2026 41st Youth Academic Annual Conference of Chinese Association of Automation (YAC) | 979-8-3195-2143-9/26/$31.00 ©2026 IEEE | DOI: 10.1109/YAC71005.2026.11161518

Blockchain-Enabled Equilibrium for Low-Carbon Computing Power Routing: A Tri-Level Stackelberg Approach

1st Zixuan Zhou
School of Advanced Manufacturing and Robotics
Peking University
Beijing, China
zxzhou@pku.edu.cn

2nd Junkai Cheng
School of Advanced Manufacturing and Robotics
Peking University
Beijing, China
2501111774@stu.pku.edu.cn

3rd Jiawei Feng
Sustainability X-Lab
The University of Hong Kong
Hong Kong, China
fengjiawei126@gmail.com

4th Wenjun Mei*
School of Advanced Manufacturing and Robotics
Peking University
Beijing, China
mei@pku.edu.cn

***Abstract*—This work proposes a game theory-based framework for jointly optimizing task routing, carbon-aware pricing, and GPU resource allocation in geographically distributed edge computing networks. We construct a three-layer Stackelberg game model where government regulators set carbon taxes, platform operators determine prices between heterogeneous data centers, and users choose computing destinations under random information. A blockchain transparency mechanism is introduced to eliminate information asymmetry regarding user perception of congestion and carbon costs. We construct a computing capacity wall model (similar to a macroscopic basic graph in transportation) to capture throughput collapses caused by excessive server load. The underlying user equilibrium is solved using a logit-based random allocation algorithm combined with continuous averaging. Numerical experiments on a 24-node network with 8 data centers show that the blockchain transparency mechanism reduces total latency by 29% by improving spatial load distribution; and carbon-differentiated pricing improves social welfare by 93% compared to uniform pricing. Sensitivity analysis reveals a severe GPU congestion trap at a task ratio of 50-60% and confirms that the carbon tax is an assessment tool, not a routing control mechanism.**

***Index Terms*—Blockchain, Stackelberg game, Carbon emission, User equilibrium, GPU scheduling**

I. Introduction

Global data centers account for approximately 1% to 1.5% of global electricity consumption, and this figure is expected to continue to rise [1]. The rapid growth of artificial intelligence workloads has transformed computing power from a commodity into an infrastructure requiring network-level coordination [2]. Data centers with high power load requirements are often isolated from the urban power grid and located in relatively energy-rich areas. Therefore, computing tasks need to be routed from various users to the data center via network transmission [3]. However, existing platforms suffer from three fundamental inefficiencies: (1) users cannot understand the congestion and energy composition of the data center in real time, resulting in suboptimal task routing; (2) uniform pricing fails to reflect the differences in carbon footprints between data centers with different energy structures; and (3) the scarcity of GPU resources leads to nonlinear congestion externalities, which traditional queuing models struggle to capture.

The intersection of computing networks and transportation science provides a powerful analytical framework. Just as urban transportation networks exhibit congestion that can be simulated using the Wardrop equilibrium model [4], computing power networks exhibit similar phenomena. Users route tasks through the network to data centers as their 'destination', experiencing both network latency and computational latency. This similarity has been applied to congestion game models for computing resource allocation [5], [6] and Stackelberg game models for network pricing in transportation [7] and communications [8], but rarely considers carbon emissions. Multi-layer Stackelberg models have also been applied to cloud computing and fog computing resource allocation [9], but typically optimize only within a single data center or use simplified demand models rather than a complete spatial network equilibrium.

Blockchain technology offers a natural solution for information asymmetry. By recording real-time DC load, energy sources, and carbon emission intensity on an immutable ledger, blockchain enables users to make fully informed routing decisions [10], [11]. Drawing on the value blockchain has demonstrated in supply chain transparency [12], [13] and energy markets [14], we design a mechanism that leverages

This work is supported by the National Natural Science Foundation of China 72595832.
* Corresponding author

this transparency to transform the game from a noisy Bayesian equilibrium to a fully informed equilibrium with significant efficiency improvements. Geographic load balancing across data centers has been shown to reduce the use of green energy [15], while carbon-sensing online control achieves a trade-off between emissions and costs [16]. However, these approaches typically adopt a centralized operator's perspective, while our model captures decentralized user routing under carbon-differentiated pricing.

In this work, we propose *computing power wall*, an analogue of the macroscopic fundamental diagram (MFD) [17], [18], which captures throughput collapse when active tasks exceed a DC's thermal-electrical threshold. This concept extends MFD theory from transportation [19] to computing systems, where thermal throttling creates an analogous throughput-density relationship [20]. First, we formulate a tri-level Stackelberg game integrating carbon taxation (government), differentiated pricing (platform), and stochastic user equilibrium (users) within a unified network framework (Section II). Second, we introduce the computing power wall model to capture DC throughput degradation under overload (Section II-C). Third, we provide numerical evidence that blockchain transparency reduces system delay by 29%, that carbon-differentiated pricing outperforms uniform pricing by 93% in welfare, and that a GPU congestion trap emerges at moderate GPU task ratios (Section IV).

II. System Model

A. Network and Data Center Architecture

Consider a directed network $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ with $|\mathcal{N}| = N$ nodes and $|\mathcal{L}| = L$ links, where a subset $\mathcal{D} \subset \mathcal{N}$ with $|\mathcal{D}| = M$ denotes DC locations. Each link $\ell \in \mathcal{L}$ has free-flow travel time t_ℓ^0 and capacity C_ℓ. Link delay follows the Bureau of Public Roads (BPR) function:

$$t_\ell(x_\ell) = t_\ell^0 \left[1 + \alpha \left(\frac{x_\ell}{C_\ell}\right)^\beta\right], \tag{1}$$

where x_ℓ is link flow and $\alpha = 0.15$, $\beta = 4$ are standard parameters [21].

Each DC $n \in \mathcal{D}$ is characterized by the tuple $(\kappa_n^{\text{cpu}}, \kappa_n^{\text{gpu}}, \gamma_n, \xi_n, N_n^{\text{jam}})$ representing CPU cores, GPU cards, carbon intensity factor (kgCO$_2$/kWh), power usage effectiveness (PUE), and thermal jamming threshold, respectively. Computing tasks are heterogeneous: a fraction $r \in [0,1]$ requires GPU acceleration, while the remainder uses CPU only.

B. Computing Delay Models

1) CPU Processing (M/M/1 Queue): The base computing delay at DC n follows an M/M/1 model with service rate $\mu_n^{\text{cpu}} = \kappa_n^{\text{cpu}} \cdot \mu^{\text{cpu}}$:

$$d_n^{\text{cpu}}(\lambda_n) = \frac{d_n^0}{1-\rho_n}, \quad \rho_n = \frac{\lambda_n}{\mu_n^{\text{cpu}}}, \tag{2}$$

where λ_n is the CPU task arrival rate and d_n^0 is the minimum processing time.

2) GPU Queue (M/M/c): GPU tasks at DC n with $c_n = \kappa_n^{\text{gpu}}$ servers and per-server rate μ^{gpu} experience waiting time:

$$w_n^{\text{gpu}} = \frac{C(c_n, A_n)}{c_n \mu^{\text{gpu}}(1-\rho_n^{\text{gpu}})}, \tag{3}$$

where $A_n = \lambda_n^{\text{gpu}}/\mu^{\text{gpu}}$ is the offered load and $C(c, A)$ is the Erlang-C probability.

C. Computing Power Wall

Drawing on the MFD from transportation science [17], [18], we model a *computing power wall* that captures throughput collapse under excessive loading. When active tasks N_n at DC n exceed the thermal-electrical threshold N_n^{jam}, frequency throttling degrades per-task processing speed:

$$v_n(N_n) = \max\left(0,\ 1 - \frac{N_n}{N_n^{\text{jam}}}\right), \tag{4}$$

yielding aggregate throughput $P_n(N_n) = N_n \cdot v_n(N_n)$, an inverted-U curve analogous to the MFD. The effective computing delay becomes:

$$d_n^{\text{eff}}(\lambda_n) = \frac{d_n^{\text{cpu}}(\lambda_n)}{v_n(N_n)}, \tag{5}$$

capturing the feedback loop between load and throughput that standard queueing models miss.

D. Carbon Emission Model

Total system carbon emission decomposes into network transport and DC computation:

$$E = \underbrace{\sum_{\ell \in \mathcal{L}} e^{\text{net}} x_\ell}_{E^{\text{net}}} + \underbrace{\sum_{n \in \mathcal{D}} \gamma_n \xi_n \left(P^{\text{cpu}} w^{\text{cpu}} \lambda_n^{\text{cpu}} + P^{\text{gpu}} w^{\text{gpu}} \lambda_n^{\text{gpu}}\right)}_{E^{\text{dc}}}, \tag{6}$$

where P^{cpu}, P^{gpu} are per-unit power consumption (W) and w^{cpu}, w^{gpu} are task durations (h). The carbon heterogeneity across DCs is substantial: a hydropower DC ($\gamma = 0.15$) produces roughly one-sixth the emissions of a coal-powered DC ($\gamma = 0.85$) per unit of computation after PUE amplification.

E. Blockchain Information Mechanism

Without blockchain, users perceive DC generalized costs with structured bias:

$$\tilde{g}_n^o = g_n + \sigma \cdot b_n^o \cdot g_n, \tag{7}$$

where σ is the opacity parameter ($\sigma = 0.4$ without blockchain, $\sigma = 0$ with blockchain) and b_n^o captures systematic misperception. We model two bias components: (i) a *capacity bias* where users underestimate congestion at large DCs, and (ii) a *carbon blindness* where users are unaware of DC emission differences, tending to select high-carbon DCs that appear cheaper:

$$b_n^o = -0.8\left(\frac{\kappa_n^{\text{cpu}}}{\bar{\kappa}^{\text{cpu}}} - 1\right) - 0.5 \cdot \hat{\gamma}_n + \varepsilon_n^o, \tag{8}$$

where $\hat{\gamma}_n$ is the normalized carbon factor and $\varepsilon_n^o \sim \mathcal{N}(0,1)$. Blockchain eliminates both biases by publishing real-time load, queue length, and certified carbon intensity on an immutable ledger.

III. Tri-Level Stackelberg Formulation

A. Three-Player Hierarchy

Upper level (Government): Sets carbon tax τ (CNY/ton) to maximize social welfare:

$$\max_{\tau \geq 0} \; W(\tau) = -\text{VOT} \cdot \text{TDT} + R - \tau \cdot E, \tag{9}$$

where $\text{TDT} = \sum_\ell x_\ell t_\ell(x_\ell)$ is total delay, R is platform revenue, and E is total emissions.

Middle level (Platform): Determines DC-specific pricing $(p_n^{\text{cpu}}, p_n^{\text{gpu}}, s_n)$ for CPU price, GPU price, and subsidy to maximize welfare subject to the user equilibrium response:

$$\max_{\boldsymbol{p}, \boldsymbol{s} \geq 0} \; W(\boldsymbol{p}, \boldsymbol{s}; \tau), \tag{10}$$

with pricing parameterized as $p_n = p^{\text{base}}(1 + \eta\hat{\gamma}_n)$, where $\eta \in [-0.3, 0.3]$ is the carbon premium/discount.

Lower level (Users): Each origin o distributes demand q_o across DCs according to a logit-based stochastic user equilibrium model [22]:

$$f_n^o = q_o \cdot \frac{\exp(-\theta\, \tilde{g}_n^o)}{\sum_{m \in \mathcal{D}} \exp(-\theta\, \tilde{g}_m^o)}, \tag{11}$$

where θ is the scale parameter and $\tilde{g}_n^o$ is the perceived generalized cost at DC n from origin o, comprising network delay, computing delay, GPU queue time, pricing, subsidy, and information bias.

B. Solution Algorithm

The lower-level stochastic user equilibrium is solved via the Method of Successive Averages (MSA) [21]: at each iteration, Dijkstra's algorithm computes shortest paths from every origin to all DC nodes, logit probabilities distribute flow, and a convex combination updates link flows with step size $1/k$ at iteration k.

The middle-level pricing optimization employs Latin Hypercube Sampling (LHS) [23] over a 4-dimensional parameter space $(p_{\text{base}}^{\text{cpu}}, p_{\text{base}}^{\text{gpu}}, s_{\text{base}}, \eta)$ with carbon-differentiated pricing structure, evaluating 50 candidate pricing vectors. This low-dimensional parameterization ensures adequate coverage compared to sampling individual DC prices in a $3M$-dimensional space. The upper-level carbon tax is analyzed parametrically.

IV. Numerical Experiments

A. Test Network Setup

We construct a computing power network based on the Sioux Falls topology [24]: 24 nodes, 76 directed links, with 8 DCs placed at nodes reflecting remote computing hub layout. Table I summarizes DC heterogeneity. Total demand is calibrated to 60% GPU bottleneck utilization (12,614 tasks/h), with a 30% GPU task ratio.

Eight scenarios are designed to answer six research questions through controlled factor isolation, as shown in Table II.

TABLE I
Data Center Parameters

DC	Node	CPU	GPU	γ_n	PUE	Energy	N_n^{jam}
A	3	10000	1416	0.25	1.25	Mix	12000
B	7	5000	692	0.85	1.45	Coal	6000
C	9	4500	642	0.55	1.35	Gas	5400
D	10	8000	1047	0.60	1.30	Wind	9600
E	16	5500	771	0.50	1.35	Mix	6600
F	17	2500	333	0.70	1.50	Coal	3000
G	20	6000	854	0.15	1.20	Hydro	7200
H	22	4000	552	0.55	1.40	Solar	4800

CPU: cores; GPU: cards; γ_n: kgCO_2/kWh; PUE: power usage effectiveness; $N_n^{\text{jam}} = 1.2\kappa_n^{\text{cpu}}$.

TABLE II
Scenario Design and Research Questions

ID	Name	GPU	BC	PW	Research Question
S1	CPU-Base	×	×	✓	Q1: BC on CPU
S2	CPU-BC	×	✓	✓	
S3	GPU-Trad	✓	×	✓	Q2: BC+GPU synergy
S4	GPU-BC	✓	✓	✓	
S5	Perfect	✓	✓	✓	Q3: Information bound
S6	No-PW	✓	✓	×	Q4: Power wall cost
S7	Fix-Price	✓	✓	✓	Q5: Pricing value
S8	Syst-Opt	✓	✓	✓	Q6: Efficiency gap

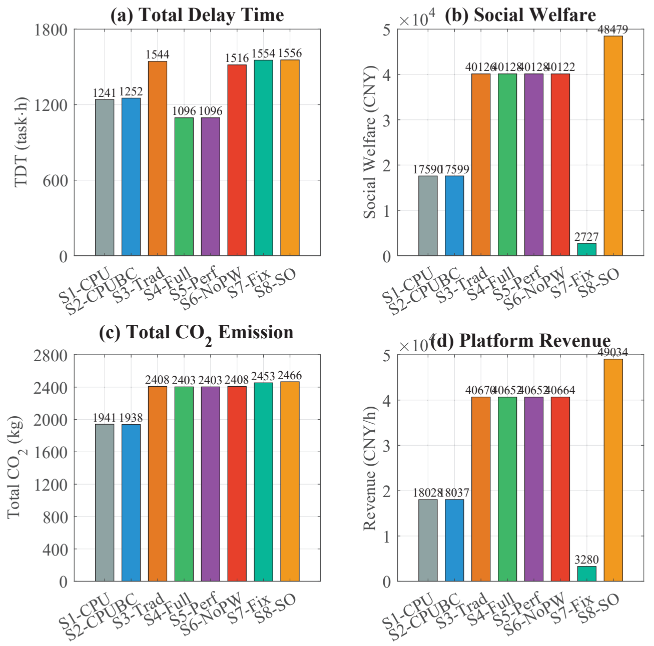

Fig. 1. Eight-scenario comparison: (a) total delay time, (b) social welfare, (c) total CO_2 emission, (d) platform revenue.

B. Main Results: Eight-Scenario Comparison

Table III and Fig. 1 present the core results.

Comparing S3 (no BC, TDT=1,544) with S4 (BC, TDT=1,096) reveals a 29% reduction in total delay time. Blockchain transparency eliminates the structured perception

TABLE III
EQUILIBRIUM RESULTS ACROSS EIGHT SCENARIOS

Scenario	TDT (task·h)	Welfare (CNY)	CO_2 (kg)	Revenue (CNY/h)
S1-CPU	1,241	17,590	1,941	18,028
S2-CPUBC	1,252	17,599	1,938	18,037
S3-Trad	1,544	40,126	2,408	40,670
S4-Full	1,096	40,128	2,403	40,652
S5-Perf	1,096	40,128	2,403	40,652
S6-NoPW	1,516	40,122	2,408	40,664
S7-Fix	1,554	2,727	2,453	3,280
S8-SO	1,556	48,479	2,466	49,934

bias that causes users to overcrowd large DCs while under-utilizing smaller ones. This delay reduction is achieved purely through information provision, without any change in physical infrastructure.

GPU scenarios (S3–S8) achieve welfare ∼40,000 CNY versus ∼17,600 for CPU-only (S1–S2), a 2.3× improvement. However, CO_2 emissions increase from 1,940 kg to 2,400 kg (24%), reflecting the higher power consumption of GPU tasks (400 Wh vs. 200 Wh per task). This quantifies the welfare–carbon tradeoff facing regulators.

S7 (uniform fixed pricing at below-market rates) achieves welfare of only 2,727 CNY—a 93% reduction compared to S4's optimized carbon-differentiated pricing (40,128 CNY). The 4-dimensional parameterized pricing structure, which adjusts prices based on each DC's carbon factor, effectively internalizes the environmental externality while maintaining market efficiency.

S8 (system optimal with expanded search) achieves welfare of 48,479 CNY, exceeding S4 by 21%. This gap represents the *price of anarchy*—the efficiency loss from decentralized user routing compared to centralized control [6].

C. Carbon Decomposition Analysis

Fig. 2 shows the carbon decomposition across scenarios. Network transport emissions account for approximately 6% of total carbon (∼140 kg), while DC computation dominates at 94%. Among GPU scenarios, S4 (BC) achieves marginally lower DC carbon than S3 (no BC): 2,403 vs. 2,408 kg. While this 0.2% reduction appears modest in aggregate, the mechanism is significant—blockchain guides users toward low-carbon DCs (e.g., DC-G with hydropower at $\gamma = 0.15$) and away from high-carbon alternatives (e.g., DC-B with coal at $\gamma = 0.85$).

D. Computing Power Wall Visualization

The power wall diagram (Fig. 3) plots each DC's operating point on its throughput-density curve. All DCs operate on the left (stable) side of their inverted-U curves, with utilization ranging from 9% (DC-F) to 16% (DC-D). DC-G, despite having the largest capacity, attracts substantial flow due to its low carbon cost and favorable pricing under the differentiated scheme.

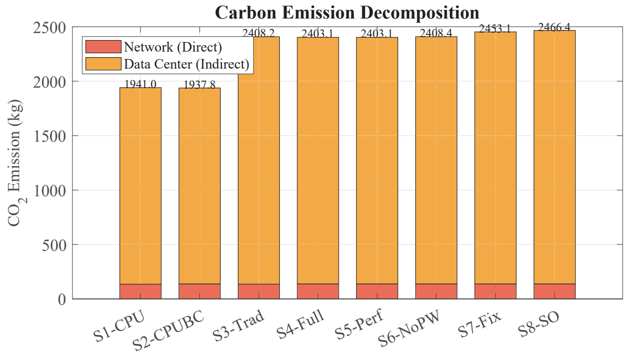

Fig. 2. Carbon emission decomposition: network transport (red) vs. data center computation (orange). DC emissions dominate at 94% of total.

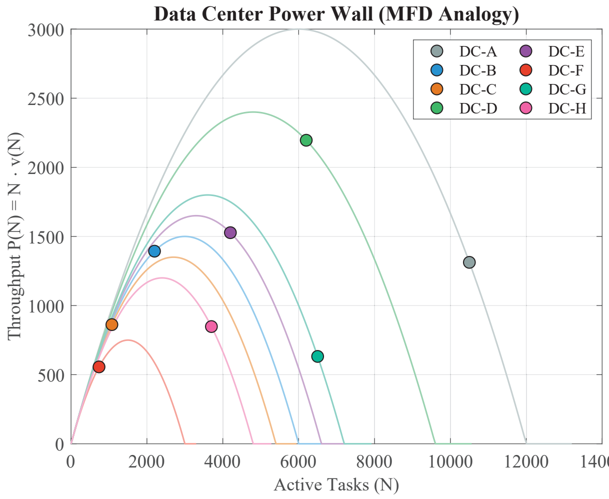

Fig. 3. Computing power wall (MFD analogy): throughput vs. active tasks. Operating points (circles) remain on the stable left side.

E. DC Load Distribution Under Blockchain

The spatial redistribution of computing tasks is a key channel through which blockchain improves system performance. Fig. 4 compares CPU and GPU task allocations across DCs between S3 (no BC) and S4 (BC). Under blockchain transparency, users shift away from high-carbon DCs (B, F) and toward the low-carbon hydropower DC-G. Although these per-DC shifts are modest in absolute terms, they represent the equilibrium response to removing systematic perception bias—the Logit model's smoothing effect distributes the information improvement across all origin-destination pairs.

The pattern is consistent: DCs with below-average carbon intensity ($\gamma < 0.52$) see increased load under blockchain, while those above average see decreased load. This carbon-sorting effect, combined with the pricing mechanism's carbon premium/discount, creates a positive feedback loop that steers the system toward lower aggregate emissions.

F. Sensitivity Analysis

1) GPU Task Ratio: Fig. 5 shows system behavior as the GPU task ratio varies from 10% to 90%. Three regimes

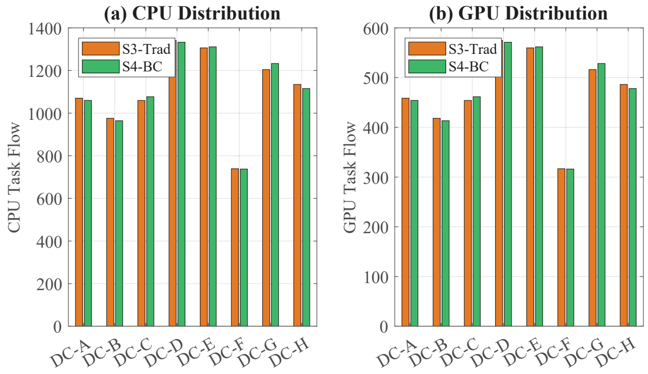

Fig. 4. DC load distribution comparison: (a) CPU task flow and (b) GPU task flow across 8 DCs under S3 (no blockchain) vs. S4 (blockchain transparency).

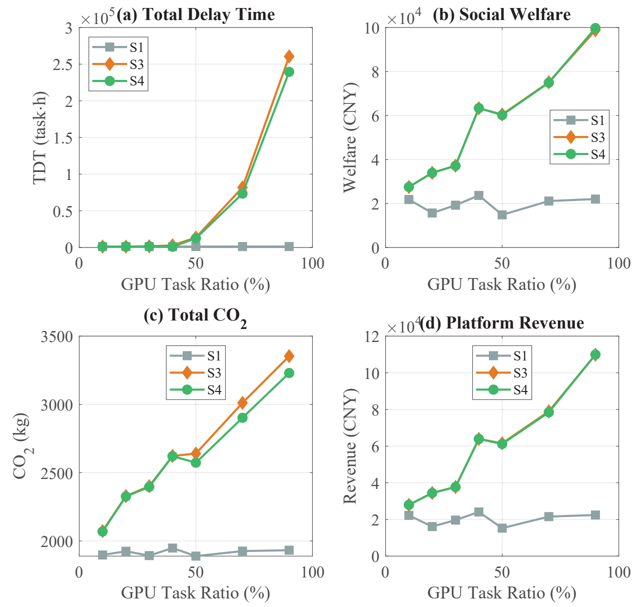

Fig. 5. Sensitivity to GPU task ratio: (a) TDT shows congestion trap at 50–60%; (b) welfare gap between S3 and S4 widens with GPU ratio; (c) CO_2 scales linearly with GPU share; (d) revenue grows monotonically.

emerge: (i) for GPU ratio below 40%, TDT remains low and scales linearly; (ii) between 50–60%, a *congestion trap* appears as GPU queues begin saturating (the M/M/c queue delay grows superlinearly as $\rho \to 1$); (iii) above 70%, TDT increases sharply to 2.5×10^5 task·h. Throughout, S4 (BC) maintains lower TDT and higher welfare than S3 (no BC), with the gap widening at higher GPU ratios—demonstrating that blockchain's benefit *amplifies under resource scarcity*. S1 (CPU-only) remains flat across all GPU ratios, confirming the correct isolation of GPU effects in the sensitivity design.

The welfare curves reveal an important finding: as GPU ratio increases, revenue grows monotonically (GPU pricing exceeds CPU pricing), but the marginal welfare gain diminishes above 50% due to rapidly increasing delay costs. This suggests an optimal GPU allocation exists that balances revenue generation against congestion externalities.

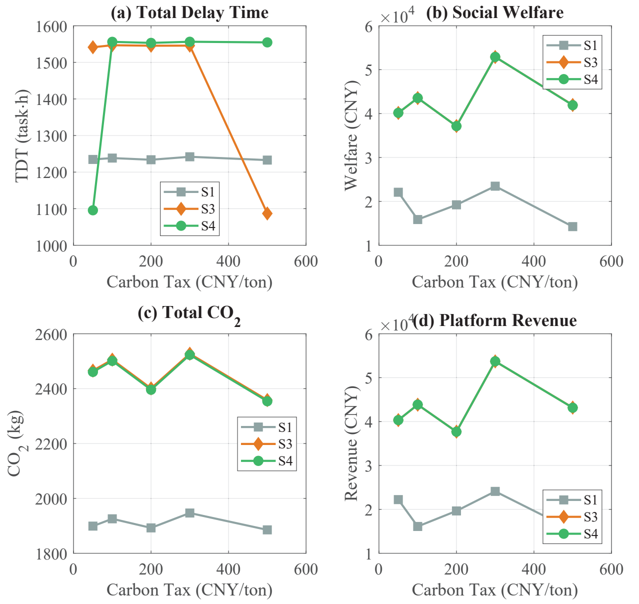

Fig. 6. Sensitivity to carbon tax: (a) TDT insensitive; (b) welfare stable; (c) CO_2 nearly constant across 50–500 CNY/ton; (d) revenue unaffected.

2) Carbon Tax: Fig. 6 displays the response to carbon tax rates from 50 to 500 CNY/ton. Both TDT and CO_2 emissions remain essentially insensitive to the tax level, with CO_2 varying by less than 5% across a 10× tax range. This insensitivity arises because the per-task carbon cost ($\tau \cdot \gamma_n \cdot \xi_n \cdot P \cdot w/10^6$ CNY) remains two orders of magnitude below the dominant cost components (computing delay valued at 40 CNY/h and DC pricing at 5–10 CNY/task). To achieve meaningful routing impact, carbon taxes would need to be at least 5,000 CNY/ton—a level far beyond current policy proposals. This confirms that in our model, carbon taxation functions primarily as a *fiscal and evaluation instrument* rather than a routing control mechanism.

V. Discussion

This work suggest three actionable insights for computing power network operators. First, investing in blockchain-based transparency infrastructure yields greater delay reduction (29%) than proportional capacity expansion would achieve at comparable cost. Second, carbon-differentiated pricing—automatically adjusting fees based on certified DC carbon intensity—is essential for both welfare maximization and emission steering. The 93% welfare gap between differentiated and uniform pricing (S4 vs. S7) underscores the inadequacy of one-size-fits-all pricing in heterogeneous networks. Third, regulators should not rely on carbon taxation alone to induce green routing; instead, direct information provision through blockchain proves more effective at current tax levels.

The computing power wall concept provides a new modeling primitive that captures throughput collapse phenomena absent from standard M/M/c queueing models. By importing the macroscopic fundamental diagram framework from transportation science [18], we enable a unified treatment of both network congestion and computing congestion within the same Stackelberg equilibrium formulation. This cross-disciplinary transfer opens new avenues for applying decades of traffic flow theory to emerging computing resource allocation problems.

The bilevel pricing mechanism with 4-dimensional parameterization demonstrates that structured low-dimensional search can dramatically outperform high-dimensional random sampling. In our experiments, the 4-dimensional approach with 50 samples significantly outperformed the naive $3M = 24$-dimensional alternative, which would require orders of magnitude more samples for comparable coverage.

The current model assumes static demand and single-period analysis. Dynamic task arrival patterns with time-varying demand, multi-period pricing adaptation using reinforcement learning [8], and strategic user behavior (e.g., task timing and splitting) represent important extensions. The bias model, while structurally motivated by information economics theory [12], uses calibrated rather than empirically estimated parameters. Field data from actual computing power exchanges would enable more precise calibration. Finally, the power wall parameters (N_n^{jam}) are currently set proportional to CPU capacity; more refined thermal-electrical modeling incorporating ambient temperature, cooling system design, and chip-level throttling characteristics would strengthen the physical foundations of the model.

VI. Conclusion

We presented a tri-level Stackelberg framework for blockchain-enabled low-carbon computing power routing. The model integrates transportation network equilibrium with computing resource allocation, introducing the power wall concept to capture throughput collapse under thermal constraints. Experiments on a 24-node, 8-DC network demonstrate that blockchain transparency reduces delay by 29%, carbon-differentiated pricing improves welfare by 93% over uniform pricing, and a GPU congestion trap emerges at 50–60% task ratios. Future work will extend the model to dynamic settings with time-varying demand and investigate learning-based pricing mechanisms.

References

[1] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984–986, 2020.

[2] S. Yukun, L. Bo, L. Junlin, H. Haonan, Z. Xing, P. Jing, and W. Wenbo, "Computing power network: A survey," *China Communications*, vol. 21, no. 9, pp. 109–145, 2024.

[3] J. Zeng and M. Yang, "Digital technology and carbon emissions: Evidence from china," *Journal of Cleaner Production*, vol. 430, p. 139765, 2023.

[4] J. G. Wardrop, "Road paper. some theoretical aspects of road traffic research." 1952.

[5] T. Roughgarden, *Selfish Routing and The Price of Anarchy*, 01 2005, vol. 74.

[6] R. Johari and J. N. Tsitsiklis, "Efficiency loss in a network resource allocation game," *Mathematics of Operations Research*, vol. 29, no. 3, pp. 407–435, 2004.

[7] L. Brotcorne, M. Labbé, P. Marcotte, and G. Savard, "A bilevel model for toll optimization on a multicommodity transportation network," *Transportation Science*, vol. 35, no. 4, pp. 345–358, 2001.

[8] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter, "A survey on networking games in telecommunications," *Computers & Operations Research*, vol. 33, no. 2, pp. 286–311, 2006, game Theory: Numerical Methods and Applications.

[9] H. Zhang, Y. Xiao, S. Bu, D. Niyato, F. R. Yu, and Z. Han, "Computing resource allocation in three-tier iot fog networks: A joint optimization approach combining stackelberg game and matching," *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1204–1215, 2017.

[10] F. Casino, T. K. Dasaklis, and C. Patsakis, "A systematic literature review of blockchain-based applications: Current status, classification and open issues," *Telematics and Informatics*, vol. 36, pp. 55–81, 2019.

[11] Y. Wang, J. H. Han, and P. Beynon-Davies, "Understanding blockchain technology for future supply chains: a systematic literature review and research agenda," *Supply Chain Management: An International Journal*, vol. 24, no. 1, pp. 62–84, 2019.

[12] V. Babich and G. Hilary, "Om forum—distributed ledgers and operations: What operations management researchers should know about blockchain technology," *Manufacturing & Service Operations Management*, vol. 22, no. 2, pp. 223–240, 2020.

[13] J. Chod, N. Trichakis, G. Tsoukalas, H. Aspegren, and M. Weber, "On the financing benefits of supply chain transparency and blockchain adoption," *Management Science*, vol. 66, no. 10, pp. 4378–4396, 2020.

[14] E. Mengelkamp, J. Gärttner, K. Rock, S. Kessler, L. Orsini, and C. Weinhardt, "Designing microgrid energy markets: A case study: The brooklyn microgrid," *Applied Energy*, vol. 210, pp. 870–880, 2018.

[15] Z. Liu, M. Lin, A. Wierman, S. Low, and L. L. H. Andrew, "Greening geographical load balancing," *IEEE/ACM Transactions on Networking*, vol. 23, no. 2, pp. 657–671, 2015.

[16] Z. Zhou, F. Liu, R. Zou, J. Liu, H. Xu, and H. Jin, "Carbon-aware online control of geo-distributed cloud services," *IEEE Transactions on Parallel and Distributed Systems*, vol. 27, no. 9, pp. 2506–2519, 2016.

[17] N. Geroliminis and C. F. Daganzo, "Existence of urban-scale macroscopic fundamental diagrams: Some experimental findings," *Transportation Research Part B: Methodological*, vol. 42, no. 9, pp. 759–770, 2008.

[18] C. F. Daganzo and N. Geroliminis, "An analytical approximation for the macroscopic fundamental diagram of urban traffic," *Transportation Research Part B: Methodological*, vol. 42, no. 9, pp. 771–781, 2008.

[19] Y. Ji and N. Geroliminis, "On the spatial partitioning of urban transportation networks," *Transportation Research Part B: Methodological*, vol. 46, no. 10, pp. 1639–1656, 2012.

[20] H. Esmaeilzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in *2011 38th Annual International Symposium on Computer Architecture (ISCA)*, 2011, pp. 365–376.

[21] M. Patriksson, *The Traffic Assignment Problem: Models and Methods*, 02 2015.

[22] Y. Sheffi, *Urban Transportation Networks: Equilibrium Analysis With Mathematical Programming Methods*, 01 1984.

[23] M. Mckay, R. Beckman, and W. Conover, "A comparison of three methods for selecting vales of input variables in the analysis of output from a computer code," *Technometrics*, vol. 21, pp. 239–245, 05 1979.

[24] L. J. LeBlanc, E. K. Morlok, and W. P. Pierskalla, "An efficient approach to solving the road network equilibrium traffic assignment problem," *Transportation Research*, vol. 9, no. 5, pp. 309–318, 1975.