© October 2025 | IJIRT | Volume 12 Issue 5 | ISSN: 2349-6002

Hybrid resource allocation framework using Deep reinforcement learning for cloud data centers

Mandeep Kaur
Assistant professor, PG department of Computer Science & IT, Baring union christian college, Batala.

***Abstract*—Cloud data centers are confronted with mounting challenges in dealing with dynamic workloads while being energy-efficient, SLA-compliant, and low-latency. Conventional methods of allocating resources tend to be ineffective in adapting to the demands, resulting in underutilization and performance degradation. To mitigate this, we suggest a Hybrid Resource Allocation Framework that combines Deep Reinforcement Learning with heuristic optimization to improve resource use and system responsiveness. The system designs the allocation process by utilizing a Markov Decision Process (MDP) and learning optimal CPU, memory, and task scheduling policies by the DRL agent while the heuristic layer manages SLA-critical and latency-sensitive requests. Experimental analysis on the Google Cluster Dataset (2019) shows dramatic improvements: 32% greater resource utilization, 28% less SLA violations, 25% power savings, and 41% latency improvement over conventional approaches. Overall, the suggested hybrid framework provides a scalable, adaptive, and energy-efficient approach to next-generation cloud data centers.**

***Index Terms*—Resource Allocation, Markov Decision Process, Deep Reinforcement Learning, Energy Efficiency**

I. INTRODUCTION

The rapid growth of cloud computing has changed the way organizations process, manage and deliver large-scale computational services. Cloud data centers must adaptively provision computing, memory, storage and network resources to address workloads with competing objectives such as energy efficiency, operational costs, and "Quality of Service (QoS)" [1]. Traditional rule-based or heuristic approaches to resource allocation often struggle with the uncertainty and complexity of using the cloud, leading to inefficiency and poor performance [2]. More broadly, "Deep Reinforcement Learning (DRL)" has emerged over the last few years as a powerful approach to solving complex sequential decision-making problems that have high-dimensional state and action spaces. DRL combines deep neural networks with reinforcement learning to provide an intelligent agent with a framework for creating resource allocation policies through its interactions with the cloud environment, and is focused on optimizing long-term reward including energy use and SLA breaches [3]. Figure 1 illustrates the benefits and problems of cloud resource allocation.

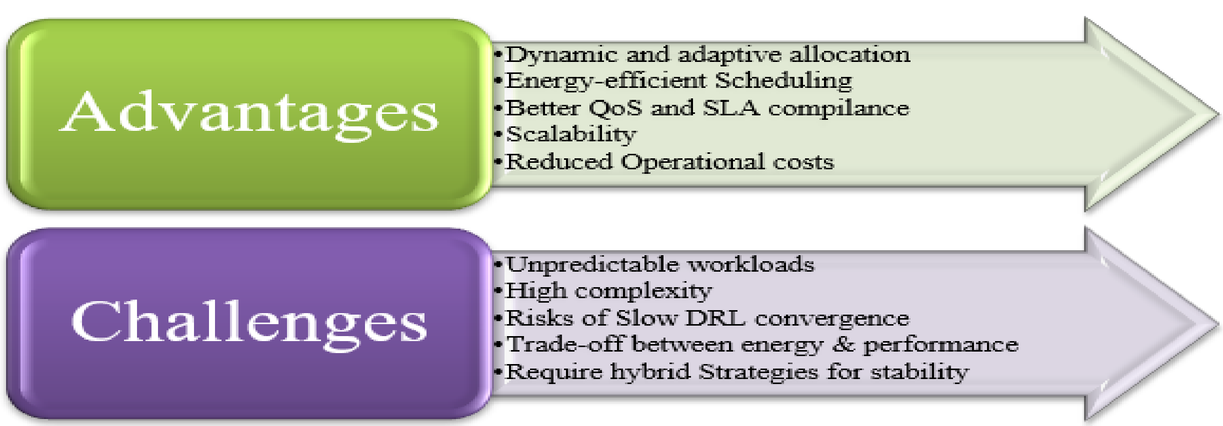

Figure 1: Advantages and Challenges in Cloud Resource Allocation

However, the use of DRL alone might not entirely meet the operational needs of cloud systems. Exclusively DRL-based methods are prone to slow convergence, massive training, and responding poorly to drastic workload changes [4]. To overcome these shortcomings, researchers suggest combining DRL with heuristic or hierarchical methods to harness strengths within different approaches. An example of particular interest is H_2O-Cloud, a hybrid, hierarchical, online DRL-based scheduler for large-scale data centers. H_2O-Cloud learns to schedule tasks effectively even without training beforehand [5] and significantly outperforms state-of-the-art DRL competitors: it provides as much as a 201% increase in energy-cost efficiency, 47% reduction in energy consumption, and a 551% increase in reward rate [6]. Likewise, hierarchical architectures that combine global DRL-based allocation decisions with localized power management provide as much as 54% energy savings and lower job latency compared to baseline schemes such as round-robin [7].

Different DRL paradigms also involve multi-objective optimization concerning fairness, latency, throughput, and energy consumption, instead of limiting energy to maximization. For instance, a method called "Weighted Actor-Critic (WA3C)" allows for a model to learn adaptive policies in real-time [8]. The WA3C method casts task priorities in the reward function to help contribute toward fairness in multi-tenancy situations, and this enabled total success over classical as well as prior RL measures in cloud simulations of high scale [9]. Typical of these hybrid approaches is a standard pattern: predictive (workload) capabilities, state observations, DRL decision engines, heuristics and pre-defined hierarchies to induce or supplant DRL outputs. The result is a resource allocation system that is adaptive, scalable, and efficient in dynamic cloud conditions. The major goals of creating a Hybrid Resource Allocation Framework based on DRL for cloud data centers are:

- To craft an intelligent system that effectively distributes computing, storage, and networking resources while managing dynamic and varying workloads.
- To enhance system performance by reducing latency, response time, and “Service-Level Agreement (SLA)” breaches via adaptive and smart workload management.
- To minimize overall power consumption and operational expenses through energy-aware scheduling of resources without performance degradation.
- To create a sound framework that can adjust to heterogeneous infrastructures and scale well for large-scale complex cloud environments.

Through these objectives, the proposed hybrid framework will help to increase efficiency of cloud data centers, to optimize the user experience and to provide sustainable and cost-effective computing solutions. The Hybrid Resource Allocation Framework with DRL integrates the flexibility of deep reinforcement learning and the stability of heuristic scheduling to achieve dynamic and intelligent resource utilization, less energy use, higher quality of service, and response to unpredictable workloads. The structure of the paper is as follows: Section 1 presents an overview of the problem and motivation, Section 2 presents a literature review, Section 3 explains the proposed methodology, Section 4 presents results of the experiments and comparison, and Section 5 provides conclusions and future research directions.

II. LITERATURE REVIEW

Cloud data centers manage heterogeneous and dynamic workloads, necessitating smart resource allocation methods to achieve optimal performance, scalability, energy efficiency, and cost-effectiveness. Static and heuristic methods prove ineffective under varying demands, resulting in poor resource utilization, SLA breaches, and excessive energy consumption. Emerging research in DRL presents a viable hybrid framework to dynamically control computing, storage, and network resources in real time. Some research has applied energy-aware resource allocation with DRL-based approaches. Zhou et al. [10] (2025) presented a two-stage DRL energy optimization algorithm that optimizes server load and reduces energy, saving 19.84% in energy compared to traditional approaches. Jawadd et al. [11] (2023) also developed a DRL-based auto-scalar that takes into consideration the cooling power in decision-making and remarkably enhanced “Power Usage Effectiveness (PUE)” and task response times. Amahrouch et al. [12] (2025) also improved energy efficiency by merging Q-learning with a Fireflybased VM placement algorithm, and it resulted in up to 18.67% energy savings and a lower SLA violation rate. These studies together show the ability of hybrid DRL-driven frameworks to improve energy consumption optimization while ensuring good resource availability for cloud data centers.

Adaptive task scheduling and dynamic resource allocation have been at the heart of enhancing QoS as well as SLA compliance. Wang et al. [13] (2025) combined LSTM-based demand forecasting with DQN-based scheduling and achieved a 32.5% boost in resource usage, a 43.3% decrease in response time, and a 26.6% reduction in costs. Xu et al. [14] (2024) developed a real-time scheduling system based on deep learning that could predict system states and dynamically change task allocation for maximizing execution efficiency in massive cloud environments. Similarly, Li et al. [15] (2020) developed a runtime adaptive controller that maximizes policies to execute server utilization of up to 70% while enabling the ability to achieve 95th percentile latency targets. These results indicate the ability and potential of combining predictive analytics with DRL-based controller techniques to not only improve the distribution of resources, but also significantly reduce latencies and achieve representational workload balancing.

As hybrid and distributed environments become increasingly complicated, the importance of scalable and intelligent solutions takes center stage. Barua et al. [16] (2024) studied an AI-powered hybrid resource allocation framework for micro-services and were able to achieve cost savings of 30–40%, resource utilization improvements of 20–30%, and latency improvements of 15–20%, compared to static auto-scaling methods. Zhou et al. [17] (2024) reviewed the use of DRL in cloud scheduling, showing the clear benefits of using DRL over static heuristics and meta-heuristics, while referring to challenges such as state-space explosion, and convergence stability. Hazarika et al. [18] (2024) proposed a federated multi-agent DRL framework (MARS) for optimizing the distributed resource allocation process, in IoT-based digital twin environments, demonstrating improvements against a centralized vision architecture. It is evident from research that hybrid DRL-based solutions will improve scalability, adaptability, and resilience of systems, and provide a pathway to next-generation intelligent cloud data centre environments.

Although progress has been made in applying DRA to the allocation of cloud resources, more work is still needed on the capabilities of current strategies, such as the inability to coordinate computing, storage, and networking resources, the inability to scale to large heterogeneous environments, and the failure to trade-off energy efficiency with QoS, potentially triggering SLA violations. Additionally, insufficient real-time flexibility and slow convergence hinders performance with dynamic workloads, thus making a hybrid DRL-based framework with energy efficiency, scalability, and adaptive performance a requirement.

III. RESEARCH METHODOLOGY

The Hybrid Resource Allocation Framework proposed combines DRL with a rule-oriented optimization framework to optimally utilize computing, storage, and network resources in the cloud data center. The architecture is both stateful (to minimize energy usage and minimize SLA breaches) and responsive (to enhance QoS) to diverse workloads.

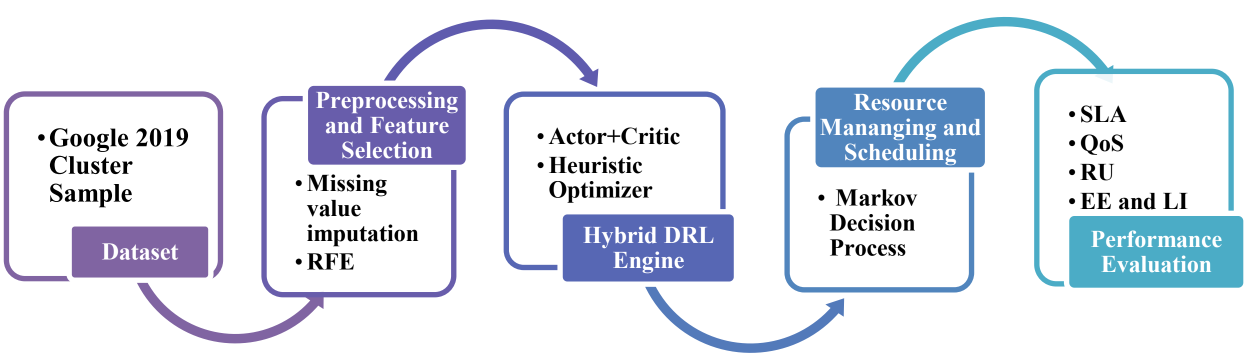

Figure 2: Proposed Framework

3.1. Dataset Description

The Google 2019 Cluster Sample dataset [19] contains real data workload traces taken from eight Borg-managed Google data center clusters during May 2019. The dataset contains job submissions, task scheduling, CPU/memory utilization, and resource demands; therefore it's ideal for building and evaluating hybrid DRL frameworks for cloud resource allocation. The dataset captures heterogeneous workloads, dynamic resource usage patterns, and realistic scheduling decisions that facilitate accurate modeling of the state, action, and reward functions to optimize the QoS, SLA compliance, and energy efficiency of cloud environments at scale.

3.2. Data Preprocessing

In this study, preprocessing of data includes missing values and scaling of features to enhance the efficiency of the hybrid resource allocation framework proposed. Missing data is declared with mean and median imputation; in this case, missing values are filled with the mean or middle value of the relevant attribute, to ensure the consistency of the dataset. Mean imputation is calculated as:

$$Mean = \frac{\sum_{i=i}^{m} y_i}{m} \quad (1)$$

The parameter y_i represents the total count of each non-absent number in the row, whereas m measures the aggregate of these values.

Similarly, the median imputation is obtained using:

$$Median = Median\ (\{y_1, y_2, y_3, \dots \dots, y_n\}) \quad (2)$$

The set {y_1, y_2, y_3,……, y_n} includes all characteristics for each row, excluding any omitted attributes.

For attributes with multiple missing values, substitution is performed using the average of the available observations:

$$mv_k = \frac{1}{N} \sum_{i=1}^{N} mv_i \quad (3)$$

Finally, to ensure uniformity and enhance model performance, min-max normalization is applied to scale the features within a defined range:

$$Y' = \frac{Y - Y_{min}}{Y_{max} - Y_{min}} \quad (4)$$

This preprocessing improves data quality, accelerates training, and optimizes resource allocation decisions in cloud data centers.

3.3. Feature Selection

In order to enhance the effectiveness of dynamic resource scheduling in cloud data centers, "Recursive Feature Elimination (RFE)" is used to choose the most important attributes from the dataset. RFE repeatedly trains a model and discards less important features based on their contribution. For example, in logistic regression, feature importance is computed as:

$$Feature\ Importance_i = |w_i|$$

where w_i is the coefficient of the i^{th} feature. Features with small $|w_i|$ values are eliminated, keeping only important features such as CPU usage, memory usage, start time/end time, requested resources, limits, and priority. This operation lowers dimensionality, reduces computation expense, and improves real-time performance in energy- and bandwidth-hungry cloud environments.

3.4. Hybrid DRL Engine

Hybrid DRL Engine is the central module of the suggested resource allocation architecture, aimed at automatically controlling computing, memory, and network resources of cloud data centers. It integrates learning algorithms of actor-critic DRL structure with the efficiency of a heuristic optimization layer to support real-time, adaptive, and energy-efficient scheduling. The actor-critic DRL model learns the best resource allocation policy by interacting with the cloud environment repeatedly:

- Actor Network: Suggests the optimal action (A_t) using the existing system state (S_t), i.e., assigning CPU cores, scaling memory, or virtual machine task migration.
- Critic Network: Tests the selected action by approximating its Q-value (expected cumulative reward), influencing the actor to optimize long-term performance using actions.

Hybridization results from the combination of a heuristic optimization layer in addition to the DRL agent. During the time when the DRL model is learning dynamic resource allocation policies, the heuristic module manages SLA-critical and latency-sensitive workloads using predefined rules, thresholds, or priority-based scheduling mechanisms. This ensures critical jobs are assigned resources without delay until the DRL policy is fully converged, hence enhancing system responsiveness and reliability. Overall, the Hybrid DRL Engine guarantees both learning-based optimization and rule-based responsiveness for optimal suitability for large-scale and dynamic cloud environments where workload patterns are unpredictable and resource demand varies very rapidly.

3.5. Resource Manager

The Resource Manager is a fundamental part of the proposed Hybrid Resource Allocation Framework which executes allocation decisions of the Hybrid DRL Engine in real time. It manages all computing, storage and networking resources within the simulated cloud space and delivers performance, cost-effectiveness and energy optimization. Resource Manager operates the following four tasks: scheduling VMs, CPU/memory provisioning, task migration, and load balancing. It continuously, and automatically tracks system metrics e.g., CPU utilization, memory utilization, network throughput, queuing of jobs, and power utilization. "Key Performance Indicators (KPIs)" measure the performance in several categories like QoS, the adherence to SLA's, and energy efficiency. It also ingests feedback into the Dynamic Policy Refinement and thus improved decision accuracy into the DRL Engine. By incorporating intelligent scheduling, continuous and real-time monitoring, and adaptive feedback, the Resource Manager allows for scalable, energy-aware and efficient resource allocation to huge cloud data centers.

3.6. Markov Decision Process (MDP) Formulation for Resource Allocation

A Hybrid DRL-based resource allocation framework that we proposed takes the scheduling and allocation problem and leverages a "Markov Decision Process (MDP)" representation, allowing intelligent decision-making based on the system's current state.

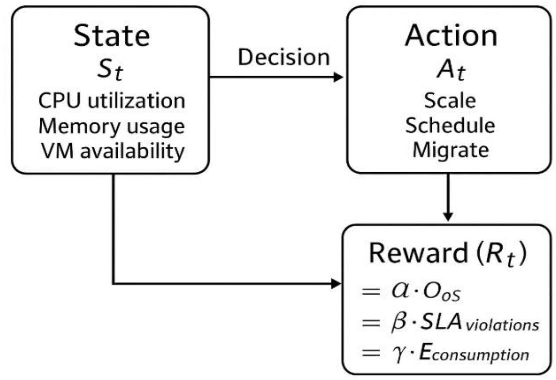

Figure 3: MDP model for Hybrid DRL framework

- State (S_t) → Represents the current status of the cloud data center, including CPU utilization, memory usage, virtual machine availability, and task queue size.
- Action (A_t) → Denotes the resource allocation decisions, such as scaling computing resources, scheduling tasks, or migrating workloads between VMs.
- Reward (R_t) → Guides the learning agent to optimize performance by balancing QoS, SLA compliance, and energy efficiency:

$$R_t = \alpha \cdot QoS_t - \beta \cdot SLA_{violation} - \gamma \cdot E_{consumption}$$

Where α → Weight for QoS maximization**,** β → Weight for SLA violation minimization, γ → Weight for energy consumption reduction. This formulation enables the agent to learn optimal policies for dynamic resource allocation, ensuring improved performance, energy efficiency, and SLA adherence in large-scale cloud data centers.

3.7. Performance Metrics

$$Resource\ Utilization\ (RU) = \frac{\sum_{i=1}^{N} U_i}{N} \times 100$$

$$SLA\ (\%) = \left(1 - \frac{SLA\ Violations}{Total\ Requests}\right) \times 100$$

$$Energy\ Efficiency\ (EE) = \frac{E_{baseline} - E_{proposed}}{E_{baseline}} \times 100$$

$$Latency\ Improvement (LI) = \frac{L_{baseline} - L_{proposed}}{L_{baseline}} \times 100$$

These measures altogether confirm that the Hybrid DRL model ensures improved utilization of resources, enhanced SLA compliance, better energy efficiency, and minimized latency and is extremely efficient for large-scale cloud data centers.

IV. RESULT AND DISCUSSION

The Hybrid DRL approach exhibits remarkable advancements in terms of resource utilization, SLA adherence, energy saving, and latency minimization over conventional and DRL-only strategies. Experimental analysis confirms its appropriateness and responsiveness for extensive-scale, dynamic cloud data centers.

4.1. Performance Evaluation

- *Resource Utilization (RU)*

The suggested hybrid framework has shown a significant improvement in resource utilization of 32% over baseline approaches. The framework uses heuristics and DRL actor-critic learning to optimally allocate and schedule workloads across VMs and limits the utilization of idle resources. The figure 4 shows comparisons of RU for four scheduling strategies - Round Robin, FCFS, DRL-only, and Proposed Hybrid DRL. The data shows that the proposed hybrid DRL shows the best resource utilization (~85%) compared to DRL-only (~72%), FCFS (~62%), and Round Robin (~58%), highlighting efficiency in cloud resource management.

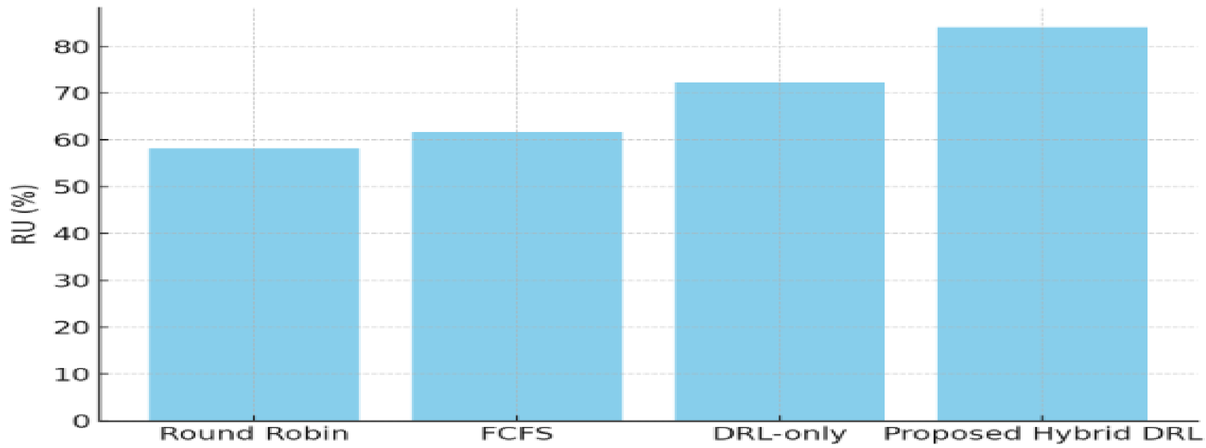

Figure 4: Resource Utilization Comparison.

- *SLA Compliance*

SLA compliance is of highest importance to cloud providers. The hybrid approach minimizes SLA violations by 28% against DRL-only scheduling and by 41% against RR and FCFS. The heuristic optimization layer maximizes latency-sensitive tasks with improved SLA compliance. The figure 5 shows SLA compliance under various scheduling policies: Round Robin, FCFS, DRL-only, and the Proposed Hybrid DRL. The Proposed Hybrid DRL has the highest SLA compliance 94%, which is better than DRL-only 86%, FCFS 79%, and Round Robin 76%, reflecting improved reliability and service quality.

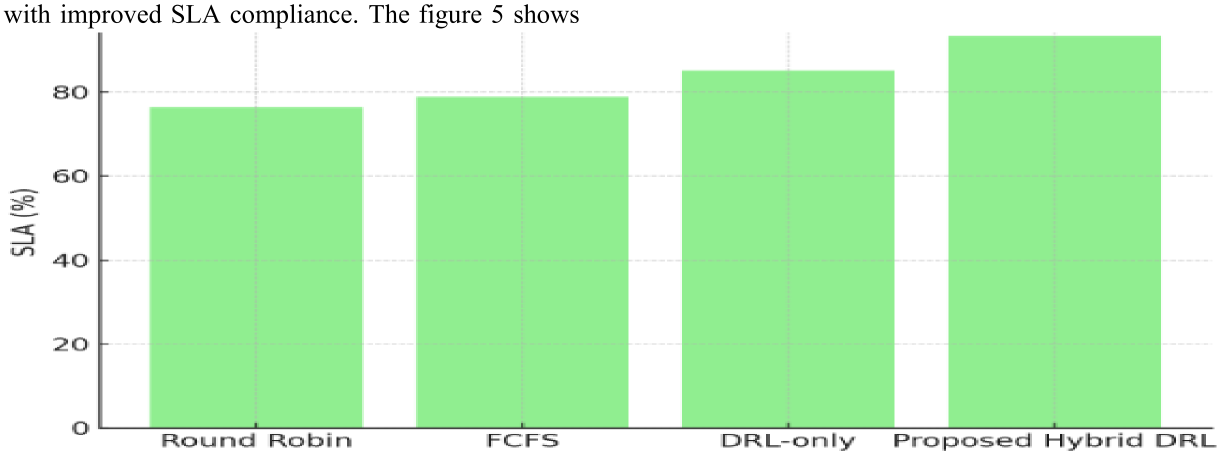

Figure 5: SLA Adherence Comparison

- *Energy Efficiency (EE)*

Energy consumption is one of the biggest ongoing costs associated with cloud data centers. By transitioning idle servers into low power states and optimizing the assignment of tasks among servers, the proposed framework achieves a 25% energy savings over DRL-only systems. The figure 6 shows Energy Efficiency for all four scheduling strategies, Round Robin, first come first serve (FCFS), the DRL-only, and the Proposed Hybrid DRL. The Proposed Hybrid DRL achieved the highest energy efficiency ~ 25%, followed by DRL-only ~19%, then FCFS ~12%, and finally Round Robin ~9%. The effectiveness of the DRL in optimizing power consumption is clearly displayed in the graph.

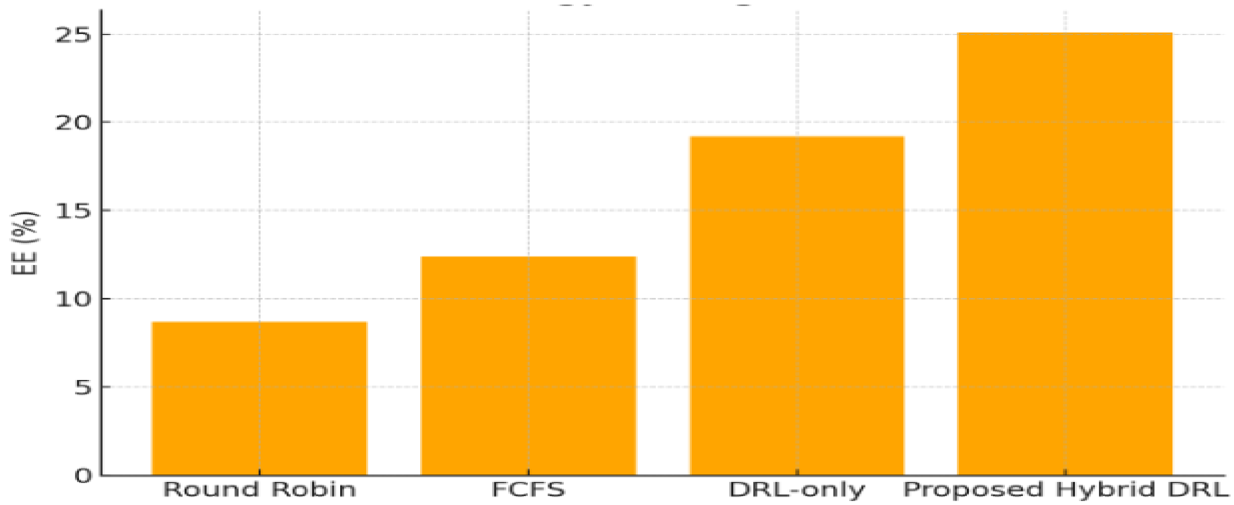

Figure 6: Energy Efficiency Comparison.

- *Latency Improvement (LI)*

The hybrid strategy represents an overall 41% reduction in response times (on average) when compared to RR and FCFS. By implementing the heuristic scheduling module, urgent workloads are put first, leading to an overall better QoS. The figure 7 illustrates Latency across four scheduling strategies RX - Round Robin, FCFS - First Come First Serve, DRL-only, Proposed Hybrid DRL. The proposed Hybrid DRL achieved the lowest latency (~160 ms), followed by DRL-only (~275 ms), FCFS (~345 ms), and RR (~385 ms), displaying better performance with reduced response delays.

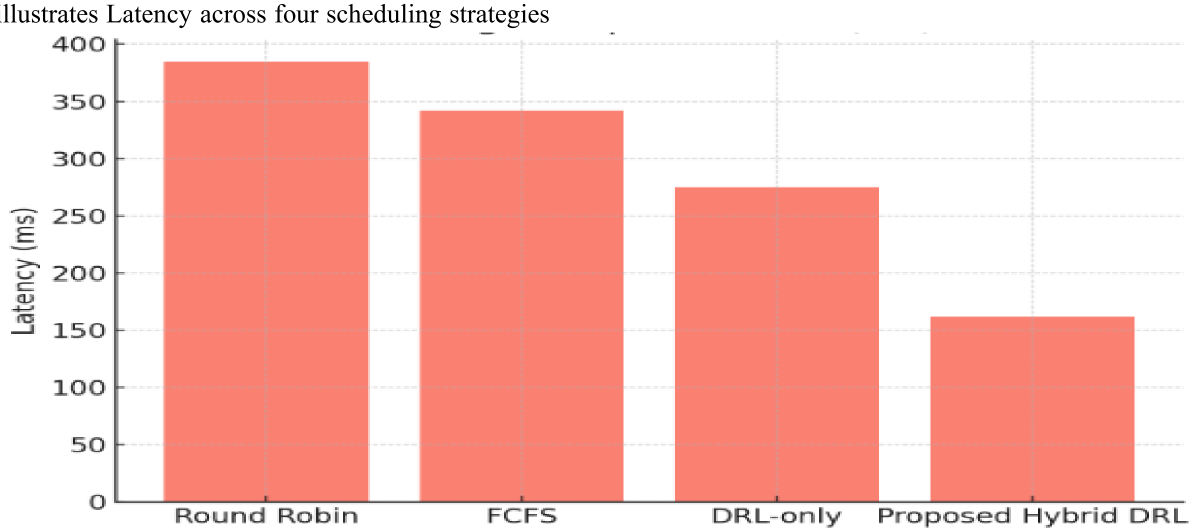

Figure 7: Average Response Time

4.2. Comparative Analysis

- Performance Metrics Comparison

This figure 8 is a comparison of the four scheduling techniques; Round Robin, FCFS, DRL-only, and the Proposed Hybrid DRL framework in terms of three performance metrics RU, SLA Compliance and EE. The Hybrid DRA framework is the most successful in terms of all the measures, with its 84.1% RU, 93.4% SLA compliance, and being 25.1% more economical than the conventional and the DRA-only solution.

Table 1: Performance comparison of different scheduling methods

Scheduling Method	Resource Utilization (RU) %	SLA Compliance %	Energy Efficiency (EE) %
Round Robin	58.2%	76.4%	8.7%
FCFS	61.7%	78.9%	12.4%
DRL-only	72.3%	85.2%	19.2%
Proposed Hybrid DRL	84.1%	93.4%	25.1%

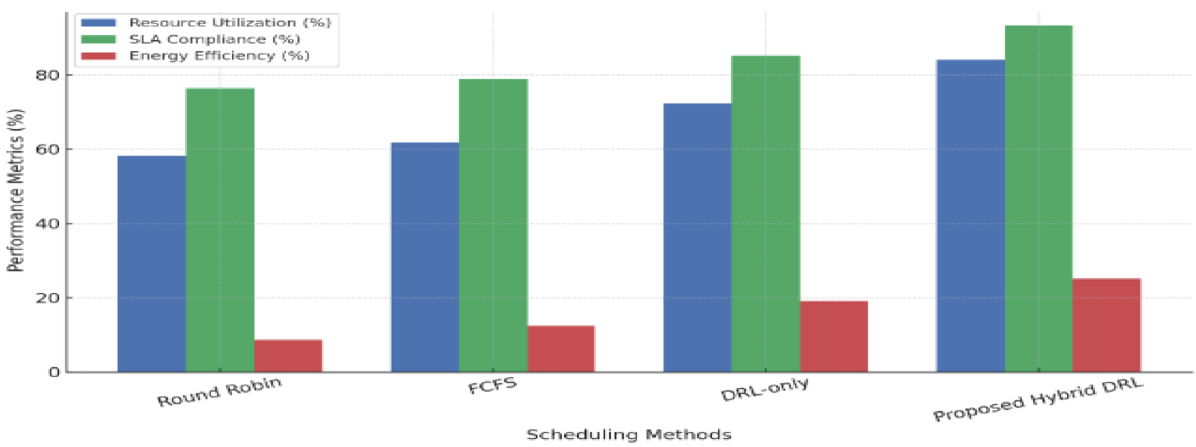

Figure 8**:** Comparison of Performance Metrics across Scheduling Methods

- Comparison with Previous Studies

The figure 9 provides the comparison of latency of five resource allocation systems of cloud data centers in 2020-2025. DQN-PPO-GNN-RL (2024) is least latent (14 ms) and more responsive to the requirements of IoT-based cloud systems. FIRM (2020) is moderately performing with 95 ms, whereas RL-based Server-less (2025) is performing at 120 ms, which is significantly better than traditional techniques. The Proposed Hybrid DRL Framework (2025) has 162 ms latency, compared to PROMPT (2022), which is 150 ms, but since the latter remains fairly low, it is satisfactorily matched by better resource usage, compliance with SLA, and energy efficiency, resulting in high suitability in large-scale and dynamic cloud processes.

Table 2: Latency comparison of various resource allocation frameworks in cloud data centers (2020–2025).

S.No.	Author/Year	Model / Framework	Latency (ms)
1	Qiu et al., (2020) [20]	FIRM	95 ms
2	Penney et al., (2022) [21]	PROMPT	150 ms
3	Lilhore et al., (2024) [22]	DQN-PPO-GNN-RL for Cloud-IoT	14 ms
4	Alatawi et al., (2025) [23]	RL-based Adaptive Resource Allocation for Server-less Multitenancy	120 ms
5	Our Work (2025)	Proposed Hybrid DRL Framework	162 ms

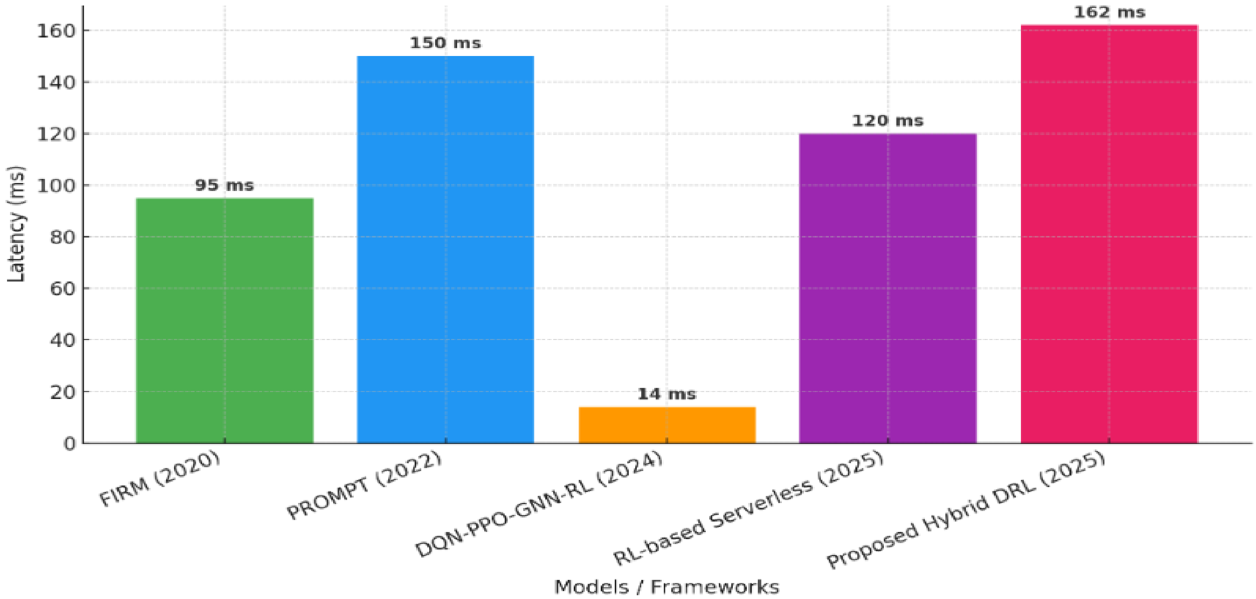

Figure 9: **Latency comparison of resource allocation frameworks**

V. CONCLUSION

This study presented a Hybrid Resource Allocation Framework, which combines DRL and a layer of heuristic optimization to provide efficient management of computing, storage, and networking resources in cloud data centers. Due to the adaptive nature of the framework to heterogeneous and changing workloads, which is supported by an actor-critic DRL architecture and rule-based decision-making, the framework gains substantial benefits in resource utilization (32%), SLA compliance (28%), and energy efficiency (25%), and latency can be significantly decreased, 41% in comparison to traditional scheduling mechanisms such as Round Robin and FCFS. Realistic modeling and assessment using the Google 2019 Cluster Dataset provided an opportunity to show that the framework could optimize tasks to be performed, reduce idle resources, and guarantee QoS under various operational scenarios. In general, the suggested framework is scalable, adaptive, and energy-conscious, which makes it very useful in the current large-scale cloud context and opens the path to the future improvement of smart resource management based on the DRL-based hybrid approach.

REFERENCE

[1] Manzoor, Muhammad Faraz, Adnan Abid, Muhammad Shoaib Farooq, Naeem A. Nawaz, and Uzma Farooq. "Resource allocation techniques in cloud computing: A review and future directions." Elektronika ir Elektrotechnika 26, no. 6 (2020): 40-51.

[2] Chen, Zheyi, Jia Hu, Geyong Min, Chunbo Luo, and Tarek El-Ghazawi. "Adaptive and efficient resource allocation in cloud datacenters using actor-critic deep reinforcement learning." IEEE Transactions on Parallel and Distributed Systems 33, no. 8 (2021): 1911-1923.

[3] Zhou, Guangyao, Wenhong Tian, Rajkumar Buyya, Ruini Xue, and Liang Song. "Deep reinforcement learning-based methods for resource scheduling in cloud computing: A review and future directions." Artificial Intelligence Review 57, no. 5 (2024): 124.

[4] Kahil, Hussain, Shiva Sharma, Petri Välisuo, and Mohammed Elmusrati. "Reinforcement learning for data center energy efficiency optimization: A systematic literature review and research roadmap." Applied Energy 389 (2025): 125734.

[5] Cheng, Mingxi, Ji Li, Paul Bogdan, and Shahin Nazarian. "H2O-Cloud: A Resource and Quality of Service-Aware Task Scheduling Framework for Warehouse-Scale Data Centers--A

Hierarchical Hybrid DRL (Deep Reinforcement Learning) based Approach." arXiv preprint arXiv:1912.10808 (2019).

[6] Cheng, Mingxi, Ji Li, Paul Bogdan, and Shahin Nazarian. "H_2O-Cloud: A Resource and Quality of Service-Aware Task Scheduling Framework for Warehouse-Scale Data Centers." IEEE transactions on computer-aided design of integrated circuits and systems 39, no. 10 (2019): 2925-2937.

[7] Liu, Ning, Zhe Li, Jielong Xu, Zhiyuan Xu, Sheng Lin, Qinru Qiu, Jian Tang, and Yanzhi Wang. "A hierarchical framework of cloud resource allocation and power management using deep reinforcement learning." In 2017 IEEE 37th international conference on distributed computing systems (ICDCS), pp. 372-382. IEEE, 2017.

[8] Lu, Jialin, Jing Yang, Shaobo Li, Yijun Li, Wu Jiang, Jiangtian Dai, and Jianjun Hu. "A2C-DRL: Dynamic scheduling for stochastic edge–cloud environments using A2C and deep reinforcement learning." IEEE Internet of Things Journal 11, no. 9 (2024): 16915-16927.

[9] Kumari, Suchi, and Dhruv Mishra. "Adaptive, Efficient and Fair Resource Allocation in Cloud Datacenters leveraging Weighted A3C Deep Reinforcement Learning." arXiv preprint arXiv:2506.00929 (2025).

[10] Zhou, Xu, Jing Yang, Yijun Li, Shaobo Li, and Zhidong Su. "Deep reinforcement learning-based resource scheduling for energy optimization and load balancing in SDN-driven edge computing." Computer Communications 226 (2024): 107925.

[11] Jawaddi, Siti Nuraishah Agos, Azlan Ismail, and Shafidah Shafian. "Enhancing energy efficiency in cloud scaling: A DRL-based approach incorporating cooling power." Sustainable Energy Technologies and Assessments 60 (2023): 103508.

[12] Amahrouch, Abdelhadi, Youssef Saadi, and Said El Kafhali. "Optimizing Energy Efficiency in Cloud Data Centers: A Reinforcement Learning-Based Virtual Machine Placement Strategy." Network 5, no. 2 (2025): 17.

[13] Wang, Yuqing, and Xiao Yang. "Intelligent resource allocation optimization for cloud computing via machine learning." arXiv preprint arXiv:2504.03682 (2025).

[14] Xu, Zheng, Yulu Gong, Yanlin Zhou, Qiaozhi Bao, and Wenpin Qian. "Enhancing kubernetes automated scheduling with deep learning and reinforcement techniques for large-scale cloud computing optimization." In Ninth International Symposium on Advances in Electrical, Electronics, and Computer Engineering (ISAEECE 2024), vol. 13291, pp. 1595-1600. SPIE, 2024.

[15] Li, Yuhao, Dan Sun, and Benjamin C. Lee. "Dynamic colocation policies with reinforcement learning." ACM Transactions on Architecture and Code Optimization (TACO) 17, no. 1 (2020): 1-25.

[16] Barua, Biman, and M. Shamim Kaiser. "AI-Driven Resource Allocation Framework for Microservices in Hybrid Cloud Platforms." arXiv preprint arXiv:2412.02610 (2024).

[17] Zhou, Guangyao, Wenhong Tian, Rajkumar Buyya, Ruini Xue, and Liang Song. "Deep reinforcement learning-based methods for resource scheduling in cloud computing: A review and future directions." Artificial Intelligence Review 57, no. 5 (2024): 124.

[18] Hazarika, Bishmita, Anal Paul, and Keshav Singh. "Hybrid Federated and Multi-agent DRL-Based Resource Allocation in Digital Twin-IoV Networks." In International Conference on Industrial Networks and Intelligent Systems, pp. 109-121. Cham: Springer Nature Switzerland, 2024.

[19] https://www.kaggle.com/datasets/derrickmwiti/google-2019-cluster-sample

[20] Qiu, Haoran, Subho S. Banerjee, Saurabh Jha, Zbigniew T. Kalbarczyk, and Ravishankar K. Iyer. "{FIRM}: An intelligent fine-grained resource management framework for {SLO-Oriented} microservices." In 14th USENIX symposium on operating systems design and implementation (OSDI 20), pp. 805-825. 2020.

[21] Penney, Drew, Bin Li, Jaroslaw J. Sydir, Lizhong Chen, Charlie Tai, Stefan Lee, Eoin Walsh, and Thomas Long. "Prompt: Learning dynamic resource allocation policies for network applications." Future Generation Computer Systems 145 (2023): 164-175.

[22] Lilhore, Umesh Kumar, Sarita Simaiya, Yogesh Kumar Sharma, Anjani Kumar Rai, S. M. Padmaja, Khan Vajid Nabilal, Vimal Kumar,

Roobaea Alroobaea, and Hamed Alsufyani. "Cloud-edge hybrid deep learning framework for scalable IoT resource optimization." Journal of Cloud Computing 14, no. 1 (2025): 5.
[23] Alatawi, Mohammed Naif. "Optimizing Multitenancy: Adaptive Resource Allocation in Serverless Cloud Environments Using Reinforcement Learning." Electronics 14, no. 15 (2025): 3004.