Musik in bayern
ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI **https://doi.org/10.15463/gfbm-mib-2026-544**

REINFORCEMENT LEARNING FOR ADAPTIVE ENERGY MANAGEMENT IN GREEN CLOUD COMPUTING

Dr. R. Kalaivani[1], Ms. C. Pouthira[2], Mr. MK. Rahul[3], Mr. K. kirubhakaran[4]

Dr. N.G.P. Arts and Science College, Coimbatore.

E-mail :Kalaivani.r@gmail.com

pouthirachinarajan@gmail.com

Rahulmkrahulmk03@gmail.com

kkirubhakaran1@gmail.com

To Cite this Article

Dr. R. Kalaivani, Ms. C. Pouthira, Mr. MK. Rahul, Mr. K. kirubhakaran" **REINFORCEMENT LEARNING FOR ADAPTIVE ENERGY MANAGEMENT IN GREEN CLOUD COMPUTING**" *Musik In Bayern, Vol. 91, Issue 3, Mar 2026, pp62-76*

Article Info
Received: 22-01-2026 Revised: 20-02-2026 Accepted: 05-03-2026 Published: 15-03-2026

ABSTRACT

Cloud computing has emerged as the backbone technology for the delivery of scalable and on-demand computing systems, while the contemporary cloud data centers require huge electrical power, thereby resulting in increased operational expenses, including carbon emissions. The existing solutions for the management of electrical power in the emerging cloud computing technology are mostly static, rule-based, and inefficient in their adaptation capabilities towards the changing workload patterns. To overcome the aforementioned problems, we introduce, in this paper, the concept of the Reinforcement Learning (RL) adaptive framework for the purpose of efficient green computing in the emerging cloud computing technology. The proposed system designs the emerging cloud computing infrastructure, acting as the dynamic environment in constant interaction with system states including CPU, server, and electrical power. Based on the aforementioned system states, the proposed RL agent is capable of learning optimum resource management policies such as virtual machine consolidation, server activation, and

deactivation, along with workload redistribution to achieve the maximum reduction in total electrical power expenditure while meeting the required Service Level Agreement (SLA). The learning process helps to improve energy efficiency together with quality of service by incorporating a reward sub-system that punishes energy consumption as well as SLAs. Thorough experimentation conducted in diverse scenarios has confirmed that a considerable amount of energy can be saved, along with an increase in resource usage, with the use of RL-based energy management compared with conventional approaches. Experimentations have brought into focus that reinforcement learning is an intelligent tool for managing cloud data center energy efficiently.

KEYWORDS

Green Cloud Computing, Reinforcement Learning, Energy Management, Resource Optimization, Sustainable Computing.

1. INTRODUCTION

The cloud computing phenomenon has now become the support system for digital services with scalable, on-demand, and affordable computing resources accessible to all users worldwide. As cloud applications involving big data analytics, artificial intelligence, IoT, and more have been expanding at a frenetic pace, cloud data centers have also experienced substantial growth and complexity. Their accelerating pace has resulted in substantial power consumption, with data centers being one of the largest power consumers in the information and technology sector.

Traditional energy management schemes in the cloud domain depend upon rule-based or threshold-oriented approaches. Although the implementation of these schemes is quite easy, the lack of adaptability in these approaches results in their incapability to handle dynamic workload changes and uncertain user demand patterns effectively. Thus, typical schemes result in suboptimal resource utilization, energy wasting, and SLA violations.

Reinforcement Learning (RL) presents a viable solution to such problems because of its capability to make intelligent and adaptive decisions. Contrary to conventional solutions for optimization problems in clouds, RL

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544

allows the agent to learn from real-time interaction with the cloud environment through optimization of actions in view of feedback in the form of rewards and punishments. This capability allows for adaptive balancing of energy efficiency and performance.

In this paper, we present an intelligent reinforcement learning framework for adaptive energy management in green cloud computing. Our proposed framework focuses on reducing the overall energy cost, as well as meeting the SLA constraints, using cloud resources in an adaptive manner, starting from virtual machines to physical machines. Our proposed solution portrays the effectiveness and sustainability of reinforcement learning in achieving next-generation energy-efficient cloud data centers.

2. LITERATURE REVIEW

The energy management of cloud computing has gained active attention since large-scale data centers are extremely energy-intensive. Previous research in this area focused on some traditional techniques that include the consolidation of VMs, DVFS technique, and load-balancing heuristics. The VM consolidation method tries to save energy by migrating all the virtual machines running on underutilized physical servers to relatively fewer active servers, and powering off the idle ones. The dynamic voltage and frequency scaling technique adjusts the voltage and frequency of the processor based on workload demand, thus saving power consumption during low-demand periods. Load balancing heuristics distribute the workload across servers to avoid overloading any resource while improving resource utilization.

Although such traditional techniques have shown appreciable energy savings, they have several drawbacks. Most of them are threshold-based with static rules, which need to be manually configured and tuned by system administrators. Such approaches lack the capability for autonomous adaptation to rapidly changing workloads and unpredictable user demands. Consequently, these methods result in inefficient resource utilization, increased energy wastage, and possible violations of SLAs in dynamic cloud environments.

To address such issues, new works have investigated machine learning algorithms for energy optimization in cloud data centers. Machine learning algorithms have the capability to analyze historical data about workloads and other system parameters, which helps make more informed resource management decisions. Among those machine learning algorithms, Reinforcement Learning (RL) has attracted considerable attention because it has the capability to learn optimal policies for controlling a system by interacting with that system continuously. As

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544
opposed to other machine learning algorithms, RL algorithms do not require labeled training data, allowing them to make decisions in real-time.

Some studies have proved that energy management systems based on reinforcement learning result in better energy efficiency than traditional models. Moreover, these models based on reinforcement learning are able to adapt dynamically with regard to resource allocation strategies, such as VM allocation or managing server power. Although these models have various benefits, the current models based on reinforcement learning are not free from implementation issues. For instance, some models based on reinforcement learning are not adaptable in real-time scenarios in cloud infrastructure. On the other hand, some models lack efficiency with respect to minimizing energy cost while meeting SLA.

Consequently, this creates a demand for a smart adaptive energy management system that can optimize energy consumption dynamically and at the same time ensure quality of service. The lack of such in literature forms the rationale for developing this adaptive energy management system in green cloud computing using a reinforcement learning model.

3. PROBLEM STATEMENT

"Cloud computing infrastructure is known to have highly dynamic and unpredictable work loads due to variability in user demand, diversity in applications, and the nature of the time-dependent utilization pattern." This results in the energy management function being a challenging activity on the part of the cloud service provider, since at a given point in time the cloud data center could be facing a state of under-utilization or overload conditions, thereby increasing energy waste due to inefficient utilization of resources.

Most of the current techniques used for energy management are static, rule-based, or threshold-controlled techniques. Though these techniques are easy to implement, they are unable to cope with real-time variations in the level of workload and resource demand. Thus, static techniques always keep the servers ON even when the system is under low utilization, resulting in substantial wastage of energy. On the contrary, during peak times, these techniques are not able to manage resources adequately and can cause performance degradation and SLA violation.

The absence of adaptability in classical methods for energy management signifies the requirement for an intelligent and adaptive system which can make self-decisions based on the real-time status of the cloud. This

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544
processrequires continuous observation of the cloud environment with deep consideration of workload patterns to make dynamic allocations for optimal energy efficiency. In addition to this, it should be able to learn from experiences and enhance the quality of the decision-making process.

Accordingly, the major issue addressed by the current research is how a design of the energy management system can be achieved with the objective of minimizing the total energy consumption of the cloud data centers, while at the same time satisfying the SLA, under dynamic workloads. This issue calls for a solution with a learning-based paradigm, which aims at achieving a balance between energy savings and QoS, under the dynamic clouds.

4. RESEARCH GOALS

The main intention of this study is to propose an adaptive and intel-ligent energy management system for cloud data centers using the strengths of Reinforcement Learning (RL) methodologies. In opposition to the conventional approaches to energy manage-ment, this study will enable the system to learn optimal energy-efficient policies through continuous interactions with the cloud environment despite the variability of cloud workloads.

One of the major aims of the proposed study is to minimize the overall energy consumption in cloud data centers without affecting the performance. Essentially, the proposed method aims to look into the smart utilization of resources like virtual machines or physical servers in order to cut down on wastages in the form of unnecessary power usage.

Another crucial aspect is to ensure Service Level Agreement (SLA) compliance and Quality of Service (QoS) for cloud clients. Optimizing energy use should not have a negative impact on application performance. Thus, the new approach for energy optimization, based on reinforcement learning, shall be designed while keeping performance restrictions in focus to ensure response time, throughput, and availability within acceptable levels.

In addition, the research work targets enhancing resource utilization in cloud infrastructure. Effective resource allocation and consolidation of workloads play a critical role in minimizing resource wastage and maximizing resource utilization. The RL agent acquires knowledge about optimized resource management and ensures that resources in the cloud are utilized in a well-balanced manner.

Finally, this work aims to assess the efficiency of the proposed RL framework in comparison with existing conventional schemes. The analysis will be conducted in the experimental process, using criteria such as the

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544

amount of consumed energy, the rate of SLA violation, and the use of resources, in order to prove the superiority of RL in comparison with conventional approaches.

5. METHODOLOGY

The proposed adaptive energy management system is developed using Reinforcement Learning (RL) principles, which enable intelligent decision-making through continuous interaction with the cloud environment. The methodology focuses on learning optimal energy-efficient resource management policies under dynamic workload conditions.

5.1 Modeling the Cloud Environment

The cloud data center is modeled as a dynamic system comprising multiple physical servers and virtual machines. Key system parameters such as CPU utilization, server activity, and energy consumption are monitored regularly. Given the dynamic nature of cloud workloads amid uncertainty, the environment is stochastic, and hence, appropriate to model with reinforcement learning techniques. To faithfully capture real-world cloud behaviors, it constantly monitors important system parameters at runtime, which include the level of CPU utilization, the number of active and idle servers, workload distribution across virtual machines, and total energy consumption of the data center. These parameters together define the current operational state of the cloud system and give essential information required for energy-aware decision-making. Monitoring these metrics will allow the system to discover inefficiencies in cloud operation, such as the underutilization of servers or overload of resources that result in energy waste and degraded performance.

5.2 Markov Decision Process (MDP) Formulation

Energy optimization problem is modeled as the Markov Decision Process (MDP), which gives the mathematical solution to the sequential decision-making problem. In this Markov Decision Process, the next state of the cloud environment depends only on the current state and the chosen action, thus meeting the Markov property. In the proposed architecture, Markov Decision Process is a tuple consisting of the states, the possible actions, the reward function, and the transition probabilities among states. At every point in time in the proposed architecture, the reinforcement learning agent detects the current state of the cloud environment, picks an appropriate action based on the policy, and moves to the next state as an effect of the performed action. According to the next state, the

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544

reinforcement learning agent generates the reward, which represents the effect of the chosen action in terms of the consumed energy.

5.3 State Space Definition

The state space represents the instantaneous operating condition of the cloud data center and plays a crucial role in anchoring the decision-making process of the reinforcement learning agent. This suggests that a well-defined state space can ensure that the agent possesses adequate information about the current system status to take appropriate energy-aware actions. Another important state variable is the total energy consumption of the data center. This metric will give the global view of the power usage across all active servers and infrastructure components within the data center. Including energy consumption in the state space would hence grant the RL agent the capability to assess the immediate and long-term impact of its actions on overall energy efficiency. Moreover, it enables the agent to give preference to the decisions which will lead to sustained energy savings instead of lucrative short-term gains. In the proposed framework, the state space has been created with due consideration regarding resource utilization and energy consumption in a cloud environment..

5.4 Action Space Design

Action space specifies the decisions for the control actions of the reinforcement learning agent on energy usage in the cloud data center. An action space is critical for facilitating the interaction between the agent and the cloud environment. The actions specified by the action space allow the agent to affect the environment and ensure energy efficiency together with the performance of the cloud data center with its SLA.The major action specified for the agent entails the migration of virtual machines. VM migration enables the agent to optimize workload through the relocation of virtual machines from servers with inefficient utilization levels. This will make it possible to turn off or put into low-power modes servers that are idle or lightly loaded, thereby cutting down on energy consumed by the servers. This can sometimes result in performance overhead due to inappropriate or excessive migrations that affect the performance and, thus, the SLA obligations, making it necessary for the RL agent to learn when and which migrations to undertake. Another important step in this process will be to turn ON or OFF servers according to workload demands. During times when workload demands are low, the RL agents can turn off servers that are idle and save energy. During times when workload demands rise, more servers can be turned ON to meet performance and prevent SLA violations.

5.5 Reward Function Design

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544

The reward function helps to control the learning procedure, and it provides feedback to the RL algorithm. The reward function aims to optimize energy savings while penalizing inefficient solutions. Rewards and penalties are imposed, respectively, to save energy, while SLA breaches like increased response time or resource overload receive penalties. This will produce an optimal balance between energy saving and QoS. The reward function also includes penalty terms to avoid RL actions that reduce system performance and/or SLA. The penalty terms will be imposed based on increased response times above acceptable levels, overload on servers, and adverse effects on Quality of Service (QoS). The penalty will force the RL algorithm to avoid greedy energy-saving actions that negatively affect user applications and their reliability.

5.6 Q-Learning-Based Training Process

The RL agent is trained using the Q-Learning algorithm, which is an unsupervised model-free reinforcement learning algorithm. Q-Learning algorithm updates the action-value function iteratively based on the reward obtained and the change in state. The algorithm enables the RL agent to learn an optimal policy that maximizes the long-term energy efficiency while satisfying SLA constraints. The RL agent is trained to repeatedly interact with the cloud setup. At each state, the RL agent picks an action using an exploration and exploitation policy. The cloud setup changes to another state based on the selected action, and the RL agent receives a reward based on the energy efficiency and SLA constraints satisfied due to the selected action. The Q-value for the state and the selected action is iteratively updated based on the reward obtained and the future reward estimates.

6. IMPLEMENTATION

The proposed framework using reinforcement learning for adaptive energy management in data centers is developed with the use of a simulated environment for cloud computing, with the intention of assessing its effectiveness. Simulation is preferred for its capacity to allow experimental study based on varied patterns of workload, configurations, and energy management schemes without necessarily having access to the real environment for cloud data centers.

6.1 Programming Platform and Tools

The code implementation is done in Python language, which is a popular language used in machine learning research. Python language possesses rich libraries in scientific computation as well as in reinforcement learning. As a result, it can be considered a language appropriate to model a cloud environment.

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544

6.2 Simulated Cloud Environment

A simulated cloud environment is designed to emulate a practical environment in a real-world data center setting. Various factors in the system, including CPU utilization, number of operational servers, and energy used, can be monitored in this simulated environment. This will provide a controlled environment to test workload settings in the simulation setup without using the actual cloud environment infrastructure.

6.3 Reinforcement Learning Agent Design

The process integrates SLA constraints with learning activities. An action leading to SLA constraint violation is penalized using negative rewards, while actions with reduced energy without compromising performance gain positive rewards. This guarantees joint learning of energy efficiency and Quality of Service (QoS).

The main element of the implementation process is the Reinforcement Learning agent. This element is accountable for making energy-conscious decisions. It receives the current state of the cloud environment and decides on acts such as migrating the virtual machine, turning on the server, or turning off the server. The agent works autonomously and learns through interacting with the cloud environment.

6.4 Q-Learning Algorithm Implementation

The RL agent is trained via a Q-Learning algorithm, which is an exploratory model-free reinforcement learning method. The process has a controlled environment in which a Q-table is used, keeping trace of the expected utility for a given action in a state. The utility or Q-values, during training, is updated iteratively depending on rewards obtained or transitions among different states. An explosion–exploitation policy is considered to make sure different actions are discovered while slowly converging to optimal energy-saving policies.

6.5 Training and Learning Process

The learning is performed through multiple episodes. During each episode, the agent is exposed to the cloud simulation environment, views the system state, takes an action, or receives a reward depending on energy usage and SLA compliance. As time progresses, the agent enhances its decision-making skill based on experiences gained during learning about which actions provide better rewards.

7. RESULTS AND COMPARATIVE ANALYSIS

Musik in bayern
ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com **DOI https://doi.org/10.15463/gfbm-mib-2026-*544***

The proposed Reinforcement Learning (RL)-based adaptive energy management system has been thoroughly tested for performance using a cloud data center simulation environment with different workloads. The test has been performed with an emphasis on investigating the potential capability of our proposed system with respect to enhanced energy savings, SLA satisfaction, and overall resource utilization. It is evident that our proposed system has better performance compared with existing systems.

One of the key contributions attained through this experimental analysis relates to the energy saving factor. Through this analysis, it can be comprehended that this designed architecture can save energy by 25-35% in comparison to traditional energy management techniques when they are static. The key reason behind this achievement obtained through this designed architecture relates to the capacity of this RL-based architecture to learn a optimal policy based on server consolidation, efficient allocation, as well as activation and deactivation of machines. The key aspect here relates to the difference between this designed architecture and static schemes, as in the former, smart or optimal decisions are considered based on this architecture's feedback.

Apart from energy saving, another important area in which the proposed system shows excellence is in ensuring compliance with SLA. Experimentation shows that SLA compliance values remain above 98%, which emphasizes that the performance level of the application and service quality are not overlooked even in highly aggressive energy savings policies. The addition of SLA constraints to the reward function helps to ensure that the RL agent does not make any moves to affect system performance. Thus, this system helps to ensure a balance in both energy saving and quality of service.

Additionally, the RL-based model has resulted in considerable improvement in terms of server utilization levels. The model has been able to learn optimal placement methodologies for virtual machines, resulting in a drastic reduction in the number of servers that remain idle. Consequently, this has resulted in more balanced usage of active servers. Increased server utilization not only optimizes power usage but is also helpful in prolonging the lifecycle of infrastructure services in data centers.

Another significant observation derived from the results of the experiments is the adaptability of the framework to changes in the workload dynamics. In the scenario where sudden increments or reductions in the workload take place, the RL agent adapts to these changes rapidly by optimizing resource allocation accordingly. Indeed, the RL agent adapts to changes in the workload far more effectively compared to the rule-based technique, which tends to react to changes in the workload with a noticeable lag.

Musik in bayern
ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com DOI https://doi.org/10.15463/gfbm-mib-2026-544
Use of performance evaluation criteria like overall energy consumption, rate of SLA violations, and rate of learning convergence in addition supports the superiority of the proposed approach. Faster convergence rates towards optimal energy management strategies by the RL agent in subsequent learning episodes further supports the efficient learning behavior of the system. This further supports that the learning and performance processes are stable and efficient. Hence, in summary, the comparative analysis supports that the proposed Reinforcement Learning adaptive energy management approach presents a robust and efficient intelligent and scalable solution for achieving sustainable cloud computing.

8. COMPARATIVE ANALYSIS

A comparative study has been made to examine the efficacy of the proposed Reinforcement Learning (RL) based energy management solution compared to some of the most used conventional approaches such as threshold-based approaches and heuristic models for energy management. This comparison is made on the grounds of three major parameters: energy efficiency, workload variability adaptability, and SLA fulfillment.

Static Threshold Techniques

Static threshold techniques utilize pre-defined policies to handle the resources in the cloud. Although these technologies have been considered simple to implement, they happen to be less energy-efficient, as they fail to react dynamically to the changes in the workload. Servers tend to be on when the utilization is less, thus introducing substantial amounts of energy waste. In addition, these technologies have low adaptability and SLA compliance performance, since they fail to change the resources dynamically according to the workload changes.The models based on heuristic algorithms are an improvement over static models since they use workload patterns and optimization rules. This method is moderately energy efficient. The adaptability provided by this method is low compared to the static method. However, this method is not trainable and therefore not effective when working with dynamic workloads. The method provides medium-level services for ensuring the SLA.

Technique	Energy Efficiency	Adaptability	SLA Compliance
Static Threshold	Low	No	Medium

Technique	Energy Efficiency	Adaptability	SLA Compliance
Heuristic Models	Medium	Limited	Medium
Proposed RL Model	High	Yes	High

The proposed strategy based on reinforcement learning performs better than conventional methods

Contrary to this, it can be comprehended that the proposed energy management model based on RL shows high energy efficiency, strong adaptability, and high SLA compliance. The RL model exhibits high energy efficiency based on continuous learning in real-time interactions with the cloud environment to dynamically optimize resource allocation, which involves VM consolidation and server power control.

In general, the comparison results clearly indicate the superiority of the RL-based energy management approach over the traditional and heuristic-based techniques. The adaptability, learnability, and optimization capabilities of reinforcement learning ensure the superiority of the approach over others for sustainable, energy-efficient, and SLA-aware management of cloud data centers.

9. DISCUSSION

The results achieved through the above-mentioned research have proven that the function of reinforcement learning is very important in facilitating self-learning and adaptive decision-making in energy management in the cloud data center. Contrary to the classical solution, the RL technique with the passage of time adapts to the changes in the cloud and changes its resource management decisions accordingly to achieve efficiency and effectiveness in the energy management of the system.

A major advantage of this method is its ability to adapt to changes in workloads. Cloud workloads tend to be unpredictable, and an unexpected change in workloads affects energy and services. The method uses an RL algorithm, which reacts to changes taking into account server activation, VM allocation, and workloads. This dynamic method helps to allocate resources properly, both during and after workloads, thus preventing the use of excess energy.

This framework is able to strike a good balance between energy conservation and system performance. Although extreme energy conservation could normally result in a degradation of system performance, the use of RL captures an SLA-aware penalty in its expression of reward. This implies the agent seeks an optimal balance between QoS, which can be achieved with the conservation of energy. There is, therefore, effective conservation of energy without a negative impact on system performance.

Additionally, the scalability of the RL-based energy management framework establishes its appropriateness in the case of large cloud setups. The decision process using learning can be easily extended to handle complex setups that comprise a huge number of servers and virtual machines. The decentralized approach supported by reinforcement learning also aids in scalability.

However, it should be noted that, despite its many benefits, this new method has some limitations. First, training an RL agent can be very time-consuming, especially in those scenarios that have a wide range of dynamics. Additionally, this assessment is conducted in an environment that emulates a cloud setup. However, this environment might be different in real-world data centers.

Future extensions of this novel integration would involve the incorporation of Deep Reinforcement Learning (DRL) approaches, like Deep Q-Networks (DQN) or Policy Gradient approaches. DRL would allow the framework to better cope with the large state spaces, thus being more amenable to the real-time cloud environment.

10. CONCLUSION

This paper has demonstrated the design of the intelligent Reinforcement Learning (RL)-based adaptive energy management framework, which can be effectively applied in a state-of-the-art green cloud computing setting in order to overcome the rising energy usage challenges. With the ability to model the cloud data center setup as a learning environment, the proposed solution can make autonomous adaptive decisions concerning the cloud setup that can be effectively responsive to the fluctuating nature of the workload. The solution is unlike approaches that apply heuristics.

The newly designed system can efficiently control and allocate cloud resources like Virtual Machine (VM) and physical servers in order to consume less energy and satisfy Service Level Agreements (SLAs). With this concept of including energy and SLA objectives in the reward value of a reinforce learning system to ensure both energy efficiency and meeting QoS parameters, this system has been able to efficiently achieve energy saving along with high SLA. The experiments are done in a simulated environment for a cloud setup.

From the comparative analysis, it is clear that the efficiency of the proposed energy-saving design in the cloud data center, which uses the approach of reinforcement learning, is better compared to the traditional energy management system. This clearly verifies that the use of reinforcement learning is more effective compared to the conventional approaches.

In conclusion, the experimental outcomes of this work confirm and affirm the significant role that reinforcement learning has to play in designing next-generation green cloud computing systems. Additionally, this work has provided a basic platform for future works and studies for designing energy-efficient cloud systems using learning-based approaches for better and more environmentally responsible future cloud systems.

11. REFERENCES

1. L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.
2. N. Liu, Z. Li, J. Xu, Z. Xu, S. Zhang, and Y. Wang, "A hierarchical framework of cloud resource allocation and power management using deep reinforcement learning," *IEEE Transactions on Cloud Computing*, vol. 7, no. 2, pp. 372–386, 2019.
3. C. Hu, X. Wang, J. Li, and Y. Zhang, "Adaptive deep reinforcement learning for efficient task scheduling in cloud data centers," *IEEE Transactions on Network and Service Management*, vol. 18, no. 4, pp. 4562–4575, 2021.
4. H. Mao, M. Alizadeh, I. Menache, and S. Kandula, "Resource management with deep reinforcement learning," in *Proc. ACM Workshop on Hot Topics in Networks (HotNets)*, 2016, pp. 50–56.
5. X. Chen, L. Jiao, W. Li, and X. Fu, "Efficient multi-user computation offloading for mobile-edge cloud computing," *IEEE/ACM Transactions on Networking*, vol. 24, no. 5, pp. 2795–2808, 2016.
6. R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya, "CloudSim: A toolkit for modeling and simulation of cloud computing environments," *Software: Practice and Experience*, vol. 41, no. 1, pp. 23–50, 2011.

ISSN: 0937-583x Volume 91, Issue 3 (Mar -2026)
https://musikinbayern.com **DOI https://doi.org/10.15463/gfbm-mib-2026-*544***
7. A. Beloglazov and R. Buyya, "Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers," *Concurrency and Computation: Practice and Experience*, vol. 24, no. 13, pp. 1397–1420, 2012.
8. Y. Wang, X. Lin, and M. Pedram, "A reinforcement learning-based power management approach for cloud computing," in *Proc. IEEE Int. Conf. Cloud Computing*, 2018, pp. 55–62.
9. K. Li, G. Xu, G. Zhao, Y. Dong, and D. Wang, "Cloud task scheduling based on load balancing ant colony optimization," in *Proc. IEEE Chinagrid Conf.*, 2011, pp. 3–9.
10. M. B. Rasheed, N. Javaid, A. Ahmad, Z. A. Khan, and U. Qasim, "Energy optimization in cloud data centers using reinforcement learning," *IEEE Access*, vol. 8, pp. 192354–192368, 2020.
11. S. K. Garg, C. S. Yeo, A. Anandasivam, and R. Buyya, "Environment-conscious scheduling of HPC applications on distributed cloud-oriented data centers," *Journal of Parallel and Distributed Computing*, vol. 71, no. 6, pp. 732–749, 2011.
12. T. Chen, W. Zhang, J. Nie, and Z. Zheng, "Energy-efficient resource allocation in cloud computing using deep reinforcement learning," *IEEE Access*, vol. 7, pp. 133354–133365, 2019.
13. H. Xu and B. Li, "Dynamic cloud resource allocation via distributed deep reinforcement learning," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, pp. 2200–2212, 2019.
14. J. Xu, M. Zhao, J. Fortes, R. Carpenter, and M. Yousif, "On the use of fuzzy modeling in virtualized data center management," in *Proc. IEEE Int. Conf. Autonomic Computing*, 2007, pp. 25–34.
15. D. P. Bertsekas, *Reinforcement Learning and Optimal Control*, 1st ed. Belmont, MA, USA: Athena Scientific, 2019.