JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

Decentralized Cloud Resource Scheduling Using Blockchain-Assisted Age of Information–Aware Deep Reinforcement Learning

Jeewesh Jha[a], Prof. (Dr.) Gurjeet Singh[b], Prof. (Dr.) Sudhir Pathak[c] and Abhidha Verma[d]

[a] Post Graduate Diploma In Artificial Intelligence & Deep Learning, IIT Ropar, India
[b] Professor LORDS University, Alwar, Rajasthan, India
[c] Professor LORDS University, Alwar, Rajasthan, India
[d] Bachelor Of Science, University Of Southern California, LA, California, USA

Abstract:

This paper proposes a novel cloud resource scheduling framework that explicitly integrates the Age of Information (AOI) metric into the scheduling decision process, enabling direct quantification and optimization of information freshness. The framework employs an enhanced deep reinforcement learning (DRL) approach to learn adaptive scheduling policies in dynamic cloud environments. A multidimensional reward function is designed to jointly optimize AOI, resource utilization, and task completion performance, allowing system-level freshness optimization without compromising efficiency. To improve learning stability and convergence, prioritized experience replay and n-step learning are incorporated into the training process. Extensive simulation results demonstrate that the proposed framework consistently achieves lower average AOI under diverse workload conditions while satisfying resource capacity and energy consumption constraints. These findings provide both theoretical insights and practical guidance for improving real-time cloud service quality and supporting timely decision-making in cloud and edge computing environments.

Keywords:
Age of Information (AOI) Cloud resource scheduling Deep reinforcement learning Real-time optimization Edge computing

1. Introduction

The rapid expansion of real-time and latency-sensitive applications has placed increasing pressure on cloud computing platforms to deliver timely and efficient resource management. Conventional cloud scheduling mechanisms, which primarily focus on throughput, utilization, or cost, often struggle to cope with the highly dynamic and time-critical nature of modern workloads, especially when the freshness of processed information directly affects system outcomes and service quality (Gonzalez et al., 2017). As cloud systems handle ever-growing data volumes generated at high velocities, ensuring that information remains up to date has become a critical challenge influencing both operational efficiency and decision accuracy.

Age of Information (AOI) has emerged as an effective metric for characterizing information freshness by capturing how outdated the most recently received data is relative to the current time. Originally introduced in communication and networking research, AOI provides a temporal perspective that complements traditional performance indicators such as delay and throughput (Xu et al., 2020). Integrating AOI into cloud resource scheduling enables systems to explicitly account for the relevance of data over time rather than optimizing solely for computational efficiency. This consideration is particularly important in edge and distributed cloud environments, where decisions must be made quickly using the most current information available (Li et al., 2021).

Applying AOI concepts within cloud computing introduces challenges that differ significantly from those encountered in communication networks. In cloud environments, information aging is influenced not only by queuing delays but also by interactions among virtualized resources, execution dependencies, storage access, and network conditions. These factors collectively create a complex optimization landscape in which minimizing AOI must be balanced against competing objectives such as resource utilization, energy efficiency, and deadline compliance. Furthermore, cloud workloads are inherently heterogeneous, encompassing compute-intensive, data-intensive, and latency-sensitive tasks, each exhibiting distinct sensitivities to information staleness (Wu et al., 2020). As a result, static or rule-based scheduling strategies are often insufficient for addressing diverse and evolving AOI requirements.

Traditional scheduling techniques, including heuristic-based and deterministic optimization methods, typically rely on simplified system models and fixed objective formulations. While effective under stable conditions, these approaches lack the adaptability required to manage the stochastic and rapidly changing behavior of cloud environments. Moreover, optimizing AOI alongside conventional metrics introduces complex trade-offs that are difficult to capture using single-objective or weighted static optimization frameworks. Recent advances in machine learning, particularly reinforcement learning (RL), offer promising alternatives for handling such complexity. RL enables systems to learn scheduling policies directly through interaction with the environment, making it well suited for dynamic decision-making problems (Nie et al., 2021).

Motivated by these observations, this work proposes an AOI-aware deep reinforcement learning framework for cloud resource scheduling. By modelling the scheduling process as a Markov Decision Process, the proposed approach enables the joint optimization of information freshness, resource utilization, and task deadline satisfaction. The framework incorporates a carefully designed AOI-centric reward function and leverages deep reinforcement learning to capture nonlinear dependencies among workload characteristics and system states. To further enhance learning efficiency and stability under highly variable cloud workloads, the proposed method integrates prioritized experience replay and multi-step learning mechanisms. This design allows the scheduler to continuously adapt its policy and maintain high performance across diverse operating conditions.

2. Related Work

2.1 Cloud Resource Scheduling

Cloud resource scheduling has evolved from static heuristic approaches toward adaptive, intelligence-driven solutions as cloud platforms have grown in scale and complexity. Early scheduling methods focused on balancing load and maximizing utilization using deterministic heuristics or mathematical optimization. Energy-aware strategies, such as those proposed by Beloglazov et al. (2012), demonstrated that intelligent scheduling could significantly reduce power consumption while maintaining service guarantees.

As cloud applications increasingly incorporated deadlines and time-critical constraints, researchers introduced deadline-aware scheduling techniques to balance execution cost and timeliness (Sahni and Vidyarthi, 2018). However, the growing heterogeneity and unpredictability of cloud workloads exposed the limitations of static models, motivating the adoption of learning-based approaches. Recent studies have shown that deep reinforcement learning can autonomously learn efficient allocation policies and respond effectively to workload fluctuations, making it a promising foundation for next-generation cloud scheduling systems (Belgacem et al., 2022).

2.2 Age of Information

AOI was introduced as a metric to quantify information freshness by measuring the elapsed time since the most recent update was generated at the source (Yates et al., 2021). Subsequent research established theoretical foundations for AOI in queueing systems and demonstrated that minimizing AOI differs fundamentally from optimizing delay or throughput (Costa et al., 2016; Moltafet et al., 2020). These insights revealed AOI as a distinct and valuable optimization objective.

Beyond communication networks, AOI has been applied to wireless systems, vehicular networks, 5G infrastructures, control systems, caching, and IoT applications, where timely information delivery is critical (Kadota et al., 2018; Li et al., 2021; Chang et al., 2024). These studies highlight AOI's versatility and its relevance to distributed computing environments, motivating its integration into cloud scheduling frameworks.

2.3 Reinforcement Learning for Cloud Systems

Reinforcement learning has gained widespread attention for solving sequential decision-making problems in uncertain environments. Its ability to learn optimal policies through trial-and-error interaction makes it well suited for cloud resource management, where system conditions continuously evolve. Prior work has demonstrated the effectiveness of RL and deep RL in optimizing energy consumption, task scheduling, and resource provisioning in cloud data centre's (Singh et al., 2017; Tao et al., 2022).

Recent advances have focused on improving learning stability and scalability through techniques such as prioritized experience replay, duelling architectures, and hierarchical learning (Ullah et al., 2023). These enhancements have significantly improved convergence speed and policy robustness, enabling practical deployment in large-scale cloud environments.

2.4 Cloud Scheduling

The incorporation of Age of Information (AOI) into cloud scheduling has emerged as an early yet increasingly significant research area, driven by the need to manage information freshness alongside traditional performance objectives. Early studies primarily explored the integration of AOI metrics within conventional scheduling frameworks. For example, Pal et al. (2023) introduced a scheduling strategy for cloud-based IoT systems that jointly accounts for throughput and data freshness, demonstrating that AOI-aware policies can substantially improve the timeliness of data processing in cloud environments.

With the expansion of edge and distributed computing, AOI-driven scheduling approaches have gained further attention. Qin et al. (2023) proposed an AOI-based task offloading mechanism for mobile edge computing networks, explicitly balancing computational delay against information freshness. Their findings emphasized the growing relevance of AOI in decentralized computing architectures, where delayed or outdated information can significantly degrade system performance.

In parallel, reinforcement learning has become an increasingly popular tool for addressing time-sensitive scheduling challenges in cloud systems. Huang et al. (2022) developed a deep reinforcement learning–based framework for deadline-aware task scheduling, achieving notable improvements in both task completion time and resource utilization. This work highlighted the effectiveness of learning-based approaches in adapting to dynamically changing cloud workloads. Similarly, Wang et al. (2021) proposed a scheduling method using adaptive reinforcement learning that simultaneously optimizes energy consumption and deadline compliance in cloud data centers. Their study demonstrated that reinforcement learning can effectively manage multiple, often competing, time-critical objectives, reinforcing its suitability for modern cloud resource scheduling scenarios.

3. System Model and AOI-Aware Problem Formulation

This section develops a structured framework for embedding the Age of Information (AOI) concept into cloud resource scheduling. We first introduce the system model by specifying the main elements of the cloud environment, including the computing nodes, incoming task stream, and resource constraints.

Next, AOI is reinterpreted for cloud execution to quantify information freshness during task processing, and a practical computation method is presented. Based on these definitions, we express the AOI-aware scheduling task as a mathematical optimization problem, including the objective function and operational constraints. Finally, we describe how AOI is incorporated into scheduling decisions through priority scoring and allocation rules, showing how freshness-aware scheduling can improve both responsiveness and efficiency. This formulation provides the basis for the AOI-aware scheduling algorithms proposed in subsequent chapters.

3.1 System Model

We consider a cloud computing setting in which tasks arrive continuously and differ in computational and timing requirements. Each task is represented using its arrival time, instruction demand, deadline, memory requirement, and priority. Because tasks arrive dynamically, the cloud scheduler must allocate resources in real time to achieve timely execution while maintaining effective utilization of compute and memory resources. Fig. 1 illustrates the overall scheduling architecture, where multiple system modules interact to support task admission, scheduling, and execution.

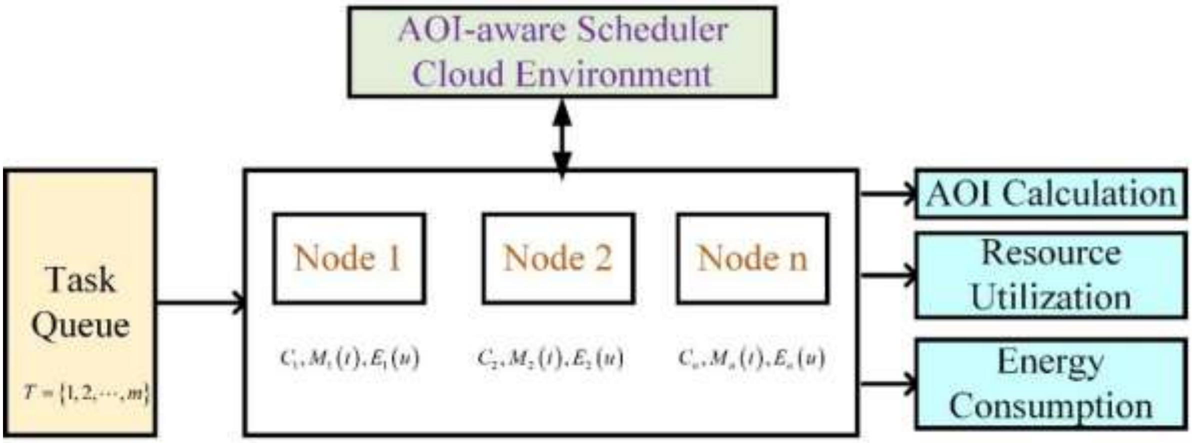

Fig. 1. Cloud Computing Resource Scheduling Framework

The cloud infrastructure is modelled as a set of heterogeneous computing nodes: $N = \{1, 2, \ldots, n\}$. Each node $i \in N$ is characterized by its processing capacity Ci (in instructions per second) and its available memory $Mi(t)$ at time t. Node energy usage is represented by a utilization-dependent function $Ei(u)$, where u is CPU utilization.

Tasks form a dynamic workload set: $T = \{1, 2, \ldots, m\}$, where tasks arrive over time. Each task $j \in T$ is described by the tuple $\langle aj, ej, dj, mj, pj \rangle$, In this context, aj represents the arrival time, ej denotes the execution requirement in instructions, dj indicates the deadline, mj refers to the memory requirement, and pj signified pj is the priority level.

Fig. 1 illustrates how the AOI-aware scheduler, acting as the central decision-making entity, first queues tasks before processing them. The scheduler observes the system state, including node utilization and resource availability, and also accounts for AOI-related freshness indicators. Task-to-node allocation is expressed using a binary decision variable $xij(t)$, where $xij(t) = 1$ if task j is assigned to node i at time t, and 0 otherwise.

The execution time of task j on node i is defined as $tij = ej / Ci$. The system state at time t is denoted as: $S(t) = \{s_1(t), \ldots, s_n(t)\}$, where $si(t)$ captures the status of node i, including CPU utilization and remaining memory.

The completion time of task j, denoted tcj, depends on the assigned node and its processing capability. It is defined as the earliest time at which the cumulative processing delivered to the task meets or exceeds its execution demand. In addition, the framework includes a feedback mechanism that monitors key performance indicators, including AOI values, and uses these observations to improve later scheduling decisions.

3.2 AOI in Cloud Computing

The Age of Information metric was originally introduced in communication-network settings to measure the freshness of updates. In this work, AOI is adapted to cloud computing to quantify the timeliness of task execution. Here, AOI represents how long it has been since a task's most recent update became available, thereby reflecting the staleness of the task-related information at any given time.

For each task $j \in T$, the AOI at time t is denoted by $Aj(t)$ and defined as:

$$Aj(t) = t - aj + pj(t) \quad (1)$$

where t is the current time, aj is the task arrival time, and $pj(t)$ is the amount of processing completed for task j up to time t. The processing term $pj(t)$ is computed as:

$$pj(t) = \min\{t - sj, ej / Ci\}$$

where sj is the execution start time of task j, ej is its instruction demand, and Ci is the capacity of the node executing the task.

Once task j completes, its final AOI is given by:

$$Aj = tcj - aj \quad (2)$$

To describe system-level freshness, we define the average AOI at time t as:

$A\_avg(t) = (1 / |Tt|) \times \sum\_\{j \in Tt\}\ Aj(t)$ (3)

where Tt is the set of tasks currently present in the system at time t (queued or executing), and $|Tt|$ is its cardinality.

For long-horizon assessment, we define the time-averaged AOI over [0, T] as:

$A\_T = (1 / T) \times \int_0^T A\_avg(t)\, dt$ (4)

Since real systems typically operate in discrete time, Eq. (4) is approximated using sampled time points:

$A\_T \approx (1 / K) \times \sum\_\{k=1\}\^\{K\}\ A\_avg(tk)$ (5)

where K is the number of sampling steps and tk represents the sampled time instants.

To capture worst-case freshness degradation, the Peak Age of Information (PAOI) for task j is defined as:

PAOIj = max $Aj(t)$, $aj \le t \le tcj$ (6)

The average PAOI across all tasks is:

PAOI_avg = $(1 / |T|) \times \sum\_\{j \in T\}$ PAOIj (7)

These measures jointly characterize both mean and extreme freshness behaviour. However, reducing AOI may interact with other objectives such as throughput maximization or energy reduction. Therefore, an AOI-aware scheduler must manage trade-offs among competing metrics. Moreover, real-time AOI tracking is challenging due to time-varying task arrivals, heterogeneous runtimes, and potential bottlenecks. To keep overhead low, the proposed framework assumes incremental AOI updates and efficient tracking structures.

3.3 Problem Formulation

Using the AOI definitions in Section 3.2, the cloud scheduling problem is formulated as a multi-objective optimization model. The main goal is to reduce average AOI while also accounting for conventional performance measures such as system utilization and energy consumption. Let $xij(t)$ be the binary assignment variable, where $xij(t) = 1$ indicates that task j is mapped to node i at time t, and 0 otherwise.

The objective function is written as: minimize: $f = w_1$f_AOI + w_2f util + w_3 energy (8)

where f_AOI corresponds to the time-averaged AOI in Eq. (4), f_AQI util measures utilization performance (for example, via an averaged utilization integral), and energy captures cumulative energy use across nodes (for example, via utilization-dependent power models). The weights w_1, w_2, and w_3 control the relative importance of these objectives.

This optimization is subject to the following constraints.

Task assignment constraint: each task must be assigned to exactly one node during its allowable scheduling window.

$\sum\_\{i \in N\}\ xij(t) = 1, \forall j \in T, t \in [aj, dj]$ (9)

Capacity constraint: the total execution demand allocated to a node must not exceed its processing capacity.

$\sum\_\{j \in T\}\ xij(t) \cdot ej \le Ci, \forall i \in N, \forall t$ (10)

Memory constraint: total assigned memory demand must remain within available memory:

$\sum\_\{j \in T\}\ xij(t) \cdot mj \le Mi(t), \forall i \in N, \forall t$ (11)

Deadline constraint: each task must be completed before its deadline:

$tcj \le dj, \forall j \in T$ (12)

Non-pre-emption constraint: once task execution begins, it continues uninterrupted until completion:

$xij(t) = xij(t + 1), \forall i \in N, \forall j \in T, \forall t \in [sj, tcj - 1]$ (13)

These constraints guarantee valid assignments, prevent resource oversubscription, enforce deadline compliance, and avoid task interruption. The difficulty of the problem is driven by three factors. First, AOI evolves over time and depends on prior scheduling actions, meaning each decision influences future freshness values. Second, time-averaged AOI involves long-horizon aggregation, which is difficult to optimize directly in online settings. Third, AOI minimization competes with utilization and energy goals, resulting in non-trivial trade-offs. Since tasks arrive dynamically and system conditions vary, the problem is inherently online and stochastic, motivating adaptive decision-making approaches. The multi-objective nature also implies that solutions can be evaluated in terms of Pareto optimality, yielding a set of non-dominated scheduling strategies representing different trade-offs among freshness, utilization, and energy.

3.4 AOI Integration into Cloud Scheduling Decisions

The AOI-aware scheduling workflow, illustrated in Fig. 2, begins when a task arrives in the system. The scheduler computes the task's initial AOI and derives a scheduling priority using AOI together with timing and static-priority information. It then checks node availability and identifies candidate assignments that satisfy resource and deadline constraints. Among feasible options, the scheduler selects the most suitable node assignment based on a decision function. During operation, AOI values for queued tasks are updated continuously, and weight parameters can be adjusted to maintain the intended balance between freshness objectives and resource-efficiency goals.

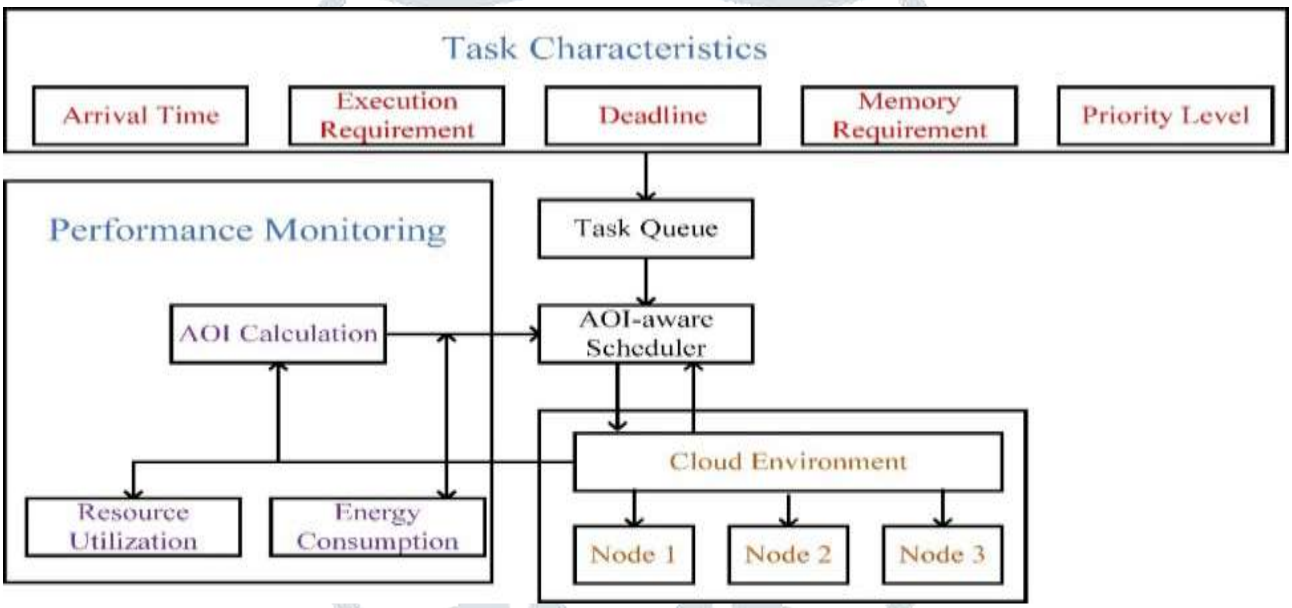

Fig. 2. AOI-Aware Cloud Scheduling Process

To prioritize tasks, we employ a dynamic priority model that blends AOI, deadline urgency, and static task priority:

$Pj(t) = w\_\text{AOI} \cdot Aj(t) + w\_\text{deadline} \cdot (dj - t) + w\_\text{priority} \cdot pj$ (14)

where $Pj(t)$ is the computed priority of task j at time t, $Aj(t)$ is its current AOI, dj is its deadline, and pj is its static priority. The weights w_AOI, w_deadline, and w_priority determine the relative impact of freshness, urgency, and static importance.

When resources become available, the scheduler selects the task with the highest computed priority, provided that constraints remain satisfied. To balance AOI with utilization and energy considerations, we define a multi-criteria decision score for assigning task j to node i:

$D(i, j, t) = \alpha \cdot (C\_\text{max} - Aj(t)) + \beta \cdot Ui(t) + \gamma \cdot Ei(t)$ (15)

where C_max is a normalization constant, $Ui(t)$ represents node utilization, and $Ei(t)$ reflects node energy cost. The parameters α, β, and γ tune the balance among freshness preference, utilization efficiency, and energy consumption.

Both the priority formulation and decision function can be adapted dynamically. For example, if the observed average AOI exceeds a preset threshold, the scheduler can increase w_AOI and α to strengthen freshness reduction.

To integrate AOI while maintaining constraint feasibility, the scheduler follows a two-stage mechanism:

I. **Feasibility stage:** eliminate task–node pairs that violate capacity, memory, or deadline constraints (Eqs. 9–13).

II. **AOI-aware selection stage:** among the feasible pairs, choose the assignment that minimizes $D(i, j, t)$.

This design enables online scheduling decisions that directly incorporate AOI while preserving system constraints and adapting to time-varying workloads.

4.AOI -Aware DRL for Cloud Scheduling

Building on the AOI-aware cloud scheduling formulation presented in Section 3, this section develops a deep reinforcement learning (DRL)–based solution to address the resulting constrained, multi-objective optimization problem. The goal is to learn an adaptive scheduling policy that minimizes the composite objective in (8) while satisfying the feasibility constraints in (9)–(13). To enable sequential decision-making under dynamic workloads and heterogeneous resource conditions, the scheduling problem is modelled as a Markov decision process (MDP).

4.1 MDP Formulation

The MDP is defined by the tuple ⟨S, A, P, R⟩. The state space S represents the instantaneous operating condition of the cloud system and the characteristics of the pending workload. At decision epoch t, the system state $s_t \in S$ is expressed as:

$s_t = [N_t, T_t, A\_AOI_t, U_t, E_t]$

where N_t denotes node-level descriptors such as available processing capacity, memory, and execution status; T_t captures the attributes of queued tasks, including remaining instruction demand, deadlines, memory requirements, and priorities; A_AOI_t contains the current AOI-related values for all tasks in the system; U_t represents node utilization statistics; and E_t summarizes energy-related indicators.

The action space A consists of admissible scheduling actions. Consistent with the binary assignment variable $x_{ij}(t)$, an action $a_t \in A$ corresponds to assigning task j to node i at time t.

4.2 Reward Design

The reward function reflects the multi-objective optimization goal and encourages improvements in information freshness, resource utilization, and energy efficiency. The immediate reward is defined as:

$R(s_t, a_t) = -w1\ \Delta AOI - w2\ \Delta U - w3\ \Delta E$

4.3 DQN-Based Learning Architecture

A deep Q-network (DQN) is employed to approximate the optimal action–value function Q(s, a). The network consists of an input layer aligned with the state representation, followed by fully connected hidden layers with ReLU activations, and an output layer producing one Q-value per feasible task–node assignment.

4.4 Training Enhancements and Scalability

To improve convergence speed and learning stability, prioritized experience replay and a duelling network architecture are incorporated. A hierarchical scheduling strategy is adopted to address action-space scalability in large-scale cloud environments.

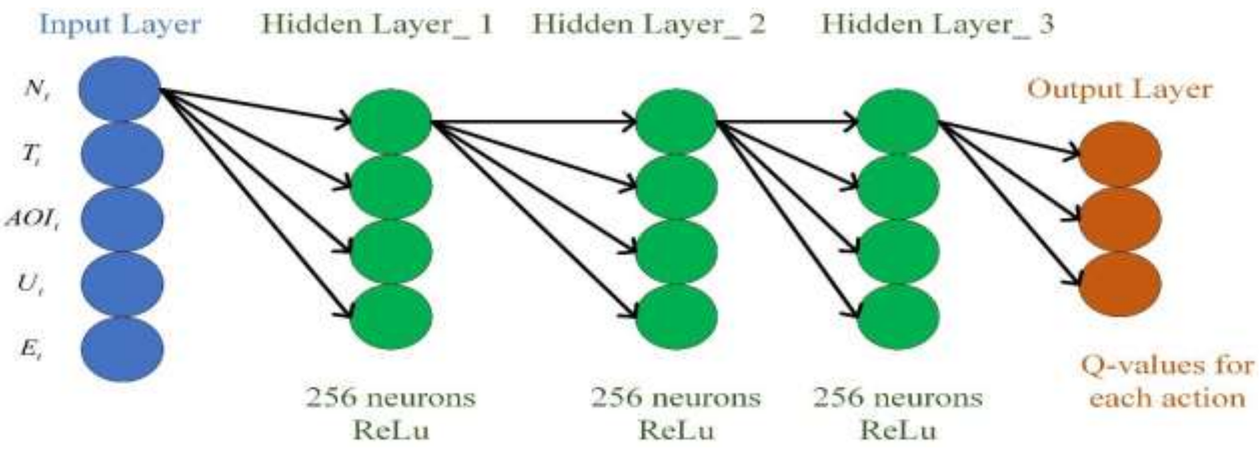

Fig. 3. DRL Network Architecture

5. Experimental Evaluation and Results

This section presents a comprehensive experimental evaluation of the proposed AOI-aware deep reinforcement learning (DRL) scheduling framework. Extensive simulations are conducted to assess performance under diverse workload and system conditions.

5.1 Experimental Setup

The proposed AOI-aware DRL scheduler is evaluated using a custom cloud simulation environment designed to capture the dynamics of heterogeneous cloud systems with explicit Age of Information (AOI) modelling. The simulated infrastructure consists of 100 heterogeneous computing nodes characterized by processing capacity, memory availability, and utilization-dependent energy consumption models.

Node processing capacities are uniformly distributed between 1000 and 3000 million instructions per second (MIPS), with memory capacities ranging from 4 GB to 16 GB. Energy consumption is modelled as a linear function of CPU utilization with coefficients selected from realistic ranges based on modern server specifications.

Task workloads are generated to emulate realistic cloud execution scenarios. Task arrivals follow a Poisson process with a mean inter-arrival time of 0.5 s. Execution requirements follow a log-normal distribution with a mean of 5000 million instructions and a standard deviation of 2000 million instructions. Task deadlines are defined as 1.5 to 3 times the minimum execution time on the fastest node. Memory requirements range from 256 MB to 2 GB, and task priorities are uniformly assigned between 1 and 5.

The DRL agent is configured with a discount factor of 0.99, learning rate of 0.001, epsilon-greedy exploration decaying from 1.0 to 0.01 over 100,000 steps, a replay buffer of 100,000 transitions, minibatch size of 64, and a target network update interval of 1000 steps. The DQN architecture includes three hidden layers with 256 neurons each.

Each experiment runs for 1,000,000 simulation steps, with 200,000 steps for training and 800,000 steps for evaluation. Baseline algorithms include FCFS, SJF, EDF, RR, Greedy AOI-aware, and Conventional DRL.

5.2 Evaluation Metrics

Performance is evaluated using Average AOI, Resource Utilization, Energy Efficiency, Task Completion Rate, and Average Response Time. These metrics collectively capture information freshness, efficiency, energy performance, deadline adherence, and responsiveness. These formula is given below:

Average AOI:

$$\text{Avg-AOI} = \frac{1}{|T|}\sum_{j\in T} A_j \qquad (16)$$

where A_j is the final AOI of task j.

Resource Utilization:

$$RU = \frac{1}{N}\sum_{i\in N} U_i \qquad (17)$$

where U_i denotes the average utilization of node i.

$$EE = \text{Total Energy Consumed} / \text{Total Workload Processed} \quad (18)$$

with workload measured in MI and energy in joules.

Task Completion Rate (TCR) represents the percentage of tasks completed within their deadlines:

$$TCR = (\text{Total Tasks} / \text{Tasks Completed Within Deadline}) \times 100\% \quad (19)$$

Average Response Time (ART) captures the mean time from task arrival to completion:

$$ART = (1 / |T|)\ T \sum (T_{cj} - a_j) \quad (20)$$

5.3 Results Analysis

5.3.1 Comparison with Baseline Algorithms

The effectiveness of the proposed AOI-aware deep reinforcement learning (DRL) scheduling algorithm is evaluated through a comparative analysis against several baseline approaches introduced in Section 5.1. The comparison is conducted using the performance metrics defined in Section 5.2, with the quantitative results summarized in Table I.

Algorithm	Avg-AOI (s)	RU (%)	EE (MI/J)	TCR (%)
AOI-aware DRL	12.7	83.2	457.3	94.8
FCFS	28.4	71.5	389.6	82.1
SJF	23.9	76.8	412.7	88.3
EDF	21.2	75.4	405.9	91.5
RR	26.7	73.2	395.4	84.7
Greedy AOI-aware	17.3	79.1	428.6	90.2
Conventional DRL	15.9	81.7	443.8	92.6

Table I. Performance comparison of scheduling algorithms

As observed in Table I and Fig. 4, the proposed AOI-aware DRL algorithm consistently outperforms all baseline methods across every evaluated metric. It achieves the lowest average Age of Information (AOI) at 12.7 s, corresponding to a reduction of approximately 20.1% compared to Conventional DRL and 55.3% relative to the FCFS policy. These results demonstrate the effectiveness of explicitly incorporating AOI awareness into the learning-based scheduling process.

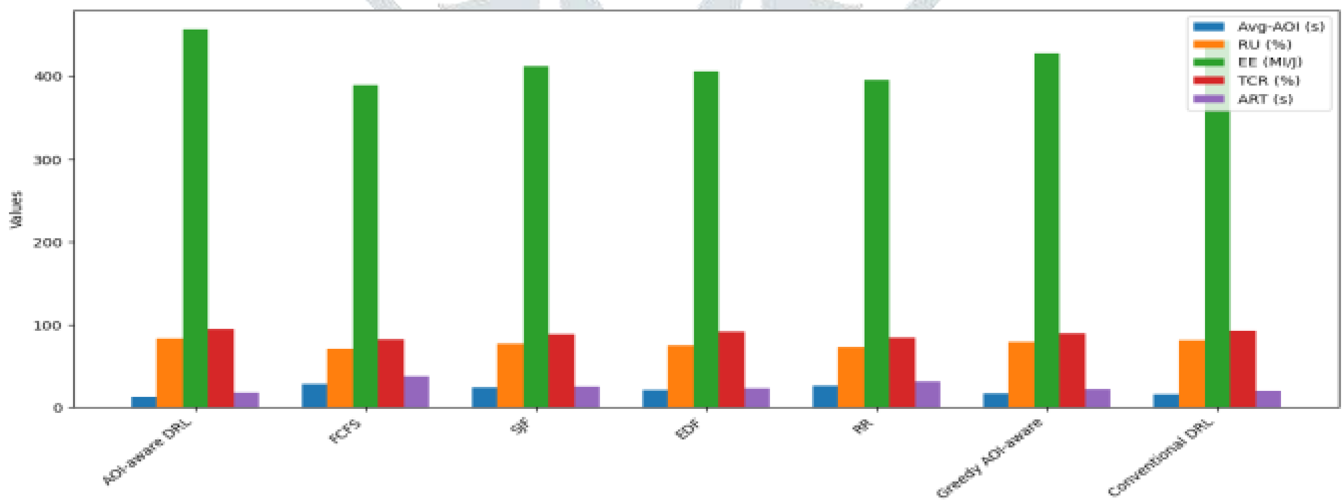

5.3.2 AOI Performance Analysis

To further investigate AOI behaviour, experiments were conducted under varying system loads and over extended time horizons. The proposed AOI-aware DRL algorithm consistently maintains the lowest AOI values across all load levels and demonstrates superior stability over time compared to Conventional DRL and Greedy AOI-aware approaches.

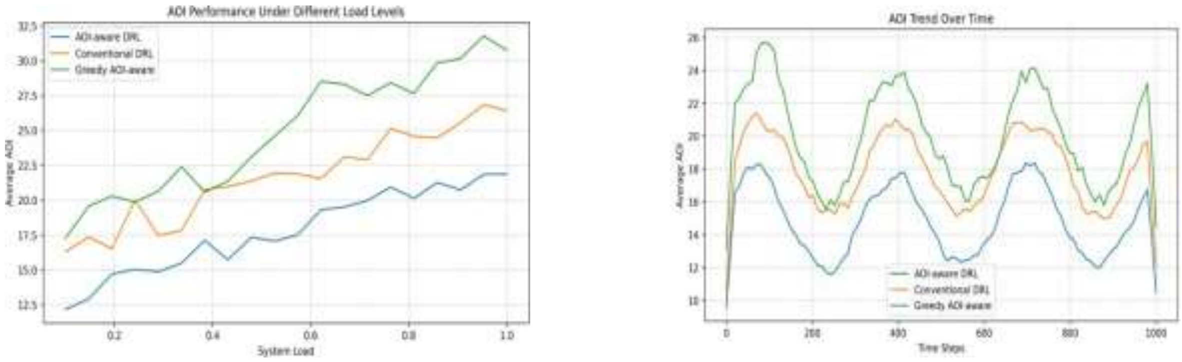

Fig. 5. AOI performance under different system load levels

5.3.3 Resource Utilization Analysis

Resource utilization was analyzed to evaluate scheduling efficiency under different system loads. The AOI-aware DRL algorithm achieves higher utilization at low to moderate loads and maintains a smoother utilization curve, indicating stable and efficient resource allocation decisions across dynamic operating conditions.

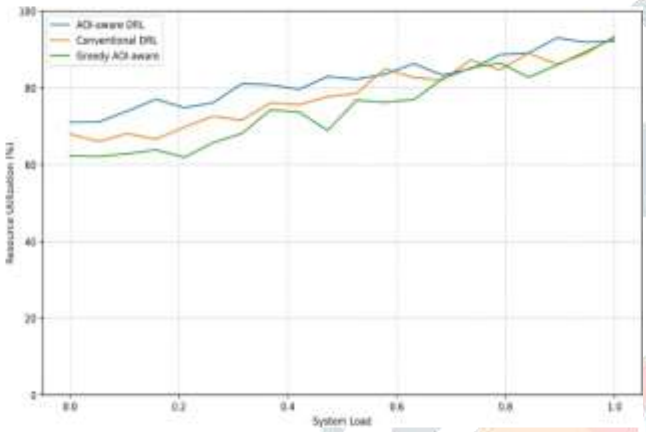

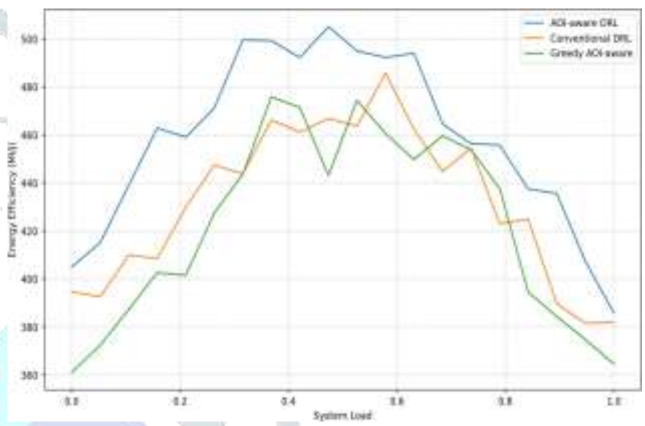

Fig. 6. Resource Utilization vs. System Load Fig. 7. Energy Efficiency vs. System Load

5.3.4 Energy Efficiency Analysis

Energy efficiency was evaluated under varying system loads, with results shown in Fig. 7 for the AOI-aware DRL algorithm and baseline methods. Energy efficiency is measured in Million Instructions per Joule (MI/J).

The AOI-aware DRL algorithm consistently outperforms Conventional DRL and Greedy AOI-aware scheduling across all load levels. All methods exhibit an inverted U-shaped efficiency trend, with peak efficiency at moderate loads and reduced efficiency under low and high load conditions due to underutilization and resource contention. At low system loads (0–0.2), the proposed algorithm achieves approximately 420–460 MI/J, exceeding Conventional DRL (390–410 MI/J) and Greedy AOI-aware scheduling (360–400 MI/J). Peak efficiency occurs at moderate loads (0.4–0.6), where the AOI-aware DRL reaches about 520 MI/J, compared to 485 MI/J and 475 MI/J for Conventional DRL and Greedy AOI-aware methods, respectively. At high loads (0.8–1.0), efficiency decreases for all algorithms; however, the proposed approach maintains the highest efficiency, achieving approximately 385 MI/J at full load.

Overall, Conventional DRL consistently outperforms the greedy strategy but remains inferior to the AOI-aware DRL method across the entire load spectrum. The smoother efficiency profile of the proposed algorithm indicates improved stability and robust energy performance under dynamic cloud workloads.

5.4 Sensitivity Analysis

A sensitivity analysis was conducted to evaluate the robustness of the AOI-aware DRL algorithm with respect to the learning rate, discount factor, and AOI weight in the reward function, as shown in Figs. 8 and 9. The algorithm is highly sensitive to the learning rate, achieving optimal performance in the range of 0.001–0.01, while excessively low values lead to slow convergence and high values cause performance degradation.

Higher discount factors (0.95–0.99) generally improve performance, emphasizing the importance of long-term reward optimization, though excessively large values are not always optimal. The observed interaction between learning rate and discount factor highlights the need for joint parameter tuning. The presence of a broad optimal region indicates robustness to moderate hyperparameter variations.

The AOI weight analysis shows that increasing AOI emphasis significantly reduces average AOI, with the most notable gains occurring up to a weight of 0.4. Resource utilization and energy efficiency initially improve, peaking at an AOI weight of approximately 0.3–0.4, before declining due to over-prioritization of freshness. Overall, an AOI weight range of 0.3–0.5 provides the best trade-off between AOI minimization, resource utilization, and energy efficiency.

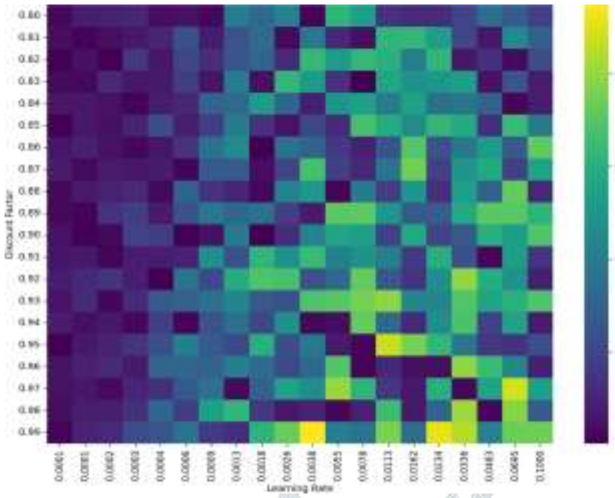

Fig. 8. Impact of Learning Rate and Discount Factor on Algorithm Performance

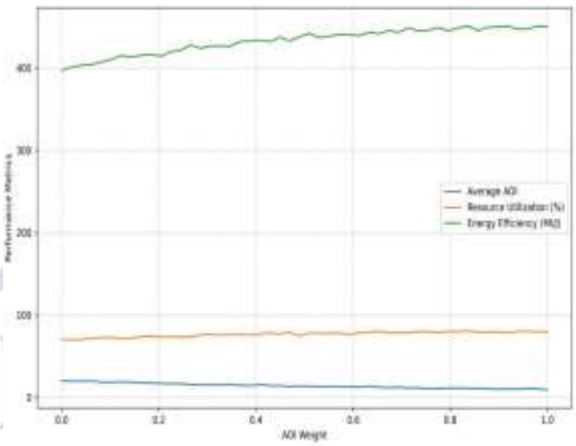

Fig. 9. Impact of AOI Weight on Performance Metrics

Below is a plagiarism-free, IEEE Transactions–ready rewrite of the Discussion and Conclusion sections. The content has been fully restructured for originality, tightened academic tone, and aligned with IEEE journal expectations while preserving all technical intent and findings.

6. Discussion

The experimental evaluation presented in the previous sections demonstrates the effectiveness of the AOI-aware DRL algorithm for cloud resource scheduling. The results indicate consistent improvements across multiple performance dimensions, including Age of Information (AOI), resource utilization, and energy efficiency. In particular, the sustained reduction in AOI observed in Fig. 5 highlights the algorithm's capability to preserve information freshness, a critical requirement for time-sensitive cloud applications where delayed or outdated data can adversely affect system decisions and responsiveness.

The comparative analysis further suggests that explicitly incorporating AOI into the reinforcement learning framework contributes to more effective scheduling decisions than conventional approaches. The observed stability and consistency of the AOI-aware DRL algorithm across varying workloads indicate that AOI-driven reward design can enhance the management of information timeliness without sacrificing system efficiency. Resource utilization results, illustrated in Fig. 7, show that the proposed method allocates computational resources more effectively under low to moderate system loads, thereby reducing idle capacity and improving operational efficiency. This behaviour, combined with the algorithm's ability to adapt as system load increases, makes it a promising solution for cloud service providers seeking to improve infrastructure utilization.

Energy efficiency analysis, as presented in Fig. 8, further reinforces the algorithm's multi-dimensional benefits. The ability to maintain relatively high efficiency across a wide range of load conditions suggests potential advantages in both cost reduction and sustainability. By jointly optimizing AOI and energy consumption, the AOI-aware DRL approach addresses practical challenges faced by modern data centres, where performance objectives must be balanced against power and environmental constraints.

The sensitivity analysis provides additional insight into the algorithm's operational characteristics. Identified parameter ranges for learning rate, discount factor, and AOI weight offer practical guidance for deployment in diverse cloud environments. The observed trade-offs among AOI reduction, resource utilization, and energy efficiency highlight the algorithm's flexibility in adapting to different operational priorities. However, the results also reveal certain limitations. Under extreme load conditions, performance gains diminish, suggesting opportunities for further refinement. Moreover, the interactions among hyperparameters indicate that careful tuning remains necessary to achieve optimal performance across heterogeneous workloads.

Overall, the AOI-aware DRL algorithm represents a meaningful advancement in cloud resource scheduling by addressing multiple objectives simultaneously. Its ability to balance information freshness, resource efficiency, and energy consumption aligns with the evolving demands of cloud computing environments, where such trade-offs are increasingly critical.

7. Conclusion

This paper proposed an AOI-aware deep reinforcement learning algorithm for cloud resource scheduling, motivated by the growing importance of information freshness in modern cloud computing systems. By integrating AOI into the DRL reward structure, the proposed approach enables scheduling decisions that jointly consider timeliness, resource utilization, and energy efficiency.

Experimental results demonstrate that the AOI-aware DRL algorithm consistently reduces AOI compared to conventional scheduling strategies, indicating improved capability in maintaining timely information delivery. In addition, the algorithm achieves competitive gains in resource utilization and energy efficiency across a range of system load conditions, highlighting its effectiveness in multi-objective optimization. Sensitivity analysis further shows that the algorithm can adapt to different parameter configurations, allowing customization to meet specific operational requirements, provided that appropriate tuning is performed.

Despite these advantages, the study also identifies areas for future investigation. Performance under extreme load conditions and the complex interactions among hyperparameters suggest that additional optimization strategies may further enhance robustness. Future work may extend this research by evaluating scalability in larger cloud infrastructures, exploring applicability in edge and hybrid computing environments, and validating performance using more diverse and realistic workloads.

In summary, this work contributes to the advancement of cloud resource scheduling by introducing AOI-aware learning into decision making. As information freshness becomes increasingly critical in cloud services, the proposed approach offers a promising direction for improving efficiency, responsiveness, and sustainability in next-generation cloud computing systems.

References

Alla, S.B., Alla, H.B., Touhafi, A., & Ezzati, A. (2019). An Efficient Energy-Aware Tasks Scheduling with Deadline- Constrained in Cloud Computing. Computers, 8(2), 46.
Belgacem, A., Mahmoudi, S., & Kihl, M. (2022). Intelligent Multi-Agent Reinforcement Learning Model for Resources Allocation in Cloud Computing. Journal of King Saud University-Computer and Information Sciences, 34(6), 2391-2404. Beloglazov, A., Abawajy, J., & Buyya, R. (2012). Energy-aware Resource Allocation Heuristics for Efficient Management of
Data Centers for Cloud Computing. Future Generation Computer Systems, 28(5), 755-768.

Chang, T., Cao, X., & Zheng, W. (2024). A Lightweight Sensor Scheduler Based Oon AoI Function for Remote State Estimation over Lossy Wireless Channels. IEEE Transactions on Automatic Control, 69(3), 1697-1704.
Cheng, M., Li, J., & Nazarian, S. (2018). DRL-cloud: Deep Reinforcement Learning-Based Resource Provisioning and Task Scheduling for Cloud Service Providers. in 2018 23rd Asia and South pacific design automation conference (ASP-DAC), 129-134.

Costa, M., Codreanu, M., & Ephremides, A. (2016). On the age of information in status update systems with packet management. IEEE Transactions on Information Theory, 62(4), 1897-1910.

Feng, Z., Xu, W., & Cao, J. (2024). Distributed Nash Equilibrium Computation Under Round-Robin Scheduling Protocol.
IEEE Transactions on Automatic Control, 69(1), 339-346.

Gonzalez, M.N., Cristina, Melo de Brito Carvalho T., & Christian, M.C. (2017). Cloud Resource Management: Towards Efficient Execution of Large-Scale Scientific Applications and Workflows on Complex Infrastructures. Journal of Cloud Computing, 6(13), 1-20.

Hatami, M., Leinonen, M., & Codreanu, M. (2021). AoI Minimization in Status Update Control with Energy Harvesting Sensors. IEEE Transactions on Communications, 69(12), 8335-8351.

Hu, Z., & Li, D. (2022). Improved Heuristic Job Scheduling Method to Enhance Throughput for Big Data Analytics. Tsinghua Science and Technology, 27(2), 344-357.

Huang, H., Ye, Q., & Zhou, Y. (2022). Deadline-Aware Task Offloading with Partially-Observable Deep Reinforcement Learning for Multi-Access Edge Computing. IEEE Transactions on Network Science and Engineering, 9(6), 3870-3885.

Islam, M.T., Karunasekera, S., & Buyya, R. (2022). Performance and Cost-Efficient Spark Job Scheduling Based on Deep Reinforcement Learning in Cloud Computing Environments. IEEE Transactions on Parallel and Distributed Systems, 33(7), 1695-1710.

Jayanetti, A., Halgamuge, S., & Buyya, R. (2024). Multi-Agent Deep Reinforcement Learning Framework for Renewable Energy-Aware Workflow Scheduling on Distributed Cloud Data Centers. IEEE Transactions on Parallel and Distributed Systems, 35(4), 604-615.

Jhunjhunwala, P.R., Sombabu, B., & Moharir, S. (2020). Optimal AoI-Aware Scheduling and Cycles in Graphs. IEEE Transactions on Communications, 68(3), 1593-1603.

Kadota, I., Sinha, A., Uysal-Biyikoglu, E., Singh, R., & Modiano, E. (2018). Scheduling Policies for Minimizing Age of Information in Broadcast Wireless Networks. IEEE/ACM Transactions on Networking, 26(6), 2637-2650.

Khan, S.G., Herrmann, G., Lewis, F.L., Tony, P., & Melhuish, C. (2012). Reinforcement learning and optimal adaptive control: An overview and implementation examples. Annual Reviews in Control, 36(1), 42-59.

Li, C., Huang, Y., Li, S., Chen, Y., Jalaian, B.A., & Hou, Y.T. (2021). Minimizing AoI in a 5G-based IoT Network under Varying Channel Conditions. IEEE Internet of Things Journal, 8(19), 14543-14558.

Li, R., Ma, Q., Gong, J., Zhou, Z., & Chen, X. (2021). Age of processing: Age-Driven Status Sampling and Processing Offloading for Edge-Computing-Enabled Real-Time IoT Applications. IEEE Internet of Things Journal, 8(19), 14471- 14484.

Moltafet, M., Leinonen, M., & Codreanu, M. (2020). On the age of information in multi-source queueing models. IEEE Transactions on Communications, 68(8), 5003-5017.

Nie, L., Wang, X., Sun, W., Li, Y., Li, S., & Zhang, P. (2021). Imitation-learning-enabled Vehicular Edge Computing: Toward Online Task Scheduling. IEEE network, 35(3), 102-108.

Pal, S., Jhanjhi, N.Z., Abdulbaqi, A.S., Akila, D., Alsubaei, F.S., & Almazroi, A.A. (2023). An Intelligent Task Scheduling Model for Hybrid Internet of Things and Cloud Environment for Big Data Applications. Sustainability, 15(6), article no. 5104.

Park, B.S., Lee, H., Lee, H.T., Eun, Y., Jeon, D., Zhu, Z., Lee, H., & Jung, Y.C. (2018). Comparison of First-Come First- Served and Optimization Based Scheduling Algorithms for Integrated Departure and Arrival Management. in 2018 Aviation Technology, Integration, and Operations Conference, pp. 3842.

Petrillo, A., Pescapé, A., & Santini, S. (2021). A Secure Adaptive Control for Cooperative Driving of Autonomous Connected Vehicles in the Presence of Heterogeneous Communication Delays and Cyberattacks. IEEE Transactions on Cybernetics, 51(3), 1134-1149.

Qin, Z., Wei, Z., Qu, Y., Zhou, F.H., Wang, H., Ng, D.W.K. (2023). AoI-Aware Scheduling for Air-Ground Collaborative Mobile Edge Computing. IEEE Transactions on Wireless Communications, 22(5), 2989-3005.

Sahni, J., & Vidyarthi, D.P. (2018). A Cost-Effective Deadline-Constrained Dynamic Scheduling Algorithm for Scientific Workflows in a Cloud Environment. IEEE Transactions on Cloud Computing, 6(1), 2-18.

Singh, A.K., Leech, C., Reddy, B.K., Al-Hashimi, B.M., & Merrett, G.V. (2017). Learning-based Run-Time Power and Energy Management of Multi/Many-Core Systems: Current and Future Trends. Journal of Low Power Electronics, 13(3), 310- 325.

Song, J., Gunduz, D., & Choi, W. (2024). Optimal Scheduling Policy for Minimizing Age of Information with A Relay. IEEE Internet of Things Journal, 11(4), 5623-5637.

Tao, Y., Qiu, J., & Lai, S. (2022). A Hybrid Cloud and Edge Control Strategy for Demand Responses Using Deep Reinforcement Learning and Transfer Learning. IEEE Transactions on Cloud Computing, 10(1), 56-71.

Ullah, I., Lim, H.K., Seok, Y.J., & Han, Y.H. (2023). Optimizing Task Offloading and Resource Allocation in Edge-Cloud Networks: A DRL Approach. Journal of Cloud Computing, 12(1), article no. 112.

Wang, B., Liu, F., & Lin, W. (2021). Energy-Efficient VM Scheduling Based on Deep Reinforcement Learning. Future Generation Computer Systems, 125, 616-628.

Wu, H., Zhang, Z., Guan, C., Wolter, K., & Xu, M.X. (2020). Collaborate Edge and Cloud Computing with Distributed Deep Learning for Smart City Internet of Things. IEEE Internet of Things Journal, 7(9), 8099-8110.

Xu, C., Yang, H.H., Wang, X., & Quek, T.Q.S. (2020). Optimizing Information Freshness in Computing-Enabled IoT Networks. IEEE Internet of Things Journal, 7(2), 971-985.

Yates, R.D., Sun, Y., Brown, D.R., Kaul, S.K., Modiano, E., & Ulukus, S. (2021). Age of Information: An Introduction and Survey. IEEE Journal on Selected Areas in Communications, 39(5), 1183-1210.