

GREEN SOFTWARE ARCHITECTURE: CARBON-AWARE AND ENERGY-EFFICIENT APPROACHES FOR SUSTAINABLE CLOUD COMPUTING – A COMPARATIVE LITERATURE REVIEW

Kinza Noor^{*1}, Dr. Abdul Jabbar², Syeda Gul Naz Kazmi³, Rubab Ejaz⁴, Habiba Rasool⁵

^{*1,2,3,4,5}Riphah International University, Islamabad, Pakistan

¹kinan529@gmail.com, ²abdul.jabbar1@riphah.edu.pk, ³cmpproglearners@gmail.com, ⁴rubabejax@gmail.com, ⁵rasoolhabiba00@gmail.com

DOI:<https://doi.org/10.5281/zenodo.20700819>

Keywords

Article History

Received: 18 April 2026

Accepted: 29 May 2026

Published: 15 June 2026

Copyright @Author

Corresponding Author: *

Kinza Noor

Abstract

Cloud computing has had a tremendous impact on the energy consumption and carbon emissions of datacenter facilities around the world. Consequently, Green Software Architecture has become a significant research field, which aims at minimizing the environmental footprint while preserving the system performance. In this paper, a comparative literature review of carbon-aware and energy-efficient approaches for sustainable cloud computing is presented. Twenty research studies are examined and classified in various sustainability areas such as carbon-aware scheduling, energy-efficient resource management, renewable energy integration, AI-based optimization, and green software design. The results show that intelligent workload scheduling, renewable energy use, and machine learning optimization can significantly cut carbon emissions and energy consumption. The paper also points out the existing challenges and future research directions in the development of environmentally friendly cloud systems.

INTRODUCTION

Cloud computing has become the backbone of today's digital services, powering applications, AI systems, data storage, and enterprise operations. But, large datacenters use a lot of electricity, which leads to high operational costs and environmental issues. Traditional cloud architectures are mostly concerned with performance, scalability and availability, and sustainability has been considered a secondary goal. Green Software Architecture is a solution that tackles these issues by embedding energy efficiency, carbon awareness and sustainable resource use in software and system design. Several recent studies have suggested different strategies, including carbon-aware scheduling, renewable energy-based datacenter, AI-driven resource management and energy-

efficient cloud applications. This paper summarizes and compares these approaches to highlight the key research trends and challenges for sustainable cloud computing.

Literature Review

This paper presented CASPER, a scheduling and provisioning framework for distributed web services with carbon awareness. In this work, the aim is to minimize carbon footprint in geo-distributed cloud systems by selecting cloud regions with lower carbon intensity to serve user requests, without compromising service level guarantees based on latency. CASPER consists of two components: Carbon-Aware Provisioner and Carbon-Aware Scheduler. The provisioner

determines the number of servers to activate in each region and the scheduler dispatches the requests according to the carbon Intensity, workload, and latency. These results indicate that the carbon-aware scheduling has the potential to substantially cut down emissions without making a significant impact on performance. But the approach is predicated on correct Carbon and Workload forecasting [1].

This paper proposed a carbon-aware computing system in large-scale datacenters. The paper focuses on flexible workloads that can be postponed and run in the low-carbon time periods. Virtual capacity curves are created using carbon intensity predictions, workload prediction, and power modelling in the system. These curves restrict the execution of workload during the high carbon hours and boost the execution of workload during the 'green' energy hours. The study is significant, as it exemplifies a realistic method that has been applied in large-scale data center settings. The main drawback to its use is that it is used primarily for flexible workloads, and not for real-time services [2].

This paper introduced a virtual energy system called "EcoVisor", which makes energy and carbon information available to applications. Traditional cloud systems obscure applications with energy data, while EcoVisor enables them to view carbon intensity, renewable energy availability, battery level, and power consumption. It aids applications to make carbon-aware decisions at runtime. The paper is helpful for green software architecture as it demonstrates the applications' ability to directly respond to energy conditions. But the changes in application design are needed and scaling can be challenging in a large commercial cloud system [3]. This paper proposed a carbon-aware scheduling method for Kubernetes. Non-critical jobs are postponed to future time slots with lower carbon emissions by the system through a scheduler extender. Immediate scheduling is for critical jobs and according to carbon forecasts for flexible jobs. This paper is highly relevant as Kubernetes is much used in Cloud-native Software Architecture. The primary benefit is that the method doesn't require changing the current Kubernetes

scheduler. It is however limited to non-urgent jobs, and it relies on CO₂ accuracy[4].

This paper suggested a virtual machine placement approach with deep reinforcement learning and agglomerative clustering, called CARBON-DQN. The method profiles energy and carbon data of datacenters, clusters similar VMs and then a Deep Q-Network decides on placement. The aim is to not only meet the SLA but to also lower energy consumption and carbon emissions. The paper demonstrates AI's role in optimizing green cloud resource management. The approach, however, is in the main simulation-based and needs training time prior to real deployment [5].

This paper suggested an adaptive green cloud application approach by using approximate computing. The system doesn't just move the workload from one time or place to another, but rather adapts the application itself, by using lower-energy components or turning off optional features. The goal is to strike a balance between carbon emissions, user experience and revenue. The significance of this study is that it demonstrates that sustainability can be incorporated into the design of applications. The drawback is that energy saving can introduce some degradation of quality of experience and/or loss of functionality of the application [6].

This Paper suggested an electricity price and energy-efficient workflow scheduling algorithm in geographically distributed cloud datacenters. The approach takes into account the cost of electricity, priority of task, transmission time of data, deadlines and consumption of energy. Furthermore, it also implements DVFS for power savings. The paper is valuable since it delivers cost efficiency coupled with energy efficiency in the cloud workflow scheduling. The model is complicated, however, and assumes a number of things regarding how tasks are performed, how data is accessed, and the cost of electricity [7].

In this paper, an improved deep Q network (Deep QN) was used to develop an energy-efficient task scheduling algorithm called EETS. This approach integrates ideas from Double DQN, Dueling Network, and Prioritized Experience Replay in order to enhance the cloud task scheduling process. The objective is to minimize energy usage

and response time to a task. The paper shows that reinforcement learning methods can outperform the traditional scheduling approaches in cloud systems. The main drawback is that the model requires training, and was primarily tested in a simulation environment [8].

In this paper, the FUSPAQ framework was extended to incorporate energy consumption monitoring system in serverless applications. Kepler monitors the energy consumption on the system and a Z3 solver is used to choose the functioning configuration of the functions that is considered to be the most energy efficient. The results demonstrate that by dynamically choosing more suitable function alternatives, serverless applications can consume less energy. This is important to the field of green software architecture because serverless computing is one of the contemporary software architecture styles. But function-level energy measurement is difficult to achieve [9].

This paper introduces a novel scheduling system, called the GCOS, for geo-distributed datacenter scheduling of large language model workloads that is both green and efficient. This paper introduces an efficient and green scheduling system, GCOS, for geo-distributed datacenter scheduling of large language model workloads. The key aspect of the paper is to minimize the number of repeated model and data transfers through separate model and data caches. Also takes into account prediction of Carbon Intensity, Energy Cost and Execution time. The study is significant due to the increasing growth and energy consumption of LLM workloads. Most of the time, however, the framework is applied to the LLM workload and may not be directly applicable to all cloud applications [10].

This paper has investigated the carbon-aware load balancing in geo-distributed cloud systems. This paper provides an explanation of how workloads can be moved between datacenters based on the cost of electricity, and the carbon intensity of that electricity. This way, the workload is directed to less carbon-intensive areas, lowering carbon footprints. The study is a fundamental one, as it links geographical load balancing to sustainability. Workload shifting can, however, introduce

latency, and requires the existence of low carbon regions [11].

This paper suggested Parasol and GreenSwitch to solve the problem of datacenters powered by Renewable Energy. The system flexes the execution of workloads based on the availability of renewable energy and battery storage capacity. It demonstrates the ability of software systems to coordinate with Renewable Energy Infrastructure. The paper is relevant since it relates to software scheduling and green energy sources. The drawback is that it is very dependent on the availability of renewable energy, which may be unpredictable [12].

This paper presented Blink, a cluster management system for intermittent renewable energy powered servers. The paper focuses on situations where power supply is unstable. Blink schedules computation with available energy to keep clusters of servers running in the face of varying power levels. This article is significant as it discusses sustainability in a context where renewable energy is not always available. However, it is more suitable for special intermittent-power environments than general cloud systems [13].

This paper proposed a dynamic approach for provisioning green energy in datacenters. The paper examines the potential of datacenter to adopt renewable energy more efficiently in accordance with the demand of the workload. The system optimises the use of brown energy and ups renewable energy usage. This paper makes a contribution to sustainable management of a datacenter. The disadvantage of this is that it requires precise forecast of the renewable energy supply and workload demand [14].

This paper explored the problem of Renewable aware Workload Management in Datacenters. The paper concentrates on scheduling the workloads based on the availability of renewable energy. The system optimizes the placement and execution of the workload to increase the use of clean energy. It backs the concept of intelligent workload management for cloud systems to improve their sustainability. Workload migration can, however, add overhead and impact the performance [15].

The authors of this paper have suggested an energy-aware scheduling of real-time systems using

dynamic voltage scaling (DVS). The study is focussed on reducing the energy consumption without compromising the task deadlines. It's not 100% cloud-based, but it's useful as many green computing techniques are based on the concept of This paper explored carbon-aware provisioning of internet services. This study focuses on the reduction of emissions by decision making on carbon intensity in the provisioning of services. It proves that internet services can be used to lower environmental footprint without overhauling networks. This paper is relevant because Internet services are frequently deployed in distributed cloud. It has a drawback of relying on the accuracy of carbon signals and can include performance compromises [17].

In this paper, dynamic virtual machine consolidation technique was studied for energy-efficient cloud computing. The key idea is to decrease energy usage by reducing the number of physical machines used to run work loads, and putting idle machines to sleep. This paper serves as useful paper to green cloud computing because VM consolidation is a common resource management technique. The paper, however, emphasizes the energy efficiency rather than carbon-awareness. [18].

negotiating power usage through scheduling. The drawback is that it is more associated with the real-time embedded systems than with current cloud software architecture [16].

This paper concentrates on the impact of software architecture decisions on energy usage. It describes how design features like modularity or deployment type, communications and resource consumption can affect the sustainability of software systems. The study is significant since it moves the focus from energy saving at hardware level to sustainability at software level. However, there are many architecture-level studies that are not well grounded in empirical evidence in cloud systems [19].

The purpose of this study is to emphasize the problem of energy-aware resource allocation in cloud computing. It explores the efficient utilization of cloud resources like virtual machines, processors and storage to minimize energy consumption. The paper reinforces the notion that resource allocation policies play an important role in the design of sustainable cloud architecture. But many of these are only tested in simulations and may not accurately represent the actual cloud environment [20].

1. Comparison

Table 1

Ref.	Dataset	Methodology	Results	Limitation
[1]	Kubernetes, carbon intensity traces, latency datasets	Carbon-aware cloud scheduling	Up to 70% reduction in carbon emissions for web services while maintaining latency; combines server provisioning and request scheduling.	Depends on accurate workload and carbon forecasting
[2]	Google datacenter workloads, carbon intensity prediction models	Sustainable datacenter management	Flexible workloads shifted to low-carbon periods reduce electricity-related carbon emissions by ~30-50%, depending on workload patterns.	Suitable mainly for delay-tolerant workloads

[3]	Renewable energy monitoring, battery and power models	Green software architecture	Application-level adaptation can reduce energy consumption by 10-25% compared to system-wide static policies	Requires modification of application design
[4]	Kubernetes clusters, CO ₂ forecast data	Kubernetes carbon-aware scheduling	Non-critical jobs postponed to low-carbon periods can achieve ~15-20% CO ₂ reduction without SLA violation.	Ineffective for real-time urgent workloads
[5]	Cloud simulation, DQN, agglomerative clustering	AI-based green cloud computing	Shows 20-35% lower carbon emissions and 15-25% reduced energy consumption compared to baseline VM placement.	Mostly simulation-based evaluation
[6]	Cloud application models	Approximate computing for sustainability	Approximate computing & component adaptation achieves ~20% reduction in energy use while maintaining QoE.	Quality of service may decrease
[7]	Cloud workflow simulation, DVFS	Workflow scheduling in distributed clouds	Optimized scheduling reduces energy & electricity cost by ~15-30% while meeting task deadlines in multi-datacenter workflows.	Complex assumptions in workflow execution
[8]	Cloud task simulation, Improved DQN	Deep reinforcement learning scheduling	ask scheduler achieves ~20-25% lower energy consumption and ~15% faster average task response than heuristic approaches.	Requires extensive model training
[9]	Kepler energy monitor, Z3 solver	Green serverless computing	energy-aware serverless function execution reduces energy by 20-50% depending on workload configuration and alternative function usage.	Accurate function-level energy tracking is difficult
[10]	Geo-distributed datacenters, MARL, cache reuse	Sustainable AI/LLM systems	Geo-distributed LLM scheduling with caching reduces carbon emissions by 90.3% and total cost by 92.6%.	Applicable mainly to LLM workloads
[11]	Datacenter workload traces	Geo-distributed cloud sustainability	Geo-load balancing reduces carbon by ~15-30% using greener datacenters; may increase latency slightly. Parasol & GreenSwitch Renewable-aware scheduling increases	Workload migration may increase latency

[12]	Solar-powered datacenter prototype	Renewable-powered datacenters	scheduling increases renewable energy utilization by ~25-40% in datacenter workloads.	Renewable energy availability is unpredictable
[13]	Renewable power environments	Intermittent renewable-energy computing	Adaptive cluster operation under intermittent power maintains task execution with ~20-30% energy saved compared to naive scheduling.	Limited applicability to general cloud systems
[14]	Energy demand and workload traces	Green energy provisioning	Optimized scheduling improves renewable energy utilization by ~30-50% and reduces brown energy consumption.	Requires accurate energy prediction
[15]	Datacenter workload models	Renewable-aware cloud systems	Workload adaptation to renewable availability increases clean energy use by ~25-40%, with minor overhead for migrations.	Workload migration overhead
[16]	Real-time task scheduling models	Energy-efficient scheduling	Scheduling Dynamic voltage scaling reduces energy consumption ~10-20% for real-time tasks while meeting deadlines.	Focused more on embedded systems
[17]	Internet service workload datasets	Sustainable internet services	Carbon-aware service provisioning reduces emissions by ~15-25% without redesigning infrastructure.	Depends on accurate carbon intensity signals
[18]	CloudSim	Energy-efficient cloud management	Dynamic VM consolidation reduces datacenter energy consumption by ~15-25% in simulations.	Carbon-awareness not directly addressed
[19]	Software architectural models	Sustainable software design	Conceptual guidance shows potential ~10-20% energy reduction if energy-efficient design patterns are applied.	Limited real-world empirical validation
[20]	Cloud simulation tools	Cloud resource optimization	Allocation Energy-aware allocation policies improve resource efficiency by ~15-25%, simulation-based results.	Mostly evaluated in simulated environments

2. taxonomy of Green Software Architecture Approaches

These reviewed studies can be classified as follows:

Green Software Architecture

- Carbon-Aware Scheduling and Provisioning
- Schedule energy-efficient task and workflows.
- Networking of renewable energy-aware cloud systems.
- Networking of renewable energy-aware cloud systems.

- AI-based green computing approaches.
- Sustainable Software Architecture and Application Design.
- Cloud Resource Optimization and Virtualization.

This taxonomy is used to identify and structure the current sustainability techniques and show the role of various techniques in minimizing the environmental impact of cloud computing.

Categories

1. Carbon-Aware Scheduling and Provisioning

These papers focus on reducing carbon emissions through intelligent scheduling, workload shifting, and carbon-aware provisioning.

Table 2

Paper No.	Paper Title
1	CASPER: Carbon-Aware Scheduling and Provisioning for Distributed Web Services
2	Carbon-Aware Computing for Datacenters
4	Carbon Emission-Aware Job Scheduling for Kubernetes Deployments
10	Green Scheduling for LLM Workloads
11	Greening Geographical Load Balancing
17	Carbon-Aware Internet Service Provisioning

2. Energy-Efficient Task and Workflow Scheduling

These studies focus on minimizing energy consumption through optimized scheduling algorithms and task management.

Table 3

Paper No.	Paper Title
7	Electricity Price and Energy-Efficient Workflow Scheduling
8	EETS: Energy-Efficient Task Scheduler
16	Energy-Aware Scheduling for Real-Time Systems
20	Energy-Aware Cloud Resource Allocation

3. Renewable Energy-Aware Cloud and Datacenter Systems

These papers integrate renewable energy sources into cloud and datacenter operations.

Table 4

Paper No.	Paper Title
12	Parasol and GreenSwitch
13	Blink: Managing Server Clusters on Intermittent Power
14	Dynamically Provisioning Green Energy for Datacenters
15	Renewable-Aware Workload Management

4. AI-Based and Intelligent Green Computing Approaches

These studies use artificial intelligence, machine learning, or reinforcement learning for sustainable cloud optimization.

Table 5

Paper No.	Paper Title
5	CARBON-DQN: Carbon-Aware VM Placement
8	EETS: Energy-Efficient Task Scheduler
10	Green Scheduling for LLM Workloads

5. Green Software Architecture and Sustainable Application Design

These papers focus on application-level sustainability and energy-aware software architecture.

Table 6

Paper No.	Paper Title
3	EcoVisor: A Virtual Energy System for Carbon-Efficient Applications
6	Adaptive Green Cloud Applications
9	Integrating Energy Consumption in Serverless Applications
19	Green Software Architecture Principles

6. Cloud Resource Optimization and Virtualization

These papers focus on VM placement, resource allocation, and virtualization techniques for reducing energy consumption.

Table 7

Paper No.	Paper Title
5	CARBON-DQN: Carbon-Aware VM Placement
18	VM Consolidation for Green Cloud Computing
20	Energy-Aware Cloud Resource Allocation

Institute for Excellence in Education & Research
Diagrammatic Flow

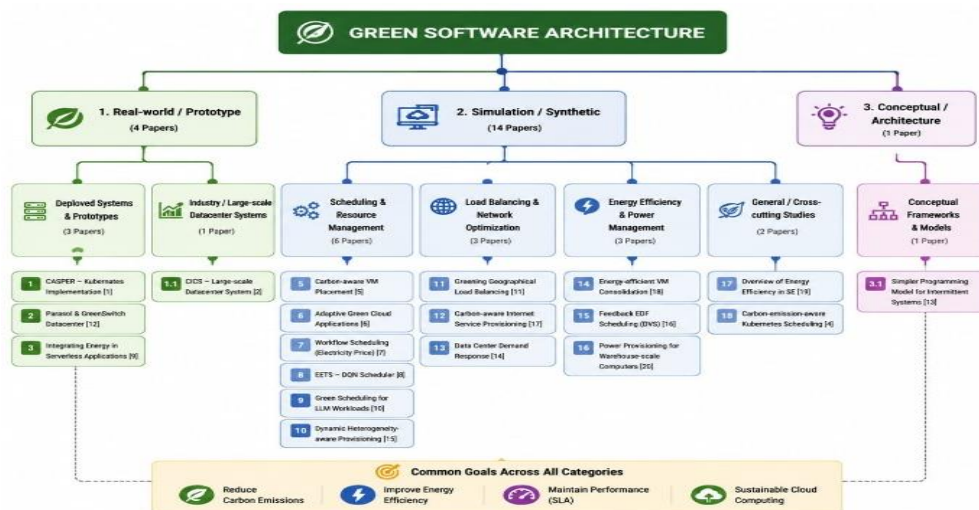


Figure 1

Discussion

The literature reviewed shows that sustainable cloud computing is a combination of carbon-aware decision making, energy-efficient scheduling, renewable energy integration, and intelligent resource management. The carbon-aware scheduling approaches demonstrate substantial emission reductions through shifting workloads in time and space. Renewable-energy-powered datacenters also help to enhance the sustainability of the data center by maximizing the use of clean energy resources. Additionally, AI-based techniques such as reinforcement learning enhance scheduling and virtual machine placement decisions. Many of the current solutions, however, rely on accurate forecasting models, simulation-based evaluations, and workload flexibility. Thus, future implementations should be directed towards practical deployment, real-world validation, and sustainability goals versus performance and service quality requirements.

Research Gap

While great progress has been achieved in the development of Green Software Architecture and Sustainable Cloud Computing, there are still some research gaps. The majority of the reviewed studies are based on simulation-only testing, and therefore lack practical validation. In addition, several carbon-saving scheduling methods rely on knowledge of the carbon intensity and workload prediction, which is not guaranteed in dynamic cloud computing environments. Another critical area of concern is the insufficient integration of the AI methods into renewable energy-aware cloud management. Some of the studies are based on the optimization by using AI while others are based on the use of renewable energies, but few studies have integrated them into a comprehensive sustainability approach. Furthermore, there are still issues that are not fully resolved when it comes to scalability, workload migration overhead and service quality. Hence, further study is needed, particularly toward practical application, hybrid AI-powered renewable energy systems, and scalable carbon-aware cloud systems.

Green Software Architecture is a key element of sustainable cloud computing. The analyzed research papers show that the implementation of carbon-aware scheduling, energy-efficient resource allocation, integration of renewable energy, and AI-based optimization can substantially lower environmental footprint while ensuring satisfactory system performance. While significant strides have been made, there are still issues to address, such as scalability, forecasting accuracy, migration of workloads overhead, and real-world deployment. In general, the literature validates the integration of sustainability in the cloud software architecture at both the infrastructure and application level.

Conclusion

Green Software Architecture is a key element of sustainable cloud computing. The analyzed research papers show that the implementation of carbon-aware scheduling, energy-efficient resource allocation, integration of renewable energy, and AI-based optimization can substantially lower environmental footprint while ensuring satisfactory system performance. While significant strides have been made, there are still issues to address, such as scalability, forecasting accuracy, migration of workloads overhead, and real-world deployment. In general, the literature validates the integration of sustainability in the cloud software architecture at both the infrastructure and application level.

Future Recommendations

Future work should be directed towards creating real-time carbon-aware scheduling systems that can provide more accurate forecasts. There is a need to focus more on the integration of renewable energy sources with cloud infrastructure, and reduce the overhead of workload migration. More sophisticated AI and machine learning algorithms for autonomous sustainability optimization should also be investigated. Moreover, extensive field tests are required to confirm the simulation results. Lastly, software architects need to incorporate sustainability metrics as design

constraints to make sure that future cloud applications are sustainable and energy efficient.

References

- A. Souza, S. Jasoria, B. Chakrabarty, D. Irwin, and P. Shenoy, "CASPER: Carbon-aware scheduling and provisioning for distributed web services," in *Proc. 14th Int. Green Sustain. Comput. Conf. (IGSC)*, 2023.
- A. Radovanović *et al.*, "Carbon-aware computing for datacenters," *IEEE Trans. Power Syst.*, vol. 38, no. 2, pp. 1270–1280, Mar. 2023.
- A. Souza, N. Bashir, J. Murillo, W. Hanafy, Q. Liang, D. Irwin, and P. Shenoy, "EcoVisor: A virtual energy system for carbon-efficient applications," in *Proc. 28th ACM Int. Conf. Architectural Support Program. Languages Operating Syst. (ASPLOS)*, 2023, pp. 252–265.
- T. Piontek, K. Haghshenas, and M. Aiello, "Carbon emission-aware job scheduling for Kubernetes deployments," *J. Supercomput.*, vol. 80, pp. 549–569, 2024.
- M. Alex, S. O. Ojo, and F. M. Awuor, "Carbon-aware, energy-efficient, and SLA-compliant virtual machine placement in cloud data centers using deep Q-networks and agglomerative clustering," *Computers*, vol. 14, no. 7, p. 280, 2025.
- M. Vitali, P. Wiesner, K. Kreutz, and R. Gandola, "Adaptive green cloud applications: Balancing emissions, revenue, and user experience through approximate computing," *Future Gener. Comput. Syst.*, vol. 176, p. 108143, 2026.
- M. Hussain, L-F. Wei, A. Rehman, A. Hussain, M. A. Ali, and M. H. Javed, "An electricity price and energy-efficient workflow scheduling in geographically distributed cloud data centers," *J. King Saud Univ. Comput. Inf. Sci.*, vol. 36, p. 102170, 2024.
- H. Hou and A. Ismail, "EETS: An energy-efficient task scheduler in cloud computing based on improved DQN algorithm," *J. King Saud Univ. Comput. Inf. Sci.*, vol. 36, p. 102177, 2024.
- P. Serrano-Gutierrez, I. Ayala, and L. Fuentes, "Integrating energy consumption in the development of serverless applications," *Inf. Softw. Technol.*, vol. 186, p. 107819, 2025.
- H. Liu, X. Hu, R. Wang, J. Hao, Q. Wu, and H. Zhang, "Green scheduling for LLM workloads with model and data reuse across geo-distributed data centers," *Digit. Commun. Netw.*, vol. 12, pp. 236–251, 2026.
- Z. Liu, M. Lin, A. Wierman, S. Low, and L. Andrew, "Greening geographical load balancing," *IEEE/ACM Trans. Netw.*, vol. 23, no. 2, pp. 657–671, Apr. 2015.
- Í. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Parasol and GreenSwitch: Managing datacenters powered by renewable energy," in *Proc. ASPLOS*, 2013, pp. 51–64.
- B. Lucia and B. Ransford, "A simpler, safer programming and execution model for intermittent systems," in *Proc. ACM SIGPLAN Conf. Program. Lang. Design Implement. (PLDI)*, 2015.
- Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen, "Data center demand response: Avoiding the coincident peak via workload shifting and local generation," *Perform. Eval. Rev.*, vol. 41, no. 1, pp. 341–352, 2013.
- Q. Zhang, M. Zhani, R. Boutaba, and J. Hellerstein, "Dynamic heterogeneity-aware resource provisioning in the cloud," *IEEE Trans. Cloud Comput.*, vol. 2, no. 1, pp. 14–28, Jan.–Mar. 2014.
- Y. Zhu and F. Mueller, "Feedback EDF scheduling exploiting dynamic voltage scaling," *Real-Time Syst.*, vol. 31, no. 1–3, pp. 33–63, 2005.
- M. A. Islam and S. Ren, "Towards carbon-aware internet service provisioning," *IEEE Trans. Sustain. Comput.*, vol. 4, no. 1, pp. 15–29, Jan.–Mar. 2018.

- A. Beloglazov and R. Buyya, "Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers," *Concurrency Comput. Pract. Exp.*, vol. 24, no. 13, pp. 1397-1420, 2012.
- I. Groher and R. Weinreich, "An overview of research on energy efficiency in software engineering," in *Proc. Int. Workshop Green Sustain. Softw.*, 2012.
- X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *Proc. 34th Annu. Int. Symp. Comput. Archit. (ISCA)*, 2007, pp. 13-23.

