Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Green Ad Tech: A Systems Approach to Carbon-Aware Digital Advertising

Nikhil Kokal
The Walt Disney Company, USA

ARTICLE INFO

Received: 04 Oct 2025

Revised: 08 Nov 2025

Accepted: 17 Nov 2025

ABSTRACT

The digital advertising ecosystem is one of the most latency-sensitive distributed systems on the Internet, fulfilling millions of real-time bidding transactions per day and producing significant greenhouse gas emissions due to server operation, network transmission, and machine learning inference. Combined with highly tight sub-100-ms latency, multifaceted multi-party economic incentives, and increasing regulatory demands that businesses focus on environmental responsibilities, these converging factors present a set of unique carbon reduction challenges to which the carbon-aware computing models have failed to provide sufficient solutions. Green Ad Tech presents a comprehensive systems architecture integrating five coordinated mechanisms: real-time carbon-aware request routing leveraging geographic and temporal variations in grid carbon intensity, lightweight machine learning model switching to reduce inference energy consumption, intelligent edge caching and creative optimization, adaptive auction reweighting through contextual bandit algorithms, and granular per-impression emission attribution aligned with industry measurement standards. The framework formalizes carbon-aware advertising as a constrained multi-objective optimization problem balancing emission minimization against latency service-level objectives, revenue preservation requirements, and data protection compliance constraints. Specialized algorithmic methods to reduce carbon and still be competitive in the auction time window of microseconds allow greedy routing with latency constraints, forecast-based time optimization of deferrable operations, and value-adaptive model selection to be the primary decision primitives. Measurement infrastructure offers per-impression carbon accounting with full provenance metadata, allowing third-party auditing and avoiding greenwashing, and incentive mechanisms can link carbon goals among advertisers, publishers, platforms, and content delivery networks using configurable preferences, eco-labeling, and billing changes. The equity considerations deal with the possible short side of publishers in carbon-intensive areas by providing temporal flexibility, fairness restrictions, and capacity-building investments that do not result in sustainability efforts contributing to global economic inequalities. Privacy and regulatory compliance are considered as hard constraints that cannot be negotiated, and carbon optimization will consider the data protection requirements, such as GDPR restrictions on cross-border processing of data. The framework demonstrates that meaningful emission reductions are achievable within the operational constraints of production advertising systems while maintaining economic viability and stakeholder trust through transparent, auditable measurement aligned with emerging industry standards for media sustainability reporting.

Keywords: Carbon-Aware Scheduling, Programmatic Advertising, Real-Time Bidding, Carbon Intensity, Energy-Aware Routing, Green AI, Sustainability Accounting

1. Introduction

The digital advertising ecosystem is a time-sensitive and geographically dispersed system of the modern Internet. Real-time bidding (RTB) systems bid on billions of impression opportunities each and every day, conducting sophisticated auctions and serving personalized content in the presence of tight latency limits, typically less than 100 milliseconds, and coordinating activities in global data centers, content delivery networks (CDNs), and billions of end-user devices. The large size of this, coupled with the growing use of computationally-intensive machine learning models in targeting and

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

personalizing, is a major source of greenhouse gas (GHG) emissions, which regulators, advertisers, and environmental stakeholders have become concerned about. The energy footprint of digital advertising infrastructure can differ widely depending on the region, with infrastructure relying on renewable energy generators using electricity at about 50-100 gCO 2 e per kilowatt-hour, and infrastructure relying on fossil fuel generation potentially approaching 800 gCO 2 e per kilowatt-hour with optimization potentials through smart geographic routing and load shifting over time.
The more recent efforts in the industry have started to institute the framework defining media emissions measurement and reporting, which may argue in favor of the increased awareness of the environmental impact of the advertising industry. Organizations that are going through the net-zero path experience significant difficulties in properly quantifying and mitigating their carbon emissions in their complex value chains, where many businesses have been unable to set strong baseline levels and effective reduction measures against aggressive timeframe undertakings [1]. The pathway to net zero requires a comprehensive assessment of Scope 1, Scope 2, and particularly challenging Scope 3 emissions, which in digital advertising encompasses not only direct operational energy consumption but also the distributed activities of multiple intermediary platforms, creative production processes, and end-user device rendering that collectively constitute the majority of the ecosystem's carbon footprint [1]. However, the unique technical constraints of programmatic advertising—ultra-low latency requirements, complex multi-party coordination, and revenue-sensitive auction dynamics—present challenges that existing carbon-aware computing research has not adequately addressed. While substantial literature exists on carbon-aware scheduling for flexible workloads and energy-efficient machine learning, these approaches typically assume temporal flexibility measured in minutes or hours, rather than the sub-100ms decision windows characteristic of RTB systems.
The Ad Net Zero Global Media Sustainability Framework provides industry-wide guidance for measuring and reducing carbon emissions across the media value chain, establishing standardized methodologies that enable consistent reporting and comparison across different advertising channels and geographic markets [2]. This framework emphasizes the importance of comprehensive measurement spanning creative production, media distribution, and consumer engagement phases, while acknowledging that digital media's carbon footprint extends beyond direct energy consumption to include embodied emissions in hardware infrastructure, network transmission energy, and the cumulative impact of billions of daily user interactions [2]. The framework's adoption by major industry bodies and advertisers creates an opportunity for technical innovations like Green Ad Tech to align with established governance structures, ensuring that carbon reduction claims can be verified through standardized attribution methodologies rather than proprietary metrics that risk greenwashing accusations [2].

Framework Component	**Current State**	**Implementation Barrier**	**Standardization Status**
Net Zero Commitments	Organizations struggle with baseline measurements	Complex value chain emissions	Framework guidance available
Scope 3 Attribution	The majority of the ecosystem footprint	Distributed intermediary platforms	Standardized methodologies emerging
Measurement Systems	Comprehensive spanning required	Creative to consumption phases	Industry-wide adoption underway
Governance Structures	Technical innovations needed	Verification requirements	Alignment opportunities present

Table 1: Industry Sustainability Framework Maturity and Implementation Challenges [1,2]

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

2. Problem Formulation and System Architecture

2.1 Formal Optimization Framework

The carbon-aware advertising optimization problem can be formalized as a constrained multi-objective program that balances environmental impact against operational performance requirements. Let U denote the set of impression opportunities arriving over a time horizon T, and R represent available geographic regions encompassing data centers or edge locations distributed across multiple continents and electrical grid zones. Each region $r \in R$ exhibits time-varying carbon intensity $C_r(t)$ measured in grams CO_2-equivalent per kilowatt-hour (gCO_2e/kWh), obtainable from grid monitoring services that provide comprehensive carbon intensity analytics and forecasting capabilities for power markets worldwide [3]. Advanced carbon intensity toolkits enable energy market participants and technology operators to access granular emission data with temporal resolutions ranging from 15-minute intervals to hourly averages, supporting decision-making processes that require understanding of both current grid conditions and anticipated future states based on forecasted renewable generation and demand patterns [3]. These monitoring systems aggregate data from transmission system operators, renewable energy forecasting services, and fuel mix reporting to generate location-specific carbon intensity values that capture the complex interplay between baseload generation from nuclear and fossil sources, variable renewable generation from wind and solar installations, and demand-responsive resources including battery storage and demand-side management programs [3]. The temporal variability of grid carbon intensity presents significant optimization opportunities, as renewable energy penetration creates predictable daily and seasonal patterns where carbon intensity fluctuations can enable strategic load shifting to minimize emissions while maintaining service quality and operational performance metrics.

Each impression request u arrive at time t_u and must receive an auction response within deadline D, typically constrained to 80-100 milliseconds in contemporary RTB ecosystems. Serving request u from region r incurs measurable energy consumption decomposable into computational energy $E^{comp}_{u,r}$ encompassing bid evaluation algorithms, machine learning model inference operations, and auction logic implementing second-price mechanisms. Comprehensive analysis of data center energy consumption patterns reveals that United States data centers consumed approximately 70 billion kilowatt-hours in 2014, representing roughly 1.8% of total national electricity consumption, with energy use distributed across IT equipment accounting for approximately 43% of consumption, cooling systems representing 38%, power distribution and conversion contributing 10%, and lighting and other facility operations consuming the remaining 9% [4]. The energy intensity of data center operations exhibits substantial variation based on facility design, with power usage effectiveness values ranging from 1.1 in state-of-the-art hyperscale facilities optimized for energy efficiency to 2.0 or higher in older enterprise data centers lacking modern cooling infrastructure and power management systems [4]. Network energy $E^{net}_{u,r}$ summarizes the communication of data to components within the system, whereas CDN energy $E^{cdn}_{u,r}$ summarizes the operations of content storage and delivery across hierarchies of distributed caches.

Routing decision x u r 01 denotes whether request u can be allocated to region r, with the requirement that one region should be selected on a request to guarantee a deterministic execution of the auction. Additionally, each routing choice induces predicted latency $L_{u,r}$ which must satisfy the bound $L_{u,r} \le L^{max}$ with high probability to meet quality-of-service requirements. The primary objective minimizes total carbon emissions: minimize $\Sigma_{u \in U_T} \Sigma_{r \in R} x_{u,r} \cdot E_{u,r} \cdot C_r(t_u)$, subject to latency constraints, revenue preservation constraints, and regulatory requirements mandating that certain requests must be processed in specific jurisdictions for data protection compliance. Revenue R_u associated with each impression depends on auction outcomes, and carbon-aware interventions must preserve advertiser value within acceptable tolerances. In practice, this multi-objective formulation is often scalarized using weighted coefficients α, β, and γ that encode stakeholder priorities regarding emissions reduction, latency penalty avoidance, and revenue preservation, respectively.

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

Infrastructure Element	Monitoring Capability	Temporal Resolution	Efficiency Metric
Grid Carbon Intensity	Comprehensive analytics toolkit	15-minute to hourly intervals	Variable by renewable penetration
Transmission System Data	Aggregated from operators	Real-time and forecasted	Generation mix dependent
Data Center IT Equipment	Energy consumption tracking	Continuous telemetry	43% of facility consumption
Cooling Infrastructure	Power usage effectiveness measurement	Facility-level monitoring	38% of facility consumption

Table 2: Carbon Intensity Data Sources and Data Center Energy Characteristics [3,4]

3. Algorithmic Techniques for Sub-100-ms Decision Making

3.1 Latency-Constrained Carbon-Minimizing Routing

The fundamental algorithmic primitive addresses single-request routing: given impression request u with deadline D, select region r* that minimizes estimated emissions $E_{u,r} \cdot C_r(t_u)$ subject to predicted latency $L_{u,r} \leq L^{max}$. The algorithm proceeds by sorting candidate regions in ascending order of estimated carbon impact and greedily selecting the first region satisfying the latency constraint. If no candidate meets the latency requirement, the algorithm falls back to selecting the minimum-latency region, prioritizing service quality over carbon optimization. Advanced forecasting approaches for high-dimensional time series demonstrate that deep neural network architectures can effectively predict complex temporal patterns by decomposing global forecasting problems into local subproblems, with recurrent neural networks and temporal convolutional networks achieving prediction accuracies that outperform traditional statistical methods by 15-25% on benchmark datasets involving thousands of interconnected time series [5]. These forecasting techniques prove particularly valuable for carbon-aware routing systems that must anticipate grid carbon intensity across multiple geographic regions simultaneously, where the "think globally, act locally" paradigm enables models to capture both cross-regional correlations in renewable generation patterns and region-specific characteristics such as local weather conditions, demand cycles, and generation portfolio composition [5]. The application of deep learning to carbon intensity forecasting supports routing decisions by providing probabilistic predictions with quantified uncertainty, enabling systems to implement risk-aware routing policies that account for forecast confidence when selecting serving regions under time-critical constraints where routing errors could result in latency violations or unnecessary carbon emissions [5].

For deployments with n candidate regions, sorting requires O(n log n) comparisons. In typical production environments, geographic routing considers 4-12 candidate regions, making this approach computationally trivial—empirical measurements demonstrate consistent sub-2ms execution times on contemporary server processors. Pre-computation further reduces online costs: routing tables mapping request characteristics to optimal regions can be periodically regenerated every 5-30 seconds based on updated carbon signals, with online logic reduced to simple table lookups requiring only microseconds. The algorithm incorporates forecast uncertainty through conservative hedging, where, rather than routing to the region with minimum expected carbon intensity, the selection criterion requires that expected carbon reduction exceed a threshold accounting for forecast variance and switching costs, reducing the risk of latency degradation from mis-routed requests.

3.2 Forecast Lookahead for Deferred Operations

While real-time auction responses cannot be delayed, several system operations exhibit temporal flexibility, including creative pre-warming, cache population, background model retraining, and log

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

processing. A lookahead optimizer exploits short-term carbon intensity forecasts typically extending 30 seconds to 2 hours ahead to schedule these deferrable tasks during predicted low-carbon windows, particularly when renewable generation is forecast to be abundant. Machine learning techniques for electricity consumption and price forecasting in smart grid environments have demonstrated significant capabilities in predicting short-term and medium-term patterns, with studies showing that hybrid approaches combining multiple algorithms can achieve mean absolute percentage errors below 5% for day-ahead electricity load forecasting and below 8% for price forecasting when trained on sufficient historical data encompassing seasonal variations and demand response events [6]. The application of ensemble methods that combine predictions from artificial neural networks, support vector machines, and traditional time series models such as ARIMA improves forecast robustness compared to single-algorithm approaches, with weighted averaging schemes or stacked generalization architectures reducing prediction variance and improving reliability for operational decision-making systems [6]. For carbon-aware workload scheduling, accurate electricity price forecasting serves as a proxy for carbon intensity in liberalized electricity markets where price often correlates with marginal generation source, as renewable energy with zero fuel costs typically sets lower market clearing prices compared to fossil generation with substantial fuel expenses, enabling systems to infer favorable low-carbon windows from price forecast signals even when direct carbon intensity data is unavailable [6].

The scheduler formulates a rolling-horizon optimization problem: given forecasted carbon intensities $C\_r(t+\tau)$ for future time steps $\tau \in [0, T\_horizon]$, a catalog of schedulable tasks with deadlines and resource requirements, and capacity constraints, determine task-to-timeslot assignments minimizing total carbon impact while respecting deadlines and resource limits. For tractability, the problem is solved using mixed-integer linear programming for small problem instances or greedy heuristics for larger deployments, with solutions recomputed every 10-60 seconds. Creative pre-warming provides a concrete application: the Edge Cache Controller analyzes impression forecasts to identify creatives likely to be requested in the coming hours. Rather than immediately fetching these assets, the controller schedules prefetch operations during upcoming low-carbon periods, potentially delaying by 15-60 minutes to align with renewable generation peaks.

3.3 Adaptive Auction Reweighting via Contextual Bandits

Carbon-aware routing alone cannot capture all optimization opportunities; auction-level policies that subtly adjust bid evaluation or creative selection can achieve additional reductions by preferring lower-emission advertising options when multiple near-equivalent choices exist. However, manual calibration of auction penalties risks revenue loss or advertiser dissatisfaction. The Green Ad Tech framework employs contextual multi-armed bandit algorithms to automatically learn advertiser-specific and campaign-specific carbon-revenue tradeoffs through online experimentation. A context is given by each advertiser or campaign; actions are given by fractional adjustments of effective bid values (or soft re-ranking weights) in the set of values: 0.0, 0.01, 0.05, 0.10. The system tracks a composite reward $r = \lambda$ revenue - $(1-\lambda)$ CO 2e after each impression, where 0 1 is a relative weight between the significance of revenue and carbon decrease. Online Thompson sampling or upper confidence bound algorithms optimize policy parameters, trade-offs between exploration and exploitation. Safety constraints prevent excessive revenue loss during exploration: hard per-campaign revenue-loss caps bound acceptable degradation, with exploration limited to campaigns that have explicitly opted into carbon-aware policies.

3.4 Model Switching and Green AI Integration

Machine learning inference represents a significant fraction of computational energy in modern advertising systems. Recent advances demonstrate that model distillation, quantization, and architecture search can reduce inference energy by 30-70% with minimal accuracy degradation. The Model Manager component integrates these techniques by maintaining multiple model variants along a quality-efficiency frontier. Online model selection treats each incoming request as a contextual decision: given predicted impression value, available computational budget, current carbon intensity, and model performance profiles, select the model variant maximizing expected utility. High-value impressions justify full-precision models despite higher energy costs; lower-value impressions are served with efficient variants. Warm pool management ensures selected models are readily available, while admission control prevents resource exhaustion during traffic surges.

Forecasting Domain	Technical Approach	Accuracy Characteristic	Application Context
Carbon Intensity Patterns	Deep neural network architectures	Outperform traditional methods	Multi-region simultaneous prediction
High-Dimensional Time Series	Global-to-local decomposition	Probabilistic with uncertainty	Risk-aware routing policies
Electricity Load Prediction	Hybrid algorithm ensemble	Mean absolute percentage error below threshold	Day-ahead scheduling
Price Forecasting Proxy	Multiple algorithm combination	Improved robustness	Carbon intensity inference

Table 3: Forecasting Technologies and Prediction Performance [5,6]

4. Measurement, Attribution, and Auditability

4.1 Per-Impression Carbon Accounting

Plausible claims of emission reduction must have an accurate, auditable measurement infrastructure, which can be challenged by examination by regulators, third-party auditors, and other stakeholders, who are becoming increasingly sensitive to the dangers of greenwashing in the sustainability reporting of corporate reporting. Carbon attribution in the Green Ad Tech framework assigns end-to-end emissions to attributing components, allowing unparalleled accountability on the impact of any advertisement at the lowest unit of delivery. For each served impression u routed to region r, the total carbon footprint is computed as: $CO_2e_{u,r} = C_r(t_u) \cdot (E^{comp}_{u,r} + E^{net}_{u,r} + E^{cdn}_{u,r}) + client_u$, where $C_r(t_u)$ represents the grid carbon intensity at serving time, energy terms capture server-side operations, and $client_u$ approximates device-side rendering energy often considered Scope 3 emissions under standard accounting frameworks. The Corporate Value Chain Accounting and Reporting Standard provides a full-fledged direction concerning the quantification and reporting of indirect greenhouse gas emissions that occur within organizational value chains, and there are fifteen different categories of Scope 3, which include purchasing goods and services, capital goods, fuel and energy related activities not covered in Scope 1 and 2, upstream and downstream transportation and distribution, waste within the organization, business travel, employee commuting, upstream and downstream leased assets, processing and usage of sold products, end of life treatment of the sold products, franchises, and investments The category of the relevant scope 3 in the case of digital advertising platforms can be classified as Category 1 purchased services offered by technology vendors and data providers, Category 3 transmission and distribution losses in electricity supply chains, Category 8 upstream leased assets in the case of co-location facility use, and Category 11 use-phase emissions in end-user devices rendering advertisement purpose, with allocation methodologies taking special account of activities, based metrics of impressions served, data transmitted, or revenue generated to apportion appropriately the emissions of shared infrastructure supporting multiple business activities [7]. The standard emphasizes that Scope 3 accounting presents significant challenges including data availability limitations where suppliers may not track or report emissions data, calculation complexity requiring lifecycle assessment expertise and access to emissions factors databases, and completeness tensions where organizations must balance comprehensive coverage against resource constraints and materiality thresholds that focus efforts on emission sources representing the largest shares of total footprints [7].

Each component requires distinct estimation methodologies: computational energy derives from CPU/GPU utilization telemetry scaled by instance power profiles and regional PUE; network energy applies per-byte coefficients accounting for data center networking and wide-area transmission.

Research examining the electricity intensity of Internet data transmission demonstrates that energy consumption estimates vary substantially across studies, with reported values spanning three orders of magnitude from 0.00043 kWh/GB to 136 kWh/GB depending on system boundaries, analytical approaches, and temporal scope of assessments [8]. Core network transmission energy ranges from approximately 0.002 to 0.01 kWh per gigabyte transferred when considering only operational energy of routers, switches, and optical transmission equipment, though comprehensive lifecycle assessments that include embodied energy in network infrastructure manufacturing and end-of-life disposal can increase total intensity estimates by factors of 1.5 to 2.5 [8]. Access network energy, representing the last-mile connection between Internet service providers and end users, contributes substantially to total transmission intensity, with digital subscriber line, cable modem, and fiber-to-the-home technologies consuming 0.003 to 0.015 kWh per gigabyte, depending on technology generation, utilization rates, and power management capabilities [8]. These transmission energy estimates exhibit declining trends over time as network equipment efficiency improves through technology upgrades, traffic growth amortizes fixed infrastructure energy across expanding data volumes, and industry adoption of energy-proportional architectures that scale power consumption with actual traffic loads rather than maintaining constant energy draw regardless of utilization levels [8].

Client-side rendering energy poses particular measurement challenges, as advertising platforms lack direct visibility into end-user device characteristics and power states. Industry frameworks increasingly recommend standardized modeling approaches: representative device profiles combined with creative complexity metrics yield reasonable emission estimates, though with greater uncertainty than server-side measurements. Each impression record includes extensive provenance metadata: carbon signal source and timestamp, energy estimator version, model variant identifier, routing decision rationale, and applied auction penalties. These fields enable retrospective auditing where third parties can verify that reported emissions reflect actual system state at serving time rather than post-hoc manipulation. Append-only logging and cryptographic commitments through hash chains further strengthen auditability by preventing retroactive modification of emission records.

Accounting Dimension	Scope Coverage	Measurement Challenge	Energy Characteristic
Value Chain Emissions	Fifteen distinct categories	Data availability limitations	Indirect throughout the chain
Digital Platform Categories	Purchased services for end-use	Allocation methodology complexity	Activity-based apportionment
Network Transmission	Core and access segments	Substantial estimate variation	Declining trends over time
Lifecycle Assessment	Infrastructure manufacturing	Embodied energy inclusion	Factor increase in totals

Table 4: Emission Accounting Standards and Network Energy Intensity [7,8]

5. Incentives, Economics, and Equity Considerations

5.1 Multi-Stakeholder Incentive Alignment

When advertisers have their carbon preference, publishers in low-carbon areas enjoy preferential routing and could raise premium prices on their green inventory. Nevertheless, this dynamic implies distributional issues: publishers in carbon-intensive areas can be at a systematic disadvantage, not their fault, because grid composition is not something they can influence themselves, but the result of government decisions affecting infrastructure and policy. The Incentive and Marketplace Layer addresses this challenge by creating explicit economic signals that align carbon objectives with existing revenue models. Industry research examining the state of readiness for sustainability in digital advertising reveals that while 86% of organizations acknowledge sustainability as a priority

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

consideration in their business strategies, only 41% have implemented concrete measures to reduce their environmental impact, indicating a significant implementation gap between stated intentions and operational reality [9]. Furthermore, 63% of surveyed organizations report lacking the necessary skills, resources, or technological infrastructure to effectively measure and report their carbon footprint, highlighting systemic barriers to widespread adoption of sustainability initiatives across the advertising ecosystem [9]. The research identifies that among organizations that have begun sustainability efforts, 52% focus primarily on measurement and reporting capabilities as foundational steps, while only 28% have progressed to implementing active carbon reduction strategies such as energy-efficient infrastructure or carbon-aware operational practices, suggesting that the industry remains in early stages of sustainability maturity with substantial opportunity for frameworks like Green Ad Tech to accelerate progress [9].

Advertisers gain control through configurable carbon preferences integrated into campaign management interfaces: hard carbon budgets establishing maximum gCO_2e per campaign, soft preferences such as "minimize emissions within 5% of optimal eCPM," or implicit carbon prices representing willingness to pay per gram CO_2e avoided. These preferences translate into algorithmic adjustments in routing and auction logic, with impacts on campaign performance and costs made transparent through reporting dashboards. Survey evidence from industry working groups suggests approximately 40% of advertisers would adopt carbon preferences if revenue-neutral, with an additional 20% willing to accept modest cost increases for demonstrable emission reductions. Publishers in low-carbon regions benefit from preferential routing when advertisers express carbon preferences, potentially commanding premium prices for "green inventory." However, this dynamic creates distributional concerns: publishers in carbon-intensive regions may face systematic disadvantage through no fault of their own, as grid composition reflects infrastructure and policy decisions beyond individual control.

5.2 Addressing Regional Inequities

Geographic carbon intensity difference is an indicator of inherent variations in electricity infrastructure, climatic conditions, and the level of economic development. Areas dependent on coal or natural gas are 400-800 gCO_2e /kWh intensities; areas with sufficient hydroelectric, solar, or wind power are much lower, usually under 100 gCO_2e /kWh. In 2020, the environmental impact of information and communication technology analyzed that as the digital technologies used around the world consumed about 2,600 terawatt-hours of electricity in 2020-at 5-9% of all the energy consumed globally-the carbon intensity of this use was many times higher than average, with data centers relying on renewable energy sources near zero carbon consumption and regions with heavy reliance on coal to power data centers approaching 1,000 grams CO_2e per kilowatt-hour in some regions [10]. Emphasizing that concentrating only on operational energy efficiency improvements without regard of the carbon intensity of electricity sources can lead to incorrect judgments about environmental impact, the study notes that a 50% efficiency gain in a coal-powered plant may nonetheless lead to more absolute emissions than an unoptimized facility driven by renewable energy, therefore underlining the need of geographical routing methods accounting for grid composition [10]. Furthermore, the study highlights that embodied emissions from manufacturing digital infrastructure equipment including servers, networking hardware, and end-user devices contribute 30-80% of total lifecycle emissions depending on usage patterns and equipment lifespan, suggesting that carbon-aware routing alone addresses only operational emissions while comprehensive sustainability requires attention to hardware procurement, utilization efficiency, and circular economy principles that extend equipment lifespans and improve recycling rates [10].

Naive carbon-aware routing that invariably prefers low-carbon areas would always harm publishers in poor countries and areas with carbon-intensive grids, therefore aggravating world economic inequities. Many methods reduce these equity issues without giving up carbon optimization. Temporal flexibility takes advantage of the wide variation in carbon intensity across hours inside any area; routing to high-carbon areas during their local low-carbon hours yields significant decreases without categorical regional exclusion. Formally incorporating fairness limitations into the optimization

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

objective guarantees that no area suffers catastrophic revenue loss while yet attaining considerable overall reductions: the system optimizes a regularized objective containing terms penalizing excessive deviation from baseline regional impression distributions instead of pure carbon reduction.

5.3 Privacy and Regulatory Compliance

Carbon-aware routing decisions incur the natural costs of information processing (user location and request properties) and privacy concerns. Strict data protection laws, especially the General Data Protection Regulation (GDPR) of the European Union and the California Consumer Privacy Act (CCPA), present limitations on cross-border data transfers and the location of data processing decisions. The Green Ad Tech framework treats these requirements as hard constraints: routing decisions must respect data residency requirements, excluding candidate regions legally prohibited from handling particular requests. This integration of constraints has practical effects: European requests falling under GDPR may be restricted to European Economic Area regions, regardless of carbon intensity elsewhere, which may limit optimization opportunities. In a wider context, considering privacy and data protection as non-negotiable limits are ASEAN measures of both necessities in the ethical and legal context. Even so, the reduction of carbs is not worth sacrificing the privacy of people or the abuse of loopholes in the regulations.

Conclusion

Green Ad Tech establishes a foundational framework for lowering carbon emissions in programmatic digital advertising while sustaining the ultra-low latency and economic performance that is crucial for real-time bidding ecosystems. By incorporating carbon-aware request routing, adaptive model switching, smart edge caching, and auction-level emission optimization, the architecture proves that significant emission reductions can be achieved within operational constraints of production systems serving billions of daily impressions with less than 100ms latency. Algorithmic techniques specifically designed for microsecond-scale decision windows enable real-time carbon optimization without violating service-level objectives, while measurement infrastructure aligned with emerging industry standards provides the auditability necessary to build stakeholder trust and avoid accusations of greenwashing. The framework's broader significance extends beyond advertising to provide a template for carbon-aware optimization in other ultra-low latency distributed systems, including financial trading platforms, real-time content recommendation services, and interactive multiplayer gaming infrastructure that face similar tensions between environmental objectives and stringent performance requirements. Economic incentive systems that match carbon targets across several independent players—advertisers, publishers, ad exchanges, and content delivery networks—convert voluntary corporate social responsibility into quantifiable business considerations integrated into operational decision-making. Equity concerns addressing possible disadvantages for publishers in carbon-intensive areas guarantee that sustainability projects do not worsen world economic disparities or unfairly handicap impoverished nations with historic fossil-dependent power systems by means of temporal flexibility, fairness restrictions, and capacity-building investments. Viewing privacy and regulatory compliance as nonnegotiable limitations shows that environmental goals have to honor more general social values and legal systems rather than trying to cut carbon in isolation from other vital factors, including individual data rights and cross-border data protection needs. The climate crisis demands transformation across all sectors, including digital infrastructure, often mistakenly assumed to be inherently clean due to its intangibility, and programmatic advertising, with its global scale and precise measurability, presents both a significant emissions source and an opportunity for innovation in carbon-aware system design that can inform sustainable computing practices across the technology industry.

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

References

[1] PwC, "Navigating Net Zero: Our Progress and Challenges on the Journey," 2025. [Online]. Available: https://www.pwc.com/mt/en/publications/sustainability/navigating-net-zero-our-progress-and-challenges-on-the-journey.html
[2] UN Global Compact, "Ad Net Zero Global Media Sustainability Framework,". [Online]. Available: https://unglobalcompact.org/cmo-blueprint-for-sustainable-growth/resources/ad-net-zero-global-media-sustainability-framework
[3] Montel, "Carbon intensity tracking and forecasting," [Online]. Available: https://montel.energy/products/analytics/carbon-intensity-toolkit
[4] Shehabi et al., "2024 United States Data Center Energy Usage Report," Lawrence Berkeley National Laboratory, 2024. [Online]. Available: https://escholarship.org/uc/item/32d6m0d1
[5] Rajat Sen, Hsiang-Fu Yu, Inderjit S. Dhillon, "Think Globally, Act Locally: A Deep Neural Network Approach to High-Dimensional Time Series Forecasting," Amazon Science, 2019. [Online]. Available: https://www.amazon.science/publications/think-globally-act-locally-a-deep-neural-network-approach-to-high-dimensional-time-series-forecasting
[6] Arooj Arif, et al., "Electricity Load and Price Forecasting Using Machine Learning Algorithms in Smart Grid: A Survey," SpringerNature Link, 2020. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-44038-1_43
[7] Greenhouse Gas Protocol, "Corporate Value Chain (Scope 3) Accounting and Reporting Standard," World Resources Institute and World Business Council for Sustainable Development, [Online]. Available: https://ghgprotocol.org/sites/default/files/standards/Corporate-Value-Chain-Accounting-Reporing-Standard_041613_2.pdf
[8] Joshua Aslan, et al., "Electricity Intensity of Internet Data Transmission: Untangling the Estimates," Wiley Online Library, 2017. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1111/jiec.12630
[9] Lauren Wakefield, "State of Readiness: Sustainability in Digital Advertising Report 2025," IAB Europe, 2025. [Online]. Available: https://iabeurope.eu/knowledge-hub/state-of-readiness-sustainability-in-digital-advertising-report-2025/
[10] Charlotte Freitag et al., "The Real Climate and Transformative Impact of ICT: A Critique of Estimates, Trends, and Regulations," ScienceDirect, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666389921001884