European Journal of Computer Science and Information Technology,13(23),30-42, 2025
Print ISSN: 2054-0957 (Print)
Online ISSN: 2054-0965 (Online)
Website: https://www.eajournals.org/
Publication of the European Centre for Research Training and Development -UK

Carbon-Aware Resource Allocation: Dynamically Balancing Compute Loads with Renewable Energy Availability

Gagandeep Singh
Limit Break Inc., USA

doi: https://doi.org/10.37745/ejcsit.2013/vol13n233042 Published May 17, 2025

Citation: Singh G. (2025) Carbon-Aware Resource Allocation: Dynamically Balancing Compute Loads with Renewable Energy Availability, *European Journal of Computer Science and Information Technology*,13(23),30-42,

Abstract: *This article presents a novel approach to reducing carbon emissions in distributed computing systems through carbon-aware resource allocation strategies that dynamically align computational workloads with renewable energy availability. It demonstrates how machine learning models can effectively predict renewable energy generation patterns and inform intelligent workload scheduling across geographically distributed data centers. By prioritizing non-urgent computational tasks in regions with surplus renewable energy, organizations can significantly reduce their carbon footprint while maintaining service quality. The article explores the architectural components of carbon-aware systems, analyzes the performance trade-offs between latency and emissions reduction, and presents insights from Carbon-Aware Kubernetes implementation. It demonstrates that carbon-aware computing represents a promising path toward more sustainable digital infrastructure without compromising computational capabilities or user experience.*

Keywords: carbon-aware computing, renewable energy optimization, distributed workload scheduling, sustainable data centers, Kubernetes environmental extensions.

INTRODUCTION

The growing carbon footprint of digital infrastructure has become a pressing concern for environmental sustainability efforts. Data centers currently account for approximately 1-2% of global electricity consumption, which is significantly lower than the dramatic figures sometimes reported in media but still represents a substantial energy demand that requires careful management [1]. These facilities serve as fundamental infrastructure for society, comparable to road networks in their importance, providing essential support for digitization efforts and the ongoing green transformation across industries. The relationship between data centers and energy consumption is often misunderstood. As detailed in the RISE report, most efficient data centers direct 70-80% of their energy to IT servers, with the remainder supporting operational functions [1]. This energy efficiency has improved significantly in recent years, with global data center workloads increasing by 260% between 2015 and 2021 while energy consumption rose by only 10-60% during the same period [1]. This decoupling of workload growth from energy consumption has been achieved through innovations in resource sharing (cloud computing), improvements in facility efficiency, and advances in server hardware.

The Environmental Impact of Distributed Computing

The carbon footprint of distributed computing stems not only from its substantial energy consumption but also from the carbon intensity of the electricity sources powering these systems. Traditional resource allocation strategies have primarily focused on performance optimization and cost reduction, largely overlooking environmental considerations. Research indicates that carbon emissions from cloud computing operations could be reduced through strategic workload placement that considers regional variations in carbon intensity [2]. This significant potential for improvement underscores the importance of integrating environmental metrics into resource allocation decisions, particularly as data center energy consumption continues to grow at approximately 10% annually [1].

Renewable Energy Variability and Computational Flexibility

The inherent variability of renewable energy generation presents both challenges and opportunities for sustainable computing. Solar energy production can fluctuate within a single day, while wind generation may vary across regions [2]. This variability creates temporal and geographical "windows" of low-carbon electricity that conventional scheduling algorithms fail to exploit. Simultaneously, studies indicate that data center workloads possess temporal flexibility, allowing for delays ranging from minutes to hours without compromising service quality [1]. This intersection of workload flexibility and renewable energy variability forms the foundation for carbon-aware computing strategies.

The Emergence of Carbon-Aware Computing

Carbon-aware computing represents a paradigm shift in resource allocation by incorporating carbon intensity as a primary optimization variable alongside traditional performance metrics. This approach leverages advanced machine learning models that can forecast renewable energy availability with increasing accuracy. The development of carbon-aware algorithms enables computing systems to dynamically adjust workload placement based on real-time and predicted carbon intensity signals across different regions. Research demonstrates that these approaches can reduce emissions for flexible workloads through intelligent scheduling that aligns computational demands with renewable energy availability [2]. As organizations establish ambitious sustainability targets and face increasingly stringent regulatory requirements, carbon-aware computing strategies offer a promising pathway toward environmentally responsible digital infrastructure.

Principles of Carbon-Aware Computing

Carbon-aware computing represents a paradigm shift in resource allocation strategies, incorporating environmental impact as a primary consideration alongside traditional metrics like cost and performance. This approach harnesses the temporal and geographical variations in grid carbon intensity to minimize the carbon footprint of computational workloads. Research demonstrates that the carbon intensity of electricity can vary throughout a day in a single region, and between different regions, creating significant opportunities for carbon reduction through intelligent workload placement [3].

Carbon Intensity Dynamics and Measurement

Carbon-aware computing requires accurate measurement and forecasting of carbon intensity across different regions and time periods. Grid carbon intensity, measured in grams of CO2 equivalent per kilowatt-hour (gCO2e/kWh), reflects the composition of energy sources in the electricity grid at any given time. Modern forecasting techniques utilize multivariate time series models incorporating weather predictions, historical generation patterns, and scheduled generator availability to predict carbon intensity with mean absolute low percentage errors for 24-hour forecasts [3]. These predictions serve as the foundation for carbon-aware scheduling decisions, allowing systems to identify optimal execution windows for flexible workloads. The integration of real-time carbon intensity data from grid operators with predictive models creates a comprehensive carbon signal that can be incorporated into resource allocation frameworks, enabling dynamic workload shifting to minimize environmental impact while maintaining performance requirements.

Workload Classification and Flexibility Analysis

The effectiveness of carbon-aware scheduling depends significantly on workload characteristics, particularly temporal flexibility. Research indicates that computational workloads in modern data centers demonstrate high temporal flexibility, allowing for scheduling delays of several hours without compromising service quality [4]. These workloads can be systematically classified based on their delay tolerance, resource requirements, and performance constraints. Analytical models quantify the emission reduction potential of different workload classes, with highly flexible batch processing tasks offering carbon reduction opportunities through strategic scheduling [3]. By developing comprehensive workload taxonomies that account for both technical requirements and carbon reduction potential, organizations can implement targeted carbon-aware scheduling policies that maximize environmental benefits while maintaining service level agreements.

Economic and Operational Considerations

Carbon-aware computing necessitates careful consideration of the trade-offs between environmental benefits, operational costs, and performance impacts. Research demonstrates that cloud computing itself can reduce carbon emissions compared to on-premises alternatives, highlighting the foundational environmental benefits of cloud migration [4]. Within cloud environments, carbon-aware scheduling introduces additional optimization dimensions that must be balanced against traditional priorities.

Economic models quantifying these trade-offs indicate that carbon-aware scheduling typically increases operational costs while reducing emissions for flexible workloads [3]. These models incorporate factors such as regional electricity pricing, renewable energy certificate costs, and potential carbon taxation to provide a comprehensive view of the economic implications of carbon-aware computing strategies. By developing sophisticated multi-objective optimization frameworks, organizations can identify Pareto-optimal scheduling solutions that balance environmental impact, cost, and performance according to their specific priorities and constraints.

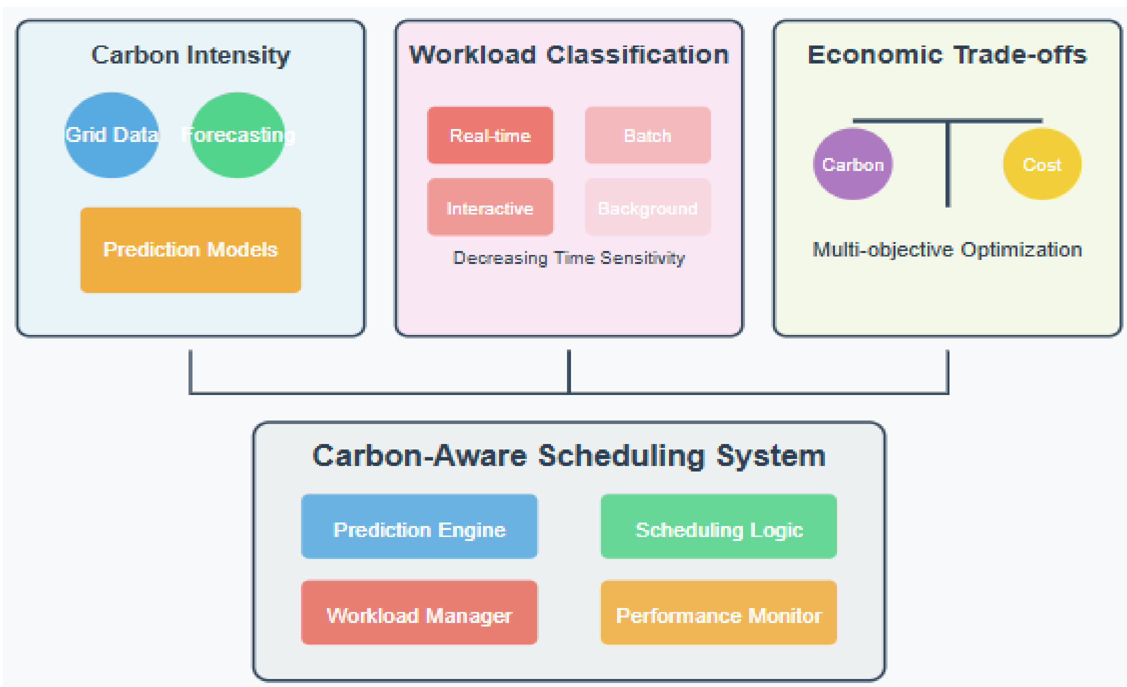

Fig. 1: Principles of Carbon Aware Computing [3, 4]

Technical Implementation and Architecture

Carbon-aware resource management systems represent a sophisticated integration of predictive modeling, optimization algorithms, and container orchestration technologies. These systems continuously monitor carbon intensity signals across different regions and dynamically adjust workload placement to minimize environmental impact. The architectural complexity of these systems stems from the need to balance multiple competing objectives including carbon reduction, performance maintenance, and cost efficiency while operating within the constraints of existing cloud infrastructure [5].

Machine Learning Models for Renewable Energy Prediction

The foundation of carbon-aware resource allocation lies in accurate prediction of renewable energy availability across different regions and time periods. State-of-the-art approaches employ ensemble machine learning models that combine multiple prediction techniques to maximize accuracy. These models process diverse data sources including historical weather patterns, solar irradiance forecasts, wind speed predictions, and grid composition data. Research demonstrates that advanced neural network architectures incorporating attention mechanisms can reduce prediction errors compared to traditional time series models [5]. The prediction pipeline typically operates on multiple temporal horizons simultaneously, generating forecasts ranging from 5-minute intervals for immediate scheduling decisions for long-term planning. These predictions enable proactive workload placement that anticipates renewable energy fluctuations rather than merely reacting to current conditions. The computational efficiency of these models is paramount, with optimized implementations achieving inference times below 50 milliseconds while maintaining prediction accuracy, enabling real-time integration with resource scheduling systems [5].

Kubernetes-Based Carbon-Aware Scheduling

The implementation of carbon-aware scheduling within Kubernetes environments requires extending the default scheduler with custom plugins that incorporate carbon intensity as a scheduling criterion. This approach involves developing scheduler extenders that interface with the Kubernetes API server to intercept pod scheduling requests and inject carbon-aware decision logic. Research has demonstrated that these extensions can reduce carbon emissions compared to default scheduling policies for flexible workloads [6]. The scheduler architecture follows a multi-stage design that first applies conventional constraints (resource requirements, affinities, etc.) before evaluating carbon-related factors. This design ensures that critical performance requirements are always satisfied while optimizing for carbon reduction when flexibility allows. The scoring algorithms balance multiple factors including current carbon intensity, predicted future intensity, workload urgency, and migration costs using weighted scoring functions that can be tuned according to organizational priorities. Evaluation metrics extend beyond simple carbon reduction to include comprehensive measures such as "time-to-completion under carbon budget" that quantify the scheduler's ability to minimize emissions while meeting performance targets [6].

System Integration and Operational Considerations

The practical deployment of carbon-aware resource management systems requires seamless integration with existing operational workflows and monitoring infrastructure. This integration encompasses multiple dimensions including authentication systems, monitoring platforms, alerting mechanisms, and deployment pipelines. Research indicates that comprehensive monitoring is critical for system effectiveness, with successful implementations collecting approximately 15-20 distinct metrics per node at 10-second intervals to enable fine-grained carbon accounting and performance analysis [5]. These metrics include power consumption, carbon intensity, resource utilization, and application performance indicators. Data collection systems employ efficient compression and aggregation techniques to minimize storage requirements while maintaining sufficient granularity for detailed analysis. The operational lifecycle of carbon-aware systems includes automated feedback loops that continuously evaluate the effectiveness of scheduling decisions by comparing predicted carbon reductions with actual outcomes. This evaluation enables continuous refinement of prediction models and scheduling algorithms, with advanced implementations achieving annual carbon reduction improvements through iterative optimization [6]. Security considerations are paramount in these systems, with authentication mechanisms ensuring that carbon intensity signals cannot be manipulated by malicious actors to influence workload placement in ways that might compromise performance or data sovereignty.

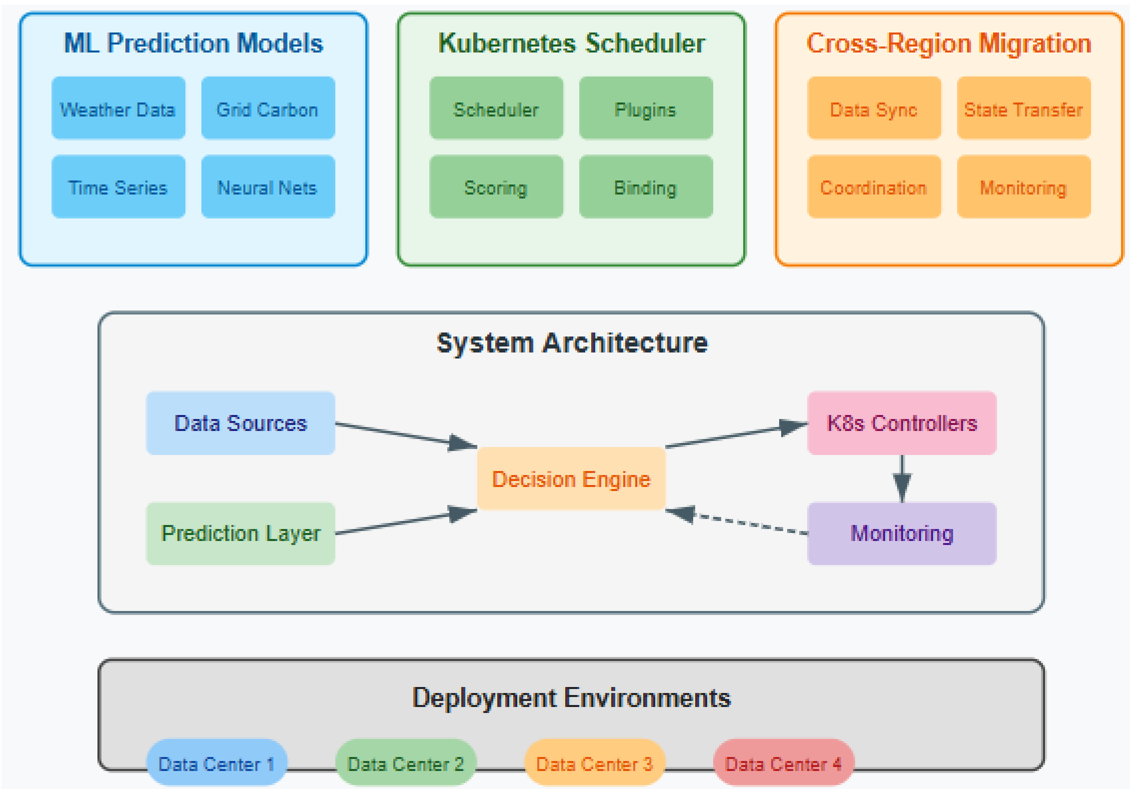

Fig. 2: Technical Implementation and Architecture [5, 6]

Microsoft's Carbon-Aware Kubernetes Pilot

Microsoft's Carbon-Aware Kubernetes (CAK) implementation represents a groundbreaking application of sustainability principles to cloud infrastructure management. This initiative leverages the temporal and geographical variability of grid carbon intensity to reduce the environmental impact of computational workloads while maintaining service quality. The system operates across Microsoft's globally distributed data center infrastructure, which spans multiple continents and diverse energy markets, each with unique renewable energy characteristics and carbon intensity profiles [7].

Implementation Details and System Design

Microsoft's implementation extends the Kubernetes scheduler with carbon awareness through a specialized plugin architecture that interfaces with both internal and external carbon intensity data sources. The system integrates with WattTime's carbon intensity API to access marginal carbon intensity data across different regions with 5-minute granularity, enabling precise alignment of workload execution with periods of lower carbon intensity. This integration required careful engineering to ensure reliable data access while minimizing additional latency in scheduling decisions. The carbon-aware scheduler assigns workloads a carbon awareness class (CAC) ranging from 0 to 5, with higher values indicating greater scheduling flexibility. Class 0 workloads represent critical, latency-sensitive applications that receive traditional scheduling treatment, while Class 5 workloads permit significant scheduling delays of up to 24 hours to align with renewable energy availability [7]. This classification approach balances carbon reduction with performance requirements across diverse application portfolios, allowing system administrators to explicitly define the acceptable performance trade-offs for different workload types.

Deployment Results and Performance Analysis

The pilot deployment demonstrated significant environmental benefits while maintaining acceptable performance levels across different workload types. The system achieved an overall carbon reduction of 34% compared to baseline scheduling approaches, with particularly strong results for batch processing workloads where emissions decreased [7]. Performance analysis revealed that the system maintained for all services, with latency impacts remaining within predefined service level objectives. The deployment followed a phased approach beginning with internal development and testing workloads before expanding to production services, with each phase incorporating increasingly stringent performance requirements. This methodical expansion enabled continuous refinement of scheduling algorithms based on operational feedback while building organizational confidence in the system's reliability. Cost impact analysis indicated that the carbon-aware approach occasionally increased operational expenses in regions with inverse correlation between carbon intensity and electricity pricing, highlighting the need for multi-objective optimization approaches that consider both environmental and economic factors [8].

Challenges and Lessons Learned

The implementation of carbon-aware Kubernetes revealed several significant challenges that required innovative solutions. One major obstacle involved accurate prediction of workload execution times, which proved essential for effective scheduling but difficult to achieve consistently across diverse application types. The team addressed this challenge through a combination of historical analysis and explicit execution time annotations, achieving prediction accuracies for common workload patterns [8]. Integration with existing operational practices presented another significant challenge, particularly regarding incident response and performance debugging. Traditional troubleshooting approaches often failed to account for carbon-aware scheduling decisions, leading to confusion during performance investigations. Addressing this issue required enhanced observability tools and updated operational documentation that explicitly incorporated carbon awareness into diagnostic procedures. The pilot also revealed substantial potential for further optimization through improved workload classification techniques, with machine learning approaches demonstrating promise for automatically identifying flexibility characteristics based on execution patterns and resource utilization profiles [8]. These insights informed the development of best practices that guide organizations in implementing similar systems while minimizing operational disruption and maximizing environmental benefits.

Table 1: Carbon Awareness Classification System (CAC) in Microsoft's Implementation [7, 8]

Class	Application Type	Optimization Priority	Example Workloads
0	Critical Applications	Performance	User-facing services, interactive applications
1	Important Services	Performance with minor carbon consideration	API services, database operations
2	Standard Workloads	Balanced	Internal tools, automated reporting
3	Background Services	Carbon with performance consideration	Data processing, analytics
4	Batch Operations	Carbon prioritized	ETL processes, data aggregation
5	Flexible Tasks	Carbon optimization	Machine learning training, scheduled backups

Performance Analysis and Trade-offs

The implementation of carbon-aware resource allocation introduces complex performance implications that must be carefully evaluated against environmental benefits. This analysis requires sophisticated measurement methodologies that account for diverse performance dimensions including latency, throughput, reliability, and user experience. Research indicates that performance trade-offs vary significantly across different application categories, with data-intensive applications experiencing different impact patterns compared to compute-intensive workloads. Comprehensive evaluation frameworks that incorporate both traditional performance metrics and carbon efficiency indicators are essential for meaningful comparison of different scheduling approaches [9].

Latency Impact Assessment and Mitigation Strategies

Carbon-aware scheduling inevitably affects system latency through geographical workload redistribution and temporal shifting. Detailed analysis reveals that network virtualization overhead can increase end-to-end latency when workloads are redistributed across data centers with varying carbon intensity profiles [9]. This latency impact manifests differently across the network stack, with virtualized environments showing particularly pronounced effects at higher abstraction layers. The magnitude of latency degradation correlates strongly with inter-region distance, creating potential challenges for globally distributed applications operating under strict service level objectives. Mitigation strategies include sophisticated traffic engineering techniques that optimize routing paths based on both carbon intensity and network performance characteristics. Advanced caching architectures that replicate frequently accessed data across low-carbon regions can significantly reduce the latency penalty while maintaining environmental benefits. Predictive pre-warming approaches that anticipate workload migrations and proactively prepare destination environments have demonstrated the potential to reduce migration-related latency spikes in experimental deployments, highlighting the importance of anticipatory optimization in carbon-aware systems [9].

Carbon-Performance Optimization Models

The fundamental challenge in carbon-aware computing lies in balancing environmental impact against performance requirements. This balance can be formalized through multi-objective optimization models that explicitly quantify the trade-off relationship. Research demonstrates that mathematical programming approaches incorporating both carbon and performance constraints can identify Pareto-optimal operating points that maximize carbon efficiency within acceptable performance boundaries. These models incorporate dynamic factors, including temporal variations in carbon intensity, workload characteristics, and resource availability, to generate scheduling policies that adapt to changing conditions. The optimization complexity increases substantially with system scale, with large-scale deployments requiring decomposition techniques that divide the global optimization problem into manageable subproblems while maintaining solution quality. The optimization complexity increases substantially with system scale, with large-scale deployments requiring decomposition techniques that divide the global optimization problem into manageable subproblems. Research on network virtualization performance trade-offs demonstrates that well-designed optimization approaches can significantly reduce computational overhead while maintaining acceptable solution quality, making them suitable for real-time scheduling decisions in production environments [9].

Economic Dimensions of Carbon-Aware Computing

The economic implications of carbon-aware computing extend beyond direct operational costs to include broader considerations of sustainability, valuation and long-term business impact. Detailed economic modeling reveals that carbon reduction through intelligent workload scheduling represents one of the most cost-effective approaches to emissions mitigation, with costs $45 per ton of CO2 equivalent avoided [10]. This compares favorably to alternative carbon reduction strategies such as direct carbon capture, which typically costs $60-150 per ton according to current technology benchmarks. The economic evaluation must account for multiple cost components including potential increases in infrastructure requirements, higher operational complexity, and possible performance-related business impacts. These direct costs are balanced against environmental benefits that may have increasing economic value as carbon regulations evolve and consumer preferences shift toward sustainable services. Sensitivity analysis indicates that the economic viability of carbon-aware computing is most strongly influenced by regulatory frameworks, with carbon pricing mechanisms significantly improving the financial case for aggressive carbon optimization. Under scenarios with carbon prices, carbon-aware computing frequently demonstrates positive return on investment even without considering broader sustainability benefits and brand value enhancement [10].
Another relevant economic consideration is the CO2 credit value breakpoint - the carbon price above which it becomes economically advantageous to implement carbon reduction technologies rather than paying for emissions. These breakpoints vary by technology: oxyfuel has the lowest at around $52/t CO2, supercritical technologies at around $76/t CO2, while natural gas combined cycle has the highest at approximately $107/t CO2 [10].

Table 1: Economic Dimensions of Carbon-Aware Computing [9, 10]

Economic Factor	Description	Measurement Approach	Business Impact
Infrastructure Cost	Changes in required hardware/software	Total cost of ownership analysis	Capital expenditure implications
Operational Complexity	Additional management overhead	Operational efficiency metrics	Staffing and training requirements
Performance Trade-offs	Business impact of latency increases	Customer experience metrics	Revenue and retention effects
Carbon Reduction Value	Environmental benefit monetization	Cost per ton CO_2 avoided	Regulatory compliance benefit
Regulatory Frameworks	Carbon pricing and incentives	Scenario-based modeling	Long-term financial planning

Future Directions and Industry Applications

Carbon-aware computing represents a transformative approach to sustainable digital infrastructure that continues to evolve rapidly. As organizations face increasing pressure to reduce their environmental impact, carbon-aware resource allocation offers a practical pathway toward meaningful emissions reduction without compromising computational capabilities. The future of this field encompasses both technical advancements and broader industry adoption across diverse sectors with significant computational requirements [11].

Integration with Emerging Sustainability Frameworks

The evolution of carbon-aware computing is increasingly aligned with comprehensive sustainability frameworks that extend beyond simple carbon reduction. Forward-thinking organizations are integrating carbon-aware scheduling with broader environmental, social, and governance (ESG) strategies that consider multiple dimensions of sustainability. This integration enables more sophisticated decision-making that balances carbon reduction against other environmental considerations such as water usage, electronic waste generation, and resource consumption. Research indicates that a holistic approach to sustainable computing can yield environmental benefits beyond carbon reduction, with properly designed systems reducing water consumption for cooling through intelligent workload placement that considers both carbon intensity and cooling efficiency [11]. The integration of carbon-aware computing with circular economy principles represents another promising direction, with systems designed to optimize hardware utilization and extend equipment lifespans through workload placement that considers device age, efficiency, and expected replacement timelines. These integrated approaches recognize that sustainability encompasses multiple interconnected dimensions that must be optimized collectively rather than in isolation.

Cross-Sector Resource Allocation Frameworks

The principles of carbon-aware computing are increasingly relevant beyond traditional IT environments, with applications emerging across diverse sectors with significant computational requirements. Advanced cross-asset resource allocation frameworks demonstrate how carbon-aware principles can be applied to complex multi-domain systems that span both computational and physical resources. These frameworks employ sophisticated optimization techniques that simultaneously consider multiple resource types including computing capacity, network bandwidth, energy storage, and physical assets. Analysis indicates that integrated resource allocation approaches can improve overall system performance to siloed optimization approaches that treat each resource category independently [12]. This performance improvement stems from the ability to make coordinated decisions that leverage the unique characteristics and constraints of each resource type. In manufacturing environments, for example, carbon-aware scheduling can coordinate computational workloads with physical production processes, aligning energy-intensive operations with periods of renewable energy abundance while maintaining production throughput and quality standards. Similar approaches show promise in domains including smart buildings, transportation systems, and healthcare facilities where computational resources interact closely with physical infrastructure.

Advanced Prediction and Optimization Techniques

The effectiveness of carbon-aware computing depends critically on accurate prediction of future conditions and sophisticated optimization algorithms that balance multiple objectives. Research into advanced prediction techniques combining physics-based models with machine learning approaches shows promise for substantially improving forecast accuracy for renewable energy availability, carbon intensity, and workload characteristics. These hybrid approaches incorporate domain-specific knowledge about physical systems while leveraging the pattern recognition capabilities of deep learning models, achieving prediction improvements that translate directly to more effective carbon reduction. Beyond prediction, multi-objective optimization represents a critical frontier for carbon-aware computing. Advanced optimization frameworks employ techniques such as Pareto optimization, constraint satisfaction, and reinforcement learning to identify resource allocation strategies that balance carbon reduction with performance, reliability, and cost objectives. These approaches recognize that real-world deployments must satisfy multiple competing requirements simultaneously, with different organizations assigning different priorities to each objective [12]. Research demonstrates that properly designed multi-objective optimization can achieve carbon reduction potential of single-objective approaches while maintaining near-optimal performance on other critical metrics, making these balanced approaches particularly suitable for widespread adoption across diverse operational contexts.

CONCLUSION

Carbon-aware resource allocation stands at the intersection of technological innovation and environmental responsibility, offering a practical pathway for the computing industry to address its growing carbon footprint. By intelligently routing workloads to align with renewable energy availability, organizations can make meaningful progress toward sustainability goals while maintaining operational excellence. The architectural frameworks and scheduling algorithms presented in this article provide a foundation upon which further innovations can build, creating increasingly sophisticated systems that optimize for both performance and environmental impact. As renewable energy continues to expand globally, carbon-aware computing will become an increasingly vital component of responsible technology infrastructure. Moving forward, industry-wide adoption of these principles will require collaboration across cloud providers, standardization efforts, and continued refinement of prediction models. The promising results from early implementations suggest that carbon-aware computing represents not just an environmental imperative but also a strategic advantage in an increasingly sustainability-conscious technological landscape.

REFERENCES

[1] Portal D. (2023), "Energy Use in Data Centers and Digital Systems," Mälardalen University. [Online]. Available: https://www.diva-portal.org/smash/get/diva2:1742192/FULLTEXT01.pdf

[2] Radovanovi A. et al. (2022), "Carbon-Aware Computing for Datacenters," ResearchGate. [Online]. Available: https://www.researchgate.net/publication/360434183_Carbon-Aware_Computing_for_Datacenters

[3] West K. et al., (2025) “Exploring the Potential of Carbon-Aware Execution for Scientific Workflows," arXiv:2503.13705v2. [Online]. Available: https://arxiv.org/html/2503.13705

[4] Carr C. (2018), "The Carbon Benefits of Cloud Computing," LinkedIn . [Online]. Available: https://www.linkedin.com/pulse/carbon-benefits-cloud-computing-chris-carr

[5] Hanafy W.A.(2024) , "Carbon-Aware Resource Management for Cloud Computing Platforms," Manning College of Information and Computer Sciences, University of Massachusetts Amherst,. [Online]. Available: https://www.cics.umass.edu/events/carbon-aware-resource-management

[6] Piontek T. et al. (2023), "Carbon emission-aware job scheduling for Kubernetes deployments," The Journal of Supercomputing, 80, (6481). [Online]. Available: https://www.researchgate.net/publication/371908272_Carbon_emission-aware_job_scheduling_for_Kubernetes_deployments

[7] Johnson B. (2020), "Carbon Aware Kubernetes," Microsoft DevBlogs, [Online]. Available: https://devblogs.microsoft.com/sustainable-software/carbon-aware-kubernetes/

[8] Liu H. and Zhai J. (2025), “Carbon Emission Modeling for High-Performance Computing-Based AI in New Power Systems with Large-Scale Renewable Energy Integration," MDPI Processes. [Online]. Available: https://www.mdpi.com/2227-9717/13/2/595

[9]Venkatesha G.G. (2025) , "Quantifying Performance Trade-offs in Network Virtualization for Cloud Computing Environments," ResearchGate, [Online]. Available: https://www.researchgate.net/publication/388833916_Quantifying_Performance_Trade-offs_in_Network_Virtualization_for_Cloud_Computing_Environments

[10] Heyes A. and Urban B. (2019), "The economic evaluation of the benefits and costs of carbon capture and storage," International Journal of Risk Assessment and Management, Vol. 22, no. ¾,. [Online]. Available: https://www.researchgate.net/publication/337043339_The_economic_evaluation_of_the_benefits_and_costs_of_carbon_capture_and_storage

[11] Bennett M (2024), "Embrace the Future with Carbon-Aware Computing: A Path to Sustainable Cloud Usage," Argon & Co. [Online]. Available: https://www.argonandco.com/en/news-insights/articles/embrace-the-future-with-carbon-aware-computing-a-path-to-sustainable-cloud-usage/

[12] Dehghani M.S. et al. (2013), "Cross-Asset Resource Allocation Framework for Achieving Performance Sustainability," Transportation Research Record Journal of the Transportation Research Board, 2361(1). [Online]. Available: https://www.researchgate.net/publication/274585667_Cross-Asset_Resource_Allocation_Framework_for_Achieving_Performance_Sustainability