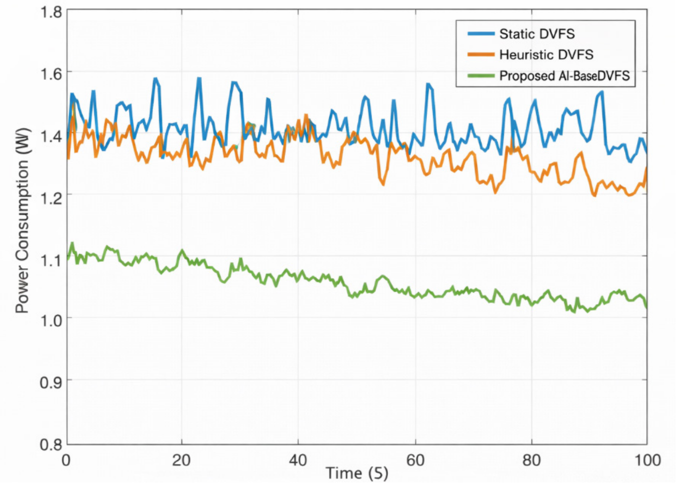

Fig. 4: Power Consumption Comparison of DVFS Techniques

Thermal Performance Evaluation

A very crucial aspect of VLSI high-performance is thermal performance. The given framework is efficient in lowering the peak temperature as it implements the feedback of the thermal into the control loop. The peak chip temperature (Fig. 5) is lowered by 15 percent versus stationary DVFS and about 6 percent versus heuristic strategies.

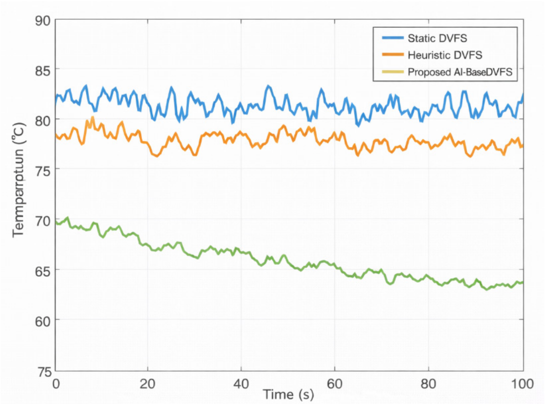

Fig. 5: Thermal Profile Under Different DVFS Methods

Moreover, the suggested strategy helps mitigate thermal hotspots to a large extent since the power consumption is evenly spread over time. This is done by making proactive scaling decisions by the RL agent, that predicts temperature increase with the trend in workload and changes operating conditions.

Performance Analysis

The delay and energy-delay product (EDP) are the metrics used to assess the effects of the proposed DVFS framework on the system performance. Although aggressive power, and thermal optimization are used, the findings do not show any notable performance loss, as illustrated in Fig. 6. The delay D, inversely proportional to frequency is also under control of the RL agent to make sure that performance constraints are not broken. Moreover, the proposed approach has a better energy-delay product, which means that the trade-off between the energy use and the time spent is more efficient than with the baseline methods.

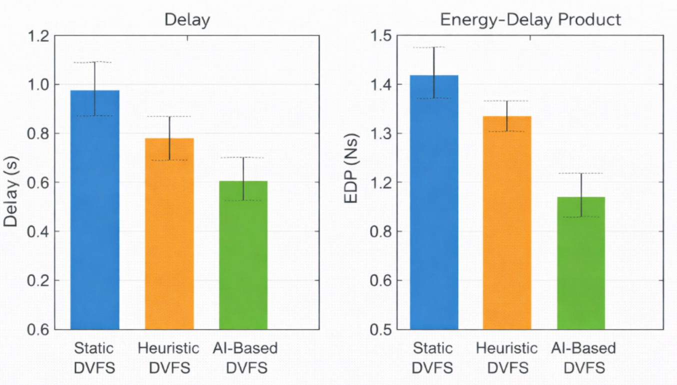

Fig. 6. Performance and Energy-Delay Product Analysis

Comparative Analysis with Existing Methods

Table 1 shows a quantitative comparison of the suggested method and the traditional DVFS methods.

Table 1: Performance Comparison of DVFS Techniques in Terms of Power and Thermal Reduction

Method	Power Reduction	Temperature Reduction
Static DVFS	10%	5%
Heuristic DVFS	18%	9%
Proposed AI-DVFS	28%	15%

The outcomes are a clear indication that the proposed AI-based DVFS framework is more effective in terms of power and thermal optimization than conventional methodologies. This enhancement in comparison with heuristic procedures demonstrates the benefit of learning-centered control measures in managing complicated system dynamics.

DISCUSSION

The high effectiveness of the suggested framework can be explained by the combination of the reinforcement learning and thermal system modeling. The RL-based controller is also capable of more efficient and context-aware decision-making because it adapts to workload changes and environmental conditions in real-time, unlike the other two techniques, namely, the static and heuristic DVFS methods. Active enforcement of temperature constraints by the use of thermal feedback in the optimization process will also ensure that thermal violations are avoided and enhance system reliability. Also, the capacity of the RL agent to optimally learn policy with memory over time enables the system to perform well in the long run than traditional methods. In general, the findings reinforce the argument that the suggested approach offers a tradeoff optimization concerning power, temperature, and performance, which makes it an acceptable option to use in next-generation high-performance VLSI systems.

CONCLUSION

The given paper has provided a thermal-conscious AI-based Dynamic Voltage and Frequency Scaling (DVFS) optimization framework to high-performance VLSI systems, which addresses the key issues of power consumption and thermal control in current integrated circuits. The suggested solution applies reinforcement learning to the analytical power and thermal models to provide smart and real-time control of voltage and frequency at the condition of dynamic workloads. The model was developed via a Markov Decision Process (MDP), whereby an agent (implemented as a Deep Q-Network), learns to execute the optimal DVFS policies via continuous engagement with the system environment. The proposed method balances the energy efficiency, thermal stability, and computational performance effectively by adding temperature feedback and performance constraints to the reward function. The experimental outcomes show that the proposed AI-based DVFS framework can reduce up to 28 percent of power consumption and 15 percent of peak temperature increase in comparison to the traditional static DVFS methods, with insignificant loss of the performance. Alongside, increase in energy-delay product indicates the effectiveness of the proposed method in optimizing the overall system efficiency. The major contributions of the work are creating the closed-loop thermal-conscious DVFS architecture, incorporating the reinforcement learning to the adaptive controllers, and demonstrating the approach in the realistic workload conditions. Future directions include developing the framework to hardware implementation on FPGA and ASIC platforms, using multi-core and heterogeneous architectures, and developing new techniques of enhanced deep reinforcement learning to achieve faster convergence and better scalability in real-time processes.

REFERENCES

1. M. A. Al-Garadi, M. S. Mohamed, A. Al-Ali, X. Du, I. Ali, and M. Guizani, “A survey on energy efficient narrow-band Internet of Things (NB-IoT): Architecture, application and challenges,” *IEEE Access*, vol. 7, pp. 16739–16776, 2018, doi: 10.1109/ACCESS.2018.2881533.
2. Y. Wu et al., “Deep reinforcement learning-based power management for chiplet-based multicore systems,” *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, 2024, doi: 10.1109/TVLSI.2024.3415487.
3. A. B. Abdallah, M. M. Fouad, and M. A. Abdelrahman, “Dynamic workload-aware DVFS for multicore systems using machine learning,” *Computing*, vol. 103, pp. 1747–1769, 2021, doi: 10.1007/s00607-020-00845-2.
4. A. Bellemare, W. Dabney, and R. Munos, “A survey on multi-agent deep reinforcement learning: From the perspective of challenges and applications,” *Artificial Intelligence Review*, vol. 54, pp. 3215–3238, 2021, doi: 10.1007/s10462-020-09938-y.
5. A. Gupta, S. Sharma, and R. K. Singh, “AI/ML algorithms and applications in VLSI design and technology,” *Integration, the VLSI Journal*, vol. 93, 2023, doi: 10.1016/j.vlsi.2023.06.002.
6. B. K. Sahu, S. K. Panda, and S. Mishra, “Reinforcement learning-based dynamic voltage and frequency scaling for energy-efficient computing,” in *Proc. Int. Conf. Distributed Computing and Electrical Circuits and Electronics (ICDCECE)*, 2024, doi: 10.1109/ICDCECE60827.2024.10549241.
7. C. H. Lim, J. Park, and S. H. Lee, “Accurate load prediction in dynamic voltage frequency scaling systems,” *J. Integrated Circuits and Systems*, vol. 20, no. 1, pp. 1–14, 2025, doi: 10.29292/jics.v20i1.977.
8. S. S. Shetty, R. S. Kumar, and K. N. Balasubramanian, “Q-learning based DVFS for multi-core real-time systems,” in *Advances in Natural Computation, Fuzzy Systems and Knowledge Discovery*, 2022, doi: 10.1007/978-3-030-89698-0_35.
9. J. Chen, X. Wang, and Y. Li, “Online power management for multi-cores: A reinforcement learning based approach,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 33, no. 3, pp. 751–764, 2021, doi: 10.1109/TPDS.2021.3092270.
10. J. Li, Y. Zhang, and Z. Wang, “Workload forecasting and energy state estimation in cloud data centres: ML-centric approach,” *Future Generation Computer Systems*, vol. 128, pp. 320–332, 2022, doi: 10.1016/j.future.2021.10.019.
11. R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
12. M. A. Khan, S. U. Rehman, and A. Ullah, “Toward intelligent resource management in dynamic fog computing-based Internet of Things environment with deep reinforcement learning: A survey,” *Int. J. Communication Systems*, vol. 36, 2023, doi: 10.1002/dac.5411.
13. M. A. Rahman, S. U. Islam, and M. Guizani, “A DRL-driven intelligent joint optimization strategy for computation offloading and resource allocation in ubiquitous edge IoT systems,” *IEEE Trans. Emerging Topics Comput. Intell.*, vol. 7, no. 1, pp. 39–54, 2022, doi: 10.1109/TETCI.2022.3193367.
14. P. K. Sharma, A. Kumar, and R. Singh, “FiDRL: Flexible invocation-based deep reinforcement learning for DVFS scheduling in embedded systems,” *IEEE Trans. Computers*, 2024, doi: 10.1109/TC.2024.3465933.