

Carbon Footprint Optimization in Cloud Data Centers Using Cloud Computing

Gudise Prathyusha¹, Dr. P.Arulmozhi²

¹Student, CSE, SRM Institute of Science and Technology

²Assistant Professor(Sr.G), CSE, SRM Institute of Science and Technology

Abstract

The rapid growth of cloud-based applications has increased the computational workload handled by modern data centers. Continuous operation of servers, storage systems, networking equipment, and supporting infrastructure results in substantial energy consumption and contributes to carbon emissions. Conventional cloud scheduling generally concentrates on performance, resource utilization, execution time, and cost, while environmental factors are often considered separately. This work presents a carbon-aware cloud resource management framework that combines machine learning-based prediction with adaptive resource allocation. The proposed framework uses cloud workload and energy-consumption datasets obtained from Kaggle. Python is used for data preparation and predictive analysis, while MATLAB is used for optimization, simulation, and performance evaluation. Two algorithms, Carbon-Aware Adaptive Resource Balancing (CAARB) and Dynamic Green Load Prediction Optimizer (DGLPO), are introduced. CAARB adjusts virtual-machine allocation according to workload requirements, resource availability, and estimated carbon impact. DGLPO uses predicted workload and energy requirements to support proactive scheduling and reduce unnecessary resource activation. Carbon emissions are estimated from energy consumption and carbon intensity. The resulting framework is evaluated using energy consumption, carbon emissions, resource utilization, operational cost, execution time, makespan, prediction accuracy, and Quality of Service (QoS). Results reported in the project indicate reduced energy usage and estimated carbon footprint, improved resource utilization, and more adaptive workload scheduling compared with conventional approaches. The framework provides a practical foundation for developing more sustainable cloud data center management systems.

Keywords: Cloud computing, carbon footprint, green computing, cloud data centers, machine learning, resource allocation, workload prediction, CAARB, DGLPO.

1. INTRODUCTION

Cloud computing has become an important part of modern computing infrastructure. Applications such as web services, artificial intelligence, online platforms, enterprise systems, and data analytics depend on cloud resources for storage and computation. The main advantage of cloud computing is its ability to provide resources according to demand. However, this flexibility also creates an environmental challenge because cloud data centers must operate large numbers of servers, storage systems, networking equipment, and cooling facilities continuously.

Energy consumption is closely related to the way cloud resources are allocated. A server may continue consuming power even when its workload is relatively low. If workloads are not distributed efficiently,

some servers may remain underutilized while others become heavily loaded. Activating additional resources unnecessarily increases energy use and can consequently increase carbon emissions. The project report identifies this relationship between workload, energy consumption, resource utilization, and carbon footprint as a major motivation for the proposed work.

Traditional scheduling techniques such as First Come First Serve (FCFS), Round Robin, priority scheduling, and basic VM consolidation generally make decisions using workload or resource availability. These techniques are useful for conventional scheduling but do not fully consider carbon impact or future workload conditions. The project therefore focuses on combining current workload information with prediction-based scheduling.

The proposed framework has two main components. The first is Carbon-Aware Adaptive Resource Balancing (CAARB), which makes resource allocation decisions according to current workload, server utilization, and carbon impact. The second is Dynamic Green Load Prediction Optimizer (DGLPO), which uses historical workload and energy information to estimate future requirements and support proactive scheduling.

The main contributions of this work are:

1. Development of a carbon-aware cloud resource management framework.
2. Introduction of CAARB for adaptive VM resource balancing.
3. Introduction of DGLPO for predictive workload and energy management.
4. Integration of machine learning with carbon-aware scheduling.
5. Evaluation using energy consumption, carbon emissions, resource utilization, operational cost, execution time, and QoS-related measures.

The objective is not simply to reduce energy usage, but to achieve a practical balance between environmental impact and cloud service performance

2. RELATED WORK

Recent research has increasingly focused on reducing energy consumption and carbon emissions in cloud and distributed computing environments. Different approaches have been proposed using multi-objective optimization, machine learning, reinforcement learning, carbon-aware scheduling, renewable-energy integration, and intelligent resource management.

De Carvalho and Sichman [1] proposed a Multi-Armed Bandit Hyper-Heuristic for cloud scheduling that jointly considers makespan, cost, and carbon emissions. Their approach dynamically selects suitable scheduling heuristics and was evaluated using real-world AWS and Teads workloads. The study demonstrates the importance of considering environmental impact along with conventional cloud scheduling objectives.

Zhao *et al.* [2] introduced CFWS, a Deep Reinforcement Learning-based framework for reducing energy cost and carbon footprint in cloud data centers. The framework supports dynamic VM migration across geographically distributed data centers and considers renewable energy availability. The study shows the potential of predictive and adaptive resource management for reducing unnecessary VM migrations and energy consumption.

Beena *et al.* [3] developed a green cloud framework that incorporates carbon-intensity information into task scheduling for heterogeneous cloud servers. The framework uses environmental information to make scheduling decisions and supports cloud platforms and containerized workloads. This work highlights the importance of including regional carbon intensity when allocating cloud resources.

Yu *et al.* [4] investigated carbon reduction in edge computing through joint task offloading and energy sharing. Their work demonstrates that coordinated computation and energy management can reduce the environmental impact of distributed computing systems. However, the approach is primarily focused on edge environments rather than general-purpose cloud data centers.

Alvarado-Alcon *et al.* [5] studied machine-learning-based carbon footprint optimization for LoRa-based IoT networks. Their work uses machine learning to determine suitable gateway positions while considering environmental impact. Although the application is different from cloud data centers, the study demonstrates how predictive models can support carbon-aware optimization.

Qi *et al.* [6] proposed SHIELD-EB, a hybrid evolutionary and XG-Boost framework for sustainable data center management. The framework considers several objectives, including operational cost, carbon emissions, wastewater, and network-related costs. Its multi-objective design demonstrates the value of combining machine learning with optimization for sustainable data center operation.

Nambi and Thanapal [7] presented EMO-TS, a multi-objective optimization approach for energy-efficient task scheduling in cloud data centers. The method focuses on improving task allocation while reducing energy requirements. The study supports the use of optimization techniques for balancing cloud performance and sustainability.

Ahmad *et al.* [8] conducted a systematic study of green cloud computing adoption from the client's perspective. Their work identifies challenges and practices associated with sustainable cloud computing. The study provides a broader understanding of the factors that influence green cloud adoption and the need for energy-conscious cloud services.

Yu, Jiang, and Zou [9] investigated real-time energy management for cloud data centers operating with smart microgrids. Their work connects cloud resource management with energy-system conditions, demonstrating how real-time energy information can support more efficient data center operation.

Shahab *et al.* [10] proposed a bi-objective policy for resilient and sustainable Service Function Chain management in telco-cloud environments. The study considers sustainability together with resilience in distributed cloud infrastructure. Although its primary application is telco-cloud networking, its multi-objective resource management concepts are relevant to sustainable cloud computing.

3. METHODOLOGY

The proposed methodology combines workload prediction, energy estimation, carbon-impact calculation, and adaptive cloud resource management to reduce unnecessary energy consumption in cloud data centers. The overall process begins with collecting cloud workload and energy-related data, followed by preprocessing and feature selection. Machine-learning models are then used to predict workload and energy requirements. These predictions are supplied to the proposed Carbon-Aware Adaptive Resource Balancing (CAARB) and Dynamic Green Load Prediction Optimizer (DGLPO) algorithms. The optimized resource allocation is finally evaluated using energy consumption, carbon emissions, resource utilization, operational cost, execution time, make span, and quality of service (QoS).

A. SYSTEM DESIGN

The proposed system combines machine learning, carbon estimation, and cloud resource optimization to reduce energy consumption and carbon emissions in cloud data centers. Workload and energy datasets are first pre-processed, and relevant features are selected. Machine learning models are then used to predict workload and energy requirements. The predicted information is provided to the resource

optimization algorithms for efficient VM allocation and scheduling.

Data Processing and Prediction

The dataset is cleaned and normalized before training the prediction models. Different machine learning techniques can be used to estimate workload and energy consumption. Model performance is evaluated using MAE, MSE, RMSE, and (R^2).

Carbon Estimation

Carbon emissions are estimated from energy consumption and carbon intensity using:

$$\text{Carbon Emission} = \text{Energy Consumption} \times \text{Carbon Intensity}$$

This estimation is used to support carbon-aware resource allocation.

B. The Architecture Diagram




Fig. 1. Architecture Diagram of Carbon Footprint optimization

Explanation of the Architecture Diagram

The proposed architecture introduces a carbon-aware optimization framework for improving the efficiency of cloud data centers. The process starts with cloud energy and workload datasets obtained from Kaggle, which are processed through machine-learning-based feature engineering to extract meaningful resource and energy-related characteristics. The engineered features are then provided to the carbon footprint optimization framework, consisting of CAARB for carbon-aware VM allocation and DGLPO for predictive workload forecasting. The forecasting component estimates future workload requirements, while the allocation component uses these insights to support suitable VM placement decisions. These outputs are subsequently used by the carbon-aware VM scheduling layer to dynamically manage computing resources across servers. Finally, the system evaluates the resulting performance based on energy consumption, carbon emissions, cost, and resource utilization. Through this workflow, the architecture aims to make cloud resource management more adaptive while reducing the environmental impact of data center operations.

C. Experimental Setup

The implementation uses Python for data preprocessing and machine learning, while MATLAB is used for optimization, simulation, and performance evaluation. The proposed CAARB algorithm performs adaptive resource balancing based on workload and server utilization. DGLPO predicts future workload and energy requirements to support proactive scheduling.

The system is evaluated using energy consumption, carbon emissions, resource utilization, operational cost, execution time, make span, and prediction accuracy. The results are compared with conventional

resource-management approaches to assess the effectiveness of the proposed methodology.

4. RESULTS

The proposed RL-based carbon-aware resource management approach was evaluated against the baseline system. The training results show that the reward gradually increases while carbon emissions, energy cost, and SLA violations decrease over the episodes. The Q-table also shows convergence, indicating that the agent progressively learns an effective resource-allocation policy.

The final carbon comparison shows 2.68 kg CO₂ for the RL agent compared with 42.64 kg CO₂ for the baseline, indicating a substantial reduction in carbon emissions. The RL approach also shows lower energy cost and fewer SLA violations than the baseline.

TABLE I
Comparison of Performance Metrics

Metric	RL Agent	Baseline	Observation
Carbon Emissions (kg CO ₂)	2.68	42.64	Reduced
Energy Cost	Lower	Higher	Reduced
SLA Violations	Lower	Higher	Reduced
Jobs Completed	Improved	Lower	Improved
Resource Management	adaptive	conventional	More efficient

Table I shows the performance comparison. The proposed method achieves the highest accuracy at 95.4% while maintaining competitive processing time.




Fig. 2. Carbon-Aware RL Training Results.

- **Reward:** Gradually increases, showing improved learning.
- **Carbon emissions:** Decrease during training.
- **Energy cost:** Shows a downward trend.
- **SLA violations:** Decrease as the policy improves.
- **Q-table:** Progressively converges.
- **Carbon reduction:** RL agent achieves 2.68 kg CO₂ compared with 42.64 kg CO₂ for the baseline.




Fig. 3. RL Agent vs Baseline Performance Comparison

5. DISCUSSION

The experimental results show that the proposed carbon-aware reinforcement learning approach improves cloud resource management compared with the baseline system. During training, the reward gradually increased, indicating that the agent learned better resource-allocation decisions. At the same time, carbon emissions, energy cost, and SLA violations showed an overall decreasing trend.

The most significant difference was observed in carbon emissions. The RL agent produced 2.68 kg CO₂, whereas the baseline produced 42.64 kg CO₂. This indicates that the learned policy was able to select resource configurations with substantially lower carbon impact. The reduction is supported by the agent's gradual learning and Q-table convergence.

Energy cost also decreased as training progressed, suggesting that the agent learned to avoid unnecessary resource usage. Similarly, the reduction in SLA violations indicates that the optimization was not limited to environmental performance but also considered service requirements.

The gradual reduction in exploration rate allowed the agent to shift from exploring different resource-allocation strategies toward using learned decisions. The convergence of the Q-table further indicates that the allocation policy became more stable over the training episodes.

Overall, the results demonstrate that integrating carbon awareness, workload management, and reinforcement learning can provide a more adaptive approach to cloud resource allocation. However, the presented results are based on the experimental setup shown in the supplied results, so further testing with larger datasets and real cloud environments would be required to establish the method's performance under different workload conditions.

6. CONCLUSION

The proposed carbon-aware reinforcement learning approach provides an adaptive method for managing cloud resources while considering environmental and operational objectives. The experimental results show that the RL agent learned improved resource-allocation decisions over successive training episodes, with reductions in carbon emissions, energy cost, and SLA violations.

The carbon comparison shows 2.68 kg CO₂ for the RL agent compared with 42.64 kg CO₂ for the baseline, demonstrating a substantial difference in the evaluated setup. The decreasing exploration rate and Q-table convergence further indicate that the agent gradually developed a more stable resource-management policy.

Overall, the study demonstrates the potential of combining reinforcement learning, carbon awareness, and cloud resource optimization to support more sustainable data-center operation. Future work can extend the approach using larger real-world datasets, real cloud platforms, dynamic carbon-intensity information, and additional workload conditions to further evaluate its scalability and practical applicability.

REFERENCES

1. V. R. De Carvalho and J. S. Sichman, "Toward Cost-Efficiency and Reduced Carbon Footprint: A Multi-Armed Bandit Hyper-Heuristic for Cloud Scheduling Problems," *IEEE Access*, vol. 13, pp. 204656–204675, 2025, doi: 10.1109/ACCESS.2025.3639273.
2. D. Zhao, J.-T. Zhou, and K. Li, "CFWS: DRL-Based Framework for Energy Cost and Carbon Footprint Optimization in Cloud Data Centers," *IEEE Transactions on Sustainable Computing*, vol. 10, no. 1, pp. 95–107, Jan–Feb. 2025, doi: 10.1109/TSUSC.2024.3391791.
3. Z. Yu, Y. Zhao, T. Deng, L. You, and D. Yuan, "Less Carbon Footprint in Edge Computing by Joint Task Offloading and Energy Sharing," *IEEE Networking Letters*, vol. 5, no. 4, pp. 245–249, Dec. 2023, doi: 10.1109/LNET.2023.3286933.
4. B. M. Beena, P. C. Ranga, T. S. S. Manideep, S. Saragadam, and G. Karthik, "A Green Cloud-Based Framework for Energy-Efficient Task Scheduling Using Carbon Intensity Data for Heterogeneous Cloud Servers," *IEEE Access*, vol. 13, pp. 73916–73938, 2025, doi: 10.1109/ACCESS.2025.3562882.
5. F.-J. Alvarado-Alcon, R. Asorey-Cacheda, A.-J. Garcia-Sanchez, and J. Garcia-Haro, "Minimizing the Carbon Footprint in LoRa-Based IoT Networks: A Machine Learning Perspective on Gateway Positioning," *IEEE Open Journal of the Computer Society*, vol. 6, pp. 531–542, 2025, doi: 10.1109/OJCS.2025.3559331.
6. S. Qi, H. Moore, D. Milojicic, C. Bash, and S. Pasricha, "SHIELD-EB: Sustainable Hybrid Evolutionary-Boosting Framework for Carbon, Wastewater, and Cost-Aware Datacenter Management," *IEEE Access*, vol. 14, pp. 40878–40898, 2026, doi: 10.1109/ACCESS.2026.3674066.
7. S. Nambi and P. Thanapal, "EMO-TS: An Enhanced Multi-Objective Optimization Algorithm for Energy-Efficient Task Scheduling in Cloud Data Centers," *IEEE Access*, vol. 13, pp. 8187–8200, 2025, doi: 10.1109/ACCESS.2025.3527031.
8. A. Ahmad, S. U. Khan, H. U. Khan, G. M. Khan, and M. Ilyas, "Challenges and Practices Identification via a Systematic Literature Review in the Adoption of Green Cloud Computing: Client's Side Approach," *IEEE Access*, vol. 9, pp. 81828–81840, 2021, doi: 10.1109/ACCESS.2021.3085707.
9. L. Yu, T. Jiang, and Y. Zou, "Real-Time Energy Management for Cloud Data Centers in Smart Microgrids," *IEEE Access*, vol. 4, pp. 941–950, 2016, doi: 10.1109/ACCESS.2016.2539369.
10. M. H. Shahab, Y. Sharma, A. Jindal, and A. Al-Dulaimy, "A Bi-Objective Policy for Resilient and Sustainable SFC Management in Telco-Cloud Environments," *IEEE Access*, vol. 13, pp. 215453–215473, 2025, doi: 10.1109/ACCESS.2025.3646975.