Hierarchical Federated MADRL for Sustainable Cloud-Edge Resource Orchestration

Saurabh Singhal1, Research Scholar, Lincoln University College, Malaysia pdf.SaurabhSinghal@lincoln.edu.my
Pawan Kumar Verma, Department of Computer Science and Engineering, Symbiosis Institute of Technology, Nagpur Campus, Symbiosis International (Deemed University), Pune, Maharashtra, India abes.pawan@gmail.com

Abstract:

Cloud edge computing is a necessity for latency-sensitive applications such as smart cities and IoT, although efficient resource management is difficult due to dynamic workload and tight SLAs. Existing approaches suffer from scalability problems with the centralisation approach for resource management or the lack of coordination with the decentralisation approach. In this paper, we discuss the literature survey for the sustainable orchestration of resources. Our goal is to propose a framework that maximises the optimisation for task allocation, resource utilisation, SLA compliance, latency, and energy consumption while minimising the carbon footprint.

Keywords: Cloud-Edge Computing; Federated Learning; Multi-Agent Reinforcement Learning; Resource Orchestration; Sustainable Computing.

Introduction

There have been tremendous changes in IoT, AI, and their influence in modern computing. As a result, cloud-edge computing has emerged to satisfy the ever-evolving needs of low-latency and bandwidth requirements for speed and efficient resource usage, especially for autonomous driving and smart cities. Resource management is, however, not an easy feat in the cloud-edge computing system, calling for intelligent solutions in resource management.

Resource orchestration in distributed cloud/edge environments involves improving the performance of applications to minimise delays, maximise resource utilisation, minimise costs, and guarantee SLA compliance [1]. Scheduling mechanisms are ineffective in dynamic environments [2], while Deep Reinforcement Learning (DRL) provides a remedy in the form of autonomous agents [3]. While centralised DRL encounters challenges of scalability and robustness, MADRL overcomes the challenges of scalability and robustness using multiple independent agents in a cloud-edge environment [4].

MASs face problems related to conflicting scheduling and privacy issues. Conflicts may occur during independent learning due to resource conflict. It will be easier to scale and protect privacy if you use federated learning and multi-agent reinforcement learning, but a hierarchical cloud-edge architecture can make things less efficient [5]. Scheduling techniques currently do not address environmental sustainability and growing energy consumption. Sustainable cloud-edge computing focuses on the optimisation of resource utilisation in order to cut down emissions through cost-effective SLA satisfaction.

Related work

Cloud-edge computing is a paradigm shift in latency-critical and highly data-driven applications through the integration of cloud and edge computing, thereby improving QoS by reducing latency and congestion. However, these distributed systems have issues like resource management, task scheduling, energy-aware assignment, SLA agreement, and sustainability. Some recent improvements on methods such as DRL, MARL, and FL have attracted much interest due to their potential in dynamic resource management.

The orchestration of resources plays an important role in cloud-edge computing because it helps in workload allocation between cloud data centers and edge nodes. Scheduling techniques like First Come First Served (FCFS), which were initially developed for centralized systems, fail in dynamic settings and hence face problems such as inefficient resource utilization, high response time, and high costs in cloud-edge computing [1].

Many metaheuristic optimization algorithms, such as Genetic Algorithm have been proposed to solve the problems in optimization. The objective of these algorithms is to optimize several objectives such as makespan, execution costs, energy consumption, and load balancing [1]. Although these algorithms perform better than traditional algorithms, they require more computations due to iterative searches in populations and thus fail to adapt to very dynamic settings. However, new studies stress the significance of AI-based orchestration mechanisms that change their scheduling strategies according to the environmental conditions [2]. The use of machine learning can be instrumental in predicting the demand and performance. Yet, all current AI schedulers are centralized in nature, ignoring the aspects of privacy concerns and distribution of decision-making.

Deep Reinforcement Learning (DRL) [3] serves as a useful technique for managing cloud resources in a smarter way by enabling autonomous agents to develop policies of optimal scheduling by interacting with the environment without the requirement of a labeled dataset. The advantage of DRL lies in its capability to deal with dynamic cloud-edge environments, which can be done through maximization of reward based on scheduling decisions. The schedulers optimize many quality of service parameters including response time, throughput, resource utilization, and energy efficiency.

Although Deep Reinforcement Learning (DRL) has many benefits in resource management, a major limitation lies in the use of a centralized framework for training models [4]. It reduces scalability since the number of edge nodes and IoT devices is increasing rapidly. In addition, centralized learning introduces the risk of failure in case of single point failure in addition to heavy communication overhead due to data transmission from many edge nodes to central servers.

MARL makes Deep Reinforcement Learning (DRL) more advanced by employing several agents that collaborate or compete to improve the overall performance of the system [5]. The decentralized nature of the decision-making process in MARL matches the architecture of the cloud-edge systems. In this way, it allows for the semi-autonomous function of the distributed edge clusters in order to provide global services. The current frameworks that use MARL for resource management have specialized agents for cloud data centers, edge servers, or pools of virtual resources for performing functions like task assignment and load balancing. Such agents perform cooperative learning and increase the system's scalability, fault tolerance, and adaptability to the changes in the load. MARL can be used for making local decisions based on the availability of the resources and applications.

There are several disadvantages of the current MARL approaches, which include the presence of conflicts in scheduling, ineffective resource usage, and unstable learning [5]. They always take into account ideal communication between the agents without accounting for possible lags and costs. Most of the models only consider traditional criteria of evaluation like latency and cost without taking environmental sustainability and CO_2 emissions into account.

Federated learning, a privacy-preserving distributed machine learning approach, allows participants to train local models without sharing raw data [6]. Model parameters are sent to a centralized server for aggregation, creating a global model that participants can utilize without compromising sensitive information. In cloud-edge computing, it is applied in various areas such as workload prediction and resource allocation, adhering to privacy regulations and minimizing communication overhead. Researchers are also integrating federated learning with reinforcement learning to enhance distributed scheduling policies [7]. However, existing frameworks typically rely on inefficient flat aggregation, leading to communication issues and slow convergence, with a lack of hierarchical aggregation strategies to improve scalability in larger infrastructures.

Sustainability has become imperative in the growing field of cloud computing owing to the huge consumption of energy by data centers [8]. Due to the increasing presence of applications based on artificial intelligence and the Internet of Things, there has been an increase in energy consumption, thus making it imperative to reduce carbon emissions in the context of resource scheduling. Recent works have presented energy-aware scheduling algorithms in which the concept of consolidation and low power modes is used, while renewable energy-aware scheduling takes into consideration data centers that make use of either solar or wind energy. On the other hand, carbon-aware scheduling makes use of the carbon intensity of regions in order to allocate workloads. But many of the existing systems lack SLA satisfaction, latency, or distributed learning [9].

Although there have been improvements in the field of intelligent cloud-edge resource orchestration, as indicated in the literature review, there is still room for more research to be done. Current algorithms used for scheduling are not flexible enough in dynamic conditions; centralized reinforcement learning approaches face scalability problems. Although multi-agent reinforcement learning approaches allow for decentralized decision-making, they lack coordination and hierarchy. The federated learning approach is a suitable method for protecting data privacy; however, the architecture employed is flat in most cases when applied to large-scale environments. In addition, current solutions focus on optimizing either SLAs or energy efficiency without a carbon-aware perspective. Table 1 summarizes the different approach highlighting their strengths, limitations and gap.

Table 1. Summary of different approaches, strengths, limitations and gap.

Approach	Strengths	Limitations	Gap
Traditional Scheduling (FCFS, Heuristics, Metaheuristics)	Simple implementation and improved resource allocation	Poor adaptability to dynamic workloads; high computational overhead; limited scalability	Intelligent adaptive scheduling for dynamic cloud-edge environments

Deep Reinforcement Learning (DRL)	Learns optimal scheduling policies and improves QoS	Centralized training causes scalability issues, communication overhead, and single-point failure	Distributed and scalable learning architecture
Multi-Agent Reinforcement Learning (MARL)	Decentralized decision-making, improved adaptability and fault tolerance	Agent conflicts, unstable coordination, and lack of hierarchical control	Coordinated hierarchical multi-agent orchestration
Federated Learning (FL)	Privacy-preserving collaborative learning without sharing raw data	Flat aggregation limits scalability and slows convergence	Hierarchical federated learning for large-scale cloud-edge systems
Sustainable Resource Scheduling	Reduces energy consumption and carbon emissions	Most methods optimize only energy or carbon without jointly considering SLA, latency, and resource utilization	Multi-objective sustainable resource orchestration

Key Contribution

Some important contributions of the article include a detailed literature review on resource orchestration approaches in cloud-edge computing, gap identification related to scalability and privacy, and the development of the Hierarchical Federated Multi-Agent Deep Reinforcement Learning (HFMADRL) approach to resource orchestration. Besides, it also provides a multi-objective approach to optimize task allocation and sustainability, and a conceptual framework for intelligent and scalable resource orchestration in cloud-edge environments for IoT and smart cities.

Conclusion

The paper presents a literature review of studies conducted in sustainable orchestration of resources in cloud-edge computing for latency sensitive services including those provided by smart cities and Internet of Things. It highlights the unsolved problems in both centralized and decentralized models of resource management. Centralized resource management suffers from issues of scalability and singleness of failure point whereas decentralized models suffer from lack of coordination. Current studies mostly concentrate on performance criteria and overlook sustainability aspects such as energy savings. In future a sustainable model of cloud-edge orchestration will be proposed that will optimize task scheduling, resource utilization, service-level agreement, latency, and energy usage, reducing carbon footprint.

References

1. Bhaskaran, Sreebha, and Supriya Muthuraman. "A comprehensive study of resource provisioning and optimization in edge computing." *Computers, Materials & Continua* 83.3 (2025): 5037-5070.

2. Zhang, Yi, et al. "Dynamic job shop scheduling based on deep reinforcement learning for multi-agent manufacturing systems." *Robotics and Computer-Integrated Manufacturing* 78 (2022): 102412.
3. Jalali Khalil Abadi, Zahra, Najme Mansouri, and Mohammad Masoud Javidi. "Deep reinforcement learning-based scheduling in distributed systems: a critical review." *Knowledge and Information Systems* 66.10 (2024): 5709-5782.
4. Hu, Fangxiang, et al. "A Multi-Agent Deep Reinforcement Learning-Based Task Offloading Method for 6G-Enabled Internet of Vehicles with Cloud-Edge-Device Collaboration." *Computers, Materials & Continua* 87.3 (2026).
5. Oroojlooy, Afshin, and Davood Hajinezhad. "A review of cooperative multi-agent deep reinforcement learning." *Applied Intelligence* 53.11 (2023): 13677-13722.
6. Patil, Rajat, Ribhav Yadav, and Kaushik Mishra. "Federated Learning in Vehicular Networks for Resource Management, Privacy, and Security: Challenges and Future Directions." *IEEE Access* 13 (2025): 190766-190785.
7. Mali, Saroj, et al. "Federated reinforcement learning-based dynamic resource allocation and task scheduling in edge for IoT applications." *Sensors* 25.7 (2025): 2197.
8. Buyya, Rajkumar, Shashikant Ilager, and Patricia Arroba. "Energy-efficiency and sustainability in new generation cloud computing: a vision and directions for integrated management of data centre resources and workloads." *Software: Practice and Experience* 54.1 (2024): 24-38.
9. Al-Shammaa, Ameer A., and Doaa H. Al-Hadrawi. "Carbon-aware federated orchestration in edge-cloud systems: A multi-objective optimization framework with dynamic spatio-temporal workload management." *Sustainable Computing: Informatics and Systems* (2026): 101398.