



# On the Implications of Choosing Average versus Marginal Carbon Intensity Signals on Carbon-aware Optimizations

Thanathorn Sukprasert      Noman Bashir      Abel Souza  
University of Massachusetts Amherst    Massachusetts Institute of Technology    University of Massachusetts Amherst

David Irwin      Prashant Shenoy  
University of Massachusetts Amherst    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 types of carbon intensity signals: *average* and *marginal*. Both signals provide distinct information about grid operations and affect the electric grid’s short- and long-term functioning in different ways. Unfortunately, there is a lack of consensus on the “right” signal for carbon-aware optimizations, and decarbonization efforts across domains have used both signals to decide when and where to shift demand. To understand the implications of signal selection on carbon-aware optimizations, this paper performs a data-driven analysis using both the average and marginal carbon intensity. Our analysis for 65 regions reveals multiple insights, including i) both signals are statistically different with very low correlation between them, ii) optimizing for one signal could lead to more carbon emissions from the other signal’s standpoint, and iii) differences in signal characteristics in each region lead to different electricity use incentives.

## CCS CONCEPTS

• **Hardware** → **Impact on the environment**; • **General and reference** → **Performance**; **Metrics**.

## KEYWORDS

Average and marginal carbon emissions, carbon intensity, societal decarbonization, sustainability

### ACM Reference Format:

Thanathorn Sukprasert, Noman Bashir, Abel Souza, David Irwin, and Prashant Shenoy. 2024. On the Implications of Choosing Average versus Marginal Carbon Intensity Signals on Carbon-aware Optimizations. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24)*, June 04–07, 2024, Singapore, Singapore. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3632775.3661953>



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

*E-Energy '24*, June 04–07, 2024, Singapore, Singapore  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0480-2/24/06  
<https://doi.org/10.1145/3632775.3661953>


## 1 INTRODUCTION

The growing concern about the climate impact of human activities has amplified the importance of assessing and reducing the carbon footprint of energy consumption across societal sectors, including datacenters [1, 2, 18, 21], buildings [13, 28], and transportation [20, 24]. 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 [16], and WattTime [29], that provide carbon intensity of electricity across regions worldwide. The carbon intensity of electricity is the grams of carbon dioxide emitted per kilowatt-hour of electricity ( $g \cdot CO_2eq/kWh$ ) at the point of consumption [17], which varies over time and across locations based on the energy sources used to satisfy the electricity demand.

Carbon information services provide carbon intensity information using two metrics: the *average* and *marginal* carbon intensity. 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 weights are the ratio of the production of each generator to the total production that serves the current demand. The marginal (or consequential) carbon intensity is the carbon emissions rate of the generator that responds to incremental changes in energy usage. Although other, more nascent, signals such as long-term marginal emission rates (LMER) [9] incorporate the effect of carbon-aware optimizations on electric grid’s capacity planning, they are unavailable for most regions globally. Thus, our work focuses on the average and marginal carbon intensity signals.

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 (Section 3.2), which holds important implications for carbon-aware optimizations. The weak correlations in the signals imply that the choice of the signal determines how much carbon savings are perceived due to carbon-aware scheduling (Section 3.3). As a result, both signals often have conflicting impacts on grid operations and long-term capacity planning as *when and where* carbon-aware execution shifts the workloads differ (Section 3.4).

Due to the vast and critical implications of choosing an electricity carbon emissions signal for carbon-aware optimizations, there



**Figure 1: The energy generation mix of a hypothetical grid ordered in terms of how various generators are engaged to serve demand. The carbon intensity, in  $g \cdot CO_2eq/kWh$ , for each generator is labeled at the top of each bar.**

is an ongoing debate as to which signal should be used for decarbonization [5, 9, 19]. Since there is no clear consensus yet, this paper aims to facilitate this discourse using a data-driven approach. In doing so, we make the following contributions.

- (1) **Large-scale data analysis.** We analyze year-long average and marginal carbon intensity traces from 65 regions worldwide to assess the statistical properties of the individual signals and quantify the correlation between them.
- (2) **Analyzing implications of carbon-aware optimizations.** We leverage state-of-the-art carbon-aware temporal workload shifting approaches, e.g., WaitAWhile [30], and spatial workload migration strategies [22], to quantify carbon savings using both signals. Our analysis reveals that mixed use of signals, one for scheduling and the other for reporting, detracts from carbon-aware optimizations. Our findings further demonstrate that the choice of signal can impact the grid in a way that discourages the adoption of renewable energy sources, such as solar.
- (3) **Future research directions.** We discuss how the data-driven insights from this work can help shape the discourse on the choice of carbon intensity signal. We enumerate numerous research questions for future research that will enable a better understanding of carbon intensity signals and their implications for holistic and societal-scale decarbonization.

## 2 BACKGROUND

This section provides an overview of the electric grid's operation and discusses how carbon intensity signals are computed.

**Electric grid operations.** The electric grid's operators must balance energy generation and energy demand in real time. The energy demand varies primarily based on weather, which dictates the energy needed for indoor heating and cooling, and human behavioral patterns, e.g., time of the day, day of the week, holidays, etc. The mix of generators used to satisfy the energy demand changes over time and is determined by the electricity markets. The characteristics of energy resources, such as fuel types, capacities, and carbon emissions, used to satisfy the energy demand vary across time and location. Figure 1 shows a hypothetical electric grid, with a mix of renewable energy, nuclear, coal, natural gas, and oil as the energy sources for a given time and location. Further details on the electricity markets can be found in a primer on energy markets [23].

**Carbon intensity of grid's electricity.** The carbon intensity ( $CI$ ) of electricity is measured in grams of carbon dioxide equivalent per kilowatt-hours ( $g \cdot CO_2eq/kWh$ ). Figure 1 depicts the two methods of estimating carbon intensity: the **average** paradigm and the **marginal** paradigm. The average carbon intensity  $CI_{avg}$  of electricity is estimated as the weighted average of carbon emissions factors  $c_i$  for all the generators satisfying the current demand, i.e.,

$$CI_{avg} = \frac{\sum_i (c_i \cdot w_i)}{\sum_i w_i}. \quad (1)$$

Here, the weight  $w_i$  for generator  $i$  is proportional to the portion of demand it satisfies. It is worth noting that when an increase in demand engages the next generator with higher emissions, increases in emissions are uniformly distributed across all units of existing demand and the new demand that triggered the next generator.

In the marginal paradigm, the carbon emission rate for the marginal generator is used as the carbon intensity signal,  $CI_{mar}$ , and can be represented mathematically as

$$CI_{mar} = \frac{\Delta CE}{\Delta D}, \quad (2)$$

where  $\Delta CE$  and  $\Delta D$  represent the change in carbon emissions (CE) and the change in electricity demand  $D$ , respectively. Equation 2 calculates the rate of change of carbon emissions if demand changes. As not all the generators serve the added demand, only the marginal generator does (see Figure 1), and the higher or lower marginal emissions are assigned to consumers whose demand triggered it.


Since both the energy generation mix and the marginal generators change over time and across locations, the average and marginal carbon intensities of the grid-supplied electricity also change over time and locations. As shown in Figure 1, the generators that cannot be turned off or ramped up/down are *always-on* and serve as the grid's base load [8], e.g., nuclear. The generators with large to medium response times, such as coal plants, serve the intermediate demand, while fast-responding generators – e.g., natural gas – serve the variable and unexpected demand. However, depending on where demand occurs, generators with low emitting rates, such as hydro, solar, and wind power plants, may also serve as marginal generators [8]. Until recently, the carbon intensity of energy was opaque to consumers since generation data was inaccessible. However, balancing authorities have begun publicly releasing information about the active generator set and their real-time energy output via web APIs. Carbon information services, such as Electricity Maps [16] and WattTime [29], combine this information with data-driven or physical models to estimate the grid energy's carbon intensity in each region, making it available via web APIs.

## 3 ANALYSIS OF CARBON INTENSITY SIGNALS

The main objective of our analysis is to quantify the statistical differences between the average and marginal carbon intensity signals and assess how the magnitude of carbon savings differs depending on the carbon intensity signal used for workload shifting.

### 3.1 Evaluation setup and methodology

In this section, we provide details on our data sources, analysis methodology, and evaluation metrics.



**Figure 2:** The mean average carbon intensity and its mean daily coefficient of variation (a), the mean marginal carbon intensity and its mean daily coefficient of variation (b), and the mean daily correlation, categorized as Strong, Moderate, and Weak, between the average and marginal signal (c).

**Carbon traces.** We collect hourly average and marginal carbon intensity data from 65 regions for the year 2022 using Electricity Maps and WattTime web API [16, 29]. The 65 regions cover most of the world’s population and also include data center locations for major cloud providers, such as Amazon Web Service (AWS), Google Cloud Platform (GCP), and Microsoft Azure.

**Workload characteristics.** While our work applies to any flexible workload, we consider a computing job of a given length  $L$  that needs to be completed within a time horizon of  $L + T$ , where  $T$  is the slack for the completion. We set the value of  $T$  as 24hrs for all the experiments. We specify the value of  $L$  with experiments.

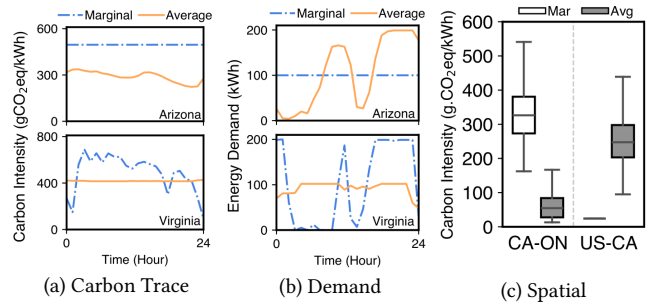
**Spatial and temporal scheduling.** As a representative temporal scheduling approach, we use a carbon-aware suspend-resume policy that assumes perfect knowledge of carbon intensity where the algorithm picks the  $L$  lowest carbon slots within the  $L + T$  horizon. This policy is based on prior work by Wiesner et al. [30] and yields the lowest carbon emissions for the job. For the spatial workload shifting, our policy migrates the job to the region with the lowest carbon intensity within all the possible geographical regions under analysis, inspired by approaches in prior work [7, 22].

**Carbon savings calculations.** To compute savings for either signal, we first compute the total carbon emissions in running the job in a carbon-agnostic manner, i.e., run the job as soon as it arrives (temporal) and at the location it arrives (spatial). We next compute the carbon emissions after the scheduling policy has determined when and where the job runs. We compute carbon savings as the reduction in emissions with respect to the carbon-agnostic baseline.

### 3.2 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. Figure 2 shows the magnitude and variability of the two carbon signals for all the regions in the trace. We quantify the variability of a carbon signal as the daily coefficient of variation (CV), computed as the standard deviation over the mean.

Figure 2a and Figure 2b show that the average carbon intensity signal has a lower mean value of  $381.4 \text{ g} \cdot \text{CO}_2\text{eq}/\text{kWh}$  as compared to  $444.29 \text{ g} \cdot \text{CO}_2\text{eq}/\text{kWh}$  for the marginal carbon intensity signal. However, the average signal exhibits a higher variability (0.11 CV)



**Figure 3:** Sample carbon intensity traces from Arizona and Virginia, with one varying and one relatively constant (a) and final demand (b) following carbon-aware temporal scheduling based on the two signals. For spatial scheduling, we show two sample regions (Ontario (CA-ON) and California (US-CA)) where average and marginal signals are opposite.

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 trace statistics are similarly spread. Figure 2c 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.


**Key takeaway.** Regions do not significantly differ in their average and marginal carbon intensity signals based on their mean and coefficient of variation values. However, in most regions, signals exhibit weak-to-moderate negative correlation with potential conflicting implications for carbon-aware optimizations.

### 3.3 Scheduling and accounting implications

The differences in the carbon intensity signal profiles have profound implications on the scheduling decisions and estimated carbon savings. In both temporal and spatial scheduling, we consider 1-hr long jobs and omit other job lengths and slacks as we focus on the relative implications rather than the actual carbon savings.

**Impact on workload scheduling.** We next illustrate how the differences in carbon intensity signals impact the scheduling decisions made by carbon-aware optimizations. We take two sample regions for each of the temporal scheduling and spatial scheduling. Figure 3a shows the average and marginal carbon intensities for Arizona and Virginia for a sample day in 2022. Both the two regions exhibit one dynamic and one comparatively stable carbon signal. For instance, Arizona has an almost constant marginal intensity signal, while it is highly variable in Virginia.

Figure 3b shows how a carbon-aware scheduler will schedule a hypothetical flexible workload of 2400 kWh in both regions. In Arizona, shifting the workload based on the marginal signal does not provide any incentive (constant signal) and the scheduler distributes the workload evenly throughout the day. We observe the



**Figure 4: 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.**

same trend for the average signal in Virginia. However, since Virginia has slight variations in its intensity, the scheduled load mirrors the variations. On the other hand, the average signal in Arizona and the marginal signal in Virginia show significant variations, and the scheduled workload follows the low-carbon slots. In both cases, most of the load is shifted to the nighttime periods when the relevant carbon intensity signal is lower.

Figure 3c shows the average and marginal carbon intensity signals for two sample locations where the spatial workload shifting between the two locations occurs. The choice of signal will determine where the load is executed. It will be executed in California if the average signal is used. It will be executed in Ontario if the marginal signal is used. Our results for both the temporal and the spatial scheduling demonstrate that neither of the signals yields consistent outcomes for the carbon-aware optimizations.

**Key takeaway.** *The negative correlation between the average and marginal carbon intensity signals results in conflicting scheduling decisions for carbon-aware temporal and spatial optimizations. The scheduling outcomes are region-dependent and are not consistent across regions for either of the signals.*

**Impact on carbon savings calculations.** Our previous analysis demonstrates that the scheduling decisions guided by both signals do not always align. Our next set of experimental results shows that they also significantly impact the perceived carbon savings as a result of carbon-aware scheduling. This is an important consideration where the enhanced focus on sustainability necessitates reporting any reductions in the carbon footprint resulting from a change in the operations of an individual or an organization.

Figure 4a shows carbon savings across all the regions when the average signal is used to schedule the workloads. The two box-plots correspond to the signal used for calculating carbon savings compared to the counterfactual of no workload shifting (carbon-agnostic execution). Figure 4b shows the same for a scenario when the marginal signal is used. Interestingly, in both scenarios, the carbon savings only manifest when the same signal is used to compute carbon savings. 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, Figure 4c and Figure 4d show mean carbon savings of ~87% when the average and marginal signals are used for spatial workload shifting. As with 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 based on the other signal. 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.

**Key observation.** *Even when the decisions of carbon-aware optimizations are fixed, the carbon savings estimated based on both signals do not yield the same results providing unclear incentives to use either of the signals for scheduling and accounting.*

### 3.4 Implications on grid operations

In addition to estimating the carbon footprint of electricity demand and quantifying the carbon savings from carbon-aware optimizations, the aim of a carbon intensity signal is to shape the electricity demand such that the use of electricity from renewable and low-carbon energy resources is maximized [3, 6]. While accurately estimating the impact on the grid without physical measurements or a high-fidelity simulator is challenging, we present key observations relating to how the choice of carbon intensity signal for carbon-aware optimizations may impact the grid operations.

**Impacts of temporal shifting.** As outlined in Section 3.3, the negative correlation between the average and the marginal signals leads to conflicting scheduling decisions made by carbon-aware optimizations. Here, we revisit Figure 3a and Figure 3b to understand how these conflicting decisions impact grid operations. If the carbon intensity signal in a given region is almost static (e.g., the marginal signal in Arizona), a scheduler may not shift the workload to avoid shifting costs. This can potentially indicate to the grid operators that the demand is inflexible.

Furthermore, if the signal is variable but has low values at night (e.g., the marginal signal in Virginia), the scheduler will shift the workloads to nighttime. The increase in demand at nighttime would require grid capacity planners to add more generators at night. As widely available low-carbon energy sources, such as solar and wind, are either unavailable or intermittent at night, they cannot be used to fulfill the added demand. This can potentially force the operators to add high-emitting generation resources such as natural gas generators [9]. Of course, using one of the signals in several regions can shift the demand to time periods with low carbon energy potential, but the same signal would not achieve the same outcome for the remaining regions.

**Impact of spatial shifting.** Similar to our analysis for the temporal shifting, we revisit Figure 3c, which shows two regions to illustrate the conflicting decisions for the spatial shifting scenario. In this scenario, using the average signal would shift any added demand from California to Ontario, as it almost always has a lower average carbon intensity than California. As most of Ontario's energy demand is fulfilled by nuclear and hydro, a significant addition of new demand may trigger the natural gas generators that serve the marginal demand in the region. This is because nuclear and hydro resources require long-term planning, are subject to public opinion, and cannot handle short-term increases in demand.

On the other hand, using the marginal signal would shift the demand to California, where the marginal carbon intensity is very



low as it is often fulfilled by wind or solar power plants that may have been curtailed before the added demand. However, if a significant demand is added to the California grid at all hours of the day, solar and wind would not be able to satisfy all the demand, and it would be satisfied by using the dominant energy resource of natural gas. As the average carbon intensity of California is 5× higher than Ontario, the migrated workload would be served by a carbon-intensive electricity supply if solar or wind do not engage as marginal generators. Furthermore, as the marginal signals for both locations do not overlap, all the additional demand in Ontario at all timeslots would migrate to California.

**Key takeaway.** *The choice of carbon intensity signal for carbon-aware optimizations determines: if the workloads shift temporally, the set of generators that are used to satisfy the added demand in the short term (both temporally and spatially), and the deployment of renewable energy as part of the long-term capacity planning.*

## 4 RELATED WORK

To the best of our knowledge, 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 [9], who look at the impact of traditional marginal signal and how it can be extended to incorporate future capacity planning implications. Recent work has focused on leveraging the flexibility in the energy demand for various workloads to reduce their carbon footprint. For computing workloads, recent work has explored leveraging the temporal and spatial flexibility of computing to reduce its carbon footprint [1, 2, 10–12, 14, 15, 18, 21, 22, 25–27, 30, 31]. While household energy demand cannot be shifted to other locations, recent work explores shifting background and non-interactive appliances to low-carbon periods to reduce their carbon footprint [4, 13]. Similarly, researchers have begun exploring decarbonization potential across other sectors, such as transportation [20, 24].

Importantly, though, there is little consensus on the choice of carbon signal, its short- and long-term impact on the grid, and its usefulness for accurate carbon accounting [5, 9, 19]. Most of the prior work uses the average signal, while some studies use the marginal carbon intensity signal for carbon-aware optimizations [12].

## 5 CONCLUSION AND FUTURE WORK

In this study, we demonstrate the fundamental differences between two commonly used carbon signals, marginal and average intensities, which yield distinct outcomes when applied to spatiotemporal scheduling. We have observed that the relative variability of these carbon intensity signals varies by location, leading to divergent results in carbon-aware spatial workload shifting. The ongoing debate in industry and research surrounding the choice of carbon signals for sustainability policies highlights the complexity and opacity of the factors influencing the relationship between these paradigms. Our goal is to facilitate the ongoing discourse on the topic. At the same time, we plan to continue further research using data-driven and analytical models for the electric grid. We also plan to explore approaches to devising either new carbon intensity signals or leveraging the novel combinations of existing signals.

## ACKNOWLEDGMENTS

We thank the e-Energy reviewers for their valuable comments, which improved the quality of this paper. We thank Electricity Maps and WattTime for access to their carbon-intensity datasets. This research was supported by NSF grants 2213636, 2105494, 2021693, 2020888, 2213636, DOE grant DE-EE0010143, and by VMware.

## REFERENCES

- [1] Bilge Acun, Benjamin Lee, Fiodar Kazhemiaka, 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] Emre Celebi and J David Fuller. 2012. Time-of-use pricing in electricity markets under different market structures. *IEEE Transactions on Power Systems* 27, 3 (2012), 1170–1181.
- [4] CityLearn. 2023. CityLearn Challenge: Using AI for Building's Energy Management. <https://www.aicrowd.com/challenges/neurips-2022-citylearn-challenge>.
- [5] 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>.
- [6] Valeria Di Cosmo, Sean Lyons, and Anne Nolanab. 2014. Estimating the impact of time-of-use pricing on Irish electricity demand. *The Energy Journal* 35, 2 (2014), 117–136.
- [7] 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*. ACM, New York, NY, USA, 1877–1894.
- [8] US EIA. 2012. Electric generator dispatch depends on system demand and the relative cost of operation.
- [9] 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.
- [10] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu. 2022. ACT: Designing Sustainable Computer Systems with an Architectural Carbon Modeling Tool. In *Proceedings of the 49th Annual International Symposium on Computer Architecture (New York, New York) (ISCA '22)*. Association for Computing Machinery, New York, NY, USA, 784–799. <https://doi.org/10.1145/3470496.3527408>
- [11] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2021. Chasing Carbon: The Elusive Environmental Footprint of Computing. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, Los Alamitos, CA, USA, 854–867.
- [12] 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.
- [13] 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>
- [14] Adam Lechowicz, Nicolas Christianson, Jinhang Zuo, Noman Bashir, Mohammad Hajiesmaili, Adam Wierman, and Prashant Shenoy. 2024. The Online Pause and Resume Problem: Optimal Algorithms and An Application to Carbon-Aware Load Shifting. In *Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS 2024)*. ACM, New York, NY, USA, 35 pages.
- [15] 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.
- [16] Electricity Maps. 2022. Electricity Map. <https://app.electricitymaps.com/map>.
- [17] Electricity Maps. 2024. Electricity Maps Methodology. <https://www.electricitymaps.com/methodology>.
- [18] 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.

- [19] 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/>.
- [20] Mahsa Sahebdeh, 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>
- [21] 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 *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, New York, NY, USA, 252–265.
- [22] 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.
- [23] FERC Staff. 2020. Energy Primer: A Handbook of Energy Market Basics.
- [24] 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>
- [25] Thanathorn Sukprasert, Abel Souza, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. Quantifying the Benefits of Carbon-Aware Temporal and Spatial Workload Shifting in the Cloud. arXiv:2306.06502 [cs.DC]
- [26] Jennifer Switzer, Gabriel Marcano, Ryan Kastner, and Pat Pannuto. 2023. Junkyard Computing: Repurposing Discarded Smartphones to Minimize Carbon. In *Proceedings of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2* (Vancouver, BC, Canada) (*ASPLOS 2023*). ACM, New York, NY, USA, 400–412. <https://doi.org/10.1145/3575693.3575710>
- [27] Seyedali Tabaeiaghdaei, Simon Scherrer, Jonghoon Kwon, and Adrian Perrig. 2023. Carbon-Aware Global Routing in Path-Aware Networks. In *Proceedings of the ACM International Conference on Future Energy Systems (e-Energy)*. ACM, New York, NY, USA, 144–158.
- [28] 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>
- [29] WattTime. 2022. WattTime. <https://www.watttime.org/>.
- [30] 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.
- [31] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, et al. 2022. Sustainable AI: Environmental Implications, Challenges and Opportunities. *Proceedings of Machine Learning and Systems (MLSys)* 4 (2022), 795–813.