How the Choice of Carbon Signal Impacts Carbon-Aware Scheduling Decisions

Thanathorn Sukprasert
University of Massachusetts Amherst

Abstract

The carbon intensity of grid-supplied electricity depends on the mix of generation sources used to satisfy its demand and varies widely over time and across locations. There are two widely used carbon intensity signals: *average* and *marginal*. Both signals provide distinct information about grid operations and affect the electric grid's short- and long-term operation differently. Unfortunately, there is a lack of consensus on the "right" signal for carbon-aware optimizations, and decarbonization efforts across domains are using both signals to guide carbon-aware workload scheduling. To understand the implications of signal selection on carbon-aware optimizations, we perform a data-driven analysis using the average and marginal carbon intensity from 65 regions to understand how the choice of carbon intensity signal impacts the carbon-aware scheduling decisions.

1 Introduction

Growing concern about climate change has elevated the importance of assessing and reducing the carbon footprint of energy consumption across societal sectors, including datacenters [1, 2, 12], buildings [8, 18], and transportation [14, 17]. Many of these decarbonization initiatives aim to shift energy demand to when and where low-carbon electricity is available. Such carbon-aware optimizations are enabled by the recent emergence of third-party carbon information services, such as Electricity Maps [11], and WattTime [19], that provide carbon intensity of electricity across regions worldwide. Carbon information services provide carbon intensity information using two metrics: the *average* and *marginal* carbon intensity, both expressed in the grams of carbon dioxide emitted per kilowatt-hour of electricity ($g{\cdot}CO_2eq/kWh$). The average (or attributional) carbon intensity is the weighted average of the carbon intensity of all the generators used to satisfy the current grid demand. The marginal (or consequential) carbon intensity is the carbon emissions rate of the generator that responds to incremental changes in energy usage.

The two carbon intensity signals express different aspects of electric grid operations to satisfy the electricity demand. The average signal provides information on the grid's overall portfolio of energy generation resources. The marginal signal derives from a smaller set of fast-responding generators that fulfill the marginal segment of electricity demand. Interestingly, the signals do not always align, which is critical for carbon-aware optimizations, for the weak correlations in the signals imply that the choice of the signal will lead to different scheduling outcomes for the same workload. Due to the vast and critical implications of carbon signal choice, there is an ongoing debate as to which signal should be used for decarbonization [3, 5, 13]. Since there is not yet a consensus, this paper aims to facilitate this discourse using a data-driven approach.

2 Overview Of Proposed Work

We plan to use several state-of-the-art carbon optimization techniques to show how the choice of carbon intensity signal impacts the workload scheduling decisions and resources.

2.1 Carbon Optimization Techniques

Temporal Workload Scheduling. The carbon-aware temporal scheduling techniques that we plan to use for the evaluation are *WaitAwhile* [20] and *CarbonScaler* [6]. WaitAwhile is a carbon-aware suspend-resume policy that splits and schedules the workloads towards low-carbon periods. Besides suspending and resuming the workload that can incur long delays in job completion times, CarbonScaler exploits the elasticity of batch workloads, where this scheduler dynamically allocates resources for a job based on variations in carbon intensity.

Spatial Workload Shifting For the spatial workload shifting, we plan to use *one-migration*, *infinite-migrations* and a few variations of load-balancing to evaluate the effects of the scheduling decisions. One-migration policy migrates a job to a region with the lowest mean overall, and the job is executed in the lowest mean region until completion. On the contrary, the infinite-migrations policy migrates the job to a region with the lowest carbon intensity of that hour, and the migration continues every hour until the job completes its execution.

2.2 Impacts on Resources Usage

Some of the resources we want to examine how they are affected by the carbon intensity signal chosen by the carbon-aware scheduler include:

Carbon Savings. We want to examine how the choice of signal affects the carbon savings based on the optimized signal and the other signal's standpoints.

Completion Time. Different signal choices could lead to different overall completion times for a job.

Locational Marginal Pricing (LMP). The marginal energy price is based on the demand and congestion of an electric grid, expressed in dollars per megawatt-hour ($/MWh). With this metric, we want to analyze how the choice of signal impacts the total cost to execute the workload.

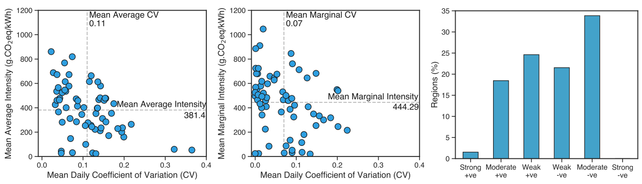

(a) Average Signal (b) Marginal Signal (c) Correlation

Figure 1. *The mean ts mean daily coefficient of variation for average and marginal signals (a)-(b), and the mean daily correlation, categorized as Strong, Moderate, and Weak, between the average and marginal signal (c).*

Energy Demand. Since the energy demand can influence capacity planning and energy use incentives, we want to analyze how different signal choices impact the distribution of energy demand across time.

Migration Overhead. Since migrating workloads to other regions incur state-transfer overheads, we want to examine the net savings from spatial migration from the average and marginal carbon intensity.

Latency Overhead We want to examine the latency incurred during the migration with respect to the different signal choices.

3 Preliminary Results

3.1 Carbon Intensity Signal Characteristics

The extent of savings from carbon-aware optimizations depends on the spatiotemporal variability in the carbon intensity signals. The larger the difference between the magnitudes of the carbon signals across regions, the higher the spatial savings. The larger the variations within a region's carbon intensity, the higher the temporal savings. In Figures 1(a)-(b), we quantify the variability of a carbon signal as the daily coefficient of variation (CV), computed as the daily standard deviation over the daily mean. The average carbon intensity signal has a lower mean value of 381.4 $g \cdot CO_2eq/kWh$as compared to 444.29 $g \cdot CO_2eq/kWh$for the marginal carbon intensity signal. However, the average signal exhibits a higher variability (0.11 CV) than the marginal signal (0.07 CV).

While these statistics provide information about overall carbon emissions and the potential for carbon savings, they do not necessarily indicate that the signals differ, as signals with different magnitudes can still be correlated. Figure 1(c) shows the distribution of the mean daily correlation between the average and marginal carbon intensity signals. We categorize the values as positively or negatively correlated with *strong*, *moderate*, and *weak* correlations, specified by the ranges of (0.7, 1], (0.2,0.7], and [0,0.2], respectively. Among 65 regions, 36 regions (55.4%) exhibit a negative correlation between their average and marginal carbon intensity signal, and only 1.5% have a strong positive correlation.

3.2 Impact on Carbon Savings Calculations

Figure 2(a) shows carbon savings across all the regions when the average signal guides the workload scheduling. The two

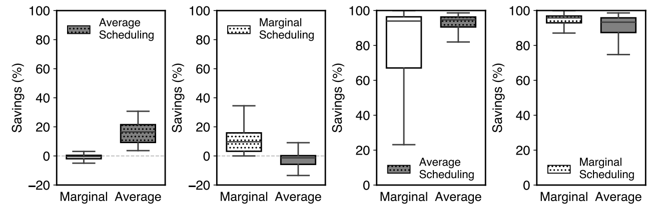

(a) Temp:Avg (b) Temp:Mar (c) Spatial:Avg (d) Spatial:Mar

Figure 2. *Scheduling and accounting implications: Carbon savings using the average signal (a) and the marginal signal (b) for temporal scheduling. Carbon savings using the average signal (a) and the marginal signal (b) for spatial scheduling.*

boxplots correspond to the signal used for calculating carbon savings compared to the counterfactual of no workload shifting (carbon-agnostic execution). Figure 2(b) shows the same for a scenario when the marginal signal is used as the guide signal. From Figure 2(a)-(b) it can be seen that based on the other signal, the carbon savings are negative, i.e., carbon emissions actually increased. Also, the estimated carbon savings based on the scheduling signal differ for both signals; scheduling and accounting based on average signal yields 18% savings while based on marginal signal yields 11% savings. Moreover, Figures 2(c)-(d) show mean carbon savings of ~87% when the average and marginal signal are used for spatial workload shifting. Like the temporal scheduling, the other signal yields less carbon savings than the scheduling signal. Generally, choosing one signal for carbon-aware optimizations for temporal workload scheduling leads to more carbon emissions from the other signal standpoint. While the opposite signal gains some savings from the decisions of the scheduling signal in spatial shifting, the savings of the opposite signal are always less than the scheduling signal.

4 Future Work

For the work to be done, we plan to analyze how the spatiotemporal scheduling policies (§2.1) based on the average and marginal carbon intensity signal impact the resources described in §2.2. We will navigate the multi-dimensional optimization configurations, which are the resource metrics to learn about their relationship with carbon savings. We will also evaluate the types of workloads that stand to benefit from different signal choices.

5 Related Work

There is no consensus on the choice of carbon signal, for some works use the average signal [2, 9, 10, 15, 16, 20], some studies use the marginal carbon intensity signal for carbon-aware optimizations [4, 7]. Moreover, there is very limited prior work on understanding the difference between the average and the marginal carbon intensity signals and their implications on grid operations. The most relevant work on this topic is done by Gagnon and Cole [5], who look at the impact of traditional marginal signal and how it can be extended to incorporate future capacity planning implications.

References

[1] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. 2023. Carbon Explorer: A Holistic Framework for Designing Carbon Aware Datacenters. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, New York, NY, USA, 118–132.

[2] Noman Bashir, Tian Guo, Mohammad Hajiesmaili, David Irwin, Prashant Shenoy, Ramesh Sitaraman, Abel Souza, and Adam Wierman. 2021. Enabling Sustainable Clouds: The Case for Virtualizing the Energy System. In *Proceedings of the ACM Symposium on Cloud Computing (SoCC '21)*. Association for Computing Machinery, New York, NY, USA, 350–358. https://doi.org/10.1145/3472883.3487009

[3] Olivier Corradi. 2023. Marginal vs Average: Which One to Use for Real-time Decisions? https://www.electricitymaps.com/blog/marginal-vs-average-real-time-decision-making.

[4] Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A Smith, Nicole DeCario, and Will Buchanan. 2022. Measuring the carbon intensity of AI in cloud instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1877–1894.

[5] Pieter Gagnon and Wesley Cole. 2022. Planning for the Evolution of the Electric Grid with a Long-Run Marginal Emission Rate. *iScience* 25, 3 (2022), 103915.

[6] Walid A. Hanafy, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. CarbonScaler: Leveraging Cloud Workload Elasticity for Optimizing Carbon-Efficiency. arXiv:2302.08681 [cs.DC]

[7] Julian Huber, Kai Lohmann, Marc Schmidt, and Christof Weinhardt. 2021. Carbon Efficient Smart Charging using Forecasts of Marginal Emission Factors. *Journal of Cleaner Production* 284 (2021), 124766.

[8] Adam Lechowicz, Noman Bashir, John Wamburu, Mohammad Hajiesmaili, and Prashant J. Shenoy. 2023. Equitable Network-Aware Decarbonization of Residential Heating at City Scale. In *ACM International Conference on Future Energy Systems (e-Energy),*. ACM, New York, NY, USA, 268–272. https://doi.org/10.1145/3575813.3576870

[9] Adam Lechowicz, Nicolas Christianson, Jinhang Zuo, Noman Bashir, Mohammad Hajiesmaili, Adam Wierman, and Prashant Shenoy. 2023. The Online Pause and Resume Problem: Optimal Algorithms and An Application to Carbon-Aware Load Shifting. *arXiv preprint arXiv:2303.17551* (2023).

[10] Liuzixuan Lin and Andrew A Chien. 2023. Adapting Datacenter Capacity for Greener Datacenters and Grid. In *Proceedings of the ACM International Conference on Future Energy Systems (e-Energy)*. ACM, New York, NY, USA, 200–213.

[11] Electricity Maps. 2022. Electricity Map. https://www.electricitymap.org/map.

[12] 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. 2023. Carbon-Aware Computing for Datacenters. *IEEE Transactions on Power Systems* 38, 2 (2023), 1270–1280.

[13] Henry Richardson. 2023. Is Your Goal Real-world Impact? Then Use Marginal Emissions. https://www.watttime.org/news/is-your-goal-real-world-impact-then-use-marginal-emissions/.

[14] Mahsa Sahebdel, Ali Zeynali, Noman Bashir, Mohammad H. Hajiesmaili, and Jimi Oke. 2023. Data-Driven Algorithms for Reducing the Carbon Footprint of Ride-Sharing Ecosystems. In *Companion Proceedings of the ACM International Conference on Future Energy Systems* (Orlando, FL, USA) *(e-Energy '23 Companion)*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3599733.3606300

[15] Abel Souza, Noman Bashir, Jorge Murillo, Walid Hanafy, Qianlin Liang, David Irwin, and Prashant Shenoy. 2023. Ecovisor: A Virtual Energy System for Carbon-Efficient Applications. In *ASPLOS*. ACM, New York, NY, USA, 252–265.

[16] Abel Souza, Shruti Jasoria, Basundhara Chakrabarty, Alexander Bridgwater, Axel Lundberg, Filip Skogh, Ahmed Ali-Eldin, David Irwin, and Prashant Shenoy. 2023. CASPER: Carbon-Aware Scheduling and Provisioning for Distributed Web Services. In *The International Green and Sustainable Computing (IGSC)*. IEEE, Piscataway, NJ, USA, 7 pages.

[17] Junyan Su, Qiulin Lin, and Minghua Chen. 2023. Follow the Sun and Go with the Wind: Carbon Footprint Optimized Timely E-Truck Transportation. In *Proceedings of the ACM International Conference on Future Energy Systems* (Orlando, FL, USA) *(e-Energy '23)*. ACM, New York, NY, USA, 159–171. https://doi.org/10.1145/3575813.3595193

[18] José R. Vázquez-Canteli, Jérôme Kämpf, Gregor Henze, and Zoltan Nagy. 2019. CityLearn v1.0: An OpenAI Gym Environment for Demand Response with Deep Reinforcement Learning. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation* (New York, NY, USA) *(BuildSys '19)*. ACM, New York, NY, USA, 356–357. https://doi.org/10.1145/3360322.3360998

[19] WattTime. 2022. WattTime. https://www.watttime.org/.

[20] Philipp Wiesner, Ilja Behnke, Dominik Scheinert, Kordian Gontarska, and Lauritz Thamsen. 2021. Let's Wait Awhile: How Temporal Workload Shifting Can Reduce Carbon Emissions in the Cloud. In *Proceedings of the 22nd International Middleware Conference (Middleware)*. ACM, New York, NY, USA, 260–272.