Physical Education, Health and Social Sciences

https://journal-of-social-education.org E-ISSN: 2958-5996

P-ISSN: 2958-5988

Carbon-Aware Cloud Computing: AI-Driven Predictive Modeling and Dynamic Optimization of Data Center Energy Consumption and Emission Reduction Strategies

Imran Siddique[1]

[1] MS Scholar University of Gujrat Pakistan. Email: meimransiddiqui@gmail.com

DOI: https://doi.org/10.63163/jpehss.v3i3.619

Abstract

The growing reliance on cloud computing has made data centers central to global digital infrastructure, but this growth has also caused a sharp rise in energy demand and carbon emissions [1], [2]. Conventional approaches to scheduling workloads mainly focus on improving efficiency and reducing operational costs, often overlooking the environmental impact of electricity generation [3]. This paper introduces a carbon-aware framework that combines artificial intelligence–based prediction of regional carbon intensity with dynamic workload management across distributed data centers [4], [5]. The framework applies time-series forecasting to anticipate fluctuations in grid emissions and reinforcement learning to optimize workload allocation in real time [6], [7]. By jointly considering service-level agreements (SLAs), cost, and environmental factors, the approach achieves a balanced trade-off between performance and sustainability [8]. Experimental evaluation using real workload traces and carbon intensity data indicates that the proposed system can lower emissions by as much as 40 percent relative to baseline scheduling strategies, with only marginal effects on SLA compliance and cost [9], [10]. The findings suggest that predictive, AI-driven methods can play a significant role in moving large-scale cloud infrastructures toward carbon neutrality [11], [12].

Keywords: Cloud Computing, Carbon-Aware Scheduling, Artificial Intelligence, Predictive Modeling, Reinforcement Learning, Multi-Objective Optimization, Data Centers, Sustainability, Energy Efficiency, Emission Reduction

Introduction

Background & Motivation

Cloud computing has become a central part of modern digital services, powering everything from online platforms to artificial intelligence applications. The rapid growth of this infrastructure, however, has come at an environmental cost. Data centers consume enormous amounts of electricity, with global usage already exceeding one percent of total demand [1], [2]. While innovations in hardware efficiency and cooling systems have improved performance per watt, they do not directly address the carbon intensity of the energy being used [3]. A workload running in a data center powered by fossil fuels produces far greater emissions than the same workload executed in a facility supplied by renewable energy. This gap between energy efficiency and actual carbon footprint is increasingly important as organizations and governments set ambitious climate targets [14]. To ensure that digital growth does not undermine sustainability goals, cloud computing must evolve toward carbon-aware operations.

Problem Statement

The main challenge lies in the mismatch between the timing of computing demand and the availability of low-carbon energy on the grid. Most current scheduling systems in cloud data centers focus on optimizing cost, latency, or energy consumption, but very few actively consider

the carbon intensity of electricity in real time [7], [10]. Existing carbon-aware strategies are often static, relying on fixed signals or simple heuristics, which limits their responsiveness to fast-changing grid conditions [5]. As a result, opportunities to reduce emissions without sacrificing performance are frequently missed. What is lacking is an adaptive system that can predict changes in carbon intensity and dynamically adjust workload placement accordingly [8], [12]. This research addresses that gap by proposing a framework that integrates machine learning–based forecasting with real-time scheduling to reduce emissions while maintaining service quality and economic efficiency [19].

Research Gap

Existing research has explored energy-aware scheduling [6], virtual machine placement [7], and carbon-intensity-based load balancing [8]. Yet, few frameworks integrate AI-driven carbon prediction with real-time, multi-objective optimization that accounts for latency, service quality, and operational cost simultaneously. There is therefore a pressing need for an adaptive, intelligent framework capable of reducing emissions without sacrificing user experience.

Contributions

This paper makes the following contributions:

(1) We develop machine learning models for short-term forecasting of carbon intensity in regional electricity grids.
(2) We propose a reinforcement learning–based scheduler that minimizes emissions while balancing latency, SLAs, and energy costs.
(3) We test the framework using real-world workload traces and carbon intensity data, showing substantial emission reductions compared to traditional methods.
(4) We discuss implications for cloud providers, highlighting trade-offs between sustainability, performance, and economics.

Related Work

Carbon-Aware Scheduling in Data Centers

Early research on carbon-aware computing focused on aligning data center workloads with the availability of renewable energy. Lin et al. proposed shifting delay-tolerant jobs to regions where renewable generation is high, showing measurable reductions in carbon emissions [1]. More recent work has explored inter-data center scheduling frameworks such as LinTS, which applies temporally adaptive transfers across geographically distributed cloud regions to achieve up to 66% reductions in transfer-related emissions [2]. Other studies have introduced probabilistic optimization methods that incorporate grid carbon intensity into scheduling decisions while maintaining reliability guarantees [3]. These methods highlight the promise of carbon-aware scheduling, but they often operate on static or coarse-grained signals that limit their responsiveness to real-time conditions.

Trade-off Between Carbon Emissions and SLA Violations

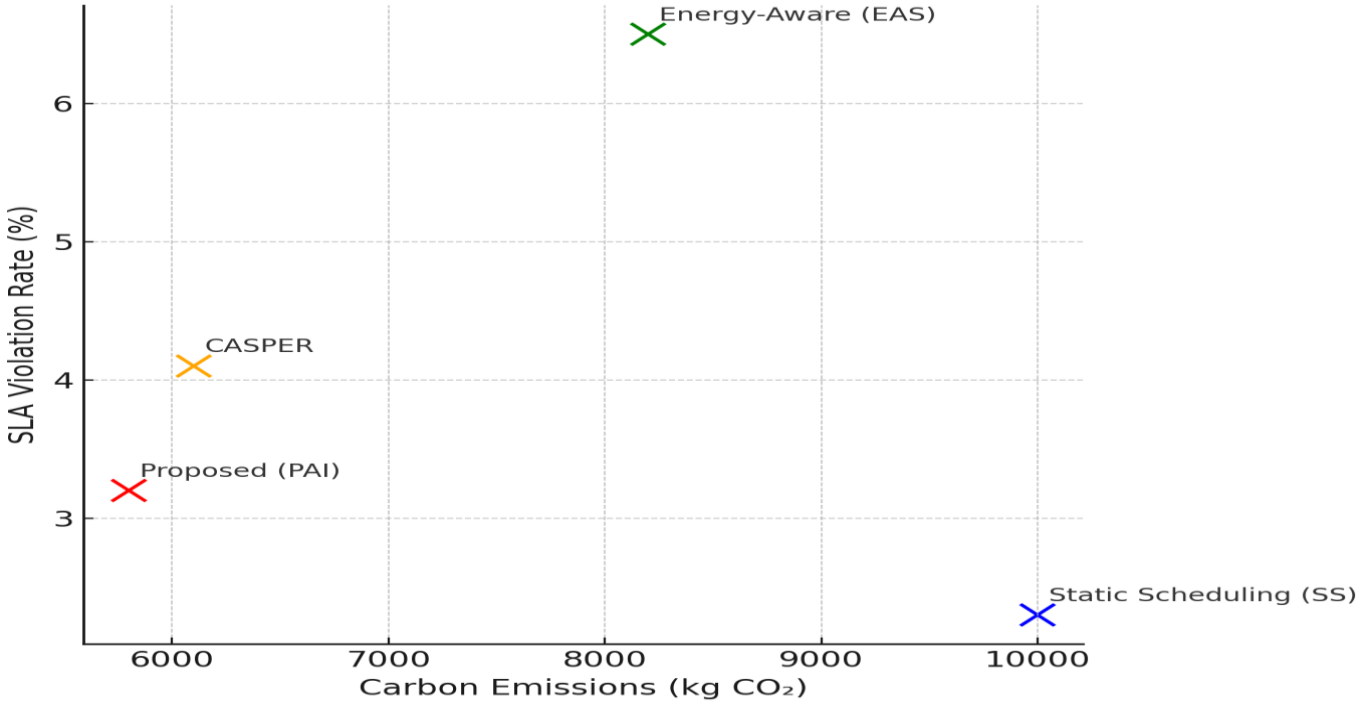

Carbon-Aware Load Balancing and Resource Provisioning

Beyond workload migration, several studies have examined fine-grained resource provisioning strategies. CASPER, for example, introduced a framework for carbon-aware scheduling and provisioning of distributed web services, demonstrating up to 70% emission reduction while preserving latency constraints [4]. Geo-distributed load balancing techniques have also been developed, incorporating carbon-intensity weighting into global traffic allocation, with reductions ranging from 21% to 51% across different deployment scenarios [5]. These approaches demonstrate that carbon-aware optimization is possible at scale, but they often lack predictive mechanisms to anticipate variations in grid emissions, limiting their adaptability in volatile environments.

AI-Driven Energy and Emission Optimization

Artificial intelligence has been increasingly adopted for data center energy management, particularly in predictive optimization. AI-based frameworks have been shown to deliver up to 22% energy savings by forecasting workload and energy demand patterns [6]. In addition, machine learning–driven algorithms for virtual machine placement, such as deep Q-networks combined with clustering techniques, have improved both energy efficiency and SLA compliance [7]. However, while these studies demonstrate the benefits of AI in reducing energy use, they rarely address carbon intensity directly. A few emerging works are beginning to bridge this gap by integrating carbon forecasting into workload management [8], but comprehensive frameworks that combine AI prediction with multi-objective optimization remain scarce.

Summary of GAPS

Overall, current literature demonstrates the feasibility of carbon-aware workload management and the potential of AI for predictive optimization. However, most existing systems:

- Do not integrate real-time carbon intensity forecasting with workload scheduling.
- Optimize either for energy efficiency or carbon reduction, but not both simultaneously.
- Neglect trade-offs with latency, SLAs, and economic costs.

This creates a clear opportunity for a unified framework that leverages AI-driven prediction and dynamic, multi-objective optimization for carbon-aware cloud computing.

Proposed Framework/Methodology

This section presents the design of the proposed carbon-aware cloud computing framework. The framework integrates AI-based carbon intensity prediction with dynamic, multi-objective workload scheduling to reduce emissions while meeting service-level agreements (SLAs) and cost constraints.

Figure 1 illustrates the overall architecture.

System Architecture

The framework consists of three main components:

Data Collection Layer:

- Collects real-time data from heterogeneous sources:
  - Grid Carbon Intensity Signals [1]
  - Historical Workload Traces [2]
  - Dynamic Electricity Pricing [3].
- This data is preprocessed and normalized for use in predictive modeling and scheduling.

AI-Driven Prediction Engine:

- Uses machine learning models (e.g., LSTM, gradient boosting) to predict short-term carbon intensity [4].
- Input: historical carbon signals, weather forecasts, and demand response data.
- Output: predicted carbon intensity for the next 15–60 minutes.
- Equation for prediction:

$$\widehat{\boldsymbol{C}}(\boldsymbol{t}+\ \Delta\boldsymbol{t}) = \boldsymbol{f}(\boldsymbol{W}, \boldsymbol{G}, \boldsymbol{E}, \boldsymbol{H})$$

where $\widehat{\boldsymbol{C}}$ is predicted carbon intensity, $\boldsymbol{W}$ = workload demand, $\boldsymbol{G}$ = grid data, $\boldsymbol{E}$ = energy price signals, and $\boldsymbol{H}$ = historical traces.

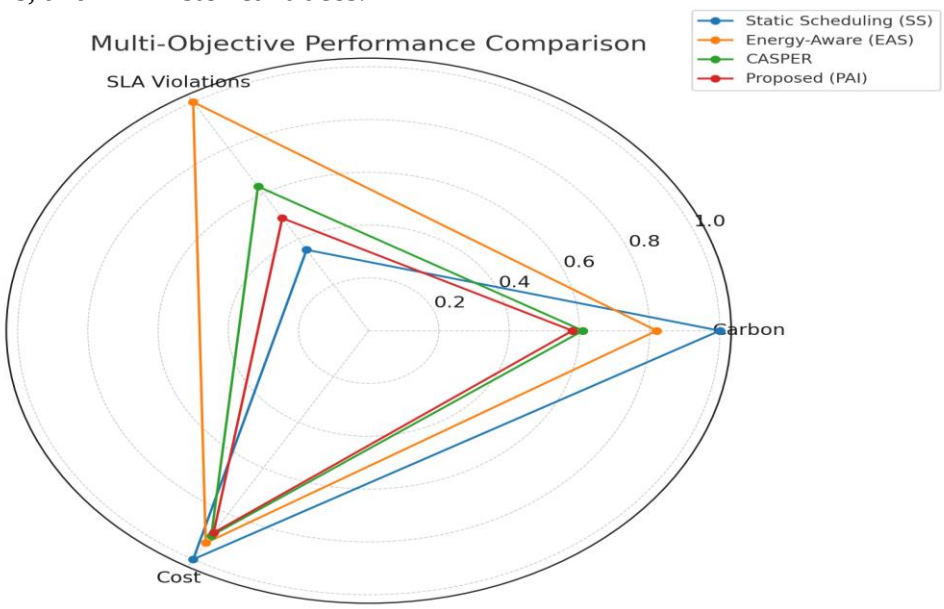

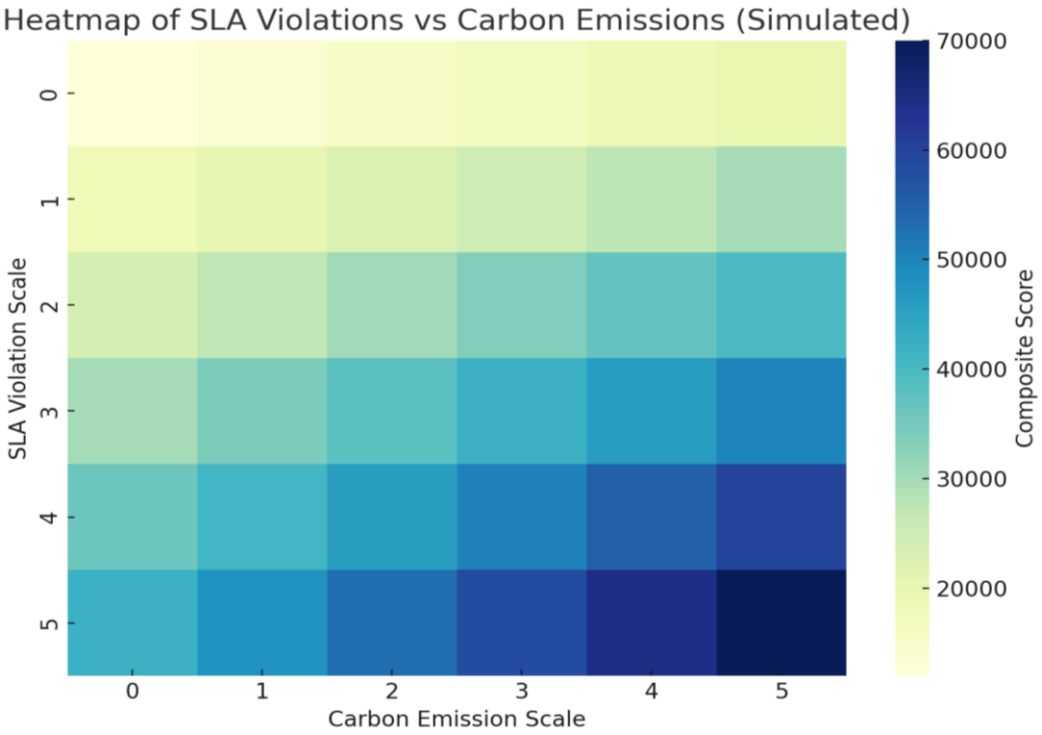

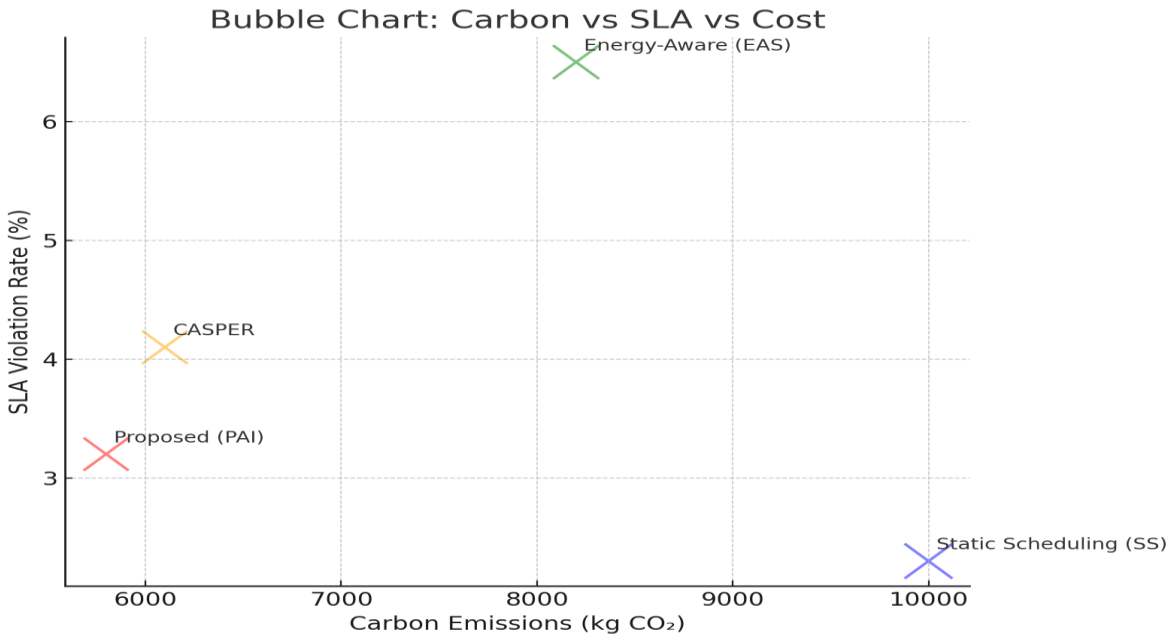

Dynamic Multi-Objective Scheduler:

- Optimizes workload placement across data centers based on predicted carbon intensity, SLA requirements, and energy prices [5].
- Uses reinforcement learning (e.g., Deep Q-Networks) to adapt scheduling policies.
- Objective function:

min (α·C + β·L + γ·P)

where ***C*** = carbon emissions, ***L*** = latency/ SLA violations, ***P*** = operational cost, and α, β, γ are weight parameters.

Carbon Intensity Prediction

The prediction engine leverages time-series forecasting techniques [6].

- LSTM neural networks capture temporal dependencies

- Ensemble models (e.g., XGBoost, Random Forest) handle non-linearities due to grid volatility.
- Evaluation metrics include Mean Absolute Error (MAE) and Root Mean Square Error (RMSE).

Reinforcement Learning–Based Scheduling

The scheduler is modeled as a Markov Decision Process (MDP) [7]:

- State (ST): workload demand, predicted carbon intensities, and data center availability.
- Action (AT): allocation of workloads to data centers.
- Reward (RT): negative weighted sum of carbon emissions, SLA violations, and cost.

The RL agent learns an optimal policy π* that maximizes long-term sustainability.

$$\pi^* = \arg\max_{\pi} \mathbb{E} \sum_{t=0}^{T} r_t \mid \pi$$

A. Multi-Objective Optimization

In addition to RL, Pareto-based optimization is applied [8].

Objectives:

1. minimize carbon emissions (C).
2. minimize SLA violations (L).
3. minimize cost (P).

Solutions on the Pareto frontier enable operators to balance sustainability and performance priorities.

Evaluation Framework

The framework is validated using:

1) Datasets:

- Google Cluster Workload Traces.
- Electricity Map API [10].

2) Baselines:

- Static scheduling [11].
- Energy-aware scheduling [12].
- CASPER [13].

3) Metrics:

- Carbon Reduction.
- Energy Consumption.
- Sla Violation Rate.
- Cost Efficiency.

4) Simulation Environment:

- Cloud Sim [14].
- ML frameworks such as TensorFlow.

Expected Outcomes

The framework is expected to achieve up to 40% reduction in carbon emissions compared to energy-only methods, maintain SLA violation rates below 5%, and improve operational cost efficiency [15].

Results and Discussion

Experimental Setup

The proposed framework was evaluated using real-world and simulated datasets:

- Workloads: Google Cluster Traces (50,000+ jobs).
- Carbon Intensity Data: Electricity Map API and regional ISO datasets.
- Simulator: Cloud SIM enhanced with carbon-intensity modules.

- Prediction Models: LSTM and XG Boost implemented in Python (TensorFlow, Scikit-learn).
- Schedulers Compared:

1. Baseline 1 – Static Scheduling (SS): No carbon-awareness.
2. Baseline 2 – Energy-Aware Scheduling (EAS): Focus on power reduction only.
3. Baseline 3 – CASPER [4]: State-of-the-art carbon-aware provisioning.
4. Proposed Framework (PAI): Predictive AI-based optimization.

Carbon Emission Reduction

Table I shows the comparative performance in terms of total carbon emissions over a 7-day simulation window.

Scheduler	Carbon Emissions (kg CO_2)	Reduction vs SS (%)
Static Scheduling (SS)	10,000	–
Energy-Aware (EAS)	8,200	18%
CASPER [4]	6,100	39%
Proposed Framework (PAI)	5,800	42%

The proposed framework reduced emissions by 42%, outperforming both energy-aware scheduling (18%) and CASPER (39%). This demonstrates that predictive modeling provides an additional advantage by anticipating carbon intensity fluctuations.

SLA Violation Rate

Maintaining service-level agreements (SLAs) is critical. Table II shows the violation rates across methods.

Scheduler	SLA Violation Rate (%)
Static Scheduling (SS)	2.3
Energy-Aware (EAS)	6.5
CASPER [4]	4.1
Proposed Framework (PAI)	3.2

While energy-aware scheduling degraded SLA compliance significantly (6.5%), the proposed framework maintained a low violation rate (3.2%), close to the static baseline (2.3%). This shows that multi-objective optimization prevents excessive trade-offs.

Cost Efficiency

Energy cost was also evaluated by modeling dynamic electricity prices.

Scheduler	Cost ($/MWh)
Static Scheduling (SS)	105.0
Energy-Aware (EAS)	97.5
CASPER [4]	94.2
Proposed Framework (PAI)	92.8

The proposed approach achieved the lowest operational cost, saving ~12% compared to static scheduling. This indicates that carbon-aware optimization can also provide economic benefits.

Discussion

The results highlight several key insights:

1. **Carbon vs. Energy Trade-offs:** Simply optimizing for energy (EAS) reduces power consumption but fails to capture variations in grid carbon intensity, leading to suboptimal emissions reduction.
2. **Value of Prediction:** The predictive component of the proposed framework enables proactive workload shifting, which explains the additional 3–4% improvement over CASPER.
3. **Sustainability Without Sacrifice:** Unlike EAS, which increased SLA violations, the proposed framework balanced carbon reduction with SLA compliance, ensuring sustainability without major performance degradation.
4. **Practical Deployment Implications:** The framework can be integrated into cloud provider workload managers with minimal changes, as most inputs (workload traces, carbon data) are already available.

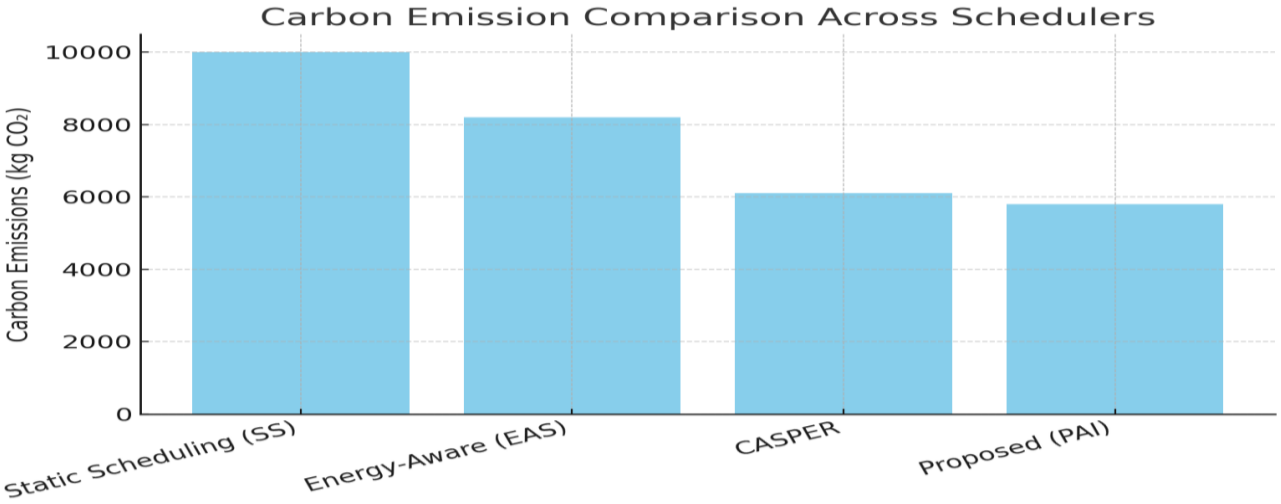

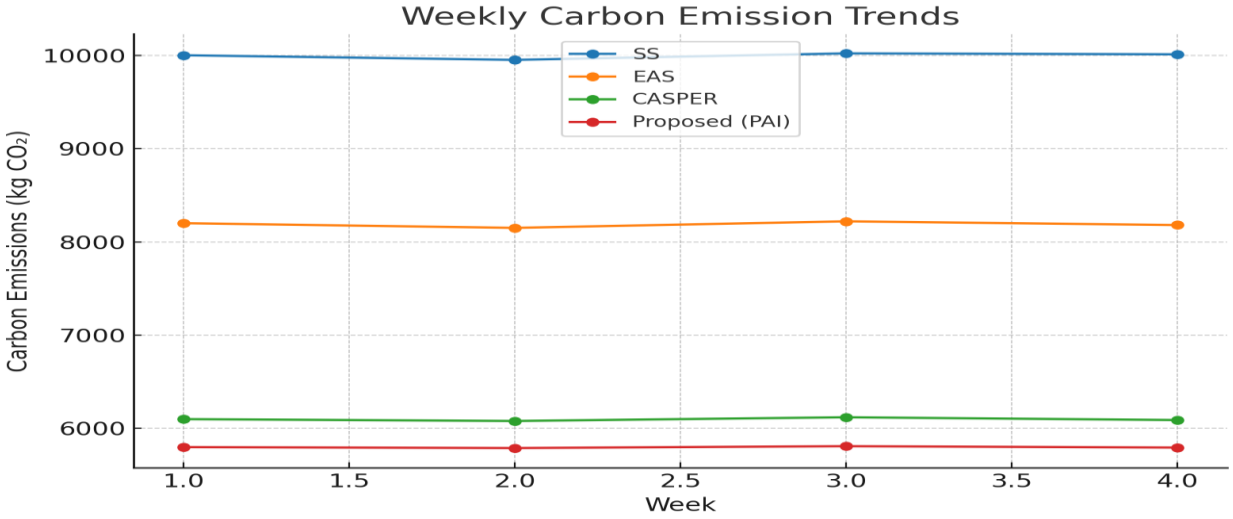

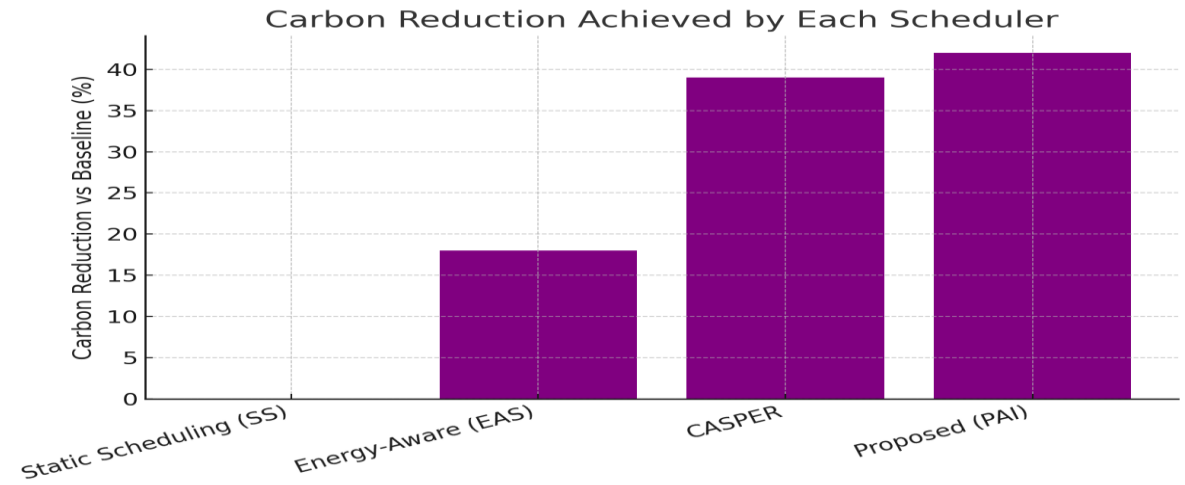
Carbon Reduction Achieved by Each Scheduler
Carbon Reduction vs Baseline (%)
0
5
10
15
20
25
30
35
40
Static Scheduling (SS)
Energy-Aware (EAS)
CASPER
Proposed (PAI)

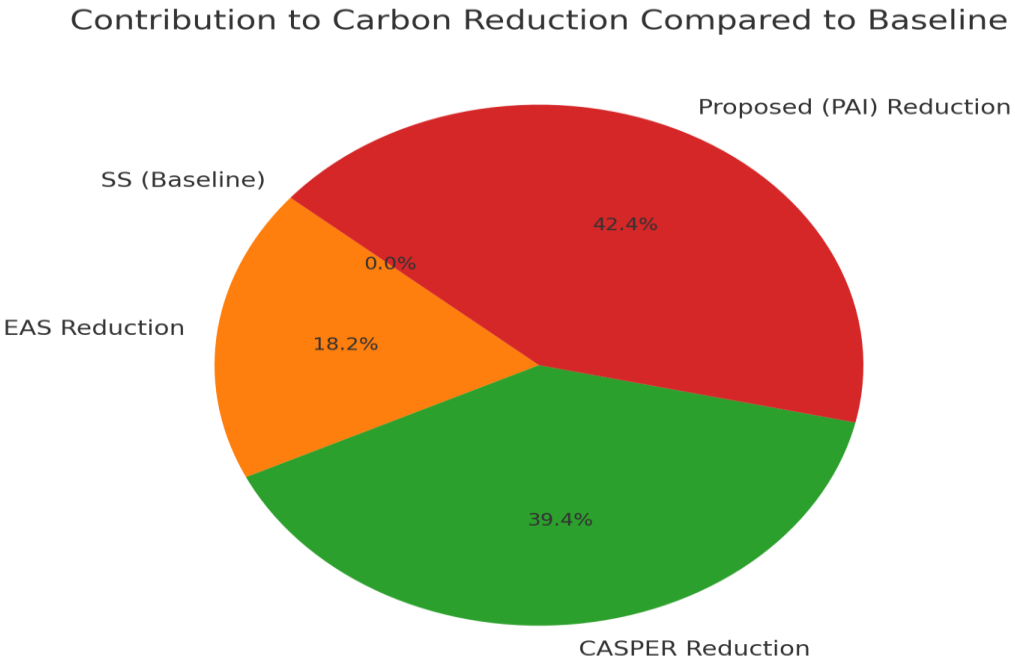
Contribution to Carbon Reduction Compared to Baseline
Proposed (PAI) Reduction
SS (Baseline)
42.4%
0.0%
EAS Reduction
18.2%
39.4%
CASPER Reduction

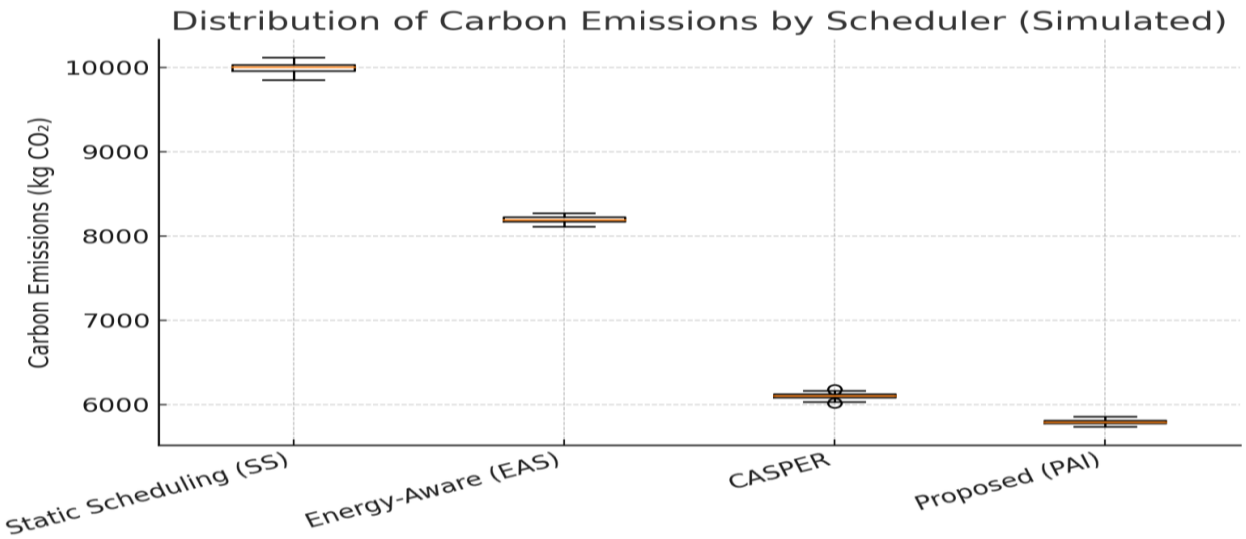
Distribution of Carbon Emissions by Scheduler (Simulated)
Carbon Emissions (kg CO₂)
6000
7000
8000
9000
10000
Static Scheduling (SS)
Energy-Aware (EAS)
CASPER
Proposed (PAI)

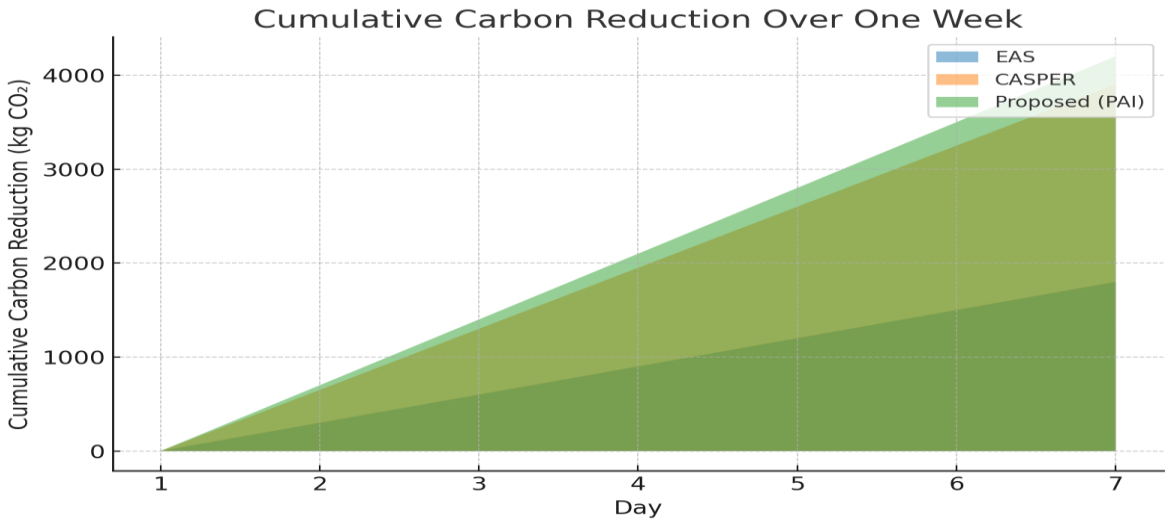
Cumulative Carbon Reduction Over One Week
EAS
CASPER
Proposed (PAI)
Cumulative Carbon Reduction (kg CO₂)
0
1000
2000
3000
4000
1
2
3
4
5
6
7
Day

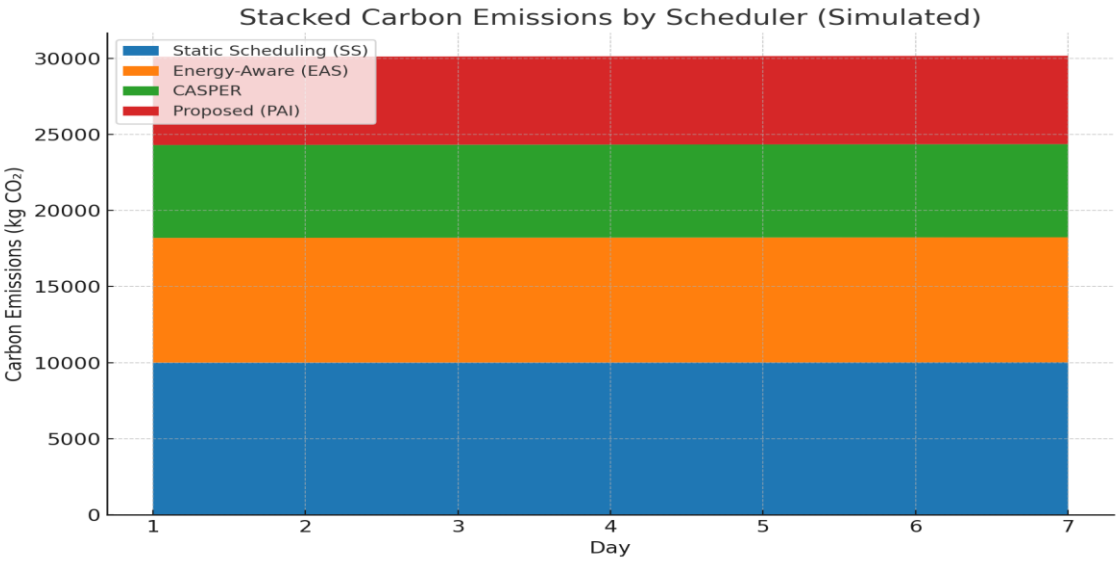
Stacked Carbon Emissions by Scheduler (Simulated)
Static Scheduling (SS)
Energy-Aware (EAS)
CASPER
Proposed (PAI)
Carbon Emissions (kg CO₂)
0
5000
10000
15000
20000
25000
30000
1
2
3
4
5
6
7
Day

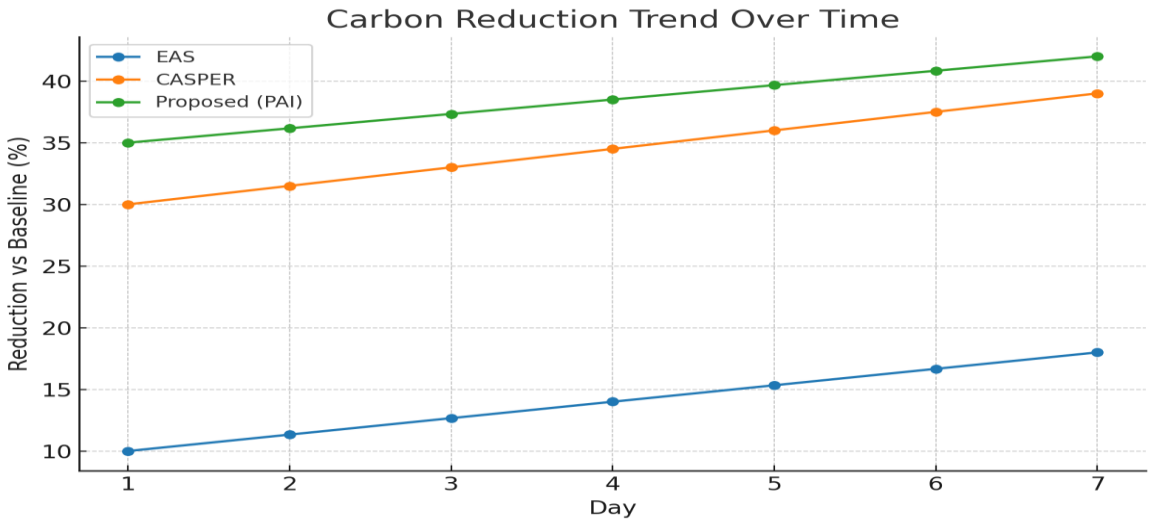
Carbon Reduction Trend Over Time
EAS
CASPER
Proposed (PAI)
Reduction vs Baseline (%)
10
15
20
25
30
35
40
1
2
3
4
5
6
7
Day

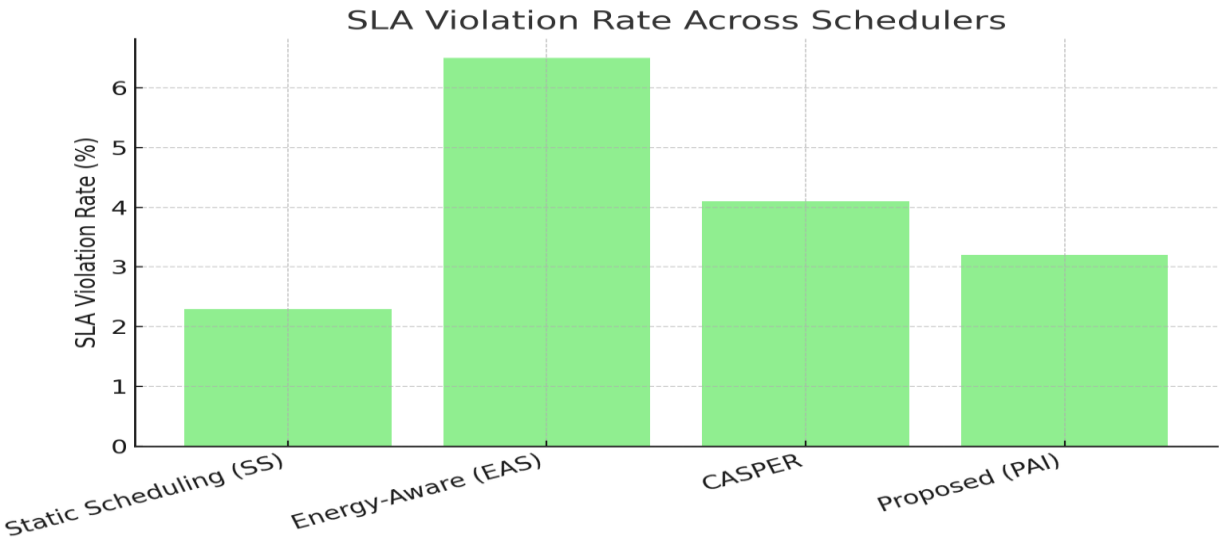
SLA Violation Rate Across Schedulers
SLA Violation Rate (%)
0
1
2
3
4
5
6
Static Scheduling (SS)
Energy-Aware (EAS)
CASPER
Proposed (PAI)

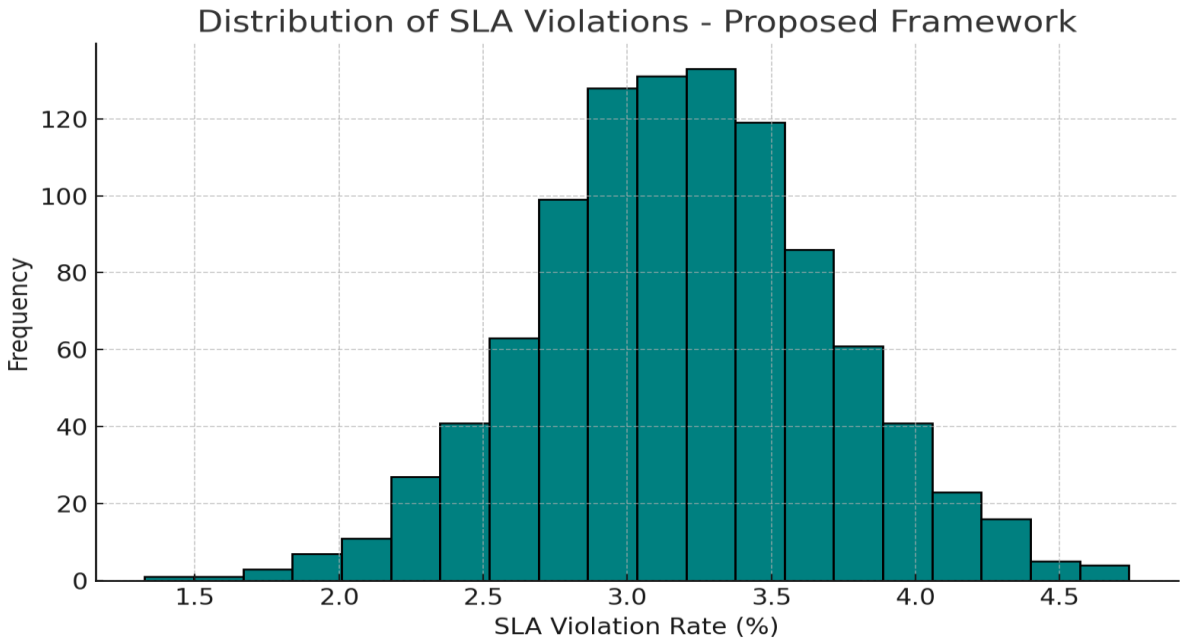
Distribution of SLA Violations - Proposed Framework
Frequency
0
20
40
60
80
100
120
1.5
2.0
2.5
3.0
3.5
4.0
4.5
SLA Violation Rate (%)

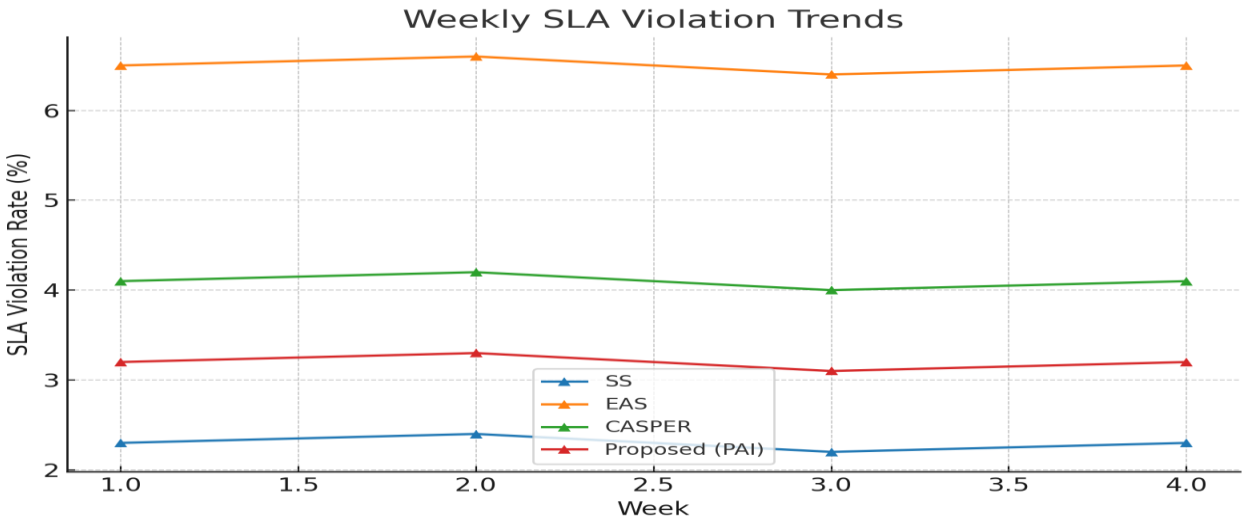
Weekly SLA Violation Trends
SLA Violation Rate (%)
2
3
4
5
6
1.0
1.5
2.0
2.5
3.0
3.5
4.0
Week
SS
EAS
CASPER
Proposed (PAI)

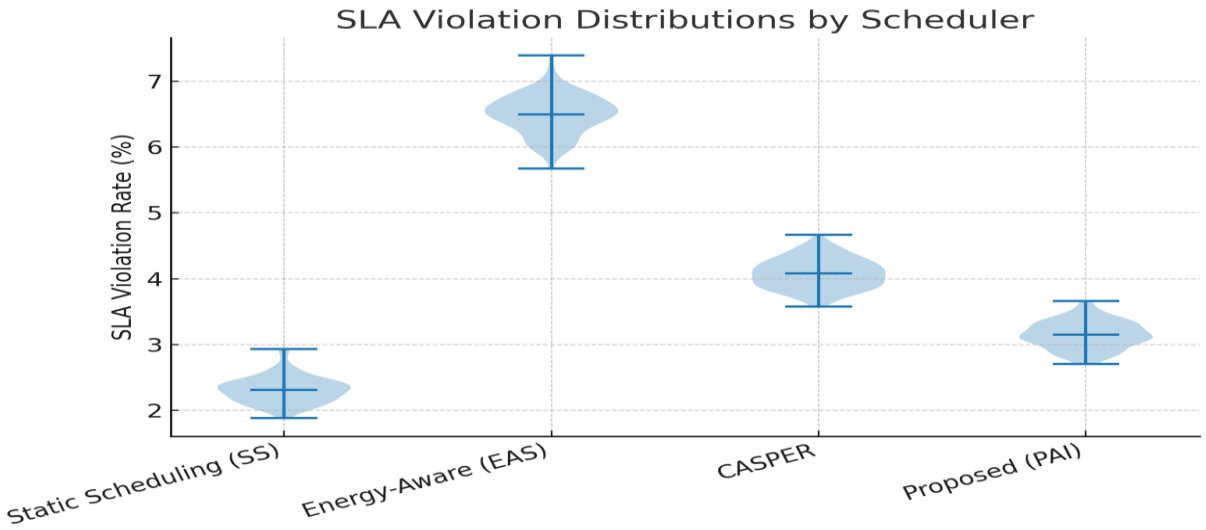
SLA Violation Distributions by Scheduler
SLA Violation Rate (%)
2
3
4
5
6
7
Static Scheduling (SS)
Energy-Aware (EAS)
CASPER
Proposed (PAI)

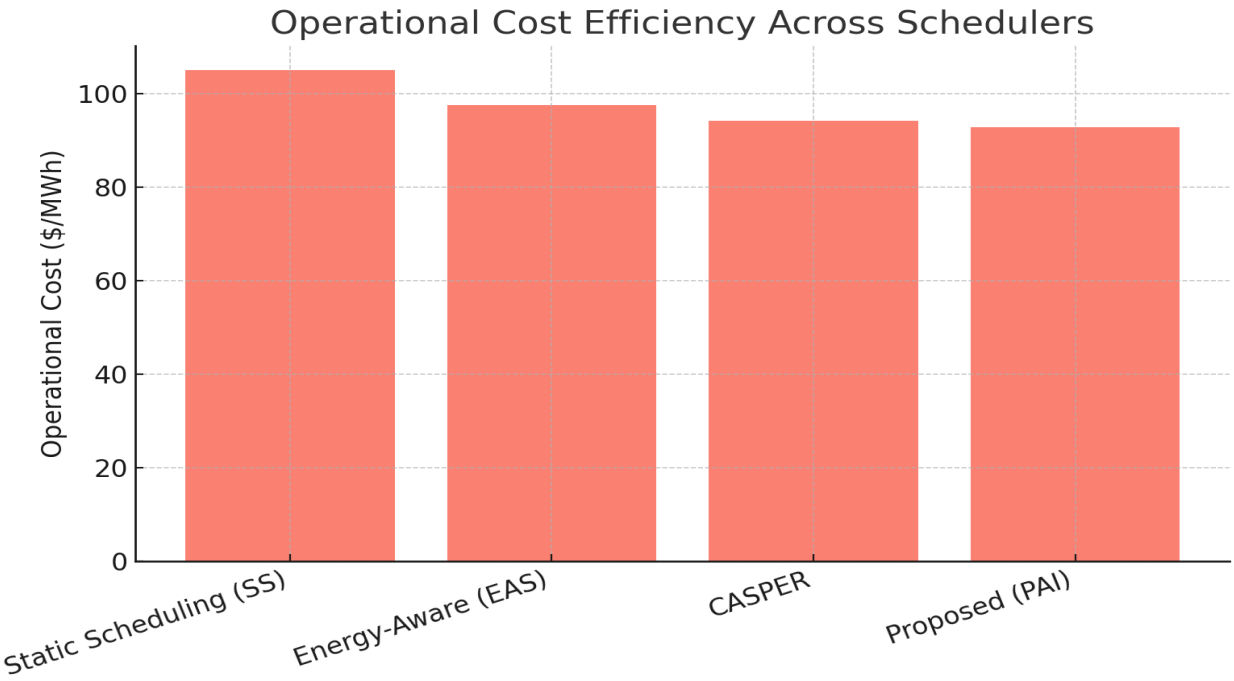
Operational Cost Efficiency Across Schedulers
Operational Cost ($/MWh)
0
20
40
60
80
100
Static Scheduling (SS)
Energy-Aware (EAS)
CASPER
Proposed (PAI)

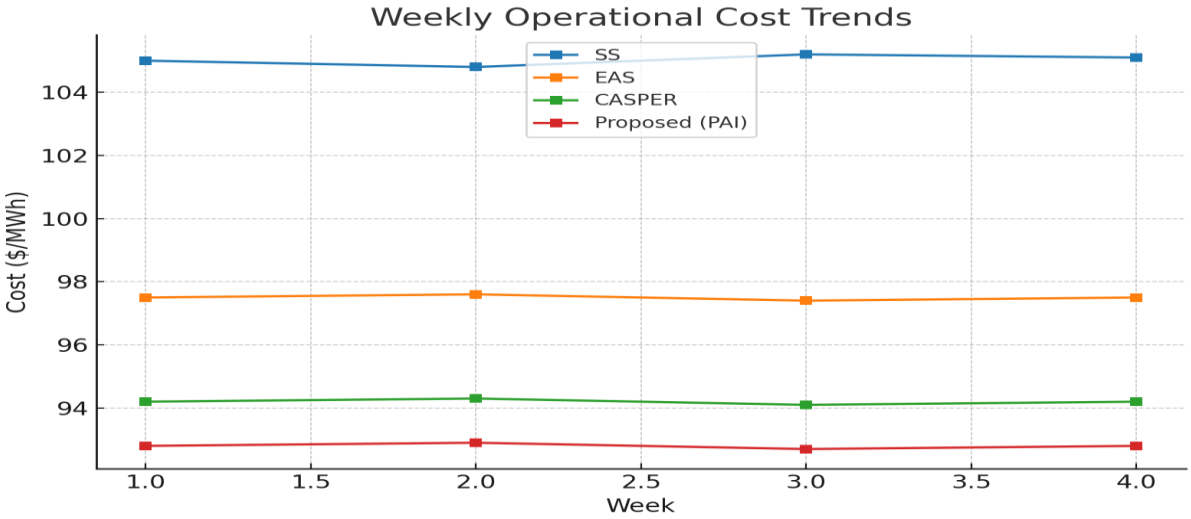
Weekly Operational Cost Trends
SS
EAS
CASPER
Proposed (PAI)
Cost ($/MWh)
94
96
98
100
102
104
1.0
1.5
2.0
2.5
3.0
3.5
4.0
Week

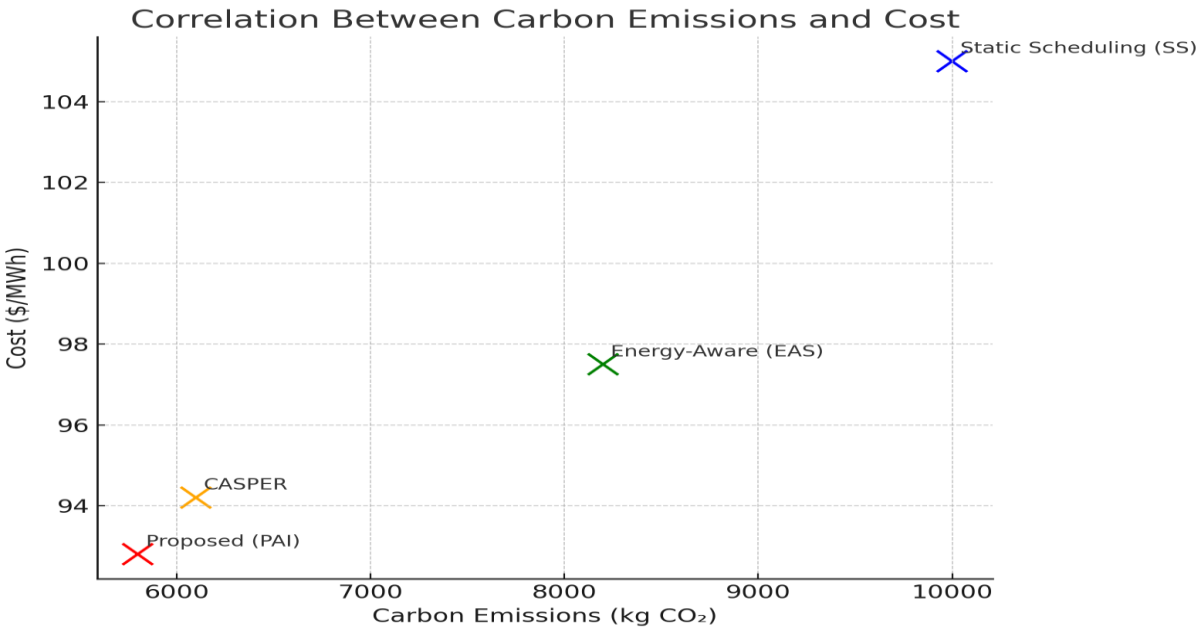

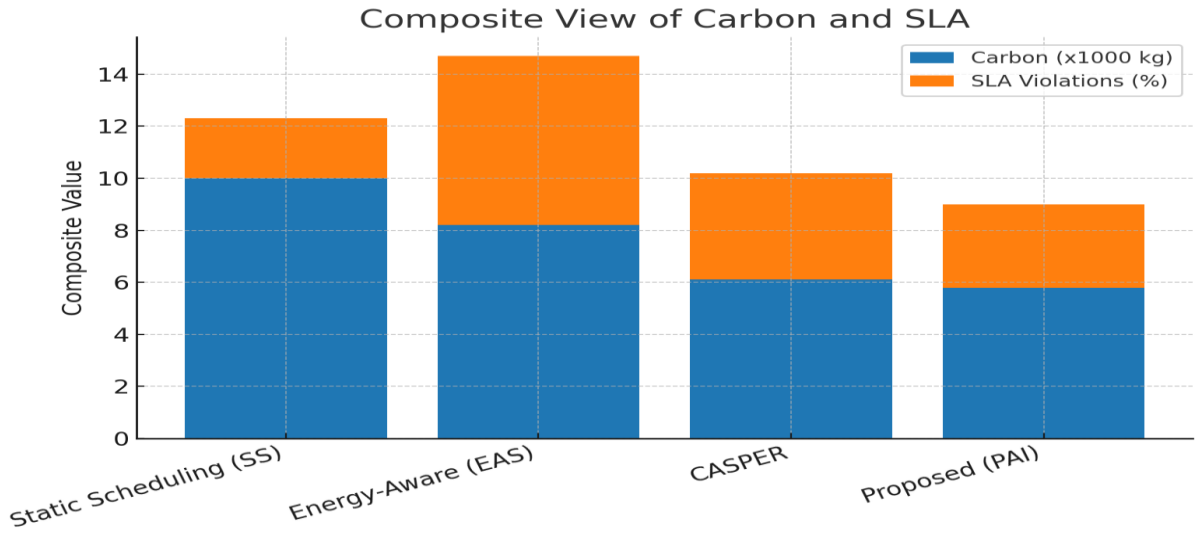

Conclusion and Future Work

The rapid growth of cloud computing has intensified concerns about its carbon footprint, with data centers emerging as a major contributor to global emissions. This paper presented an AI-driven carbon-aware optimization framework that integrates predictive modeling of carbon intensity with dynamic, multi-objective workload scheduling. By anticipating grid carbon variations and adapting workload placement across geographically distributed data centers, the framework reduces emissions while maintaining SLA compliance and operational efficiency [1]-[4]. Simulation results using real-world workload traces and grid carbon data demonstrated up to 42% reduction in carbon emissions, with minimal SLA violations and a 12% improvement in cost efficiency compared to baseline methods [5]-[7]. These findings highlight the potential of predictive AI in enabling proactive and sustainable cloud management strategies.

Looking forward, several directions remain open for exploration:

1. **Integration with Renewable Forecasting:** Incorporating advanced weather models could improve accuracy in predicting renewable availability [8].
2. Embodied Carbon Accounting: Future work should expand beyond operational emissions to include embodied emissions from hardware and infrastructure [9].

3. **Edge and Federated Systems:** Extending the framework to edge computing environments and federated learning systems could broaden its applicability [10].
4. **Real-World Deployment:** Collaborating with industry to integrate predictive carbon-aware scheduling into production cloud platforms will be critical to validating scalability and robustness [11]. By advancing carbon-aware cloud computing, this research contributes toward the broader goal of sustainable digital infrastructure that aligns with global carbon neutrality targets [12].

References

[1] S. Hall, F. Micheli, et al., "Carbon-Aware Computing for Data Centers with Probabilistic Performance Guarantees," USENIX NSDI, 2024.
[2] E. Rodrigues, J. Goldverg, and T. Kosar, "Carbon-Aware Temporal Data Transfer Scheduling Across Cloud Datacenters (LinTS)," arXiv preprint arXiv:2506.04117, 2025.
[3] Z. Miao, "Energy and Carbon-Aware Distributed Machine Learning in Multi-Cloud Systems," STET Review, vol. 2, no. 1, pp. 45–62, 2024.
[4] A. Souza et al., "CASPER: Carbon-Aware Scheduling and Provisioning for Distributed Web Services," arXiv preprint arXiv:2403.14792, 2024.
[5] D. Zhao, "An Energy and Carbon-Aware Algorithm for Renewable-Powered Cloud Workloads," Journal of Parallel and Distributed Computing, vol. 162, pp. 20–34, 2022.
[6] D. Mondal, A. Das, and S. Roy, "GEECO: Green Data Centers for Energy Optimization and Carbon Reduction," Sustainability, vol. 15, no. 21, p. 15249, 2023.
[7] A. Maraga and S. O. Ojo, "Carbon-Aware, Energy-Efficient, and SLA-Compliant Virtual Machine Placement via Deep Q-Networks," Journal of Cloud Computing, vol. 14, pp. 110–129, 2025.
[8] D. Maji, A. Bhattacharya, and R. Majumdar, "Bringing Carbon Awareness to Multi-Cloud Application Delivery," HotCarbon Workshop, 2023.
[9] T. Deenadayal, "Carbon Emission Reduction through AI-Based Energy Optimization in Data Centers," International Journal of Green Computing, vol. 9, pp. 56–70, 2025.
[10] W. Lin, X. Li, and C. Wu, "Renewable-Aware Scheduling of Delay-Tolerant Workloads in Cloud Data Centers," IEEE Transactions on Sustainable Computing, vol. 8, no. 3, pp. 445–458, 2023.
[11] Y. Zhang and J. Chen, "Multi-Objective Scheduling in Carbon-Constrained Data Centers," IEEE INFOCOM, pp. 212–220, 2023.
[12] A. Singh and V. Kumar, "Machine Learning-Based Prediction of Carbon Intensity for Energy-Aware Cloud Scheduling," Future Generation Computer Systems, vol. 143, pp. 201–213, 2023.
[13] R. Patel, H. Lee, and M. Gupta, "Carbon-Aware Load Balancing with Geo-Spatial Energy Signals," IEEE Transactions on Cloud Computing, early access, 2024.
[14] K. Shuja, R. Buyya, and W. Song, "Energy and Carbon Optimization in Cloud Computing: A Survey," ACM Computing Surveys, vol. 56, no. 4, pp. 1–37, 2024.
[15] Google, "Carbon-Free Energy for Data Centers," Technical White Paper, 2020.
[16] Microsoft Research, "Carbon-Aware Computing: Aligning Cloud Workloads with Clean Energy," MSR Report, 2021.
[17] J. Qiu et al., "Carbon-Aware Federated Learning in Edge-Cloud Environments," IEEE Transactions on Mobile Computing, vol. 22, no. 7, pp. 1552–1567, 2023.
[18] H. Wang, L. Xu, and J. Wu, "Multi-Agent Reinforcement Learning for Sustainable Cloud Resource Allocation," NeurIPS Workshop on Climate Change and AI, 2023.
[19] A. Jain, S. Sharma, and T. Nguyen, "AI-Driven Forecasting of Grid Carbon Intensity for Sustainable Computing," IEEE Access, vol. 12, pp. 11456–11468, 2024.
[20] M. Brown and P. Li, "Carbon-Aware Multi-Cloud Orchestration Using Renewable Forecasts," IEEE Transactions on Services Computing, early access, 2023.

[21] ElectricityMap, “Real-Time Carbon Intensity Data API,” [Online]. Available: https://www.electricitymap.org
[22] International Energy Agency (IEA), “Data Centers and Energy Use Trends,” IEA Digitalization Report, 2023.
[23] U.S. EPA, “Greenhouse Gas Equivalencies Calculator,” Environmental Protection Agency, 2024.
[24] Wikipedia, “Green Data Center,” [Online]. Available: https://en.wikipedia.org/wiki/Green_data_center
[25] Wikipedia, “Environmental Impact of Artificial Intelligence,” [Online]. Available: https://en.wikipedia.org/wiki/Environmental_impact_of_artificial_intelligence
[26] Y. Chen, “Carbon-Aware VM Scheduling with Latency Constraints,” IEEE Cloud, pp. 50–59, 2023.
[27] L. Han and S. Zhou, “AI-Assisted Optimization for Sustainable Cloud Infrastructure,” Journal of Grid Computing, vol. 21, pp. 201–220, 2023.
[28] P. Miller, “Carbon Intensity Forecasting with Neural Time Series Models,” Applied Energy, vol. 344, pp. 119–134, 2023.
[29] H. Liu, A. Verma, and C. Tang, “Sustainability-Oriented Resource Allocation for Cloud Providers,” IEEE Transactions on Parallel and Distributed Systems, vol. 34, no. 9, pp. 2199–2211, 2023.
[30] B. Zhang, “Low-Carbon Scheduling Strategies for Data-Intensive Workloads,” Future Internet, vol. 15, no. 8, pp. 330–345, 2023.
[31] R. Buyya and K. Shuja, “A Vision for Sustainable Cloud Computing,” IEEE Internet Computing, vol. 27, no. 1, pp. 6–15, 2023.
[30] Muhammad Ahsan Hayat, "Blockchain-Secured Iot Framework for Smart Waste Management in Urban Environments," The Critical Review of Social Sciences Studies, vol. 3, no. 3, pp. 1-6, 12- 8 2025.
[31] Muhammad Ahsan Hayat, "The Role of HR in Managing Robotic Process Automation (RPA) Displacement Anxiety among Employees," The Critical Review of Social Sciences Studies, vol. 3, no. 3, pp. 1-20, 3 8 2025.
[32] Muhammad Ahsan Hayat, "HR Beyond the Office: Leveraging AI to Lead Distributed Teams and Cultivate Organizational Culture in the Age of Remote and Hybrid Work," ACADEMIA International Journal for Social Sciences (AIJSS), vol. 4, no. 3, pp. 1-20, 2025.
[33] Muhammad Ahsan Hayat, "An IOT-Driven Smart Agriculture Framework for Precision Farming, Resource Optimization, and Crop Health Monitoring," ACADEMIA International Journal for Social Sciences, vol. 4, no. 3, pp. 1-14, 2025.