Data Centers Carbon Emissions at Crossroads: An Empirical Study

DIPTYAROOP MAJI, University of Massachusetts Amherst, USA
WALID A. HANAFY, University of Massachusetts Amherst, USA
LI WU, University of Massachusetts Amherst, USA
DAVID IRWIN, University of Massachusetts Amherst, USA
PRASHANT SHENOY, University of Massachusetts Amherst, USA
RAMESH K. SITARAMAN, University of Massachusetts Amherst, USA

Data centers' energy consumption is expected to rise significantly over the next five years due to the accelerated growth of AI and its computational demands. Simultaneously, initiatives are underway to mitigate the environmental impact of the energy usage by replacing brown energy sources (e.g., coal) with green energy sources (e.g., solar), resulting in a projected decrease in grid carbon intensity. Yet, the interplay between these two contrasting forces on data centers' carbon emissions — a function of both energy consumption and carbon intensity — has not received much attention.

In this paper, we analyze the operational carbon emissions of data centers at this crossroads, by considering the increasing data center energy consumption and the decarbonization of the electricity grid. In particular, we integrate publicly available current and future energy projections, consider multiple future scenarios, and provide a 5-year projection of datacenter carbon emissions at the global, US, and state levels. Our analysis shows that over the next five years, the rate of data center demand growth will overshadow the rate of grid decarbonization, with global (resp. US) emissions projected to rise by 4.2× (resp. 4.1×) in the worst case by 2030. Moreover, we observe considerable regional differences within the US, where emissions in some states could increase by up to 3.4× by 2030.

CCS Concepts: • **Social and professional topics → Sustainability**; • **General and reference → Estimation**; • **Applied computing → Data centers**.

Additional Key Words and Phrases: Operational Carbon Emissions, Energy Consumption, 5-year Projections.

1 Introduction

Data centers' energy consumption has been a global concern for the past 20 years, during which researchers have analyzed and projected the energy consumption of data centers and the internet. For example, researchers [32, 36, 43, 44, 50] have analyzed the energy consumption of data centers between 2000 and 2020, with a common conclusion that data centers' energy demand will be stable, accounting for 1-2% of the global energy consumption. Despite the increase in computing demand, such stability was attributed to energy efficiency gains from advances in integrated circuit (IC) designs, where energy efficiency have doubled every 1.57-2.6 years (Koomey's law) [33, 35] along with advances in cooling (e.g., open-air cooling [2]), operating data centers at higher temperatures [47, 62], and efficient power distribution approaches [2]. As a result, between 2010 and 2018, the energy efficiency of computing and storage increased by 4.1× and 9×, respectively [43]. Moreover, the International Energy Agency (IEA) estimates that from 2010 to 2020, global data centers' energy consumption increased by only 10%, despite a 9.4× increase in computing demand [28].

However, recent studies highlight that the era of significant energy efficiency improvements in data centers is coming to an end, and foresee a substantial increase in data centers' energy demands [25, 51]. For example, according to the recent report by Lawrence Berkeley National Laboratory (LBNL) [51], the Compound Annual Growth Rate (CAGR) — which represents the average annual growth rate over a specified period — of data centers' energy consumption has increased from 7% between 2014–2018 to 18% between 2018–2023, and expected to increase by 13–27% between 2023–2028. Similarly, hyperscale data centers have reported 2-3× increase in energy consumption between 2017 and 2022 [60]. Such increase in energy consumption is attributed to: 1) slowdowns in energy efficiency gains due to approaching the end of Moore's law and Landauer's limit [4, 35, 38] as well as nearing the optimal PUE with little room for further improvements, and 2) the rise of AI and its computing demand [25, 60], where data centers are expanding their compute fleets with AI accelerators, that typically consume more energy [51].

In contrast to the energy consumption of data centers, which have been extensively analyzed [18, 25, 32, 36, 37, 43, 44, 50, 51, 53], the emissions of data centers, which reflects the true environmental and health impacts of data centers [20, 27], have just begun to receive attention. In this study, we focus on the operational carbon emissions of data centers, which are the emissions directly associated with the electricity consumption of data centers, excluding the emissions from the supply chain, often denoted as embodied emissions [3, 24]. The operational carbon emissions of data centers, measured in carbon dioxide equivalent (CO_2eq), is defined by the equation $C = E \times I$, where E indicates the data center's total energy consumption and I denotes the carbon intensity of the energy feeding this data center. The carbon intensity of electricity, in CO_2eq per unit of energy (e.g., g·CO_2eq/kWh), reflects the cleanliness of the electricity grid, where grids that predominantly use fossil-based sources (e.g., coal) exhibit high carbon intensity. Conversely, electricity grids with a significant share of renewable energy have low carbon intensity. For example, countries such as Sweden and Norway meet a large fraction of their energy demand through hydroelectric power, resulting in a low carbon intensity.

In recent years, electricity grids, which are responsible for 26% of the global emissions [58] and 25% of US emissions [59], have been undergoing a significant transformation. Motivated by the

Authors' Contact Information: Diptyaroop Maji, University of Massachusetts Amherst, Amherst, MA, USA, dmaji@cs.umass.edu; Walid A. Hanafy, University of Massachusetts Amherst, Amherst, MA, USA, whanafy@cs.umass.edu; Li Wu, University of Massachusetts Amherst, Amherst, MA, USA, liwu@cs.umass.edu; David Irwin, University of Massachusetts Amherst, Amherst, MA, USA, irwin@ecs.umass.edu; Prashant Shenoy, University of Massachusetts Amherst, Amherst, MA, USA, shenoy@cs.umass.edu; Ramesh K. Sitaraman, University of Massachusetts Amherst, Amherst, MA, USA, ramesh@cs.umass.edu.

Data	Source	Scope			Year Range	
		Global	US	State	Estimate	Projection
Electricity Consumption	IEA [25]	✓	–	–	2020–2024	2025–2030
	IEA w/ LBNL CAGR	✓	–	–	2024	**2025–2030**
	Koot et al. [37]	✓	–	–	Pre-2021	2021–2030
	EPRI [18]	–	✓	✓	2016-2023	**2024–2030**
	LBNL [51]	–	✓	–	2016-2024	2025–**2030***
Carbon Intensity	Ember [17]	✓	✓	–	2016–2024	**2025–2030**
	Electricity Maps [16]	–	✓	✓	2017-2024	**2025-2030**

* we extrapolated the energy consumption for 2029 and 2030, based on the reported best and worst case CAGR.

Table 1. *Data sources used in our analysis, their scope, and the year range. Our projections are highlighted in bold.*

environmental benefits of renewables and their cost competitiveness [27], the adoption of renewable energy (e.g., solar and wind) has experienced exponential growth over the past 20 years. Consequently, the carbon intensity of electricity has seen significant decreases, with many grids targeting 2050 to achieve carbon-free electricity [10]. For example, Global, US, and Denmark carbon intensities have decreased by 0.31, 36.8%, and 77.1% between 2000 and 2024 [17].

In this paper, we investigate the operational carbon emissions of data centers at this pivotal crossroads by analyzing two competing trends: (1) energy consumption in data centers, which is experiencing (or expected to experience) an unprecedented increase [51], and (2) the grid carbon intensity, which is undergoing (or is expected to undergo) a decarbonization phase. Unlike the previous reports, which focus on energy consumption [25, 51], or global carbon emissions [25], our analysis characterizes the operational carbon emissions of data centers at global, US, and state granularities. In particular, we integrate current and future energy projections during this period under multiple scenarios that consider the grid's ability to transition to carbon-free sources, as well as scenarios where global load projections follow the worst-case analysis reported by [51].

Our Key Findings. The key findings are as follows:

(1) Over the next five years, the growth in data center demand will likely overshadow the grid decarbonization. Consequently, by 2030, the global and US data center emissions could increase by up to 4.2× (27% CAGR) and 4.1× (26.3% CAGR), respectively.
(2) Emissions computed based on different demand projections and decarbonization scenarios vary significantly. For example, 2030 emissions in the US could differ by 6× between the best and the worst case.
(3) We find significant regional differences in data center emissions within the US. Virginia is expected to see the highest emission increase within the US due to high demand growth and a low rate of decarbonization. Data center emissions in Virginia could increase by 3.4× (22.6% CAGR) by 2030 in the worst case.
(4) In 2030, emissions from the top 10 US states with the highest data center loads could vary by 24.8%, depending on whether the most or the least carbon-intensive states see the most demand growth.

2 Research Questions & Methodology

This section presents our research questions, data sources, and methodology.

2.1 Research Questions

This paper analyzes the *location-based* [15] *operational carbon emissions* of data centers from electricity use, calculated as a product of electricity consumption and carbon intensity. Specifically, we address the following research questions:

Q1 *What are the current and projected carbon emissions in 2030, from data centers globally, in the US, and across US states?*
Q2 *How do the different data center energy demand projections and grid decarbonization scenarios affect data center emissions?*
Q3 *How would the US data center emissions in 2030 vary if the data center energy demand shifted away from the current hotspots?*

2.2 Data Sources

Table 1 summarizes the data sources used in our analysis. While we recognize that our findings depend on the accuracy of the data sources, our work uses the most recent publicly available data center energy consumption estimates and projections, along with estimates of carbon intensity at the global, US, and state levels. Additionally, as shown later, we enhance these data with new scenarios that further support our conclusions. In this paper, we refer to the data prior to the publication year of a report as *estimates* and future predictions as *projections*. All the data used in this paper are available at **https://github.com/codecexp/dc-emissions-2030**.

Energy Consumption. Our analysis is based on estimates and projections of data centers' global, US, and state levels (see Figure 1). Our global analysis relies on the data from the International Energy Agency (IEA) and Koot et al. [37] (see Figure 1a). The IEA released the "Energy and AI" report [25] in 2025. Their report projects global data center electricity usage until 2035, given the accelerated growth of AI. However, we limit our analysis to their 2030 projections. Our analysis also considers the projections between 2021 and 2030 by Koot et al. [37]. Although their projections did not explicitly consider the rise of AI, their 99th percentile worst-case analysis, which we use, highly matches the witnessed trends and IEA projections.

Our US analysis is based on the Lawrence Berkeley National Laboratory (LBNL) [51] and the Electric Power Research Institute (EPRI) [18] reports (see Figure 1b). The LBNL report projects data center energy consumption between 2024 and 2028 in the US. The report projects a CAGR of 13%–27%, attributing this to the accelerated growth of generative AI in the last year. In contrast to the LBNL report, which estimates the aggregate data center energy consumption across the US, the EPRI report provides aggregated and per-state data center energy consumption in the US while considering a more conservative CAGR of 3.7%–15%. We note that higher

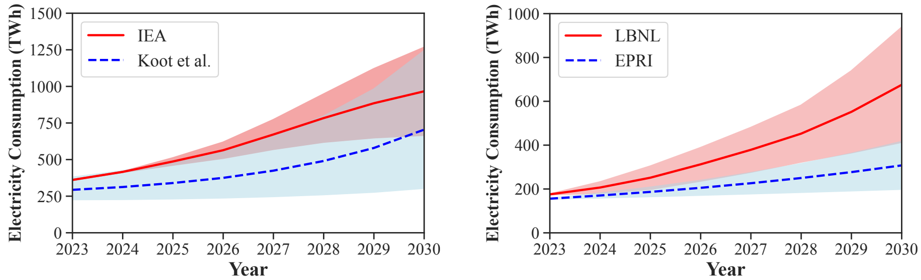

(a) *Global* (b) *US*

Fig. 1. *Data center energy consumption projections till 2030.*

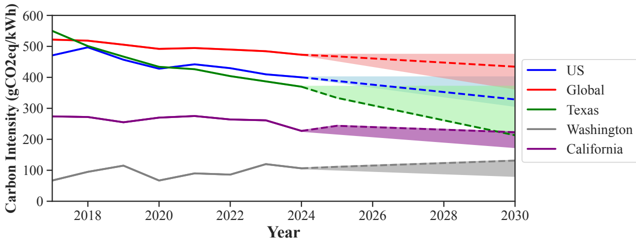

Fig. 2. *CI estimates and projections between 2017 and 2030 Projections based on current trends are shown by dashed lines.*

growth estimates in the LBNL report lead to non-overlapping estimates, highlighted in Figure 1b, which can highly affect the future carbon emissions. Per various reports, the consumption is projected to increase by 1.3–3.03× (resp. 1.2–4.05×) at the global (resp. US) level between 2024 and 2030.

Lastly, our state-level analysis considers the data from EPRI. However, in contrast to EPRI, which assumes that all states have the same CAGR, we consider that states have different CAGRs as highlighted in recent reports [7, 14, 39]. Specifically, Virginia, California, and Texas, which constitute 33.3% of the total US data centers [12], exhibit 23%, 23%, and 20% CAGRs, respectively. For other states, we assume a CAGR of 15%, as projected by EPRI.

Carbon Intensity (CI). We use the life cycle carbon intensity estimates provided by Electricity Maps [16] and Ember [17] as the other input to our analysis. Electricity Maps reports carbon intensities at the regional Independent System Operator (ISO) and the US levels, while Ember reports at the global and per-country levels. We compute the state-level carbon intensity by mapping the state to the corresponding grid ISO. For states spanning multiple ISOs, we use the carbon intensity of the ISO that covers most of the state data centers as per the US Data Center Map [12].

Figure 2 shows the historical carbon intensity estimates for different levels. The carbon intensity is generally seeing a consistent decrease, where the global and US energy's carbon intensity have decreased by 9.3% and 15% between 2017 and 2024 with an average yearly decline of 1.4% and 2.2%, respectively. Figure 2 also shows the projections till 2030 based on these trends, which we discuss in detail in Section 2.3.

2.3 Methodology

We now discuss our methodology in detail. Our analysis computes the future emissions of data centers globally, in the US, and across US states with high data center capacities using an extrapolation-based approach. For clarity, we convert all emission values from g·CO_2eq to MtonsCO_2eq, where 1 MtonsCO_2eq = 10^{12} g·CO_2eq. Our work looks at data center energy demand estimates and projections from recent governmental and academic reports, and combines them with multiple demand increases and grid decarbonization scenarios. In particular, we examine the potential permutations of energy and carbon intensity projections and report the emissions in the worst, best, and average cases. Next, we detail our growth scenarios for the grid's carbon intensity and data center demand.

Carbon Intensity Analysis. Despite the ambitious goals and significant progress in integrating renewable energy sources, fully decarbonizing the grid often encounters three primary challenges. First, consumers often prefer the stability in electricity supply that only conventional energy sources can provide. For example, the newly announced Stargate data center will be powered by a 360.5 MW local natural gas power plant [54]. Second, due to the intermittent nature of renewables, decarbonizing the grid may incur exponentially increasing marginal costs, posing a challenge in the whole grid decarbonization [1, 11]. Third, strained supply chains may impede governmental efforts to acquire the necessary generation and transmission infrastructure [25]. Considering these uncertainties, we employ three scenarios in our analysis:

(1) ZeroCarbon 2050 (`ZC50`): This method presents an optimistic scenario, where the electricity grids undergo a full decarbonization by 2050. We assume that the carbon intensities will start decreasing linearly from 2024 and become zero in 2050.
(2) Current Decarbonization (`CD`): In this scenario, we consider the historical decarbonization rate between 2017 and 2024 and employ a regression-based approach to compute the carbon intensities until 2030. We note that in some cases (e.g., Washington, see Figure 2), the carbon intensity may see minor increases.
(3) No Change (`NC`): Lastly, we consider a pessimistic scenario where carbon intensity remains unchanged from 2024 to 2030. Note that although reports highlight that new demands may be met by fossil fuels [54], we do not include scenarios where increases in data centers' energy demand lead to increases in carbon intensity.

These future trends are highlighted in Figure 2. We see that the trends vary across regions. As shown, global (resp. US) carbon intensities may decline steadily, decreasing by 10.4% (resp. 13.6%) by 2030 when averaged across the three methods. More importantly,

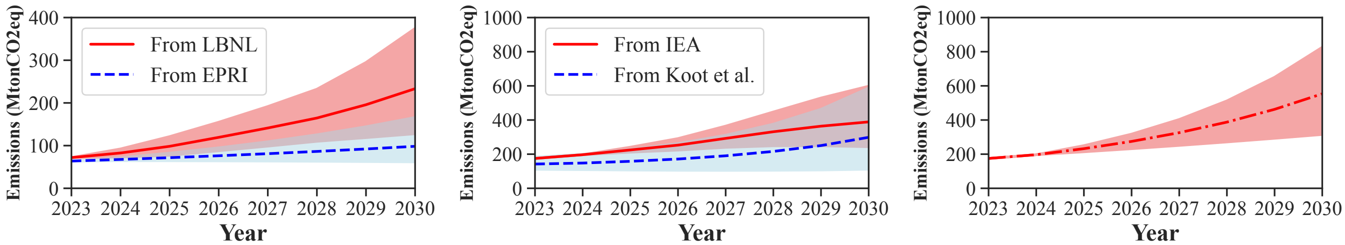

(a) *US* (b) *Global* (c) *Global (IEA w/ LBNL CAGR).*

Fig. 3. *Yearly projections of the US and global data center carbon emissions till 2030. Data center emissions in the US (resp. global) could increase by up to* 4.1× *(resp.* 4.2×*) by 2030 in the worst case.*

the figure highlights that in some locations (e.g., Texas), the current decarbonization trends outpace the ZC50 scenario. Lastly, the figure highlights situations where the carbon intensity increases, such as the state of Washington, where the carbon intensity may increase by up to 1.2× compared to 2024.

Data Center Energy Consumption Projections. In addition to the energy consumption projected by different sources, we consider two hypothetical scenarios as follows:

Scenario 1: Since the US hosts almost half of the world's data centers, it is possible that the US data center emissions will increase at a much faster rate than the global average. For instance, the CAGR outlined by LBNL [51] and IEA [25] reports reflect this difference, wherein the LBNL's lowest growth rate is comparable to the worst-case scenario considered by the IEA, leading to situations where US energy consumption sometimes exceeds global energy consumption (see Figure 1). However, it is also possible that other countries will see a similar growth rate to the US over the next five years. To explore this, we model a scenario in which the global growth rate adheres to both the best and worst-case CAGRs outlined by LBNL (referred to as *"IEA w/ LBNL CAGR"*).

Scenario 2: Data center demand growth in a region depends on several factors, such as land availability, grid generation and transmission constraints, etc. While Virginia currently has the most data centers in the US, grid strain is slowing down newer deployments [30]. Meanwhile, states with enough grid capacity are embracing data center growth [54]. Hence, the future growth across states may deviate from the previous trends. To explore this, we consider a scenario where data center demand increases differ from historical trends.

Specifically, we consider two alternate scenarios and analyze the variance in emissions in 2030 between them. The first scenario (referred to as "Brown") assumes that the states projected to be the most carbon-intensive in 2030 see the highest data center growth. The second scenario (referred to as "Green") assumes that the greenest states see the highest growth. We consider these contrasting scenarios to show the two ends of the emission spectrum. Our analysis has two assumptions: (1) the maximum demand growth in any state is 30%, since any higher CAGR may not be practical, and (2) we assume that only the top 10 states in terms of data center demand can experience higher growth rates than today, since demand increases in other states have minimal impact on the results as their capacity and demand increase abilities are limited.

3 Data Center Emissions Analysis

In this section, we analyze the changes in data center operational emissions from 2024 to 2030, starting with projections for the US. We then analyze global data center emissions and compare the two projections. Finally, we analyze regional emissions within the US and how the US emissions may be affected by where data center demand grows over the next few years.

3.1 US Data Center Emissions

Figure 3a shows the US data center emission estimates and projections from 2023 to 2030 using LBNL, EPRI, and Electricity Maps' data. US data center emissions in 2024 range between 64.6 and 92.2 MtonsCO_2eq, increasing from 61.6 MtonsCO_2eq in 2023 when considering a carbon intensity of 350 g·CO_2eq/kWh[1], as reported in [51]. US emissions calculated using the LBNL data are projected to increase between 1.8× and 4.1× by 2030, with an average increase of 2.8×. This corresponds to an average CAGR of 18.8%, and up to 26.3% in the worst case. In comparison, emissions calculated using the EPRI estimates are projected to increase by 1.5× on average and 2.3× in the worst case. This corresponds to an average CAGR of 6.4%, and up to 14.9% in the worst case. Interestingly, our analysis shows that emissions could decrease by 8% in the best case, where emissions could be 6× lower compared to the worst case emissions based on LBNL projections. However, given current trends, such an outcome will likely require a significantly faster rate of AI algorithmic and efficiency gains, as well as higher grid decarbonization rates, than predicted today.

Finally, since the emission range using the carbon intensity projections from Ember [17] falls within the range obtained using the projections from Electricity Maps, we only show the emission projections using the Electricity Maps data.

> ***Key Takeaways:*** *US data center emissions could increase by* 2.8× *in the average case and up to* 4.1× *by 2030 based on LBNL, which corresponds to CAGR of 18.8% on average, and 26.3% in the worst cases. Moreover, our analysis highlights that emissions computed based on different demand projections and decarbonization scenarios vary significantly, where the difference between emissions based on LBNL's worst case and EPRI's best case can vary by 6×.*

[1] Note that the 2023 carbon emissions, based on the 409.8 g·CO_2eq/kWh carbon intensity reported by Electricity Maps, are estimated to be 72.1 MtonsCO_2eq.

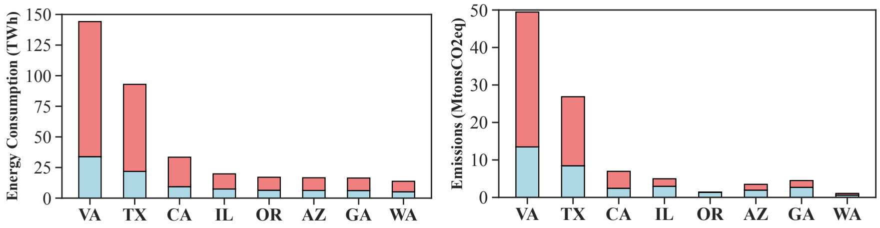

(a) *Energy Consumption* (b) *Average Emissions*

Fig. 4. *Data center demand and respective emissions (2023 and 2030) for the top eight US states.*

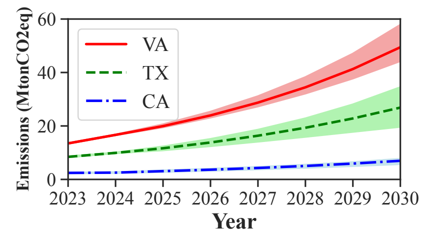

Fig. 5. *Yearly change in emissions in top three states (dark lines represent the average case).*

3.2 Global Data Center Emissions

We next analyze how the global data center emissions may change from 2024 to 2030, and compare them with the emissions from the US data centers. Figure 3b shows the projected emissions till 2030. The current global data center emissions (in 2024) are estimated to be between 109.2 and 198.8 MtonsCO$_2$eq, with estimates calculated using the data from recently published sources leaning towards 198.8 MtonsCO$_2$eq. This is equivalent to 0.26%–0.47% of the global carbon emissions across all sectors[2].

When calculated using the IEA data (resp. data from Koot et al.), global data center emissions are projected to range from 243 to 597 MtonsCO$_2$eq (resp. 111.7 – 587 MtonsCO$_2$eq) by 2030[3]. This implies an increase of 1.2–3× (resp. 1.02–2.9×) from 2024, with an average increase of 2×. This is equivalent to a 12% CAGR on average, up to 20.3% (resp. 19.7%) in the worst case. Although the best case projects that emissions will remain almost stagnant, it is based on the data from Koot et al. [37], which does not account for the current growth of AI. Hence, similar to the best case for US data center emissions, such an outcome will likely require different efficiency and grid decarbonization trends than predicted today.

Emissions if Global Demand Growth Follows the US. Figure 3a and Figure 3b show different trends of emission increase between the US and the global data center emissions, especially when calculated using the recent LBNL and IEA reports. The US emissions based on the LBNL data have an 18.8% average CAGR, much higher than the 12% CAGR projected using IEA data. We also observe that the global CAGR decreases from 13.9% between 2024 and 2028 to 8.3% beyond 2028, indicating a much slower growth in emissions in the later years. However, if the global data center demand growth follows the same rate as the US (*"IEA w/ LBNL CAGR"* scenario), global emissions from data centers could range between 315.1 – 825.7 MtonsCO$_2$eq in 2030. This would mean a 1.6–4.2× increase from 2024, with an average of 2.8× — equivalent to having an 18.8% CAGR (see Figure 3c). While data center emissions globally may increase slightly faster than in the US, as the global decarbonization rate is slightly slower, the world and the US will see similar emission growth rates on average.

[2]Global emissions across all sectors is estimated to be 41.6 $GtonsCO_2eq$ [61] in 2024.
[3]Our calculations encompass the recent IEA projections [25], which estimate the values to be 215 – 475 MtonsCO$_2$eq.

Key Takeaways: *Recent demand estimates suggest that global data center emissions could increase by up to 3× by 2030, with the rate slowing down in the later years. However, if the global demand follows the current US demand growth rate, such emissions could increase by up to 4.2× by 2030, having CAGR of 18.8% in the average case, and up to 27% in the worst case.*

3.3 US States Emissions

Next, we analyze data center emissions within different US states. State-wise emissions are highly non-uniform, with the top five states currently accounting for 39%–50% of the US data center emissions. Figure 4 shows the current data center demand and emissions (in 2024), and the projected increase by 2030, which lies within the best and worst case projections in LBNL [51]. The states are listed in order of projected demand in 2030. As expected, Virginia currently has the highest data center demand and emissions, followed by Texas. Interestingly, some states (e.g., California) have lower data center emissions than other states (e.g., Illinois) in 2024, despite having more demand. Virginia is expected to have the most data center emissions in 2030 since it has the highest CAGR for data center demand, and a slow rate of grid decarbonization. On the other hand, Oregon is projected to have only a 5.1% increase in emissions by 2030 due to a very high rate of grid decarbonization. Consequently, Oregon is expected to have lower carbon emissions than many states (such as Georgia), even though it is projected to have a similar or higher data center demand than those states.

Figure 5 focuses on Virginia, Texas, and California — the three states expected to have the most emissions in 2030 — and shows the range of increase in emissions Year-on-Year (YoY). Emissions in Virginia (resp. Texas) are expected to grow between 2.6× and 3.4× (resp. 2× and 3.4×) from 2024 to 2030, with the average rate being 3× (resp. 2.7×). This corresponds to an average CAGR of 19.8% and 18% in Virginia and Texas, respectively. California is also projected to have a 2.7× increase in emissions (18.2% CAGR) on average over the next five years, similar to Texas. Note that our state analysis only considers the *expected* increases in energy and does not account for worst-case energy demand; therefore, our state analysis shows lower emission increases than the US-level.

Interestingly, Virginia and Texas have different CAGRs for data center emissions, even though they have the same data center demand growth rate. This is because Texas is projected to have an average grid decarbonization rate of 4.6% per year, higher than the 2.7% for Virginia. Moreover, although California and Texas have similar emission growth rates, the amount of emissions from data

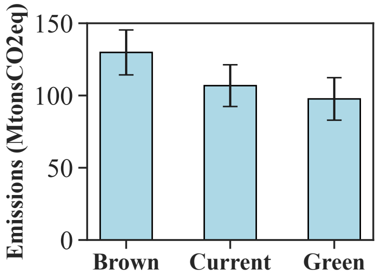

Fig. 6. *2030 data center emissions (top 10 states) under various scenarios.*

centers (in *$MtonsCO_2eq$*) in California is 3.9× less than that in Texas since the California grid is much greener than the Texas grid [16].

2030 Emissions from Alternate Demand Scenarios. Finally, we analyze the "Brown" and "Green" scenarios, where future demand growth deviates from current trends. Figure 6 shows the projected 2030 emissions of the top 10 states for the two scenarios, and also compares them to the projected emissions if the demand growth trend does not change in the future. Compared to the current trend, emissions in the top 10 states can increase by 21.5% or decrease by 8.6% on average, depending on how the trend changes over the next five years. This corresponds to an increase of 23 $MtonsCO_2eq$ or a decrease of 9.2 $MtonsCO_2eq$ from the data center emissions projected today based on the current trends. Overall, emissions in the top 10 states can vary by 24.8%, or 32.2 $MtonsCO_2eq$, on average between the two scenarios.

Key Takeaways: *Emission growth rates vary by state, with Virginia and Texas seeing up to 3.4× increase in the worst case. States with similar growth rates can have different environmental impacts depending on how carbon-intensive the respective states are. Data center emissions in 2030 in the top 10 states can vary by 24.8% in the average case, depending on whether the most or the least carbon-intensive states see the most demand growth.*

4 Other Considerations

Analyzing data center emissions has several challenges, ranging from a lack of information to the validity of the estimates. While we try to include multiple estimates and approaches in our analysis, more needs to be done to analyze such emissions holistically. In this section, we discuss some considerations and limitations that are important but beyond the scope of this paper.

Location- versus Market-Based Emissions. Accounting for carbon emissions can follow location-based or market-based approaches. Location-based approaches directly follow their local grid energy mixture and report their carbon emissions. However, the market-based approach assumes that users (or data centers) can choose their sources based on Power Purchase Agreements (PPAs). However, due to the lack of such information and the challenges associated with market-based analysis [3, 5, 24, 40, 41], we limit our work to location-based carbon emissions only.

Grid Decarbonozation Trends. While we only consider linear decarbonization trends, we acknowledge that grid decarbonization may undergo different trends. For instance, grid decarbonization may see diminishing improvements as the ratio of renewables increase [1, 11]. In addition, we do not consider the impact of nuclear energy, which may significantly reduce the grid's carbon intensity [31, 52, 57], with expectations that some of these projects will come online by 2030. Lastly, while the increasing energy demand may lead to higher carbon intensity as data centers seek more stable energy sources [26], our worst-case analysis only considers a stable carbon intensity.

Embodied Emissions. The life cycle emissions of data centers go beyond operational emissions to include complex supply chains in buildings and computing hardware (*embodied emissions*). Although recent research has made significant strides in estimating the embodied carbon emissions of hardware, noticeable uncertainties still remain in the estimates [19, 49]. Additionally, the heterogeneous nature of equipment across data centers would make it highly complicated to project embodied emissions reasonably. Therefore, this work focuses only on operational emissions, with an extension to include embodied emissions as future work.

Effects of Carbon-Aware Demand-Response. Researchers have highlighted that demand-side adjustments leveraging the temporal and spatial variability of carbon intensity [1, 6, 8, 9, 13, 21–23, 29, 42, 45, 48, 55, 56, 63] can reduce the carbon emissions of data centers. However, due to the nascent nature of this domain and the lack of information about such adjustments in practice, analyzing its holistic impact is challenging and beyond the scope of this work.

Emissions Beyond 2030. Some experts predict that data center emissions may plateau or decline after 2030 [25, 46, 60], especially as AI gains further significance and aids grid decarbonization. However, such long-term predictions have many uncertain variables and hence can often be significantly over- or under-estimated [34]. *Hence, while hoping for a better future, we present a conservative perspective and limit our projections to the next five years.*

5 Conclusions

In this paper, we project the data center carbon emissions till 2030 at the global, US, and state levels. Our analysis shows that the current demand growth will overshadow the current rate of grid decarbonization, causing the emissions to increase by up to 4.2× (resp. 4.1×) globally (resp. in the US) from 2024. Within the US, Virginia and Texas could see up to 3.4× emission increase by 2030. However, the emissions may vary widely depending on which states see the most demand growth in the next few years. We aim to extend our analysis to other regions outside the US and analyze embodied emissions in the data centers as future work.

Acknowledgments

This research is supported by National Science Foundation (NSF) grants 2213636, 2105494, 2211302, 2211888, 2325956, the U.S. Department of Energy Award DE-EE0010143, and support from VMware. This work used Amazon Web Services through the CloudBank, which is supported by NSF grant 19250001.

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 *Proceedings of*

the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (Vancouver, BC, Canada) *(ASPLOS 2023)*. 118—-132. doi:10.1145/3575693.3575754
[2] Luiz André Barroso and Urs Hölzle. 2009. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan & Claypool Publishers. http://dx.doi.org/10.2200/S00193ED1V01Y200905CAC006
[3] Noman Bashir, David Irwin, and Prashant Shenoy. 2024. On the Promise and Pitfalls of Optimizing Embodied Carbon. *SIGENERGY Energy Inform. Rev.* 4, 3 (Sept. 2024), 94–99. doi:10.1145/3698365.3698380
[4] Noman Bashir, David Irwin, Prashant Shenoy, and Abel Souza. 2022. Sustainable Computing – Without the Hot Air. In *Proceedings of the First Workshop on Sustainable Computer Systems Design and Implementation (HotCarbon)*. Association for Computing Machinery, New York, NY, USA.
[5] Anders Bjørn, Shannon M Lloyd, Matthew Brander, and H Damon Matthews. 2022. Renewable energy certificates threaten the integrity of corporate science-based targets. *Nature Climate Change* 12, 6 (2022), 539–546.
[6] Roozbeh Bostandoost, Adam Lechowicz, Walid A. Hanafy, Noman Bashir, Prashant Shenoy, and Mohammad Hajiesmaili. 2024. LACS: Learning-Augmented Algorithms for Carbon-Aware Resource Scaling with Uncertain Demand. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems* (Singapore, Singapore) *(e-Energy '24)*. 27–45. doi:10.1145/3632775.3661942
[7] California Energy Commission. 2024. Data Center Load Forecasts, 2024 - 2040. Retrieved April 29, 2025 from https://www.energy.ca.gov/filebrowser/download/6686?fid=6686
[8] Zhiwei Cao, Xin Zhou, Han Hu, Zhi Wang, and Yonggang Wen. 2022. Toward a Systematic Survey for Carbon Neutral Data Centers. *IEEE Communications Surveys & Tutorials* 24, 2 (2022), 895–936. doi:10.1109/COMST.2022.3161275
[9] Andrew Chien. 2021. Driving the Cloud to True Zero Carbon. *Communication of the ACM* 64, 2 (February 2021).
[10] Clean Energy States Alliance. [n. d.]. Table of 100% Clean Energy States. https://www.cesa.org/projects/100-clean-energy-collaborative/guide/table-of-100-clean-energy-states/ Accessed: 2025-05-03.
[11] Wesley J Cole, Danny Greer, Paul Denholm, A Will Frazier, Scott Machen, Trieu Mai, Nina Vincent, and Samuel F Baldwin. 2021. Quantifying the challenge of reaching a 100% renewable energy power system for the United States. *Joule* 5, 7 (2021), 1732–1748.
[12] DataCenterMap. 2025. USA Data Centers. Retrieved April 28, 2025 from https://www.datacentermap.com/usa/
[13] 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* (Seoul, Republic of Korea) *(FAccT '22)*. 1877–1894. doi:10.1145/3531146.3533234
[14] Dominion Energy Virginia. 2024. 20-Year Data Center Forecast. Retrieved April 29, 2025 from https://www.pjm.com/-/media/DotCom/committees-groups/subcommittees/las/2024/20241025/20241025-item-03ai---dominion-data-center-large-load-request.pdf
[15] Electricity Maps. 2022. Understanding Electricity Scope 2 Attribution Rules. Retrieved June 26, 2025 from https://ww2.electricitymaps.com/reports-and-guides/accounting-guide
[16] Electricity Maps. 2025. Granular historical electricity data. Retrieved April 28, 2025 from https://portal.electricitymaps.com/datasets
[17] Ember. 2025. Yearly Electricity Data. Retrieved April 28, 2025 from https://ember-energy.org/data/yearly-electricity-data/
[18] EPRI. 2024. Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption. Retrieved April 28, 2025 from https://www.epri.com/research/products/000000003002028905
[19] 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. doi:10.1145/3470496.3527408
[20] Yuelin Han, Zhifeng Wu, Pengfei Li, Adam Wierman, and Shaolei Ren. 2024. The Unpaid Toll: Quantifying the Public Health Impact of AI. *Preprint* (2024). https://arxiv.org/abs/2412.06288
[21] Walid A. Hanafy, Roozbeh Bostandoost, Noman Bashir, David Irwin, Mohammad Hajiesmaili, and Prashant Shenoy. 2023. The War of the Efficiencies: Understanding the Tension between Carbon and Energy Optimization. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems* (Boston, MA, USA) *(HotCarbon '23)*. Article 19, 7 pages. doi:10.1145/3604930.3605709
[22] Walid A. Hanafy, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. CarbonScaler: Leveraging Cloud Workload Elasticity for Optimizing Carbon-Efficiency. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7, 3, Article 57 (December 2023), 28 pages. doi:10.1145/3626788
[23] Walid A. Hanafy, Qianlin Liang, Noman Bashir, Abel Souza, David Irwin, and Prashant Shenoy. 2024. Going Green for Less Green: Optimizing the Cost of Reducing Cloud Carbon Emissions. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3 (ASPLOS'24)*. 479—-496. doi:10.1145/3620666.3651374
[24] Peter Holzapfel, Vanessa Bach, and Matthias Finkbeiner. 2023. Electricity accounting in life cycle assessment: the challenge of double counting. *The International Journal of Life Cycle Assessment* (2023), 1–17.
[25] IEA. 2025. Energy and AI. Retrieved April 28, 2025 from https://www.iea.org/reports/energy-and-ai
[26] Industry Insider. 2025. OpenAI Chooses Abilene for Stargate Data Center. Retrieved May 9, 2025 from https://insider.govtech.com/texas/news/openai-chooses-abilene-for-stargate-data-center
[27] Intergovernmental Panel on Climate Change (IPCC). 2023. *Climate Change 2022 - Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.
[28] International Energy Agency. 2021. Global trends in internet traffic, data centres workloads and data centre energy use, 2010–2020. https://www.iea.org/data-and-statistics/charts/global-trends-in-internet-traffic-data-centres-workloads-and-data-centre-energy-use-2010-2020. IEA, Paris. Licence: CC BY 4.0.
[29] David Irwin, Prashant Shenoy, Mohammad Hajiesmaili, Walid A. Hanafy, Jimi Oke, Ramesh Sitaraman, Yuvraj Agarwal, Geoff Gordon, Zico Kolter, Deepak Rajagopal, Mani Srivastava, Vivienne Sze, Priya Donti, Andrew Chien, John Birge, Ali Hortacsu, and Line Roald. 2025. A Vision for Computational Decarbonization of Societal Infrastructure. *IEEE Internet Computing* (2025), 1–7. doi:10.1109/MIC.2025.3575016
[30] Josh Saul. 2024. Data Centers Face Seven-Year Wait for Dominion Power Hookups. Retrieved May 9, 2025 from https://www.bloomberg.com/news/articles/2024-08-29/data-centers-face-seven-year-wait-for-power-hookups-in-virginia?embedded-checkout=true
[31] João da Silva. 2024. Google partners with Kairos Power to deploy small nuclear reactors for AI data centres. *BBC News* (2024). https://www.bbc.com/news/articles/c748gn94k95o Accessed: 2025-01-14.
[32] Jonathan Koomey et al. 2011. Growth in data center electricity use 2005 to 2010. *A report by Analytical Press, completed at the request of The New York Times* 9, 2011 (2011), 161.
[33] Jonathan Koomey, Stephen Berard, Marla Sanchez, and Henry Wong. 2011. Implications of Historical Trends in the Electrical Efficiency of Computing. *IEEE Annals of the History of Computing* 33, 3 (2011), 46–54. doi:10.1109/MAHC.2010.28
[34] Jonathan Koomey and Eric Masanet. 2021. Does not compute: Avoiding pitfalls assessing the Internet's energy and carbon impacts. *Joule* 5, 7 (2021), 1625–1628. doi:10.1016/j.joule.2021.05.007
[35] Jonathan Koomey and Samuel Naffziger. 2015. Moore's Law might be slowing down, but not energy efficiency. *IEEE spectrum* 52, 4 (2015), 35.
[36] Jonathan G Koomey. 2008. Worldwide electricity used in data centers. *Environmental research letters* 3, 3 (2008), 034008.
[37] Martijn Koot and Fons Wijnhoven. 2021. Usage impact on data center electricity needs: A system dynamic forecasting model. *Applied Energy* 291 (2021), 116798.
[38] Rolf Landauer. 1961. Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development* 5, 3 (July 1961).
[39] Liuzixuan Lin, Rajini Wijayawardana, Varsha Rao, Hai Nguyen, Emmanuel Wedan GNIBGA, and Andrew A Chien. 2024. Exploding ai power use: an opportunity to rethink grid planning and management. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 434–441.
[40] Diptyaroop Maji, Noman Bashir, David Irwin, Prashant Shenoy, and Ramesh K Sitaraman. 2024. The Green Mirage: Impact of Location- and Market-based Carbon Intensity Estimation on Carbon Optimization Efficacy. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems* (Singapore, Singapore) *(e-Energy '24)*. 256–267. doi:10.1145/3632775.3639587
[41] Diptyaroop Maji, Noman Bashir, David Irwin, Prashant Shenoy, and Ramesh K Sitaraman. 2024. Untangling Carbon-free Energy Attribution and Carbon Intensity Estimation for Carbon-aware Computing. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 580–588.
[42] Diptyaroop Maji, Ben Pfaff, Vipin PR, Rajagopal Sreenivasan, Victor Firoiu, Sreeram Iyer, Colleen Josephson, Zhelong Pan, and Ramesh K Sitaraman. 2023. Bringing carbon awareness to multi-cloud application delivery. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. 1–6.
[43] Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. 2020. Recalibrating global data center energy-use estimates. *Science* 367, 6481 (2020), 984–986. doi:10.1126/science.aba3758
[44] Eric R. Masanet, Richard E. Brown, Arman Shehabi, Jonathan G. Koomey, and Bruce Nordman. 2011. Estimating the Energy Use and Efficiency Potential of U.S. Data Centers. *Proc. IEEE* 99, 8 (2011), 1440–1453. doi:10.1109/JPROC.2011.2155610

[45] Jorge Murillo, Walid A. Hanafy, David Irwin, Ramesh Sitaraman, and Prashant Shenoy. 2024. CDN-Shifter: Leveraging Spatial Workload Shifting to Decarbonize Content Delivery Networks. In *Proceedings of the 2024 ACM Symposium on Cloud Computing* (Redmond, WA, USA) *(SoCC '24)*. Association for Computing Machinery, New York, NY, USA, 505–521. doi:10.1145/3698038.3698516

[46] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer* 55, 7 (2022), 18–28. doi:10.1109/MC.2022.3148714

[47] David Quirk, Tom Davidson, and Roger Schmidt. 2022. Thermal Guidelines for Data Processing Environments, Fifth Edition. *ASHRAE Journal* 64, 7 (2022), 24.

[48] 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. doi:10.1109/TPWRS.2022.3173250

[49] Ian Schneider, Hui Xu, Stephan Benecke, David Patterson, Keguo Huang, Parthasarathy Ranganathan, and Cooper Elsworth. 2025. Life-Cycle Emissions of AI Hardware: A Cradle-To-Grave Approach and Generational Trends. arXiv:2502.01671 [cs.AR] https://arxiv.org/abs/2502.01671

[50] Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. 2016. United states data center energy usage report. (2016).

[51] Arman Shehabi, Sarah J. Smith, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor. 2024. *2024 United States Data Center Energy Usage Report*. Technical Report LBNL-2001637. Lawrence Berkeley National Lab (LBL) - Energy Analysis & Environmental Impacts Division. https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf

[52] Natalie Sherman. 2024. Microsoft chooses infamous nuclear site for AI power. *BBC News* (20 Sept. 2024). https://www.bbc.com/news/articles/cx25v2d7zexo Accessed: 2025-06-23.

[53] Md Abu Bakar Siddik, Arman Shehabi, and Landon Marston. 2021. The environmental footprint of data centers in the United States. *Environmental Research Letters* 16, 6 (may 2021), 064017. doi:10.1088/1748-9326/abfba1

[54] Zachary Skidmore. 2025. *Natural gas plant planned for Stargate AI data center campus*. https://www.datacenterdynamics.com/en/news/natural-gas-plant-planned-for-stargate-ai-data-center-campus-report/ Accessed: 2025-05-05.

[55] 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 *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2* (Vancouver, BC, Canada) *(ASPLOS 2023)*. 252–265. doi:10.1145/3575693.3575709

[56] Thanathorn Sukprasert, Abel Souza, Noman Bashir, David Irwin, and Prashant Shenoy. 2024. On the Limitations of Carbon-Aware Temporal and Spatial Workload Shifting in the Cloud. In *Proceedings of the Nineteenth European Conference on Computer Systems* (Athens, Greece) *(EuroSys '24)*. 924–941. doi:10.1145/3627703.3650079

[57] Timothy Gardner. 2024. Meta seeks nuclear power developers for reactors to start in early 2030s. *Reuters* (2024). https://www.reuters.com/business/energy/meta-seeks-nuclear-power-developers-reactors-start-early-2030s-2024-12-03/ Accessed: 2025-06-23.

[58] United Nations Environment Programme. 2024. Emissions Gap Report 2024. https://wedocs.unep.org/20.500.11822/46404

[59] United States Environmental Protection Agency. 2024. Inventory of U.S. Greenhouse Gas Emissions and Sinks. https://www.epa.gov/ghgemissions/inventory-us-greenhouse-gas-emissions-and-sinks. Accessed: 2025-05-03.

[60] Carole-Jean Wu, Bilge Acun, Ramya Raghavendra, and Kim Hazelwood. 2024. Beyond Efficiency: Scaling AI Sustainably. *IEEE Micro* 44, 5 (2024), 37–46. doi:10.1109/MM.2024.3409275

[61] Zeke Hausfather et al. 2025. Analysis: Global CO2 emissions will reach new high in 2024 despite slower growth. Retrieved May 9, 2025 from https://www.carbonbrief.org/analysis-global-co2-emissions-will-reach-new-high-in-2024-despite-slower-growth/

[62] Yingbo Zhang, Hangxin Li, and Shengwei Wang. 2023. The global energy impact of raising the space temperature for high-temperature data centers. *Cell Reports Physical Science* 4, 10 (2023), 101624. doi:10.1016/j.xcrp.2023.101624

[63] Jiajia Zheng, Andrew A. Chien, and Sangwon Suh. 2020. Mitigating Curtailment and Carbon Emissions through Load Migration between Data Centers. *Joule* 4, 10 (2020), 2208–2222. doi:10.1016/j.joule.2020.08.001