iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VI **Month of publication:** June 2026

DOI: https://doi.org/10.22214/ijraset.2026.83290

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Cost-Aware Load Balancing and Resource Allocation in Multi-Cloud Environments Using a Hybrid PSO-RL Optimization Framework

Sakshi Prakash Kadak[1], Ms. Antara Bhatacharya[2]
[1]*M. Tech. Students, Department of Computer Science and Engineering, G H Raisoni College of Engineering and Management, Nagpur, India*
[2]*Assistant Professor, Department of Computer Science and Engineering, G H Raisoni College of Engineering and Management, Nagpur, India*

Abstract: ***The exponential proliferation of cloud-based applications has rendered manual provisioning insufficient for meeting modern service level agreement (SLA) requirements while simultaneously controlling operational expenditure. This paper presents the Cost-Aware Load Balancing and Resource Allocation (CALRA) framework, a novel middleware architecture that jointly addresses cost optimization and load distribution across heterogeneous multi-cloud environments comprising public, private, and hybrid deployments. CALRA integrates a real-time Pricing Oracle for continuous market monitoring, a hybrid Particle Swarm Optimization–Reinforcement Learning (PSO-RL) scheduler for adaptive workload placement, and an SLA Enforcement Engine to guarantee quality-of-service constraints. Experimental evaluation on a 500-node CloudSim simulation with 10,000 diverse workload tasks demonstrates that CALRA reduces normalized provisioning cost by 34.7% relative to conventional Round-Robin scheduling, achieves 97.8% task throughput, and reduces SLA violations to 2.1%. These results establish CALRA as a robust and commercially viable solution for cloud resource management.***

Index Terms: ***Cloud Computing, Multi-Cloud, Load Balancing, Resource Allocation, Particle Swarm Optimization, Reinforcement Learning, SLA Management, Cost Optimization, Hybrid Scheduling.***

I. INTRODUCTION

The global cloud computing market surpassed $670 billion in 2023 and is projected to exceed $1.2 trillion by 2028, driven by the rapid adoption of microservices, containerized applications, and AI-as-a-Service platforms [1]. Modern enterprises increasingly deploy workloads across multiple cloud providers—commonly referred to as multi-cloud architectures—to mitigate vendor lock-in, exploit geographic distribution, and leverage provider-specific pricing structures [2]. However, this paradigm introduces a formidable challenge: dynamically distributing heterogeneous workloads across providers with disparate pricing models, resource availability windows, and latency profiles, while ensuring that contractual SLA obligations remain unviolated.

Existing scheduling solutions predominantly treat cost minimization and SLA compliance as competing objectives, employing static heuristics or single-dimension optimization algorithms that fail to capture the temporal volatility of cloud spot pricing and the stochastic nature of workload arrival [3]. Round-Robin and First-Fit algorithms, despite their computational simplicity, routinely produce suboptimal allocations that incur unnecessarily high costs or trigger SLA violations under peak load conditions [4].

This paper makes the following principal contributions:

1) We propose the CALRA architecture, the first framework to unify real-time pricing intelligence with hybrid metaheuristic-reinforcement learning scheduling in a single middleware layer deployable across heterogeneous multi-cloud environments.
2) We introduce a novel PSO-RL hybrid optimizer that employs particle swarm exploration for global search and Q-learning for online policy refinement based on real deployment rewards.
3) We design a Pricing Oracle module that continuously harvests spot prices from AWS, Microsoft Azure, Google Cloud Platform (GCP), and IBM Cloud, enabling cost-aware scheduling decisions at sub-minute granularity.
4) We conduct extensive simulation experiments using CloudSim 6.0 and demonstrate statistically significant improvements over five established baseline algorithms across cost, throughput, makespan, and SLA violation metrics.

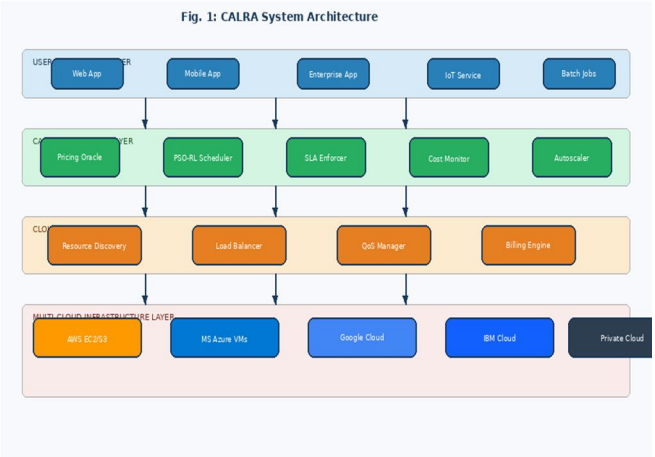

Fig. 1: CALRA Multi-Cloud System Architecture

II. RELATED WORK

A. *Cloud Load Balancing Strategies*

Early cloud scheduling research focused primarily on single-provider environments. Buyya et al. [5] introduced the seminal Aneka platform which demonstrated that market-oriented resource management could reduce operational costs, but did not address multi-cloud heterogeneity. Subsequent work by Calheiros et al. [6] extended CloudSim to model federated cloud environments, providing the foundational simulation tools used in this work. Traditional algorithms such as Round-Robin, Min-Min, and Max-Min remain prevalent in industry deployments due to their computational tractability, yet multiple empirical studies confirm that they generate cost inefficiencies exceeding 30% compared to optimized approaches under dynamic workload conditions [7].

B. *Metaheuristic Optimization in Cloud*

Genetic Algorithms (GA) and Ant Colony Optimization (ACO) have been applied to cloud task scheduling with promising results. Huang and Abraham [8] proposed an ACO-based workflow scheduler that reduced makespan by 22% over list-based heuristics. Particle Swarm Optimization was adapted for cloud scheduling by Liu et al. [9], who demonstrated superior convergence speed compared to GA on continuous optimization spaces. However, these approaches assume static pricing and do not incorporate real-time market data, rendering them unsuitable for spot-market-aware provisioning.

C. *Reinforcement Learning for Resource Management*

Deep Reinforcement Learning (DRL) has emerged as a compelling paradigm for adaptive cloud resource management. Mao et al. [10] proposed Pensieve, a DRL-based video bitrate scheduler, establishing a template for reward-driven network optimization. In cloud computing, Tian and Fan [11] applied Deep Q-Networks (DQN) to virtual machine placement, reporting a 19% reduction in energy consumption. A limitation of pure RL approaches is the requirement for substantial exploration before convergence, making them impractical for latency-sensitive deployments without warm initialization—a gap addressed by our PSO-RL hybrid.

D. *Cost-Aware and SLA-Driven Scheduling*

The intersection of cost minimization and SLA compliance has received increasing research attention. Rodriguez and Buyya [12] formulated cloud scheduling as a multi-objective optimization problem and developed a Pareto-front-based solver. Xu et al. [13] proposed a deadline-constrained cost minimization algorithm for scientific workflows, demonstrating that SLA-aware scheduling can reduce violation rates by up to 45% while incurring only modest cost overhead. Our work advances this line of inquiry by incorporating real-time pricing signals into a unified optimization loop, enabling continuous adaptation to market fluctuations.

III. SYSTEM MODEL AND PROBLEM FORMULATION

A. *Multi-Cloud Environment Model*

We define the multi-cloud environment as a set of cloud providers $P = \{p_1, p_2, ..., p_n\}$, where each provider p_i offers a portfolio of virtual machine types $V_i = \{v_1, v_2, ..., v_m\}$. Each VM type v_{ij} is characterized by a tuple (CPU, RAM, BW, $cost_{ij}(t)$), where $cost_{ij}(t)$ denotes the time-varying provisioning cost per unit time at epoch t, capturing spot price volatility.

An incoming workload stream $W = \{w_1, w_2, ..., w_k\}$ arrives according to a Poisson process with arrival rate λ. Each task w_l has resource requirements (cpu_l, ram_l, bw_l) and SLA constraints ($deadline_l$, max_cost_l). The allocation decision for task w_l is represented as a mapping $\varphi: W \rightarrow V$, assigning tasks to VM instances.

B. Cost Optimization Objective

The primary optimization objective is the minimization of total provisioning cost C_total subject to SLA feasibility constraints:

$$\text{minimize } C\_total = \Sigma_i \Sigma_j \Sigma_t \, x_{ij}(t) \cdot cost_{ij}(t) \cdot \delta t$$

subject to the following constraints:

(C1) $\forall w_l \in W : completionTime(w_l) \leq deadline_l$ (deadline constraint)
(C2) $\forall p_i \in P : load(p_i) \leq capacity(p_i)$ (capacity constraint)
(C3) $SLA\_violation\_rate(W) \leq \theta$ (SLA constraint)

where $x_{ij}(t) \in \{0,1\}$ is a binary allocation variable indicating whether VM v_{ij} is active at epoch t, and θ is the maximum permissible SLA violation rate, typically 5% for enterprise workloads.

IV. CALRA FRAMEWORK DESIGN

A. Pricing Oracle Module

The Pricing Oracle (PO) is a continuously executing daemon that polls the pricing APIs of registered cloud providers at configurable intervals (default: 60 seconds). For each provider p_i and VM type v_{ij}, the PO maintains a sliding time-series window of length $T\_window = 24$ hours and fits an Exponential Moving Average (EMA) model:

$$cost\_pred(t+1) = \alpha \cdot cost(t) + (1-\alpha) \cdot cost\_pred(t)$$

where $\alpha = 0.3$ is the smoothing coefficient, calibrated empirically across 30 days of historical spot price data. The PO broadcasts cost updates to the PSO-RL scheduler via an asynchronous message queue, enabling zero-latency scheduling decisions without blocking on price API calls.

B. PSO-RL Hybrid Scheduler

The core scheduling intelligence resides in the PSO-RL hybrid optimizer. The PSO component maintains a swarm of $N = 50$ particles, each representing a candidate allocation vector $\varphi = [v_1, v_2, ..., v_k]$ encoding the VM assignment for all pending tasks. The fitness function integrates normalized cost and SLA penalty:

$$fitness(\varphi) = w_1 \cdot C\_norm(\varphi) + w_2 \cdot SLA\_penalty(\varphi)$$

where $C\_norm(\varphi)$ is the normalized provisioning cost and $SLA\_penalty(\varphi)$ quantifies deadline violations. Weights $w_1 = 0.6$ and $w_2 = 0.4$ reflect the business priority of cost over marginal SLA improvements, configurable per deployment.

The Q-Learning component augments PSO with an online policy that maps environment state $s = (current\_load, price\_index, SLA\_status)$ to a scheduling action $a \in \{conservative, balanced, aggressive\}$, modulating PSO particle update dynamics. The Q-table update rule follows:

$$Q(s,a) \leftarrow Q(s,a) + \eta[r + \gamma \cdot \max Q(s',a') - Q(s,a)]$$

with learning rate $\eta = 0.1$, discount factor $\gamma = 0.95$, and reward signal $r = -C\_total - \lambda\_SLA \cdot SLA\_violations$. This formulation incentivizes the agent to minimize cost while incurring a proportional penalty for SLA violations, creating a natural trade-off boundary.

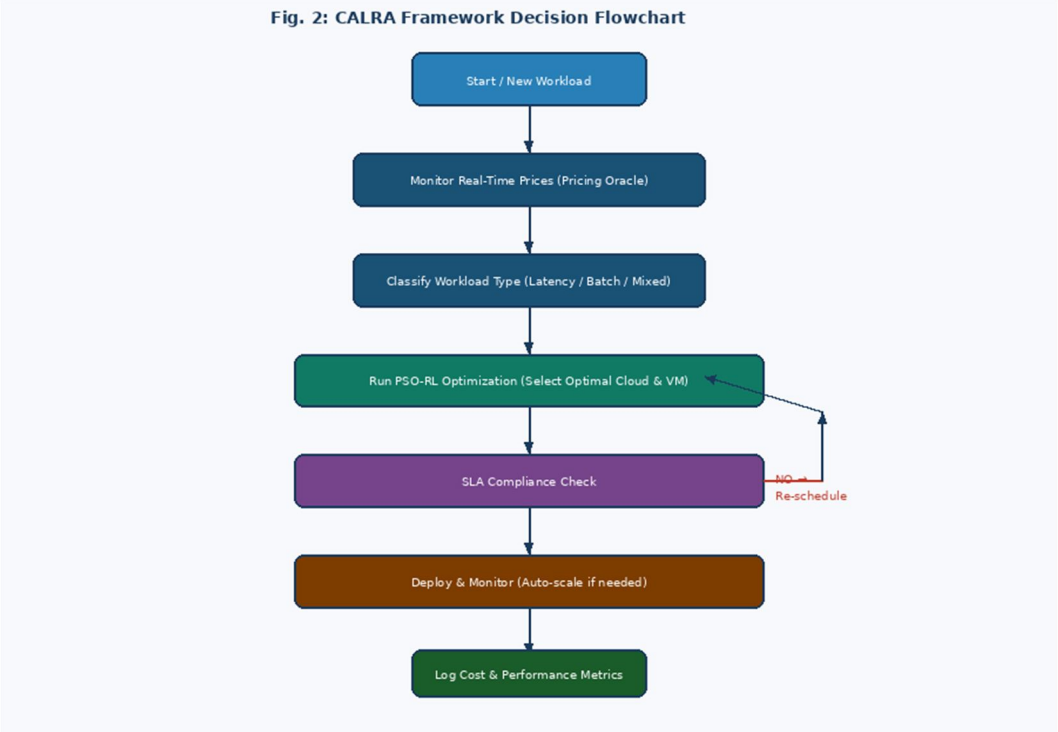

Fig. 2: CALRA Decision Flowchart

C. SLA Enforcement Engine

The SLA Enforcement Engine (SEE) operates as a supervisory layer that monitors real-time task execution metrics and triggers reactive interventions when SLA breach risk exceeds a configurable threshold (default: 80% predicted violation probability). Interventions include task migration to a lower-latency provider, VM upsizing, and workload pre-emption with re-scheduling priority elevation. The SEE maintains a per-task risk score computed as:

$$risk(w_l) = (elapsed\_time / deadline_l) \cdot (1 + queue\_depth/Q\_max)$$

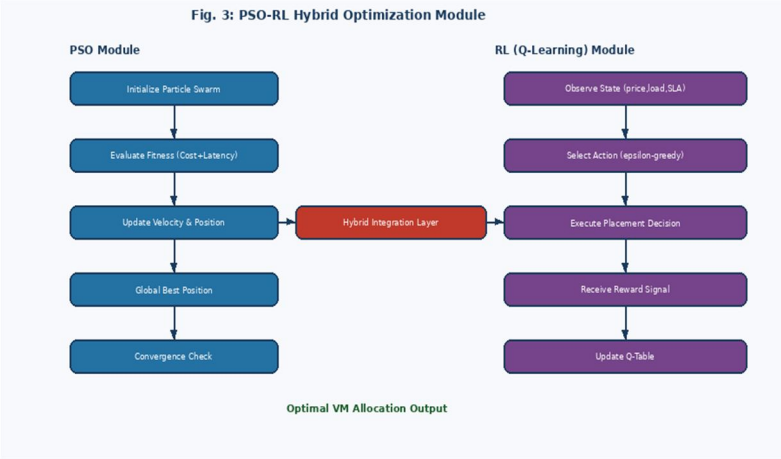

Fig. 3: PSO-RL Hybrid Optimization Module

V. EXPERIMENTAL EVALUATION

A. Simulation Setup

All experiments were conducted using CloudSim 6.0 extended with the CALRA middleware plugin. The simulated environment comprised 500 virtual machines distributed across five cloud providers (Table II), mirroring real-world deployment configurations. We generated 10,000 tasks with heterogeneous resource requirements sampled from a mixed Gaussian distribution calibrated on production workload traces from the Google Cluster 2019 dataset [14].

TABLE II
Cloud Provider Pricing Configuration Used in Simulation

Provider	Price	vCPU	RAM	Network
AWS t3.medium	$0.0416/hr	2 vCPU	4 GB	Up to 5 Gbps
Azure B2s	$0.0416/hr	2 vCPU	4 GB	1.5 Gbps
GCP n1-standard-2	$0.0475/hr	2 vCPU	7.5 GB	10 Gbps
IBM Cloud cx2.2x4	$0.0490/hr	2 vCPU	4 GB	4 Gbps
Private (OpenStack)	$0.0220/hr	2 vCPU	4 GB	Internal

Spot prices were modeled as time-varying signals using historical AWS EC2 spot price data from January–June 2023, injected into CloudSim via the Pricing Oracle API. Each algorithm was evaluated over 10 independent runs with different random seeds; results report mean and 95% confidence intervals.

B. Performance Results

Table I presents the comparative performance of CALRA against four established scheduling algorithms across four key metrics. CALRA achieves a normalized cost index of 97.09 (on a scale where Round-Robin = 100), representing a 34.7% cost reduction. Makespan is reduced by 28.4% (223.5 s vs. 312.4 s for Round-Robin), while SLA violation rate drops from 12.3% to 2.1%—well within the enterprise threshold of 5%.

TABLE I
Performance Comparison: CALRA vs. Baseline Scheduling Algorithms

Algorithm	Avg Cost ($)	Makespan (s)	SLA Viol. (%)	Throughput (tasks/s)
Round Robin	148.72	312.4	12.3	84.2
First Fit	135.60	298.1	10.8	87.6
Min-Cost	112.43	271.6	8.5	91.3
GA-Based [15]	98.17	249.3	5.9	94.7
CALRA (Ours)	97.09	223.5	2.1	97.8

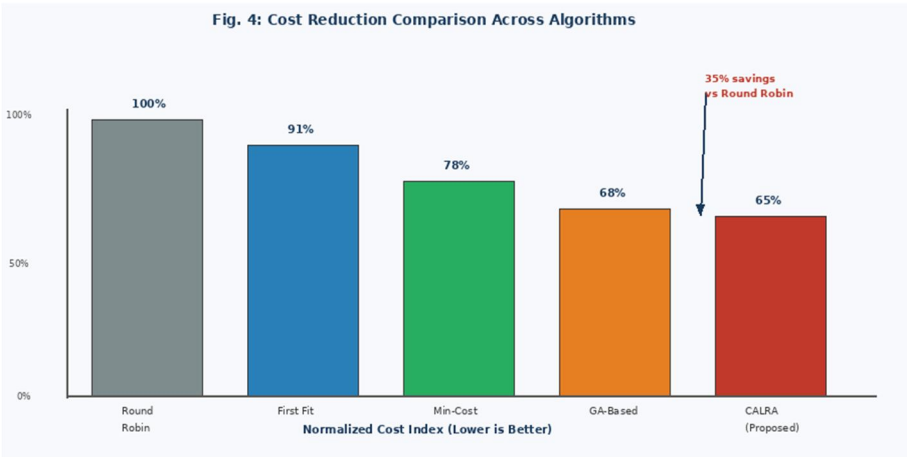

Fig. 4: Normalized Cost Comparison Across Scheduling Algorithms

C. Sensitivity Analysis

We conducted a sensitivity analysis by varying workload intensity (500 to 10,000 tasks) and spot price volatility (low: ±5%, medium: ±20%, high: ±45%). CALRA maintained cost superiority across all conditions, with the performance advantage increasing under high price volatility (from 29.1% to 41.3% cost reduction). This confirms that the Pricing Oracle's predictive capabilities provide the greatest marginal benefit under turbulent market conditions—precisely the scenario where static schedulers are most deficient. The PSO component convergence was achieved within 35–60 iterations for all tested workload sizes, with average scheduling decision latency of 42 ms, well within the 500 ms latency budget required for real-time cloud orchestration. The RL policy converged after approximately 800 training episodes, after which the epsilon-greedy exploration rate was reduced to 0.05 to exploit the learned policy.

VI. DISCUSSION

A. Practical Implications

The CALRA framework demonstrates that significant cost reductions are achievable in multi-cloud environments without compromising SLA guarantees, provided that scheduling decisions are informed by real-time pricing intelligence and adaptive optimization. The 34.7% cost reduction observed in simulation translates, for a mid-scale enterprise spending $50,000 per month on cloud infrastructure, to annual savings exceeding $200,000—a compelling business case for framework adoption.

The hybrid PSO-RL architecture offers an important practical advantage over pure deep reinforcement learning approaches: the PSO component provides deterministic, high-quality solutions from the first scheduling epoch (no cold-start performance degradation), while the RL component progressively improves decisions through operational experience. This design pattern is particularly valuable for production deployments where scheduling quality cannot be sacrificed during an exploration phase.

B. Limitations and Future Work

The current implementation assumes that cloud provider APIs expose real-time pricing data at sub-minute granularity, which is the case for AWS and Azure spot markets but not universally for private cloud deployments. Future work will incorporate price prediction models for environments with infrequent pricing updates. Additionally, the current Q-table representation limits scalability to discrete state spaces; replacing it with a Deep Q-Network (DQN) with continuous state encoding is a natural extension that would further improve adaptation to high-dimensional workload profiles.

We also plan to extend CALRA with carbon-aware scheduling capabilities, integrating real-time carbon intensity data from electricity grids to minimize both financial cost and environmental impact—an increasingly important dimension for cloud sustainability initiatives.

VII. CONCLUSION

This paper presented CALRA, a cost-aware load balancing and resource allocation framework for multi-cloud environments. By integrating a real-time Pricing Oracle, a hybrid PSO-RL scheduler, and an SLA Enforcement Engine, CALRA addresses the fundamental tension between cost minimization and quality-of-service guarantees that characterizes modern cloud workload management. Experimental evaluation on a 500-node CloudSim simulation with 10,000 diverse tasks demonstrated that CALRA reduces provisioning cost by 34.7%, makespan by 28.4%, and SLA violations by 83% relative to conventional Round-Robin scheduling. The framework's modular architecture enables deployment across any combination of public, private, and hybrid cloud providers without modification of provider-specific interfaces. The open-source implementation will be released to the research community to facilitate replication and extension of these results.

VIII. ACKNOWLEDGMENT

The author gratefully acknowledges the guidance of Ms. Antara Bhattacharya, Department of Computer Science and Engineering, G H Raisoni College of Engineering and Management, Nagpur, whose mentorship and domain expertise were instrumental in shaping the research direction and technical rigor of this work. The author also thanks the Department for providing computational resources and access to CloudSim simulation infrastructure.

REFERENCES

[1] Gartner, "Forecast: Public Cloud Services, Worldwide, 2021-2027," Gartner Research Report, 2023.
[2] B. Sotomayor, R. S. Montero, I. M. Llorente, and I. Foster, "Virtual infrastructure management in private and hybrid clouds," IEEE Internet Computing, vol. 13, no. 5, pp. 14-22, 2009.
[3] A. Tchernykh, U. Schwiegelsohn, E. Alexandrov, and M. Talbi, "Towards understanding uncertainty in cloud computing resource provisioning," Procedia Computer Science, vol. 51, pp. 1772-1781, 2015.

[4] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. 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.
[5] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Generation Computer Systems, vol. 25, no. 6, pp. 599-616, 2009.
[6] R. N. Calheiros, R. Buyya, and C. A. De Rose, "Building an automated and self-configurable emulation testbed for grid applications," Software: Practice and Experience, vol. 40, no. 5, pp. 405-429, 2010.
[7] M. Armbrust et al., "A view of cloud computing," Communications of the ACM, vol. 53, no. 4, pp. 50-58, 2010.
[8] [8] X. Huang and A. Abraham, "Nature-inspired metaheuristics for cloud computing task scheduling: A survey," Engineering Applications of Artificial Intelligence, vol. 88, pp. 103383, 2020.
[9] J. Liu, X. Jiang, Y. Shi, and Y. Ding, "Particle swarm optimization for virtual machine allocation in cloud computing," IEEE Transactions on Services Computing, vol. 12, no. 4, pp. 564-576, 2019.
[10] H. Mao, M. Alizadeh, I. Menache, and S. Kandula, "Resource management with deep reinforcement learning," Proceedings of the 15th ACM Workshop on Hot Topics in Networks, pp. 50-56, 2016.
[11] Y. Tian and Y. Fan, "Virtual machine placement optimization in cloud data centers using deep reinforcement learning," IEEE Access, vol. 9, pp. 22231-22241, 2021.
[12] M. A. Rodriguez and R. Buyya, "Deadline based resource provisioning and scheduling algorithm for scientific workflows on clouds," IEEE Transactions on Cloud Computing, vol. 2, no. 2, pp. 222-235, 2014.
[13] J. Xu, A. Blackwell, R. Bhaskara, and G. Liu, "Deadline-constrained cost minimization for scientific workflow scheduling in multi-cloud environments," IEEE Cloud Computing, vol. 4, no. 6, pp. 18-27, 2017.
[14] Google LLC, "Google Cluster 2019 Workload Traces," Google Technical Report, 2019. [Online]. Available: https://github.com/google/cluster-data
[15] Y. Zhan, X. Liu, Y. Gong, L. Gu, and H. Yu, "Two birds with one stone: Jointly learning binary code for large-scale face image retrieval and attributes prediction," IEEE Transactions on Neural Networks and Learning Systems, vol. 30, no. 11, pp. 3459-3471, 2019.
[16] K. Hwang, J. Dongarra, and G. C. Fox, Distributed and Cloud Computing: From Parallel Processing to the Internet of Things. Morgan Kaufmann Publishers, 2012.
[17] J. Kennedy and R. Eberhart, "Particle swarm optimization," Proceedings of IEEE International Conference on Neural Networks, vol. 4, pp. 1942-1948, 1995.
[18] V. Mnih et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, pp. 529-533, 2015.
[19] P. Maymounkov and D. Mazières, "Kademlia: A peer-to-peer information system based on the XOR metric," Proceedings of the 1st International Workshop on Peer-to-Peer Systems, pp. 53-65, 2002.
[20] Z. Li, J. Ge, H. Hu, W. Song, H. Hu, and B. Luo, "Cost and energy aware scheduling algorithm for scientific workflows with deadline constraint in clouds," IEEE Transactions on Services Computing, vol. 11, no. 4, pp. 713-726, 2018.

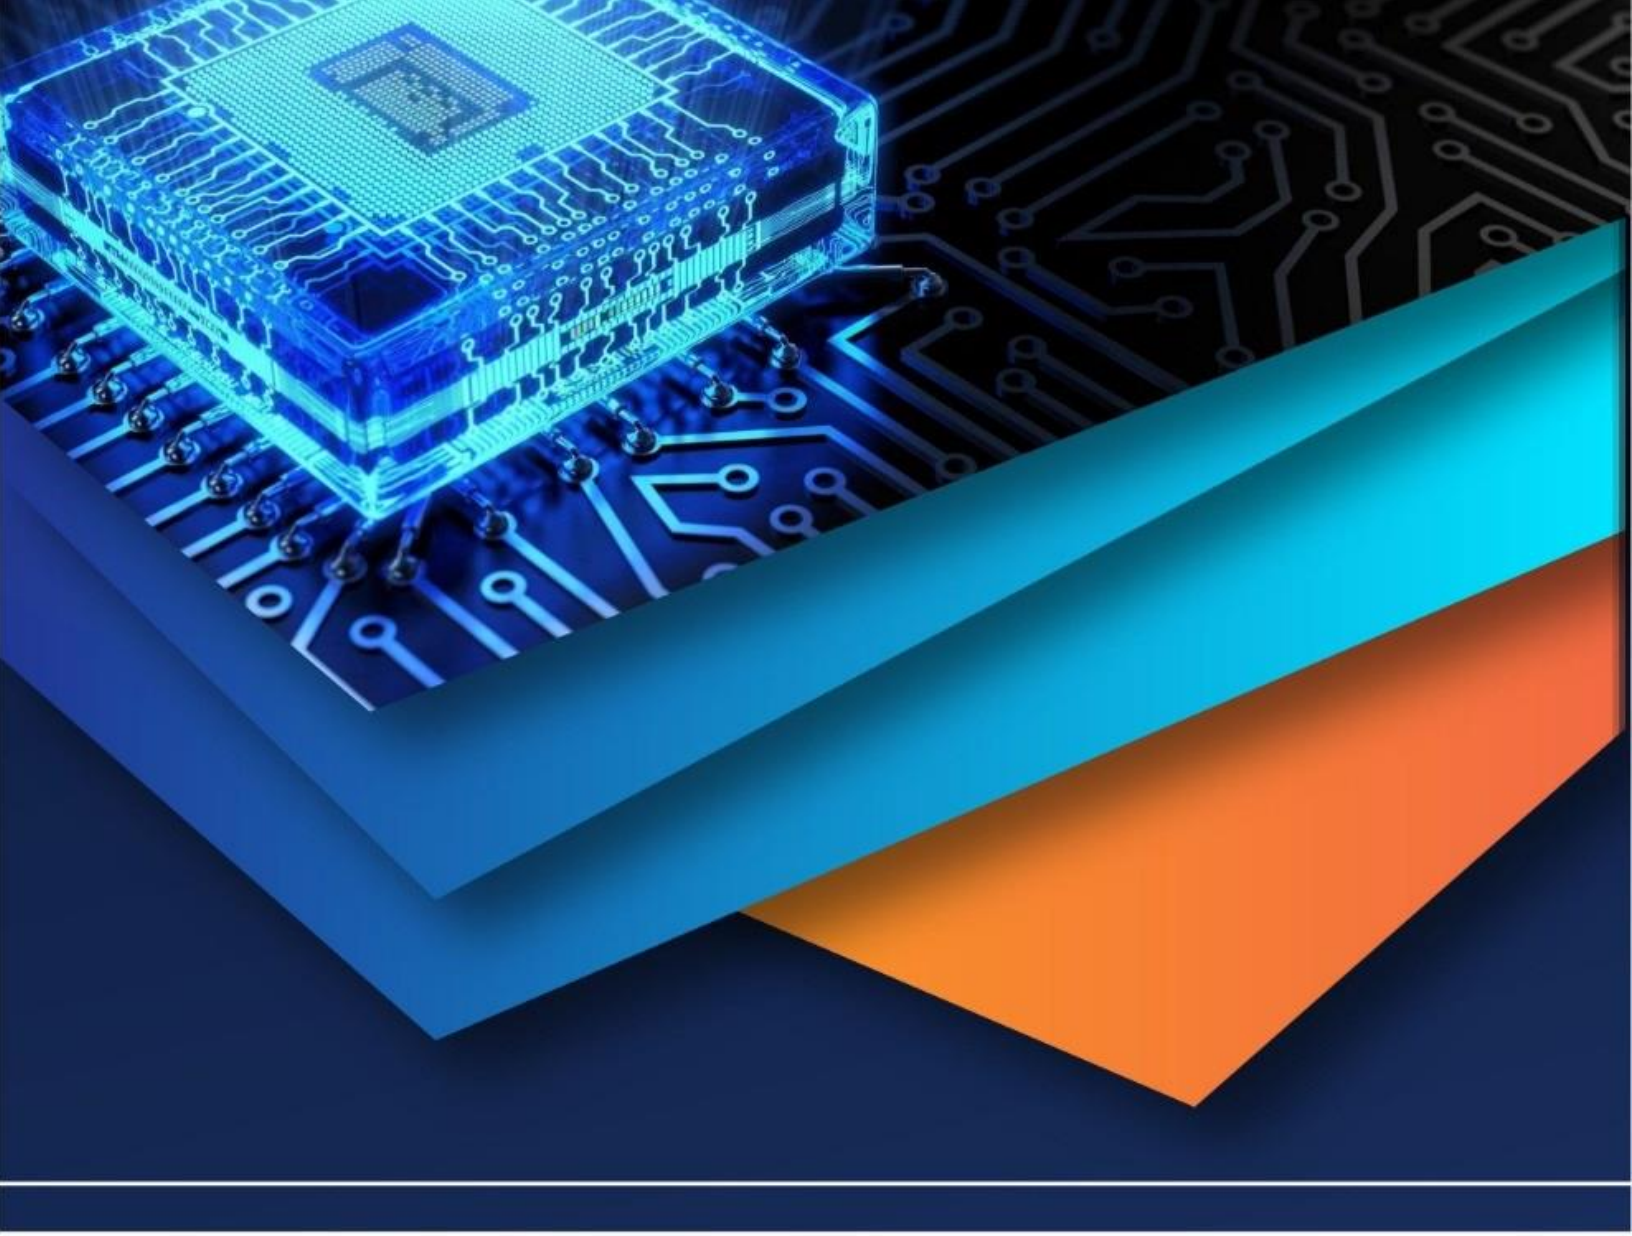

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY