Equity-Aware Spatial-Temporal Workload Shifting for Sustainable AI Data Centers

Mohammad Jaminur Islam
University of California, Riverside

Shaolei Ren
University of California, Riverside

Abstract

The escalated demand for hyperscale data centers due to generative AI has intensified the operational load, leading to increased energy consumption, water usage, and carbon emissions. In this paper, we propose `EquiShift`, a novel equitable spatial-temporal workload balancing algorithm that shifts workloads spatially and temporarily across geographically different data centers to minimize the overall energy costs while ensuring fair distribution of water and carbon footprints. Concretely, `EquiShift` introduces a model predictive control (MPC) framework to solve the equitable load balancing problem, leveraging the predictive capabilities of MPC to optimize load distribution in real-time. Finally, we present comparative evaluations against state-of-the-art load-balancing algorithms to demonstrate the performance of `EquiShift` which underscores the potential of equitable load balancing as a key strategy for enhancing the sustainability of data centers while achieving fairness in the face of growing computational demands.

1 Introduction

The rapid digitization and growth of generative AI have significantly increased demands on data centers, leading to heightened energy consumption. In 2022, data centers consumed approximately 460 TWh, accounting for around 2% of the global total electricity usage [1]. This large-scale energy consumption contributes to a significant portion of global greenhouse gas (GHG) emissions [2]. Additionally, training large language models such as GPT-3 in data centers can directly evaporate $700,000$ liters of freshwater, let alone the additional water usage for inference to serve millions of user requests [3].

To mitigate the environmental impact, industries have implemented various measures, including the adoption of green infrastructure such as wide-scale renewable energy, climate-aware cooling, hardware reuse, geographical workload balancing, and carbon-aware computing [4–11]. However, these solutions often lead to disproportionate benefits, causing certain regions to bear a heavier environmental burden due to an uneven distribution of computational loads. For example, a Google data center in Finland benefits from 23 times more green energy (and less fossil-fuel air pollution) compared to a facility in Singapore [12, 13], creating a regional disparity in carbon emissions. Similarly, Nordic countries exhibit better water efficiency (<1.0 L/kWh) compared to hotter regions (9.0 L/kWh), making water consumption in areas like Arizona disproportionately high [14].

Additionally, existing load shifting techniques such as geographical load balancing (GLB, which shifts workloads across different data centers) create unintentional disparities since their focus on minimizing the *total* energy consumption and carbon emissions can inadvertently increase water footprints putting severe stresses on limited freshwater resources and further exacerbate the uneven distribution of water usage. To address these regional disparities in environmental impact, an equitable GLB approach offers a promising solution. By prioritizing areas with greater environmental disadvantages and providing flexibility in spatial and temporal workload scheduling, this approach can help ensure a more balanced distribution of environmental burdens.

Tackling Climate Change with Machine Learning: Workshop at NeurIPS 2024.

Some recent studies [14] have also tackled equitable GLB using a dual mirror descent approach, but they only consider delay-sensitive AI inference workloads without exploiting deferrable workloads (e.g., AI model training and back-end processing). In contrast, our algorithm `EquiShift` leverages the spatial-temporal scheduling flexibility of data center workloads and employs Model Predictive Control (**MPC**) [15], enabling dynamic and adaptable workload shifting in an equitable manner.

2 Problem Formulation

We model delay-tolerant energy-intensive workloads, such as ML training tasks, that can be flexibly distributed among N geographically dispersed data centers connected via a common network. These workloads are suitable for deferral because they are not often time-sensitive, allowing them to be rerouted to data centers that are less environmentally impacted in a dynamic manner. Each data center $i \in [1, N]$ receives a portion $y_i(t)$ of the total incoming workload $\lambda(t)$ at time $t \in [1, T]$, where T is the time horizon. The associated energy cost, water usage, and carbon emissions at each data center are denoted by $f_i(y_i(t))$, $w_i(y_i(t))$, and $c_i(y_i(t))$, respectively, and are also time-varying functions for which we omit the time index for brevity.

`EquiShift` minimizes the total energy cost while also reducing the maximum water and carbon footprints among the data centers (i.e., minimax). By prioritizing the regions most affected by water usage and carbon emissions, `EquiShift` ensures an equitable and sustainable distribution of these impacts over the entire time horizon T.

$$\min_{y} \left[\sum_{t=1}^{T} \sum_{i=1}^{N} f_i(y_i(t)) + \kappa_w \left(\max_i \sum_{t=1}^{T} w_i(y_i(t)) \right) + \kappa_c \left(\max_i \sum_{t=1}^{T} c_i(y_i(t)) \right) \right] \tag{1a}$$

$$\text{s.t.} \quad \sum_{i=1}^{N} y_i(t) \leq \lambda(t) + \delta_{(t-1)}, \ \forall t \tag{1b}$$

$$\sum_{t=1}^{T} \sum_{i=1}^{N} y_i(t) = \sum_{t=1}^{T} \lambda(t), \tag{1c}$$

$$y_i(t) \leq M_i, \ \forall i, t \tag{1d}$$

In Eqn. (1), κ_w and κ_c are positive constants representing the relative importance of minimizing water and carbon footprints, respectively. The constraint in Eqn. (1b) limits the total distributed workload by the available workload at any time $t \in [1, T]$, plus any left-over workload from the previous time step $t-1$. The constraint in Eqn. (1c) ensures that the total workload distributed across all data centers matches the total incoming workload by the end of the time horizon T. Additionally, Eqn. (1d) ensures that no data center is overloaded beyond its capacity M_i. Details on the computation of energy, water, and carbon footprints are provided in Appendix A.

3 `EquiShift`: Equity-Aware Spatial-Temporal Workload Shifting using MPC

The formulation in (1) assumes complete knowledge of future contexts (i.e., an offline method) such as workloads, prices of electricity, water, and carbon efficiencies over the entire time horizon T. Therefore, it is ill-suited for the dynamic reality of data center operations. In response, we adopt a model predictive control (MPC) [15] approach to solve our equitable load balancing problem , which has also been implemented in real production-systems for sustainability without addressing equity [16]. Formally, for a prediction horizon K (length of the forecast future context), the following objective function can be formulated for the MPC approach.

$$\begin{aligned} \min_{y[t:t+K]} \Big[\Big(\sum_{k=0}^{K} \sum_{i=1}^{N} f_i(y(t+k)) + \sum_{\tau=1}^{t-1} \sum_{i=1}^{N} f_i(y(\tau)) \Big) + \kappa_w \max_i \Big(\sum_{k=0}^{K} w_i(y_i(t+k)) \\ + \sum_{\tau=1}^{t-1} w_i(y_i(\tau)) \Big) + \kappa_c \max_i \Big(\sum_{k=0}^{K} c_i(y_i(t+k)) + \sum_{\tau=1}^{t-1} c_i(y_i(\tau)) \Big) \Big] \end{aligned} \tag{2a}$$

$$s.t., \quad \text{(1b), (1d)} \tag{2b}$$

$$\sum_{k=0}^{K} \sum_{i=1}^{N} y_i(t+k) = \sum_{k=0}^{K} \lambda(t+k) + \delta_{(t-1)} \tag{2c}$$

In Eqn. (2a), the accumulated cost up to time $t-1$ represents the historical cost, serving as a penalizing factor in the optimization. This penalization influences the current workload distribution to avoid further increases in energy costs, water usage, and carbon footprints. The MPC optimization adheres to the constraints defined in Eqns. (1b) and (1d), along with the additional constraint in Eqn. (2c). The latter constraint imposes a limit on the total distributed workload across the N data centers over the predicted horizon K, ensuring it does not exceed the sum of the predicted incoming workload and the residual workload from previous time steps ($\delta_{(t-1)}$).

MPC iteratively solves this optimization problem, refining its decisions based on feedback from past performance. Although the MPC solution typically deviates from the offline solution defined in Eqn. (1) due to incomplete context information and prediction errors, it has the advantage of continuous adjustment and offers good robustness in practice. By dynamically incorporating evolving contextual information, such as electricity prices, water efficiency, and carbon efficiency, the MPC solution often remains close to the optimal offline solution.

4 Setup

We simulate a network of $N = 10$ geographically distributed data centers with homogeneous capacity, interconnected through a single gateway, located in the United States (Virginia, Georgia, Texas, and Nevada), Europe (Belgium, the Netherlands, Germany, and Denmark), and Asia (Singapore and Japan). To model the workload, we use GPU power usage data used in [14] over 18 days from September 23rd to October 11th, 2022 which is later augmented to match the demand of 10 datacenters. Although this data is from the inference phase, we will also test other workload traces over a multi-month length and defer the results to Appendix C. Electricity price data for this period is sourced from [17] and [18], with country-level data for Europe and Asia, and state-level data for the U.S. We compute hourly water efficiency for each region following the methodologies in [19] and [8], using hourly fuel mix data from [18], while carbon efficiency data is obtained from [20] (more details in Appendix B). Our `EquiShift` approach is compared in terms of cost savings and equity against six representative algorithms: (1) Energy and Water GLB [21] (**EWShift**) minimizes total energy cost and water footprint, (2) Energy and Carbon GLB (**ECShift**) [8] minimizes total energy cost and carbon footprint, (3) Energy GLB (**EnShift**) minimizing total energy costs [9–11], (4) Water GLB (**WaShift**) minimizes the total water footprint, (5) Carbon GLB (**CaShift**) minimizes the total carbon footprint, and (6) Total GLB (**ToShift**) [22] jointly minimizes the total energy cost, water footprint, and carbon footprint. We compute the cost savings of our equitable algorithm compared to other algorithms based on $\eta = \frac{\mathcal{C}_{Eq} - \mathcal{C}_{Ot}}{\mathcal{C}_{Ot}} \cdot 100\%$, where $\mathcal{C}_{Eq}, \mathcal{C}_{Ot}$ denote the cost for equitable and other algorithms respectively and η indicates the percent of cost savings over the other algorithms. In addition, we use Jain's fairness index [23] as equity metric for each algorithm that utilizes the coefficient of variance $\nu = \frac{\sigma}{\mu}$, where μ is the average regional environmental footprint and σ is the standard deviation. The following equation is used to compute the Jain's Index $\zeta = \frac{1}{1+\nu^2}$, where ν is the coefficient of variance.

5 Result

We compare the cost savings of the `EquiShift` approach against other GLBs, using $\kappa_w = 60\$/\mathrm{m}^3$ and $\kappa_c = 1500\$/\mathrm{Ton}$ assuming no prediction error. We also highlight the deviation of the `EquiShift` from the `EquiShift-Offline` approach. `EquiShift` algorithm's primary objective is to reduce regional disparities in water and carbon footprints. To demonstrate the comprehensive efficiency of our solution, we focus our comparisons on average energy cost, maximum regional water, and carbon footprints. While these metrics highlight our approach's effectiveness in addressing extreme cases, we also provide average case water and carbon footprint analysis for a more comprehensive understanding. *Further results, including the effects of varying κ_w and κ_c, prediction errors, and prediction windows, are provided in Appendix B due to space constraints.*

Table 1 illustrates the cost savings achieved by `EquiShift` compared to other GLBs. In most cases, `EquiShift` results in reduced costs and footprints, therefore we focus on the scenarios where cost increases are observed. Specifically, `EquiShift` shows a 6.01%, 6.20%, 21.82%, and 10.38% increase in energy cost compared to the EWShift, ECShift, EnShift, and ToShift methods, respectively. This increase is justified by substantial reductions in the worst-case water and carbon

Table 1: Cost/footprint savings ($\downarrow$ $\equiv$ decreased and $\uparrow$ $\equiv$ increased cost/footprint) of equitable MPC load balancing algorithm (`EquiShift`) over other algorithms and `EquiShift-Offline`

	EWShift	ECShift	EnShift	WaShift	CaShift	ToShift	`EquiShift Offline`
Avg Energy	$\uparrow$6.01%	$\uparrow$6.20%	$\uparrow$21.82%	$\downarrow$45.51%	$\downarrow$34.25%	$\uparrow$10.38%	$\uparrow$16.21%
Max Water	$\downarrow$20.58%	$\downarrow$24.21%	$\downarrow$25.83%	$\downarrow$8.57%	$\downarrow$30.56%	$\downarrow$18.45%	$\downarrow$6.60%
Avg Water	$\uparrow$5.22%	$\downarrow$3.06%	$\downarrow$4.51%	$\uparrow$18.75%	$\uparrow$1.52%	$\uparrow$1.67%	$\downarrow$1.51%
Max Carbon	$\downarrow$39.64%	$\downarrow$21.69%	$\downarrow$31.13%	$\downarrow$43.46%	$\downarrow$15.89%	$\downarrow$23.51%	$\downarrow$7.10%
Avg Carbon	$\downarrow$8.43%	$\uparrow$3.15%	$\downarrow$11.64%	$\uparrow$0.46%	$\uparrow$20.48%	$\downarrow$2.77%	$\downarrow$2.93%
Total	$\downarrow$17.38%	$\downarrow$9.93%	$\downarrow$9.40%	$\downarrow$40.35%	$\downarrow$29.03%	$\downarrow$7.67%	$\uparrow$3.85%

footprints across regions. Moreover, `EquiShift` exhibits a 16.21% and 3.85% rise in energy and total cost, respectively, compared to the `EquiShift-Offline` algorithms. However, `EquiShift` also achieves a 6% reduction in the worst-case water footprint and a 7% reduction in the worst-case carbon footprint compared to `EquiShift-Offline`. These deviations are primarily due to contextual differences, as the MPC approach does not have access to complete information, unlike the offline method. Additionally, while the average water and carbon footprints for `EquiShift` show slight increases in some cases compared to the state-of-the-art methods, these changes are justifiable given the significant improvements in addressing regional disparities in water and carbon footprints. In conclusion, the `EquiShift` algorithm offers significant reductions in water and carbon footprints, with a modest increase in energy cost, and provides the flexibility to adjust preferences based on specific goals.

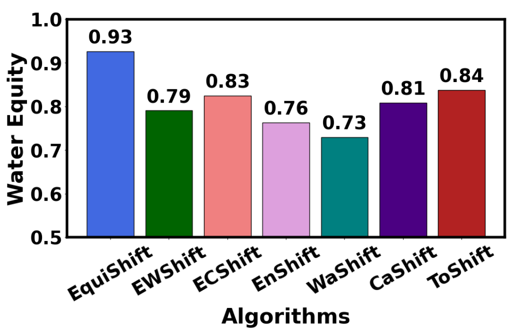

(a) Water equity index;

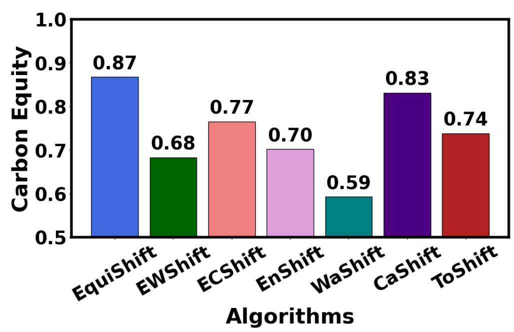

(b) Carbon equity index;

Figure 1: Equity index of water and carbon across algorithms.

Fig 1a and Fig 1b illustrate the fairness of the `EquiShift` compared to other load-balancing methods, using Jain's fairness index [23]. The results show that the `EquiShift` algorithm consistently achieves the highest fairness scores for both water and carbon footprints. This indicates that our method is more effective at reducing regional disparities across geographically distributed data centers, promoting more sustainable and equitable data center operations. However, the equity controlled by κ_w and κ_c impacts the energy consumption and thus requires careful selection to achieve the operational sustainability goal.

6 Conclusion

In this paper, we propose `EquiShift` to address a critical issue in existing sustainable load-shifting methods: the lack of awareness of regional disparities in environmental sustainability. Moreover, `EquiShift` offers greater adaptability by exploiting both spatial and temporal scheduling flexibilities. However, achieving true sustainability in data center operation may require considering factors beyond carbon emissions and water usage, such as local ecological impacts, energy source variability, socio-economic conditions, etc. Future work should aim to incorporate these factors, ensuring a holistic view that aligns operational efficiency with long-term environmental and societal goals in an equitable manner.

Acknowledgement

This work was supported in part by the NSF under grant CCF-2324941.

References

[1] Eren Çam, Zoe Hungerford, Niklas Schoch, Francys Pinto Miranda, Carlos David Yáñez de León, Carlos Fernández Álvarez, International Energy Agency, et al. *Electricity 2024*. 2024.

[2] Hessam Lavi. Measuring greenhouse gas emissions in data centres: the environmental impact of cloud computing. 2022.

[3] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. Making ai less "thirsty": Uncovering and addressing the secret water footprint of ai models, 2023.

[4] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu. Act: designing sustainable computer systems with an architectural carbon modeling tool. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, ISCA '22, page 784–799, New York, NY, USA, 2022. Association for Computing Machinery.

[5] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga Behram, James Huang, Charles Bai, Michael Gschwind, Anurag Gupta, Myle Ott, Anastasia Melnikov, Salvatore Candido, David Brooks, Geeta Chauhan, Benjamin Lee, Hsien-Hsin S. Lee, Bugra Akyildiz, Maximilian Balandat, Joe Spisak, Ravi Jain, Mike Rabbat, and Kim Hazelwood. Sustainable ai: Environmental implications, challenges and opportunities, 2022.

[6] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. Carbon explorer: A holistic framework for designing carbon aware datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, ASPLOS '23. ACM, January 2023.

[7] Minghong Lin, Zhenhua Liu, Adam Wierman, and Lachlan L. H. Andrew. Online algorithms for geographical load balancing. In *2012 International Green Computing Conference (IGCC)*, pages 1–10, 2012.

[8] Peter Xiang Gao, Andrew R Curtis, Bernard Wong, and Srinivasan Keshav. It's not easy being green. *ACM SIGCOMM Computer Communication Review*, 42(4):211–222, 2012.

[9] Zhenhua Liu, Minghong Lin, Adam Wierman, Steven H Low, and Lachlan LH Andrew. Greening geographical load balancing. *ACM SIGMETRICS Performance Evaluation Review*, 39(1):193–204, 2011.

[10] Lei Rao, Xue Liu, Le Xie, and Wenyu Liu. Minimizing electricity cost: Optimization of distributed internet data centers in a multi-electricity-market environment. In *2010 Proceedings IEEE INFOCOM*, pages 1–9, 2010.

[11] Asfandyar Qureshi, Rick Weber, Hari Balakrishnan, John Guttag, and Bruce Maggs. Cutting the electric bill for internet-scale systems. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication*, SIGCOMM '09, page 123–134, New York, NY, USA, 2009. Association for Computing Machinery.

[12] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. The carbon footprint of machine learning training will plateau, then shrink, 2022.

[13] Google 2023 environmental report. `https://www.gstatic.com/gumdrop/sustainability/google-2023-environmental-report.pdf`, 2023.

[14] Pengfei Li, Jianyi Yang, Adam Wierman, and Shaolei Ren. Towards environmentally equitable AI via geographical load balancing. In *e-Energy*, 2024.

[15] Model predictive control. `https://en.wikipedia.org/wiki/Model_predictive_control`, 2024.

[16] Ana Radovanović, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyue Xiao, Maya Haridasan, Patrick Hung, Nick Care, Saurav Talukdar, Eric Mullen, Kendal Smith, MariEllen Cottman, and Walfredo Cirne. Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems*, 38(2):1270–1280, 2023.

[17] International Energy Agency. Data and statistics. `https://www.iea.org/data-and-statistics`.

[18] U.S. Energy Information Administration. Open data. `https://www.eia.gov/opendata/`.

[19] Mohammad A. Islam, Kishwar Ahmed, Hong Xu, Nguyen H. Tran, Gang Quan, and Shaolei Ren. Exploiting spatio-temporal diversity for water saving in geo-distributed data centers. *IEEE Transactions on Cloud Computing*, 6(3):734–746, 2018.

[20] Electricity emission data. `https://www.electricitymaps.com/data-portal/`.

[21] Shaolei Ren. Optimizing water efficiency in distributed data centers. In *2013 International Conference on Cloud and Green Computing*, pages 68–75, 2013.

[22] Mohammad A. Islam, Kishwar Ahmed, Hong Xu, Nguyen H. Tran, Gang Quan, and Shaolei Ren. Exploiting spatio-temporal diversity for water saving in geo-distributed data centers. *IEEE Transactions on Cloud Computing*, 6(3):734–746, 2018.

[23] Rajendra K Jain, Dah-Ming W Chiu, William R Hawe, et al. A quantitative measure of fairness and discrimination. *Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA*, 21:1, 1984.

[24] Mohammad A. Islam, Shaolei Ren, Gang Quan, Muhammad Z. Shakir, and Athanasios V. Vasilakos. Water-constrained geographic load balancing in data centers. *IEEE Transactions on Cloud Computing*, 5(2):208–220, 2017.

[25] Miyuru Dayarathna, Yonggang Wen, and Rui Fan. Data center energy consumption modeling: A survey. *IEEE Communications Surveys & Tutorials*, 18(1):732–794, 2016.

[26] Yuanxiong Guo and Yuguang Fang. Electricity cost saving strategy in data centers by using energy storage. *IEEE Transactions on Parallel and Distributed Systems*, 24(6):1149–1160, 2013.

[27] Loïc Lannelongue, Jason Grealey, and Michael Inouye. Green algorithms: Quantifying the carbon footprint of computation, 2020.

[28] Qizhen Weng, Wencong Xiao, Yinghao Yu, Wei Wang, Cheng Wang, Jian He, Yong Li, Liping Zhang, Wei Lin, and Yu Ding. MLaaS in the wild: Workload analysis and scheduling in large-scale heterogeneous GPU clusters. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 945–960, 2022.

[29] Alibaba cloud products. `https://www.alibabacloud.com/product/`, 2024.

A Quantifying Energy Cost and Environmental Impact

In the main experiment, we used $f_i(y_i(t)$,$w_i(y_i(t)$ and $c_i(y_i(t)$ as the energy cost, water, and carbon footprint respectively for data center i at time $t \in [1, T]$. Here we present the details of computing each of the cost and footprints for the distributed workload $y_i(t)$.

Energy Cost: Energy cost is a fundamental metric for assessing the operational efficiency and sustainability of data centers, as it is directly tied to electricity consumption. Understanding and managing this cost is crucial for optimizing data center operations. Drawing from established research [24–26], we model the electricity consumption of a data center as a function of the workload it handles and its server utilization, which in turn depends on the number of active servers.

In a typical data center, not all servers are active simultaneously; however, each server still consumes a baseline level of energy, known as static power ($e_{i,s}$), even when idle. Beyond this, processing the incoming workload demands additional energy, referred to as dynamic power ($e_{i,d}$). The combined energy consumption of the data center is thus influenced by both these components and can be expressed by the following equation:

$$e_i(y_i(t)) = \gamma_i(t) \cdot M_i(t) \left[e_{i,s} + e_{i,d} \cdot \frac{y_i(t)}{M_i(t)} \right] \tag{3}$$

In Eqn. (3), $y_i(t)$ represents the workload dispatched to data center i, while $M_i(t)$ denotes the number of active servers. The static power ($e_{i,s}$) accounts for the baseline energy consumption, and dynamic power ($e_{i,d}$) scales with server utilization, represented by $\frac{y_i(t)}{M_i(t)}$. The factor $\gamma_i(t)$ captures the Power Usage Effectiveness (PUE), reflecting the efficiency of the data center in utilizing its electrical energy.

To translate energy consumption into a monetary cost, we define the power cost function ($f_i(y_i(t))$), which incorporates the price of electricity ($p_i(t)$) at data center i during time t:

$$f_i(y_i(t)) = p_i(t) \cdot e_i(y_i(t)) \tag{4}$$

In Eqn. (4), the electricity price $p_i(t)$ can vary significantly depending on the regional energy mix, including the availability of fossil fuels (such as coal and gas) and renewable energy sources. This variability highlights the importance of strategic workload distribution and timing, as these factors can greatly influence the overall energy cost. This energy cost model provides a comprehensive view covering both static and dynamic energy consumption of a data center.

Water Footprint: The operation of data centers involves substantial water consumption, primarily for cooling purposes, to maintain optimal operating conditions and mitigate the risks associated with overheating. Cooling towers, which are widely used in data centers, rely heavily on water-based cooling systems due to their efficiency compared to air cooling, especially in regions with high ambient temperatures where air cooling may be insufficient. However, the process of cooling leads to water evaporation, contributing to water loss and increasing the overall water footprint of the data center.

In quantifying the water footprint, it is crucial to consider both direct and indirect water usage. Direct water usage refers to the water consumed onsite for cooling and other operational processes within the data center. Indirect water usage, on the other hand, pertains to the water consumed offsite in the production of electricity that powers the data center. Previous research, such as [3] and [24], have explored these aspects extensively, offering methodologies for calculating the total water footprint by integrating both direct and indirect contributions. This comprehensive approach can be expressed by the following equation:

$$w_i(y_i(t)) = \left[\frac{\epsilon_{i,D}(t)}{\gamma_i(t)} + \epsilon_{i,I}(t) \right] \cdot e_i(y_i(t)) \tag{5}$$

In Eqn. (5), $\epsilon_{i,D}(t)$ represents the direct Water Usage Effectiveness (WUE), which is a measure of how efficiently water is used onsite for cooling purposes. This includes factors such as the design of the cooling system, the ambient temperature, and the efficiency of water recycling processes within the data center [3]. The term $\gamma_i(t)$ denotes the Power Usage Effectiveness (PUE), a metric that reflects the overall energy efficiency of the data center, impacting the direct water usage due to its influence on the amount of heat generated.

The second term, $\epsilon_{i,I}(t)$, is the indirect Water Usage Effectiveness, also known as the Electricity Water Intensity Factor (EWIF). This factor accounts for the water used in the generation of electricity consumed by the data center. Since the energy mix varies by region, with different proportions of fossil fuels, hydroelectric power, and other renewable sources, the indirect water usage can vary significantly. For example, regions reliant on thermal power plants, particularly those using water-intensive cooling methods, may have a higher EWIF, thereby increasing the overall water footprint.

By combining these factors, the equation $w_i(y_i(t))$ provides a comprehensive measure of the water footprint associated with the data center's operations, encompassing both the onsite and offsite water consumption.

Carbon Footprint: The carbon footprint [27] of a data center is a critical metric that quantifies the amount of carbon dioxide and other greenhouse gases (GHGs) emitted as a result of the energy consumed by the data center. This metric provides insights into how efficiently the data center utilizes energy and serves as a key indicator of its environmental impact. Understanding the carbon footprint is essential for informed decision-making aimed at reducing emissions and promoting sustainability within the data center industry. The carbon footprint can be computed using the following equation:

$$c_i(y_i(t)) = \Gamma_i(t) \cdot e_i(y_i(t)) \tag{6}$$

In Eqn. (6), $\Gamma_i(t)$ represents the Carbon Usage Effectiveness (CUE) for the i^{th} data center at time t. The CUE is a pivotal metric that defines the ratio of the greenhouse gas (GHG) emissions associated with the data center's operations to the total energy consumed by its IT equipment. Essentially, a lower CUE value indicates a more energy-efficient and environmentally friendly data center, as it implies that less carbon is emitted per unit of energy consumed.

The CUE is influenced by various factors, including the energy mix used by the data center. For example, a data center powered primarily by renewable energy sources, such as wind or solar, will have a lower CUE compared to one that relies heavily on fossil fuels like coal or natural gas. Additionally, advancements in energy efficiency technologies, such as more efficient cooling systems or optimized server utilization, can further reduce the CUE.

B Additional Experimental Details

B.1 Data Preperation

In this section, we describe the underlying pre-processing conducted for the experiment, as we do not have access to hourly energy fuel mix information for non-U.S. data center locations. Consequently, we utilize U.S.-based hourly energy fuel mix data to generate synthetic hourly fuel mixes for European and Asian locations. This synthesis is based on adhering to the average percentage of fuel mix information within the time range of September 23 and October 11, 2022, as reported by [17]. To match the average energy fuel mix of European and Asian countries, we scale the U.S. data by the factors provided in Table 2.

Different 18-day periods are intentionally selected to ensure that the data from U.S. locations does not correlate with the European and Asian energy fuel mix data.

B.2 Impact of Weight Parameters

The weight parameters κ_w and κ_c are crucial in shaping the regional balance of water and carbon footprints. Adjusting these parameters alters the focus on either water or carbon footprints, which in turn affects the prioritization of energy costs. Fig. 2 illustrates how varying κ_w and κ_c impacts energy costs, water footprints, and carbon footprints for ou `EquiShift`, as well as for EWShift, ECShift, and ToShift. Other GLBs are not included in this analysis since they remain unaffected by these weight parameters.

Table 2: Scaling factors applied to U.S. energy fuel mix data to align with the average energy fuel mix of European and Asian countries.

Country	Source (U.S. State)	Time Period	Scaling Factor
Germany	Texas	June 1 – June 19, 2022	0.8503
Belgium	Georgia	June 1 – June 19, 2022	1.5319
Netherlands	Georgia	March 1 – March 19, 2022	1.2759
Denmark	Oregon	July 1 – July 19, 2022	0.2657
Japan	Nevada	March 1 – March 19, 2022	3.2374
Singapore	Georgia	May 1 – May 19, 2022	4.4875

In Fig. 2(a), we observe that increasing κ_w and κ_c shifts the focus from minimizing energy costs to prioritizing water and carbon footprints, resulting in higher energy costs. Conversely, Figs. 2(b) and 2(d) show that with our `EquiShift` approach, raising these parameters effectively reduces water and carbon footprints, as intended. This trend may not be apparent in other algorithms, which primarily focus on minimizing overall costs or footprints.

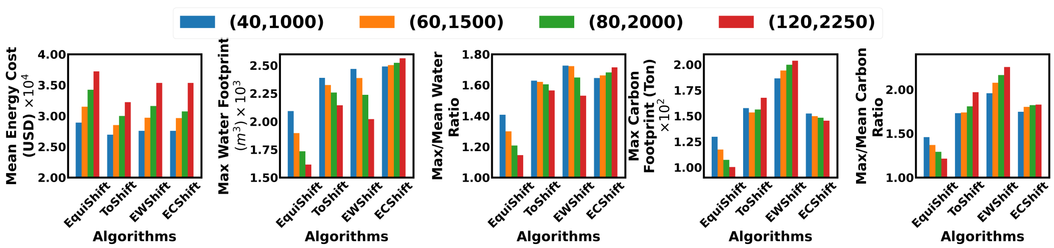

(a) Average Energy Cost (b) Maximum Water Footprint (c) Ratio of Maximum and Mean Water Footprint (d) Maximum Carbon Footprint (e) Ratio of Maximum and Mean Carbon Footprint

Figure 2: Impact of the weighting parameters κ_w and κ_c on the environmental footprint and energy cost in `EquiShift`.

To further analyze the impact, Figs. 2(c) and 2(e) present the ratio of maximum to mean footprint for water and carbon, respectively. These figures demonstrate that as the weight parameters increase, `EquiShift` achieves a more equitable distribution of footprints, with the max/mean ratio approaching 1, indicating fairness across regions. However, this increased equity in environmental footprints comes at the cost of significantly higher energy expenditure, highlighting the need to balance fairness with energy efficiency. In contrast, other GLBs may also experience higher energy costs, but they do not achieve equitable distribution of footprints, leading to regional disparities.

B.3 Impact of Prediction Error and Window Length

Fig 3 illustrates how prediction error and prediction window length influence energy costs and environmental footprints in data centers using our proposed MPC-based algorithm, `EquiShift`. The heat maps display the relationship across three metrics: (a) energy cost, (b) water footprint, and (c) carbon footprint.

In Fig 3(a), we see that higher prediction errors (0.8) combined with shorter prediction windows (2) lead to increased energy costs. However, under the same conditions, Figs. 3(b) and 3(c) show that both the water and carbon footprints are reduced. This reciprocal relationship indicates that while energy costs decrease with improved prediction accuracy and longer prediction windows, the environmental footprints tend to increase.

Importantly, these changes in energy costs and environmental footprints are not drastic. The strength of the MPC-based algorithm lies in its ability to adapt smoothly to varying conditions. Even when dealing with high prediction errors and short prediction windows, the algorithm can adjust load distribution to stay near an optimal trajectory. This adaptability stems from its use of historical data in decision-making. For instance, if a past decision results in a suboptimal load distribution, the algorithm compensates in subsequent steps.

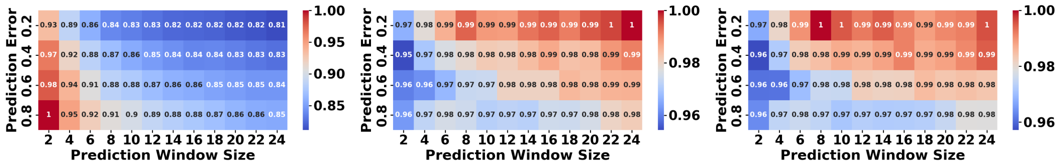

(a) Normalized Average Energy Cost **(b) Normalized Maximum Water Footprint** **(c)Normalized Maximum Carbon Footprint**

Figure 3: Heatmaps of normalized energy cost and environmental footprint for `EquiShift` under different prediction window lengths and prediction error magnitudes.

As a result, Fig. 3 demonstrates that the MPC-based approach maintains stability and remains close to the optimal solution, even under challenging conditions of high prediction errors and short prediction windows.

C Extended Experiments

C.1 Workload & Data Preparation

For the extended experiment, we use a machine learning workload trace from July and August 2020, as detailed in [28]. This trace, collected from a production cluster with 6,742 GPUs in Alibaba's Platform for Artificial Intelligence (PAI) [29], includes a mixture of training and inference jobs submitted by over 1,300 users. The workloads span a wide range of ML algorithms executed on heterogeneous GPUs. Additional details are available in [28]. For our study, we process and scale the original trace to fit the 10 locations described in Section 4.

As in the main experiment, we use U.S. fuel mix data for various locations and apply scaling factors to align the fuel mix with the average compositions reported for Europe and Asia during the same period by [17]. The scaling factors are summarized in Table 3.

Table 3: Scaling factors applied to U.S. energy fuel mix data to align with the average energy fuel mix of European and Asian countries.

Country	Source (U.S. State)	Time Period	Scaling Factor
Belgium	Texas	March 1 - April 30, 2020	1.3331
Denmark	Georgia	March 1 - April 30, 2020	0.1594
Germany	Virginia	March 1 - April 30, 2020	2.4486
Singapore	Nevada	March 1 - April 30, 2020	1.1982
Japan	Virginia	June 1 - July 31, 2020	6.7678
Netherlands	Virginia	September 1 - August 31, 2020	0.9596

C.2 Extended Results

Similar to the main experiment here we present the cost-saving and equity-based comparative analysis of `EquiShift`, with $\kappa_w = 15\,\$/\text{m}^3$ and $\kappa_c = 100\,\$/\text{Ton}$.

Table 4 shows that `EquiShift` generally achieves a reduction in cost/footprints over other GLBs. However, `EquiShift` shows 16.91%,16.91%,18.40%,17.38% increase in energy cost compared to EWShift, ECShift, EnShift, and ToShift. This rise is compensated by significant reductions in water usage and carbon emissions, consistent with the findings from the main experiment. Additionally, as in the main experiment, the average-case analysis reveals a slight increase in water and carbon footprints for `EquiShift` compared to state-of-the-art algorithms, though the trade-offs are justifiable in light of the improvements in regional equity.

An important observation from the extended experiment with `EquiShift` is the increase in energy cost, water usage, and carbon footprint compared to `EquiShift-Offline`, whereas in the main experiment, `EquiShift` achieved reductions in water and carbon footprints. This highlights the

Table 4: Cost/footprint savings (↓ ≡ decreased and ↑ ≡ increased cost/footprint) of equitable MPC load balancing (`EquiShift`) algorithm over other algorithms and `EquiShift-Offline` for extended experiment.

	EWShift	ECShift	EnShift	WaShift	CaShift	ToShift	`EquiShift Offline`
Avg Energy	↑16.91%	↑16.91%	↑18.40%	↓23.79%	↓38.71%	↑17.38%	↑10.57%
Max Water	↓39.82%	↓35.99%	↓42.78%	↓27.74%	↓32.34%	↓38.24%	↓0.65%
Avg Water	↓0.50%	↓1.91%	↓2.70%	↑10.46%	↓0.46%	↓1.17%	↑0.80%
Max Carbon	↓42.75%	↓38.93%	↓45.48%	↓32.46%	↓11.69%	↓41.19%	↓0.23%
Avg Carbon	↓3.05%	↓1.81%	↓3.67%	↑3.75%	↑11.12%	↓2.47%	↑0.30%
Total	↓9.80%	↓7.19%	↓11.29%	↓25.64%	↓35.32%	↓8.50%	↑6.98%

critical trade-off between energy costs and environmental impacts. The main experiment shows a sharper increase in energy cost compared to the extended experiment, which limits its ability to reduce water and carbon footprints. This underscores the importance of tuning κ_w and κ_c in the optimization process to balance these objectives effectively.

As demonstrated earlier, the weighting parameters κ_w and κ_c play a crucial role in managing disparities in water and carbon footprints. Prioritizing greater equity can lead to higher energy costs. However, seeking greater equity may not always be the best approach for operational efficiency for data centers, so a careful balance is required to manage these trade-offs.

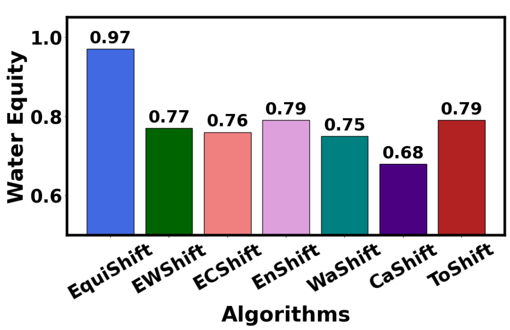

(a) Water equity index;

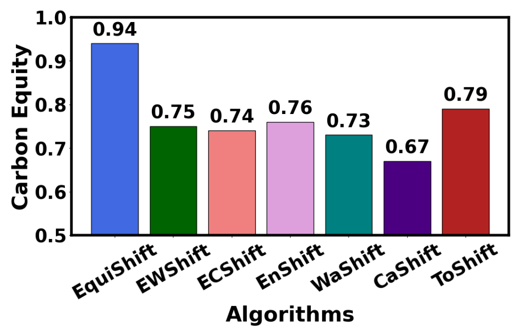

(b) Carbon equity index;

Figure 4: Equity of water and carbon across algorithms for the extended experiment.

We also compare the equity of `EquiShift` against other GLBs in Figs. 4a (for water) and 4b (for carbon) for this extended experiment. According to Jain's fairness index, [23], our proposed `EquiShift` achieves a higher score than any other GLBs. This consistency with the main experiment further underscores the advantages of `EquiShift` in promoting sustainable data center operations.