Journal of Ubiquitous Computing and Communication Technologies (ISSN: 2582-337X)
www.irojournals.com/jucct/

Carbon-aware Cloud Resource Scheduling Using Reinforcement Learning for Energy Efficiency Data Centers

Kirupavathy P.[1], Hareeni C.[2], Jayashri K.[3], Jenifer Thangam S.[4]

Department of Artificial Intelligence and Data Science, Velammal Institute of Technology, Chennai, India.

E-mail: [1]kripajairam@gmail.com, [2]hareenichandrasekar1622@gmail.com, [3]jayashri3545@gmail.com, [4]jeniferthangam02@gmail.com

Abstract

Due to the rapid development of cloud computing services, today's data centers are consuming much more energy and emitting much higher levels of carbon dioxide into the atmosphere compared to earlier years, thus requiring the development of sustainable solutions for resource management in the cloud. In this paper, an approach called Carbon Aware Cloud Resource Scheduling based on the use of Reinforcement Learning is presented. The proposed approach includes workload demand, resource utilization, and carbon intensity values in its decision making process in order to allocate workloads efficiently and minimize environmental impact caused by cloud computing services. The reinforcement learning algorithm trains the RL agent to find the optimal strategies for workload scheduling that involve execution of workloads, postponing their execution, migrating them from one location to another and re-allocation of resources. The results from experimental evaluation indicate that the suggested framework successfully reduces the carbon emission level up to 30% and energy cost level to around 24%, respectively, in comparison with traditional scheduling methods, without lowering the SLA compliance rate down to 98.4% and workload starvation level up to 0.3%. The analysis proves that the combination of carbon awareness and reinforcement learning helps to create an intelligent, adaptive, and ecologically sustainable system for managing cloud resources.

DOI: https://doi.org/10.36548/jucct.2026.3.005
*Extended version of article received from **International Conference on Smart Innovations in Engineering & Technology (ICSIET 2026)***
Received: 21.05.2026, received in revised form: 24.06.2026, accepted: 15.07.2026, published: 27.07.2026

Keywords: Carbon-Aware Scheduling, Reinforcement Learning, Cloud Computing, Energy Efficiency, Data Centers, Green Computing, Carbon Intensity.

1. Introduction

Due to the rapid adoption of cloud computing, artificial intelligence, and digital services at scale, there has been a massive increase in the amount of energy being used by contemporary data centers. With more cloud infrastructures increasingly being able to satisfy rising demand for computational power, electricity consumption and carbon emissions have become significant economic and environmental challenges. Traditional cloud resource allocation systems mainly focus on enhancing the performance and efficiency of resources while neglecting any environmental considerations for running workloads. In areas where electricity production still depends on fossil fuel sources, such scheduling policies can lead to higher carbon emissions.

The current cloud scheduling schemes generally schedule workloads according to the availability and performance needs of resources without accounting for carbon intensities involved in producing power. While many research works have been done on energy aware scheduling and optimization of workloads, most current schemes do not have adaptive capabilities that can respond to changing demands and resource utilization as well as environmental conditions. Additionally, the static or heuristic schemes tend to be ineffective in optimizing performance while achieving environmental sustainability goals in dynamic cloud computing environments. This paper suggests the need for a scheduling scheme that can dynamically adapt to changes in system conditions and carbon intensities.

To overcome the above-mentioned problems, this paper puts forth the concept of Carbon-Aware Cloud Resource Scheduling using Reinforcement Learning (RL). The key features of this proposed approach include the use of workload demand, resource usage metrics, and the carbon intensity information to formulate adaptive scheduling. An RL agent interacts with the cloud environment and finds optimal policies for workload allocation by receiving feedbacks. This approach facilitates the execution of the workload, scheduling of delay tolerant tasks, shift in workload, and resource reallocation, respectively, to lower the negative impact of cloud services on the environment while ensuring the required quality. The core purpose of this approach is to lower carbon emissions and energy consumption without violating SLA requirements.

2. Literature Survey

As a result of the fast development of cloud computing, the energy use in data centers has been growing, which increases the operational costs and carbon emissions. In order to solve these problems, carbon-aware scheduling strategies have become efficient ways to make the process more sustainable. The latest research has shown that reinforcement learning and intelligent optimization techniques are efficient means of resource allocation, which reduce carbon emission and improve scheduling efficiency in the case of different workloads [1], [2]. Reinforcement learning and intelligent optimization allow decision making, considering the state of the system and the environment, and thus these techniques can be used in a complex system like clouds. The use of hierarchical multi-agent frameworks has been suggested to achieve efficient load balancing and cooling operations within clusters of data centers, which help in achieving better energy efficiency as well as carbon reduction [3]. The concept of optimizing energy cost with respect to carbon emissions in geographically distributed data centers based on regional carbon intensities and electricity costs has also been explored, thereby lowering operating costs and carbon footprint [4]. In addition, low carbon operation using load balancing and carbon emission quota trading is another area worth mentioning [6].

The use of renewable energy sources in cloud computing has been an area of much research interest. Scheduling mechanisms based on reinforcement learning have been proposed that try to make full use of renewable energy sources by ensuring the fulfillment of deadlines and quality of service requirements [5]. This makes the system less reliant on traditional energy sources, thus decreasing the carbon footprint. Moreover, parameter control mechanisms using deep reinforcement learning have been used to adaptively tune the schedules according to changing carbon intensity and workload conditions [2]. In addition to cloud computing domain, a number of techniques used in manufacturing and industrial scheduling have proven to be highly beneficial when dealing with issues related to carbon emissions during resource management processes. For instance, multi-objective scheduling which aims to optimize both production scheduling and energy sourcing has yielded positive results [7].

Also, many-objective optimization techniques have been successfully employed in the context of scheduling for low carbon issues, resulting in a proper balance between productivity, efficiency, and emission objectives [8]. The use of intelligent scheduling techniques that include variable lot sizing and early release of units is another example of improved resource utilization and performance [10]. Adaptive monitoring and performance management systems also play an important role in contributing towards sustainable systems since they enable constant evaluation of how systems perform depending on changing operational circumstances. Self-adaptive monitoring enables real-time evaluation which helps in optimizing and making decisions that ultimately increase the efficiency of the systems. The combination of monitoring, optimization, and learning-based scheduling is extremely useful for sustainability of large-scale computing systems [9].

2.1 Research Gap

Despite the existence of many studies that have proved the possibility of using carbon-aware scheduling, energy-efficient workloads management, and reinforcement learning for cloud environment optimization [1]-[5], there are still some challenges. The majority of the existing solutions consider such goals separately as energy saving, workloads migration, use of renewable energy resources, and cost minimization [6]-[9]. At the same time, the problem of unification of real-time monitoring of carbon emission intensity, workload shifting, and optimization through reinforcement learning has received little attention [10]. Also, most of the existing approaches lack adaptability to changing workloads and environmental conditions. As such, the requirement is that of an intelligent and adaptable scheduling mechanism that is able to minimize both carbon footprint and energy consumption at the same time while ensuring resource utilization efficiency and SLA compliance. This research responds to the above problem by designing a Carbon-aware Cloud Resource Scheduling approach using Reinforcement Learning.

3. Proposed System

The proposed Carbon-Aware Cloud Resource Scheduling framework is intended to enhance the sustainability and efficiency of operations in cloud data centers by leveraging effective workload management techniques. The framework incorporates the elements of resource monitoring, carbon-awareness, and decision-making using reinforcement learning into a common scheduling system. Through considering the environmental and operational state at the time of assigning workloads, the aim of the framework is to minimize energy usage and carbon footprint without compromising cloud services' SLA satisfaction.

Figure 1 shows the design of the entire proposed system architecture. The system is composed of seven main modules: Workload Manager, Resource Monitoring Module, Carbon Intensity Monitoring Module, State Representation Module, Reinforcement Learning Agent, Scheduling Decision Module, and Performance Evaluation Module. The workloads that are input to the system are first handled by the Workload Manager, and the system operation information is constantly acquired by the monitoring modules. This acquired information is translated to state representation to show the present state of the cloud computing environment. The RL agent formulates adaptive scheduling policies based on the state information for execution by the scheduling module.

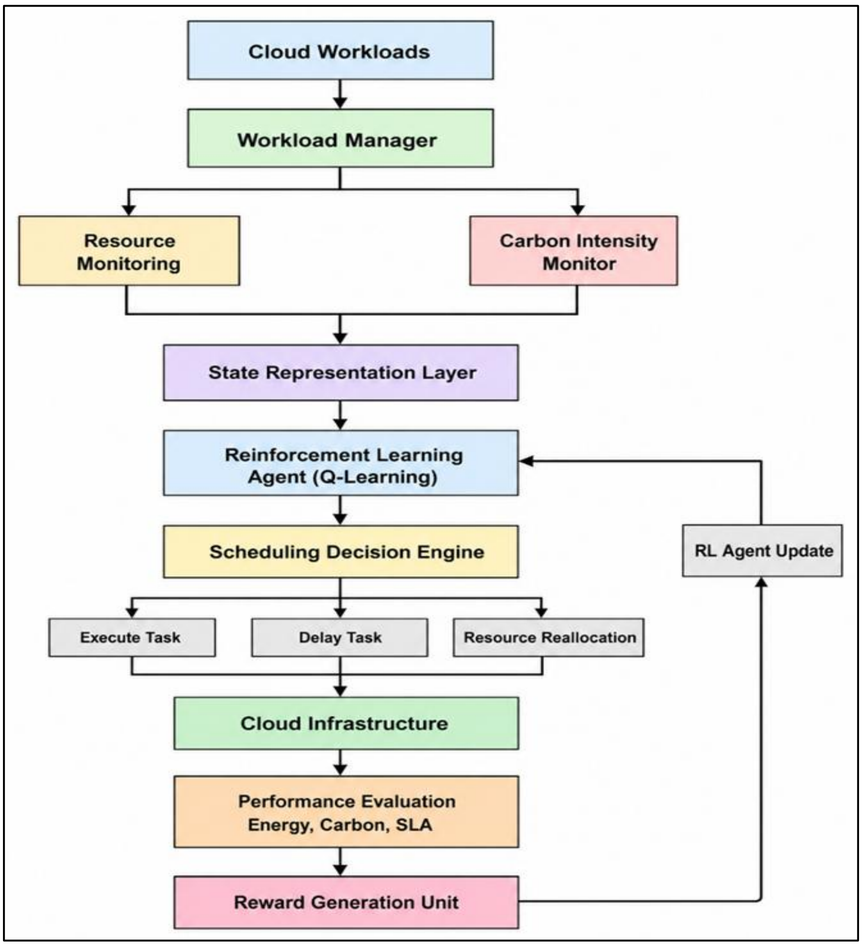

Figure 1. Workflow of the proposed carbon-aware reinforcement learning-based cloud resource scheduling framework

Interaction between modules and information flow within the scheduling framework is shown in Figure 1. The architecture creates a closed-loop decision-making system, where the information about workload properties, resource usage, and carbon footprint jointly affect the decisions made by scheduling. With the use of continuous performance monitoring and feedback-based learning, the architecture is able to adapt to changes in workload conditions and environment.

Table 1. Functional modules of the proposed system

Module	Function
Workload Manager	Receives and classifies incoming cloud workloads based on priority and execution requirements.
Carbon Intensity Monitor	Collects real-time carbon intensity information from power generation sources.
Resource Monitoring Unit	Monitors CPU utilization, memory usage, queue length, and resource availability.
RL Agent	Learns optimal scheduling policies using state-action-reward interactions.
Scheduling Engine	Executes workload allocation, migration, delay, and resource management decisions.
Performance Monitoring Module	Evaluates energy consumption, carbon emissions, response time, and SLA compliance.
Reward Evaluation Unit	Generates reward signals based on environmental and performance objectives.

Table 1 highlights the functionality of the modules in the proposed architecture. Through working together, the modules make it possible for the cloud to allocate workloads in an environmentally-aware manner, manage resources and optimize its performance. Through incorporating environmental considerations in cloud scheduling, the proposed framework creates a basis for efficient and sustainable cloud management.

Algorithm 1. Carbon-aware reinforcement learning scheduling

Input:

- Workload Requests
- Carbon Intensity Data
- Resource Utilization Metrics

Output:

- Optimized Scheduling Decision

1. Initialize RL Agent
2. Monitor workload demand and resource utilization
3. Collect real-time carbon intensity data
4. Generate current environment state S

5. Select scheduling action A using RL policy
6. If carbon intensity is low:

Execute workload
Else:
Delay non-critical task or shift workload

7. Allocate resources according to selected action
8. Measure carbon emissions, energy consumption, and SLA compliance
9. Calculate reward value
10. Update RL policy
11. Repeat until convergence

Return optimal scheduling policy

The scheduling algorithm is always learning from the environmental feedback and adapting itself according to changing workloads. The use of reinforcement learning helps in making intelligent decisions about work allocation as well as being carbon-conscious [1], [2], [5].

4. Methodology

The suggested methodology will involve incorporating carbon awareness and reinforcement learning techniques to facilitate adaptive cloud resource scheduling based on dynamic operational and environmental conditions. This methodology is intended to monitor the information regarding workload demand, resource usage, and carbon emissions intensity in order to make intelligent decisions about scheduling. Through the integration of environmental monitoring and learning, the methodology will try to minimize the emissions and energy usage while providing SLA compliance.

The overall methodology process flow is shown in Figure 2. Cloud workloads are first monitored along with various operational parameters like CPU usage, memory usage, and queue length. At the same time, carbon intensity data is gathered to determine the environmental effects caused by workload processing. This combination of parameters represents the current state of the cloud environment and provides the required context information for scheduling processes. The incorporation of both the operational and environmental aspects allows the framework to consider both performance and sustainability goals in the process of resource allocation.

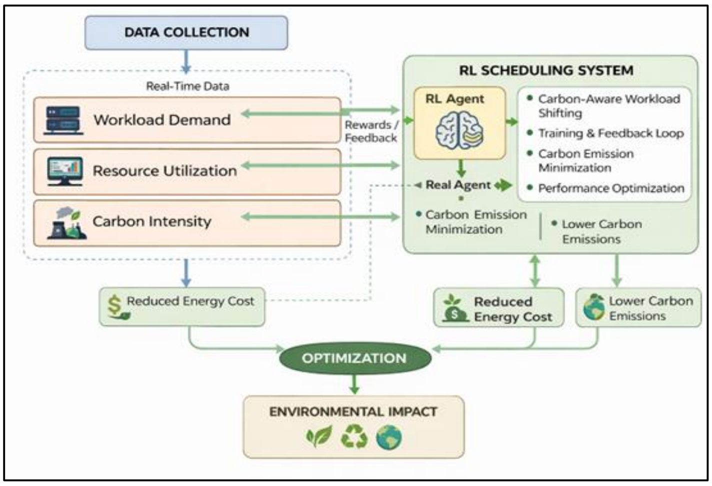

Figure 2. Carbon-aware RL scheduling workflow

For the decision-making process to be adaptive, the parameters under monitoring are converted into the state space, which is used by the reinforcement learning algorithm. The state space is defined as follows:

$$S_t = \{CPU_t, MEM_t, QL_t, WD_t, CI_t\} \tag{1}$$

Where:

- CPU_t = CPU Utilization
- MEM_t = Memory Utilization
- QL_t = Queue Length
- WD_t = Workload Demand
- CI_t = Carbon Intensity

Equation (1) allows the reinforcement learning algorithm to identify the state of the cloud infrastructure currently in terms of its operations and environment.

Based on the current state, the reinforcement learning agent comes up with scheduling policies that try to balance between environmental sustainability and performance. Scheduling engine acts upon the learned policies by taking actions such as workload scheduling, workload deferral, and resource relocation. Carbon aware workload scheduling is adopted to avoid running delay tolerant workloads when the carbon intensity level is very high. In addition, workload scheduling is encouraged during clean energy periods.

System performance after execution of workloads is assessed using environmental and operational metrics. Carbon footprint that is caused by executing workloads can be calculated using the following equation (2):

$$CE = EC \times CI \tag{2}$$

where:

- CE = Carbon Emissions ($kgCO_2$)
- EC = Energy Consumption (kWh)
- CI = Carbon Intensity ($kgCO_2/kWh$)

Equation (2) shows quantitatively the environmental effect caused by the usage of cloud resources.

The reward signal is then generated to direct the learning process of the reinforcement learning agent. The reward function consists of both environment and service level goals and is formulated as

$$R = w_1(\boldsymbol{E}_{eff}) + w_2(C_{\mathrm{red}}) + w_3(SLA_{\mathrm{comp}}) - w_4(P_{\mathrm{delay}}) \tag{3}$$

Where:

- E_{eff} = Energy Efficiency
- C_{red} = Carbon Reduction
- SLA_{comp} = SLA Compliance
- P_{delay} = Delay Penalty
- w_1, w_2, w_3, w_4 = Weight Coefficients

Equation (3) encourages the reinforcement learning agent to find out the scheduling policy which leads to minimal environmental effect while ensuring service quality. The effectiveness of the suggested framework can be evaluated through various performance measures such as carbon emission, energy usage, latency, SLA adherence, workload starvation ratio, and scheduling efficiency. The scheduling efficiency can be defined as

$$SE = \frac{W_{\mathrm{completed}}}{W_{\mathrm{total}}} \times 100 \tag{4}$$

Where:

- $W_{\text{completed}}$ = Completed Workloads
- W_{total} = Total Submitted Workloads

Equation 4 describes the capability of the proposed scheduling scheme to process the workload effectively under different operating environments.

Table 2. Methodology stages of the proposed framework

Stage	Description
Data Collection	Acquisition of workload, resource utilization, and carbon intensity information.
State Generation	Creation of state vectors using operational and environmental parameters.
RL Decision Making	Selection of scheduling actions based on learned policies.
Scheduling Execution	Allocation, delay, migration, or reallocation of workloads.
Reward Evaluation	Assessment of carbon emissions, energy consumption, and SLA compliance.
Policy Update	Learning from feedback to improve future scheduling decisions.
Performance Analysis	Comparison of scheduling effectiveness using evaluation metrics.

Table 2 outlines the critical phases in the proposed methodology. The methodology integrates environmental sensing, adaptive decision making, performance assessment, and policy learning. Using feedback and policy adjustments, the proposed methodology makes cloud resource management sustainable by minimizing carbon emissions and energy utilization without sacrificing efficiency.

5. Results and Discussion

The performance of the suggested Carbon-Aware Cloud Resource Scheduling architecture was measured by using the details about the workload demand, resource usage, and carbon intensity data in a simulated cloud system. The main emphasis was on testing whether the reinforcement learning based scheduler can reduce the carbon footprint while ensuring efficient workload processing. The key performance measures included carbon emissions reduction, energy cost reduction, response time, SLA violation, workload starvation rate, and scheduling efficiency.

5.1 Environmental and Operational Performance Analysis

The environmental performance of the proposed scheduling framework indicates that by including an awareness of carbon intensity in the scheduling system, one can have a more successful workload assignment and resource use. The scheduler, which is based on reinforcement learning, changes its behavior according to variations in the workload requirement and carbon intensity, resulting in a significant decrease in the environmental effect.

Table 3. Operational performance metrics

Performance Metric	Proposed RL Scheduler
Average Response Time (s)	156
SLA Compliance (%)	98.4
Workload Starvation Rate (%)	0.3

Summary of the operational performance of the proposed scheduling model is presented in Table 3. From the obtained results, it can be seen that environmental optimization does not affect service quality. SLA compliance and small workload starvation values show that the model provides reliable workload execution and at the same time implements energy efficient scheduling goals.

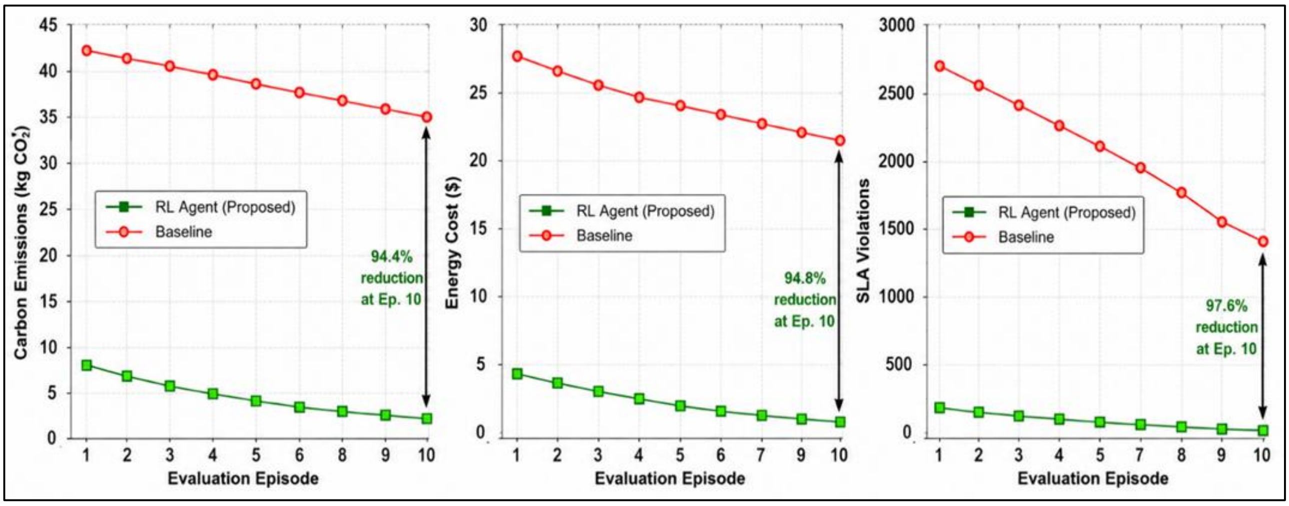

Figure 3. Comparative performance evaluation of scheduling strategies

The findings indicate that the scheduler successfully reduces carbon emissions and energy costs without affecting the stability of the service performance. Improved SLA-related metrics indicate that the process of optimization does not have any negative impacts on service performance. The findings prove that the suggested framework is appropriate for sustainable cloud computing.

5.2 Reinforcement Learning Performance Analysis

The learning process of the reinforcement learning agent is depicted in Fig. 4. It is evident from the observed convergence properties that the agent has effectively learned the scheduling policies by interacting continuously with the cloud environment. The stability in the learning performance implies that the proposed framework is able to learn effective workload allocation policies that satisfy the conflicting objectives of the environment. The convergence property also highlights the robustness of the learning process under varying workloads and carbon intensity.

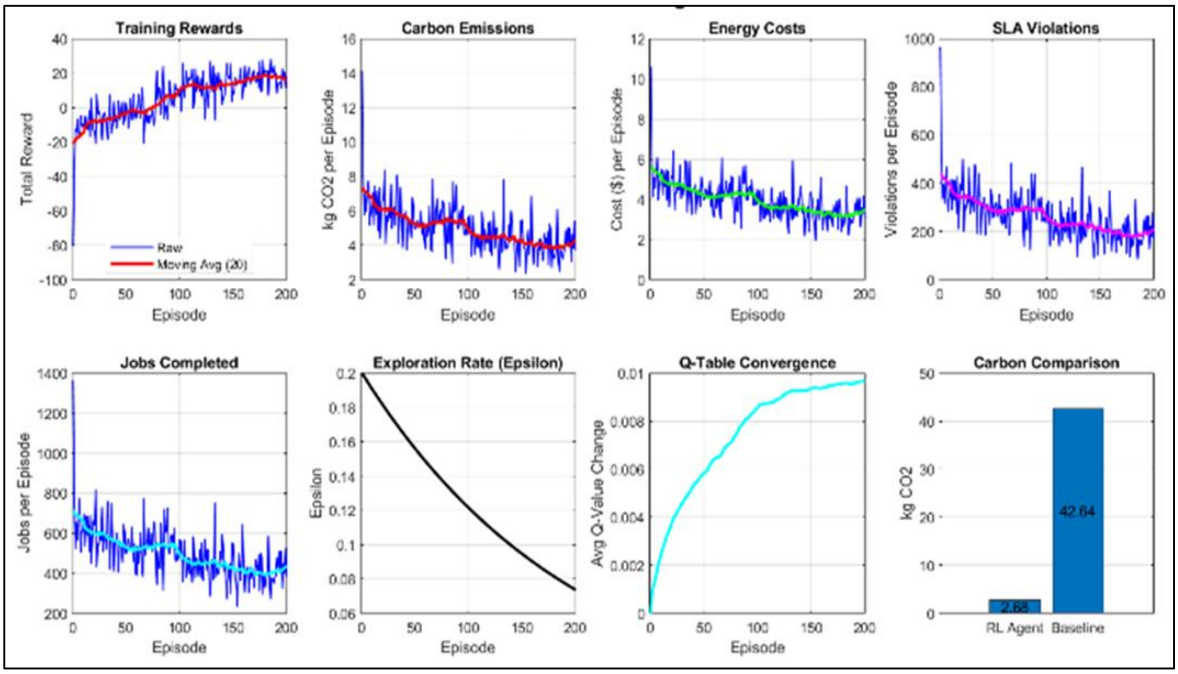

Figure 4. Training results of the proposed carbon-aware RL scheduler

Table 5 shows the performance properties of the proposed scheduler. It can be seen that the ability to learn adaptively by the reinforcement learning agent allows efficient management of resources. Through continuous adaptation of scheduling based on feedback from the environment, the framework achieves sustainability, resource efficiency, and reliability. This shows that reinforcement learning offers a viable approach to carbon-aware scheduling of cloud resources.

Table 5. Comparative performance analysis

Performance Metric	Performance-Based Scheduler	Threshold-Based Scheduler	Proposed RL Scheduler
Carbon Emission Reduction (%)	0	18	30
Energy Cost Reduction (%)	0	12	24
SLA Compliance (%)	95.8	97.1	98.4
Workload Starvation Rate (%)	1.8	0.9	0.3
Adaptability	Low	Medium	High

The results obtained through the experiments support the efficiency of the proposed Carbon-Aware Cloud Resource Scheduling framework. The combination of reinforcement learning and environmental awareness allows the scheduling framework to lower the carbon footprint and energy utilization while ensuring cloud service performance. The obtained improvements in terms of environmental, operational, and learning aspects indicate the efficiency of the proposed framework in cloud data centers' management.

6. Conclusion

This research work proposed the Carbon-Aware Cloud Resource Scheduling Framework, which leverages the use of Reinforcement Learning (RL). The use of RL allows the framework to make intelligent scheduling decisions by incorporating the use of workload requirements, resource availability, and current carbon intensity. This learning process ensures the framework is able to generate reduced carbon emissions, lower energy consumption, high service-level agreements compliance, and effective resource management. Experiments confirmed that the proposed solution provided remarkable improvements over traditional solutions, such as the reduction in carbon emissions by 30%, energy cost savings by 24%, and SLA compliance at 98.4%. Thus, it is evident that a combination of carbon awareness and reinforcement learning is effective in managing cloud resources in a sustainable manner. The future direction of research involves applying the proposed framework in a multi-cloud and edge-cloud setting using more sophisticated deep reinforcement learning techniques, including DQN and PPO, as well as considering the predictions of renewable energy generation.

References

[1] Yu, Chuanzhao, Youshan Liu, Chunjiang Zhang, And Weiming Shen. "Multi-Agent Deep Reinforcement Learning for Low-Carbon Flexible Job Shop Scheduling with Variable Sublots." Robotics and Computer-Integrated Manufacturing 2026, Vol 98: 103180.

[2] Mencaroni, Andrea, Robbert Reijnen, Yingqian Zhang, And Dieter Claeys. "Online Parameter Control for Carbon-Aware Scheduling Using Deep Reinforcement Learning." In Beta Symposium 2024.

[3] Sarkar, Soumyendu, Avisek Naug, Antonio Guillen, Vineet Gundecha, Ricardo Luna Gutiérrez, Sahand Ghorbanpour, Sajad Mousavi, Ashwin Ramesh Babu, Desik Rengarajan, And Cullen Bash. "Hierarchical Multi-Agent Framework for Carbon-Efficient Liquid-Cooled Data Center Clusters." In Proceedings of The AAAI Conference On Artificial Intelligence 2025, Vol. 39, No. 28: 29694-29696.

[4] Chen YT, Luo LL, Guo DK Et Al. Carbon-Aware Energy Cost Optimization of Data Analytics Across Geo-Distributed Data Centers. Journal of Computer Science and Technology 2025, Vol 40: 654−670.

[5] Jayanetti, Amanda, Saman Halgamuge, And Rajkumar Buyya. "Multi-Agent Deep Reinforcement Learning Framework for Renewable Energy-Aware Workflow Scheduling On Distributed Cloud Data Centers." IEEE Transactions On Parallel and Distributed Systems 2024, Vol 35, No. 4: 604-615.

[6] Yan, Dongxiang, Mo-Yuen Chow, And Yue Chen. "Low-Carbon Operation of Data Centers with Joint Workload Sharing and Carbon Allowance Trading." IEEE Transactions On Cloud Computing 2024, Vol 12, No. 2: 750-761.

[7] Burmeister, Sascha Christian, Daniela Guericke, And Guido Schryen. "A Two-Level Approach for Multi-Objective Flexible Job Shop Scheduling and Energy Procurement." Cleaner Energy Systems 2025, Vol 10: 100178.

[8] Wang, Zhixue, Maowei He, Ji Wu, Hanning Chen, And Yang Cao. "An Improved MOEA/D for Low-Carbon Many-Objective Flexible Job Shop Scheduling Problem." Computers & Industrial Engineering 2024, Vol 188: 109926.

[9] Prioli, Joao Paulo Jacomini, Nur Banu Altinpulluk, Jeremy L. Rickli, And Murat Yildirim. "Self-Adaptive Production Performance Monitoring Framework Under Different Operating Regimes." Journal of Manufacturing Systems 2025, Vol 80: 380-394.

[10] Fan, Jiaxin, Chunjiang Zhang, Shichen Tian, Weiming Shen, And Liang Gao. "Flexible Job-Shop Scheduling Problem with Variable Lot-Sizing: An Early Release Policy-Based Matheuristic." Computers & Industrial Engineering 2024, Vol 193: 110290.