

Balancing Bits and Drops: Stress-Adjusted Water Management for Data Centers

Zahidur Talukder*

Department of Math and Computer
Science
Texas Lutheran University
Seguin, Texas, USA
ztalukder@tlu.edu

Imtiaz Bin Rahim

Department of Computer Science and
Engineering
University of Texas at Arlington
Arlington, Texas, USA
ixb6394@mavs.uta.edu

Pranjol Sen Gupta

Department of Computer Science
Kennesaw State University
Kennesaw, Georgia, USA
pgupta10@kennesaw.edu

Shaolei Ren

Department of Electrical and
Computer Engineering
University of California, Riverside
Riverside, CA, USA
shaolei@ucr.edu

Mohammad A. Islam

Department of Computer Science and
Engineering
University of Texas at Arlington
Arlington, TX, USA
mislam@uta.edu

Abstract

Data centers are critical to today's digital economy, but are also among the largest industrial consumers of freshwater. Beyond the sheer volume of water use, the environmental impact of data center water consumption varies significantly across locations and seasons, depending on local and regional water stress. However, prior research has largely focused on reducing total water use, overlooking that the same unit of water can have drastically different environmental consequences depending on when and where it is consumed. In this paper, we introduce a *stress-adjusted water* framework that quantifies the true sustainability impact of data center water consumption by incorporating both spatial and temporal water stress. Using the AWARE-US model, we capture county-level monthly variations in water availability and extend this framework to account for the off-site water footprint of electricity generation. Based on this stress-aware accounting, we analyze stress-adjusted water-computing strategies spanning both the software and infrastructure layers. Specifically, we study workload scheduling policies that jointly optimize water and carbon efficiency, evaluate the potential of rainwater harvesting as a supplemental water source, and investigate the feasibility of dry cooling as a water-free alternative to evaporative cooling. Our evaluation across major U.S. data center markets shows that stress-aware workload scheduling can reduce stress-adjusted water consumption by up to 25% while preserving performance and balancing carbon emissions. We further show that rainwater harvesting can offset up to 100% of on-site cooling water in regions with sufficient precipitation, while offering diminishing returns in arid locations. Finally, our analysis of dry cooling reveals that its effectiveness depends critically on energy efficiency and off-site water stress, highlighting important trade-offs between water

savings, energy use, and carbon emissions. Together, these results demonstrate that incorporating water stress as a first-class sustainability metric enables more informed and effective data center design and operation.

CCS Concepts

• **Hardware** → **Impact on the environment**; • **General and reference** → *Empirical studies*.

Keywords

Data centers, water sustainability, water stress, workload scheduling, rainwater harvesting, dry cooling

ACM Reference Format:

Zahidur Talukder, Imtiaz Bin Rahim, Pranjol Sen Gupta, Shaolei Ren, and Mohammad A. Islam. 2026. Balancing Bits and Drops: Stress-Adjusted Water Management for Data Centers. In *The 17th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '26)*, June 22–25, 2026, Banff, AB, Canada. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3744255.3811726>

1 Introduction

Water is one of the most abundant natural resources on Earth, yet access to freshwater remains limited and unevenly distributed. The global water crisis is further intensified by climate change, rapid population growth, and aging infrastructure. The United Nations (UN) highlights water scarcity as one of the most pressing consequences of climate change, underscoring the urgency of global action [39]. Since water is a shared societal resource, industries across all sectors must actively contribute to sustainability efforts. Among them, data centers—given their immense scale and critical role in the digital economy—have both the opportunity and the responsibility to lead by example in advancing sustainable water management [36].

Water consumption in data centers. Data centers are widely recognized for their massive energy consumption, but they are also huge water guzzlers [16, 33]. Due to their high server cooling loads, most large-scale data centers rely on highly efficient cooling tower systems that use water evaporation to expel heat. However, this

*Part of this work was done at the University of Texas at Arlington.



This work is licensed under a Creative Commons Attribution 4.0 International License. *E-Energy '26, Banff, AB, Canada*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2011-6/26/06

<https://doi.org/10.1145/3744255.3811726>

evaporation leads to continuous water loss, requiring constant replenishment [44]. In 2023 alone, Google’s self-operated data centers withdrew approximately 29 billion liters of water for on-site cooling, with over 23 billion liters evaporated—nearly 80% of which was potable water [48]. As the demand for digital services surges, driven largely by advances in artificial intelligence (AI) and machine learning (ML), data center water usage is rising at an unprecedented rate. From 2021 to 2022, Google’s water consumption increased by around 20%, followed by another 17% increase from 2022 to 2023. Similarly, Microsoft reported a 34% increase between 2021 and 2022, and a further 22% rise from 2022 to 2023 [8]. The U.S. Department of Energy projects that, by 2028, total annual on-site water consumption by U.S. data centers could double or even quadruple from 2023 levels [40]. Beyond direct water use for cooling, data centers are also responsible for the water consumption associated with their grid electricity use [27, 45]. Electricity generation from nuclear and thermal energy sources also requires water in their cooling systems. Excluding hydroelectricity, which itself is a major water consumer, the national average water consumption by electricity power plants in the U.S. is as high as 1.8 L/kWh [20].

Limitations of prior works. Over the past decade, data centers have made substantial progress in improving energy efficiency and reducing carbon emissions through carbon-aware scheduling, placement, and infrastructure design [22, 32, 46]. In contrast, data center water sustainability has largely been studied through the lens of *total water consumption*, often mirroring carbon-centric optimization approaches [26–28]. While these efforts provide valuable insights into improving water efficiency, they overlook a fundamental distinction between the environmental impacts of water consumption and carbon emissions. Carbon emissions contribute to climate change at a global scale, and their marginal impact is largely independent of where or when they occur [41]. In contrast, the environmental impact of water consumption is inherently *local and time dependent* [30, 31, 37]. Consuming one liter of water in a water-stressed region during a dry season can have orders-of-magnitude greater ecological and societal consequences than consuming the same amount of water in a water-abundant region or during a wet season. As a result, treating all water consumption as environmentally equivalent fundamentally misrepresents its true impact.

Recent work has begun to move beyond purely volumetric accounting by explicitly valuing water based on its scarcity. In a recent presentation, the Open Compute Project (OCP) introduced an updated water-efficiency metric that adjusts for regional water scarcity, expressed as the ratio of water withdrawal to water availability [50]. Wu et al. proposed SCARF, which introduces an Adjusted Water Impact (AWI) metric to reflect regional water scarcity when evaluating computing workloads [52]. SCARF demonstrates that treating water as location-invariant can substantially alter sustainability conclusions. However, it applies regional adjustment at an aggregate level and does not explicitly distinguish between *on-site* and *off-site* water pathways or model the provenance of electricity generation. More closely related in spirit, Jiang et al. proposed *WaterWise*, a scheduler that incorporates water stress into workload scheduling and shows that optimizing for water and carbon independently can yield conflicting outcomes [29]. *WaterWise* models off-site water impact using region-level abstractions

and does not explicitly account for the geographic locations of water-consuming power plants, which can misestimate the water stress borne by electricity generation—especially when data centers draw power from stressed generation sites located far from the data center itself.

Our key insight and contributions. This paper builds on and extends these efforts by explicitly modeling the *provenance* of both on-site and off-site water consumption. Our key insight is that sustainable data center operation must reason about the *value of water*¹, not merely its volume. Specifically, the environmental cost of water consumption depends on *where* and *when* the water is consumed, as well as whether it occurs directly at the data center or indirectly through electricity generation.

To capture these effects, we leverage the **Available Water RE** maining for the **United States (AWARE-US)** model, which provides county-level monthly water-stress characterization across the U.S. [31]. Using AWARE-US, we introduce the notion of *stress-adjusted water consumption*, which weights water usage by local water scarcity to reflect its true environmental impact. We apply this metric to both direct (on-site) water consumption and indirect (off-site) water consumption embedded in electricity use.

While AWARE-US provides fine-grained characterization of *local* water stress, it does not capture the off-site water impacts embedded in electricity generation. To address this limitation, we explicitly trace data center electricity consumption to upstream power plants within the U.S. eGRID sub-regions [12] and estimate off-site water stress by weighting plant-level water stress according to electricity generation. This enables a more accurate estimation of off-site water stress than region-level or national-average abstractions. In practice, water intensity varies substantially across locations even for the same generation technology. As shown in Table 2, relying on fixed national water-intensity factors can lead to substantial misestimation of off-site water use. Because stress-adjusted water combines both volumetric consumption and local water stress, such inaccuracies directly propagate into errors in the overall water-impact assessment. Notably, minimizing volumetric water can misrank decisions under geographic flexibility and may even increase stress-adjusted water by shifting demand toward a more water-stressed electricity supply.

Using our framework, we analyze more than 4,000 U.S. data center locations obtained from datacentermap.com [10]. Our analysis reveals that nearly 75% of U.S. data centers experience at least one month of elevated water stress annually, and that off-site water stress can dominate total water impact in many major markets.

Stress-adjusted water computing strategies. Building on this framework, we investigate three complementary approaches for reducing stress-adjusted water impact in data center operations, spanning software-level optimizations and infrastructure-level interventions.

Workload scheduling. We study how spatial and temporal workload flexibility can be leveraged to reduce stress-adjusted water consumption by preferentially executing workloads at locations and times with lower water stress. Our analysis considers multiple

¹We interpret the value of water based on availability; its value rises with scarcity and falls with abundance.

workload classes with varying degrees of flexibility and jointly optimizes for water and carbon objectives. We show that stress-aware scheduling can reduce stress-adjusted water consumption by up to 25% while simultaneously reducing carbon emissions by up to 25%.

Rainwater harvesting. We evaluate rainwater harvesting as a location-dependent mitigation strategy that reduces reliance on stressed freshwater sources. Using year-long precipitation data across major data center markets, we analyze the feasibility of rainwater harvesting under realistic storage constraints. Our results demonstrate that, in suitable regions, rainwater harvesting can offset up to 100% of on-site cooling water demand.

Dry cooling. Finally, we examine the water-energy trade-offs introduced by dry-cooling technologies, which eliminate on-site water use but incur higher electricity consumption [44]. While dry cooling reduces direct water withdrawal, the resulting increase in electricity demand shifts water consumption off-site to power plants, where water stress may be significantly higher. We quantify this trade-off across major U.S. data center markets and show that dry cooling is beneficial primarily in highly water-stressed regions when the associated energy efficiency penalty remains sufficiently low.

Dataset. We construct a dataset comprising 4,147 U.S. data center locations annotated with on-site and off-site water stress metrics, electricity provenance, regional precipitation profiles, cooling efficiencies, and electricity-generation water efficiencies. To support transparency and future research, our dataset is made publicly available on the Open Science Framework (OSF) [47].

Limitations. Our study has several limitations that point to future research directions. We rely on AWARE-US for water stress characterization; alternative hydrological models may yield different absolute stress values. Our off-site analysis uses publicly available EPA power plant data, which may omit small or recently commissioned plants. Rainwater harvesting is evaluated using precipitation data from a single year (2023), and inter-annual climate variability may affect long-term outcomes. Our workload scheduling analysis provides an offline upper bound rather than an online scheduling algorithm. Finally, data center location data represents a snapshot in time and may not capture newly deployed facilities.

Looking forward, we believe that combining stress-aware water accounting with demographic and equity analyses—such as assessing the co-location of data centers with vulnerable or underserved communities—represents an important and largely unexplored direction for sustainable computing research.

2 Preliminaries

2.1 Water Usage

The data center community often discusses “water usage” without carefully distinguishing among *withdrawal*, *discharge*, and *consumption*. These terms are not interchangeable and lead to very different sustainability conclusions [38]. *Water withdrawal* refers to the total volume of water taken from a source (e.g., a municipal utility, river intake, or groundwater well). *Water discharge* is the portion returned to the environment after use (e.g., cooling-water blowdown or once-through return flow). *Water consumption* is the fraction of withdrawn water that is not returned to the original watershed, typically because it evaporates, is incorporated into products, or




Figure 1: Water usage in data center heat rejection.

is transported to another watershed [38]. Thus, a facility’s water consumption can be expressed as:

$$\text{consumption} = \text{withdrawal} - \text{discharge}. \quad (1)$$

Fig. 1 illustrates these concepts for cooling-tower-based data center operation. In evaporative cooling, a portion of circulating water evaporates to reject heat, which directly contributes to *consumption*. Another portion is periodically discharged (“blowdown”) to control mineral concentration, which contributes to *discharge*. Cooling systems typically cannot achieve both low withdrawal and zero consumption simultaneously. For example, *once-through* cooling can achieve near-zero consumption by returning most of the withdrawn water, but it requires far larger withdrawal volumes to remove the same heat load [30]. These trade-offs matter because they affect different stakeholders: withdrawal stresses water *supply capacity* and infrastructure (treatment, pumping, distribution), while consumption reduces *availability* for competing societal and ecological demands.

Water usage is often categorized by color to distinguish its source, quality, and suitability for various uses. *Blue Water* refers to surface and groundwater (e.g., lakes, rivers, and aquifers) that are withdrawn for human use, such as cooling data centers and industrial processes. Once consumed (e.g., through evaporation or incorporation into products), it is no longer available in the original source [24]. *Green Water* represents rainwater and soil moisture that is naturally stored in the unsaturated zone of soil and used by plants for growth. It is crucial for agriculture and forestry, but is generally not part of direct industrial or urban water consumption [14]. *Grey Water* is wastewater from domestic, commercial, or industrial activities that has been used but is not heavily contaminated (e.g., from sinks, showers, or cooling processes). It can be treated and reused for non-potable purposes such as irrigation or industrial cooling [53].

2.2 Water Consumption in Data Centers

To systematically reason about water impacts, we adapt the well-known sustainability scope framework from the Greenhouse Gas Protocol [19]. Although originally designed for carbon accounting, the scope decomposition is also useful for water because it separates operational impacts from supply-chain impacts and clarifies which levers are controllable during operation.

Scope 1 (direct on-site water consumption). This includes water consumed within the data center boundary, primarily for cooling (e.g., evaporation in cooling towers), humidification, and other facility processes. In Fig. 1, evaporated cooling water constitutes the dominant Scope 1 component in many large-scale facilities.

Scope 2 (indirect off-site water consumption from electricity). This includes water consumed upstream by electricity generation needed to power IT and cooling loads. Thermal and nuclear power plants often consume water through evaporative cooling, and this consumption occurs at the plants rather than at the data center. As a result, the off-site water footprint can be geographically decoupled from the facility and may be subject to very different water scarcity conditions.

Scope 3 (other indirect water consumption). This includes embodied water from manufacturing IT equipment, constructing facilities, and broader supply chain processes. While Scope 3 can be significant, it is difficult to track at high resolution and harder to mitigate through operational decisions. Accordingly, we focus on **operational water consumption** (Scopes 1 and 2) that can be influenced by water-aware computing and infrastructure choices.

2.3 Water Usage Effectiveness (WUE)

WUE is a widely used metric proposed by The Green Grid to quantify water efficiency [5]. For on-site operations, WUE captures facility water consumption per unit of IT energy (L/kWh), making it suitable for comparing cooling and facility designs. For off-site electricity, the analogous concept is sometimes referred to as the Electricity Water Intensity Factor (EWIF) [5, 28]. Because grid electricity comes from a portfolio of generation sources with different water intensities, off-site WUE is often computed as a generation-weighted average:

$$WUE_{offsite} = \frac{\sum_i G_i \cdot WUE_i}{\sum_i G_i}, \quad (2)$$

where G_i is the electricity generated by source i and WUE_i is its water intensity. This formulation is attractive because hourly generation data by source is often publicly available [21]. However, the approach is only as accurate as the assumed WUE_i values. A common simplification is to treat fuel-specific WUE values as fixed across large geographies (e.g., a single U.S.-wide WUE for natural gas), which, as shown in Table 2, can introduce regional errors.

OCP has recently introduced an updated WUE metric, WUE^+ , that adds two multiplicative factors to the original WUE to capture regional water availability and reuse efficiency [50]. It is defined as follows

$$WUE^+ = WUE \times \text{Regional Adjustment} \times \text{Reuse Efficiency}, \quad (3)$$

where regional adjustment is defined as the ratio of total water withdrawal and available renewable water, and reuse efficiency is defined as the ratio of potable water used and total water input.

2.4 Variation in WUE

On-site and off-site water efficiency varies across geography and time. On-site WUE depends on local weather (temperature and humidity), cooling tower characteristics, and facility design [21]. For the same IT load, warmer or more humid conditions can lead to higher evaporation and, in turn, higher on-site water consumption. Off-site WUE depends on the regional generation mix, power plant cooling technologies, and local operating conditions.

Fig. 2 illustrates spatial variation in average on-site and off-site WUE across the U.S. We observe substantial regional diversity, with some states exhibiting markedly higher off-site WUE due to reliance




Figure 2: Spatial variation in on-site and off-site water efficiency across U.S. states.

on water-intensive generation technologies. This heterogeneity motivates software strategies (e.g., workload shifting) that exploit spatiotemporal differences in WUE.

While WUE is useful, it measures *liters per kWh*, not the environmental cost of consuming those liters. Two regions can have the same WUE but very different water scarcity conditions, meaning equal volumes of water consumption may have dramatically different societal and ecological consequences. This motivates stress-aware accounting, which we introduce next.

3 Stress-Adjusted Water Footprint

3.1 Water versus Carbon

A major reason water has been historically underemphasized in computing sustainability is the community's success with carbon-aware optimization. Carbon emissions are an appropriate global sustainability metric because greenhouse gases mix in the atmosphere and persist over long time horizons; therefore, marginal damages are often treated as largely independent of where and when emissions occur [41]. This justifies objectives such as minimizing total operational emissions or shifting workloads to regions with low-carbon electricity.

Water consumption is fundamentally different. Water is locally sourced, locally constrained, and shared among households, agriculture, ecosystems, and industry. Thus, the environmental impact of consuming one liter of freshwater depends strongly on local scarcity [30, 31, 37]. Moreover, water availability is time-varying: seasonal patterns and drought conditions can sharply alter scarcity and competition for water. As a result, minimizing total water volume or maximizing WUE is insufficient to assess the sustainability impact.

3.2 Water Stress

Water stress refers to the pressure on local water resources resulting from the balance between availability and demand. Many tools quantify water stress at different spatial and temporal scales [1, 2, 31]. We use the AWARE-US model because it provides county-level monthly characterization factors across the U.S. and has been adopted in prior assessments of electricity-related water impacts [31].

AWARE-US defines a unit-less characterization factor (CF) that converts each liter of water consumption into a *stress-adjusted* quantity. CF is derived from availability-minus-demand (AMD), where AMD is computed from hydrological runoff and social/environmental




Figure 3: County-level water stress across the United States based on the AWARE-US model, illustrating strong spatial heterogeneity in water stress [31].




Figure 4: Monthly water stress across major U.S. data center markets, illustrating both spatial variation across locations and temporal variation throughout the year.

water demands. The CF of county i in month j is:

$$CF_{i,j} = \frac{AMD_{US}}{AMD_{i,j}}, \quad (4)$$

where AMD_{US} is the U.S. reference value [31]. CF is capped between 0.1 and 100 to avoid extreme values dominating assessments. In our interpretation, $CF > 1$ indicates above-average water stress.

Spatiotemporal variation. Water stress exhibits pronounced heterogeneity across both geography and time due to seasonal hydrology and regional demand patterns. Fig. 3 illustrates county-level water stress during the summer, revealing strong spatial disparities in water availability across the United States at a single point in time. Fig. 4 further highlights this variability by showing monthly water stress across the top 20 U.S. data center markets, where stress levels fluctuate substantially throughout the year and differ markedly between locations. Together, these results demonstrate that water stress is inherently spatiotemporal, implying that the environmental impact or *value* of a unit of water consumption can vary significantly depending on when and where it occurs. This variability motivates the need for stress-adjusted water accounting rather than volume-based metrics.




Figure 5: Off-site water stress across EPA eGRID sub-regions, computed using generation-weighted AWARE-US characterization factors.

3.3 Stress-Adjusted Water Consumption

We define **stress-adjusted water** as water consumption weighted by local stress conditions.

On-site stress-adjusted water. For a data center located in county c and month m , on-site stress-adjusted water is computed as:

$$W_{SA}^{on}(m) = W^{on}(m) \cdot CF_{c,m}, \quad (5)$$

where $W^{on}(m)$ is the on-site water consumption in month m .

Off-site stress-adjusted water. Off-site water consumption occurs at power plants due to electricity generation and is therefore governed by the water stress at plant locations rather than at the data center site. Estimating off-site stress-adjusted water thus requires mapping a data center's electricity consumption to the upstream generation portfolio supplying its electricity.

3.4 Estimating Off-Site Water Stress

Tracing the exact power flow from generators to a specific data center is challenging due to the physics of power networks, dispatch dynamics, and market operations. Instead, we approximate electricity provenance using the EPA eGRID framework, which partitions the U.S. grid into 27 sub-regions [12, 49]. Within an eGRID sub-region, electricity is shared among consumers, and therefore, the regional generation portfolio provides a practical approximation for attributing off-site impacts.

We compute a region-level off-site stress factor by taking a generation-weighted average of county-level CF values of water-consuming power plants in that region:

$$\overline{CF}_r = \frac{\sum_i G_i \cdot CF_i}{\sum_i G_i}, \quad (6)$$

where G_i is the annual net generation of plant i and CF_i is the CF of the county containing that plant. We align monthly AWARE-US stress factors with monthly averaged generation profiles; annual generation shares are used only to approximate spatial provenance. Fig. 5 visualizes off-site stress across eGRID regions based on this metric.

3.5 U.S. Data Centers' Exposure to Water Stress


Many major U.S. data center markets (e.g., California, Arizona, Texas) operate in regions experiencing substantial seasonal water scarcity. To quantify the prevalence of water stress across the U.S. data center market, we analyze 4,147 data center locations from

Table 1: Water efficiency in L/kWh of different fuel sources across the U.S. eGRID regions. “-” indicates no electricity generation from the corresponding source. “0” indicates no water consumption (e.g., once-through cooling) during electricity generation.

Primary Fuel	AZNM	CAMX	ERCT	FRCC	HIOA	MROE	MROW	NEWI	NWPP	NYCW	NYLI	NYUP	RFCE	RFCM	RFCW	RMPA	SPNO	SPSO	SRMV	SRMW	SRSO	SRTV	SRVC
COAL	1.82	1.35	1.68	1.42	2.04	-	1.23	0	2.19	-	-	-	2.21	1.12	1.35	1.83	1.44	2.37	1.27	2.42	2.24	1.57	1.29
NATURAL GAS	2.47	1.25	2.4	2.58	-	0	1.63	1.01	1.74	0	0	0.69	1.71	2.21	1.56	1.46	2.13	3.23	1.86	1.78	2.48	1.38	1.67
NUCLEAR	2.83	0	0.88	0	-	-	1.95	0	2.67	-	-	-	1.29	3.14	1.33	-	2.16	-	1.86	2.01	2.72	0.21	0.52
OTHER	6.49	2.59	1.82	4.8	-	0	-	3.87	2.13	-	-	-	17.2	0	49.9	1.77	-	1.35	0.71	-	0.64	3.35	1.26
PETROLEUM	-	-	-	-	0	-	-	0.23	-	-	-	0	0.41	-	0	-	-	0	1.04	-	-	-	0
SOLAR	-	0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 2: Error introduced by using fixed national water-intensity factors instead of region-specific values. Min/Max are computed across eGRID sub-regions using Table 1.

Fuel	Fixed WUE [43] (L/kWh)	Minimum Regional (L/kWh)	Maximum Regional (L/kWh)	Maximum Overestimation (%)	Maximum Underestimation (%)	Regional Variability (×)
Coal	1.82	1.12	2.42	+63%	-25%	2.2×
Natural Gas	0.80	0.69	3.23	+15%	-75%	4.7×
Nuclear	2.31	0.21	3.14	+999%	-26%	15.0×
Petroleum	1.36	0.23	1.04	+492%	+31%	4.5×
Other	0.76	0.64	49.9	+18%	-98%	78.0×

**Figure 6: Percentage of U.S. data centers experiencing at least N months of moderate or extreme on-site water stress per year.**


datacentermap.com. Fig. 6 shows the distribution of data centers across stress levels. Here, we consider only the on-site water stress. Nearly 75% of data centers experience at least one month of water stress annually; approximately 30% experience extreme stress (CF=100) in at least one month; and 12% operate under persistent stress year-round.

3.6 Limitation of On-Site Stress Alone

While on-site CF captures local scarcity at the data center, the indirect footprint from electricity consumption can be influenced by very different water-stress conditions. Fig. 7 compares on-site and off-site water stress for the top U.S. data center markets. We observe substantially higher off-site stress for several markets (e.g., Dallas, Boardman, Las Vegas, Columbus, Tulsa), indicating that a purely local assessment may underestimate the true stress burden of data center operation. Conversely, some markets exhibit lower off-site stress, suggesting potential opportunities for water-aware workload distribution.

3.7 Regional Off-Site Water Efficiency

A second source of error in off-site water accounting arises from assuming fixed water intensities for generation sources across large regions. In practice, even within the same fuel category (e.g., natural gas), water consumption per kWh varies across regions due to differences in cooling technology, plant design, local conditions,

**Figure 7: Comparison of on-site and off-site water stress across major U.S. data center markets.**

and operational practices. To more accurately estimate off-site water consumption, we derive region-specific WUE values for each eGRID sub-region using 2023 power-plant-level water consumption data [11]. Table 1 summarizes the resulting water efficiencies by generation source and region.

Table 2 quantifies the error introduced by using fixed national water-intensity factors, as is common in prior work. Across all fuel types, regional water intensity varies substantially, leading to systematic misestimation of off-site water consumption. For example, natural gas generation exhibits a 4.7× regional spread, leading to fixed values that underestimate water use by up to 75% in some regions. For nuclear and “other” fuels, the error exceeds an order of magnitude. These results motivate our use of region-specific water efficiencies when estimating off-site, stress-adjusted water consumption.

4 Stress-Adjusted Water Management

Having established *stress-adjusted water* as a metric that captures the *value* of water consumption—accounting for where and when water is consumed and explicitly separating on-site and off-site impacts—we now examine how data centers can reduce their stress-adjusted water footprint in practice. Rather than proposing a single mechanism, we analyze three complementary approaches that

Table 3: Qualitative impact of *stress-adjusted water* techniques on stress-adjusted water and carbon. Arrows indicate direction and relative magnitude.

Technique	On-site Water	Off-site Water	Carbon
Workload Scheduling	↓	↓	↓
Rainwater Harvesting	↓↓	–	–
Dry Cooling	↓↓	↑	↑

↓↓: strong reduction, ↓: moderate reduction, ↑: moderate increase, –: negligible change.

operate at different layers of the data center stack and involve distinct implementation costs and time horizons: (i) **software-driven workload scheduling** that shifts computation across time and location, (ii) **rainwater harvesting** that supplements on-site cooling demand with locally available precipitation, and (iii) **dry cooling** that eliminates on-site water use by replacing evaporative cooling with air-based heat rejection.

Crucially, these approaches affect different components of the data center footprint. Workload scheduling influences both *on-site* (Scope 1) and *off-site* (Scope 2) water consumption by changing when and where electricity is consumed, and cooling is performed. Rainwater harvesting directly offsets *on-site* (Scope 1) withdrawals without changing electricity demand. In contrast, dry cooling eliminates on-site water use but increases electricity consumption, thereby shifting water impact from *on-site* (Scope 1) to *off-site* (Scope 2) through upstream power generation. Table 3 summarizes the overall contribution of three different techniques on stress-adjusted water and carbon.


4.1 Workload Scheduling

4.1.1 Motivation: exploiting spatiotemporal heterogeneity. On-site and off-site water impacts vary significantly across locations and over time due to weather-driven cooling dynamics, differences in power-grid composition, and geographic heterogeneity in water stress. This spatiotemporal heterogeneity creates opportunities to reduce stress-adjusted water by shifting workloads toward *low-impact* time–location pairs. Modern cloud platforms already offer non-trivial scheduling flexibility, and prior work commonly categorizes workloads by their spatial and temporal freedom [46]. We adopt three representative classes:

Spatially flexible workloads. These workloads must be served immediately (tight response-time SLOs) but may be routed to one of several geographically distributed data centers. Examples include web services, interactive applications, and ML inference. Execution times typically range from milliseconds to seconds.

Temporally flexible workloads. These workloads are delay-tolerant: they have a completion deadline (e.g., hours) and can be shifted in time within that window. Examples include batch analytics, offline log processing, checkpointable HPC jobs, and many ML training pipelines.

Spatio-temporally flexible workloads. These workloads can be shifted both in time and across locations, e.g., distributed training or migratable HPC jobs that can restart at alternative sites. This class offers the largest optimization potential because it can exploit both temporal and geographic heterogeneity.

**Figure 8: Hourly variation in water and carbon efficiencies (Ashburn, VA), illustrating weak correlation.**

4.1.2 Scheduling overhead and performance assurance. Scheduling flexibility is not free. For spatial routing, sending an interactive request to a distant site increases network round-trip latency; satisfying the same SLO may therefore require additional provisioning or higher power. For temporal or migratory jobs, overheads may arise from data transfer, checkpointing, and restart costs. In our evaluation, we conservatively assume that **no performance degradation is permitted for sustainability**: any routing or deferral overhead is absorbed as additional power, ensuring SLO compliance. Operationally, this converts scheduling into a **power-allocation** problem: for each workload and SLO, the required power demand is determined, and the scheduler decides *when* and *where* to consume that power.

4.1.3 Jointly optimizing water and carbon. Although this paper focuses on water, carbon emissions remain a core sustainability objective. Importantly, water and carbon are not necessarily aligned. On-site water depends on weather and cooling dynamics, while off-site water depends on generation water intensity and water stress, which can be high even for low-carbon electricity sources. Fig. 8 illustrates that hourly water and carbon efficiencies can vary independently. Therefore, optimizing only water or only carbon can unintentionally worsen the other, motivating a joint objective that accounts for both carbon and stress-adjusted water.

4.1.4 Stress-adjusted water scheduling model. We model workload scheduling as a time-slotted power-allocation problem over a horizon of T slots. At each slot, workload arrives at one or more gateways and must be assigned to a data center, either immediately or within a limited deferral window (for delay-tolerant workloads). Each assignment decision induces IT power consumption at a specific location and time; the resulting total facility power then determines (i) carbon emissions, (ii) on-site cooling water consumption (weighted by local water stress), and (iii) off-site water consumption from electricity generation (weighted by the water stress at power-plant locations).

We jointly minimize total carbon and total stress-adjusted water over the horizon. A tunable weight parameter controls the emphasis between water and carbon, allowing exploration of the trade-off space. All workloads must be scheduled, and each data center must respect a maximum power capacity. Full formal definitions and the complete optimization formulation are provided in Appendix A.

Offline upper bound. Optimally scheduling workloads online requires forecasting future carbon intensity, water efficiency, and water stress signals. Rather than proposing a new online algorithm, our goal is to quantify the *maximum achievable benefit* of workload




Figure 9: Savings from workload scheduling under different benchmark strategies, illustrating that minimizing volumetric water can misestimate stress-adjusted water savings—particularly for spatial and spatio-temporal workloads.




Figure 10: Impact of temporal delay tolerance on sustainability savings under temporal and spatio-temporal workload.




Figure 11: Impact of the number of available data centers.

shifting under idealized conditions. Accordingly, we compute an offline optimal schedule assuming perfect future knowledge. This offline solution serves as an upper bound on potential savings and isolates the value of spatiotemporal flexibility independent of forecasting or control errors.

4.1.5 Evaluation setup. We evaluate stress-adjusted water scheduling using long-term trace-driven simulations. We consider the top five U.S. data center markets by concentration and extend the analysis to the top twenty. We use hourly on-site/off-site water efficiencies and carbon efficiencies from [21]. We use county-level monthly water stress factors from AWARE-US [31], and off-site stress factors derived from our eGRID-based methodology (Fig. 5). For workload, we use Google search workload traces from 2023 [18] and scale them to represent a 10 MW data center. We emulate temporal, spatial, and spatio-temporal flexibility by varying the deferral window and the set of reachable markets. Geographic distance between markets is used to model routing overhead for interactive workloads.

Baselines. We compare our joint optimization WWS (Water Wise Scheduling) that minimizes stress-adjusted water plus carbon against: (i) **NoScheduling** that processes workload immediately in the local data center, (ii) **CARBON** that minimizes only the carbon emission, (iii) **WATER** that minimizes only the water *volume*, and (iv) **STRESS** that minimizes only the stress-adjusted water.

4.1.6 Results. Overall savings and misestimation under different workload flexibilities. Fig. 9 summarizes savings under temporal, spatial, and spatio-temporal workloads across different optimization objectives. Temporal flexibility primarily exploits intra-market variation, and consequently volumetric water optimization provides a reasonable—but still incomplete—approximation of stress-adjusted outcomes.

In contrast, for *spatial* and *spatio-temporal* workloads, optimizing for water volume alone can substantially misrepresent true water impact. As shown in Fig. 9(b) and (c), strategies that minimize volumetric water often achieve modest water savings while simultaneously increasing stress-adjusted water. This occurs because spatial routing shifts workloads toward regions with lower cooling water usage but significantly higher water stress or stressed electricity supply, amplifying off-site stress-adjusted water consumption.

By jointly accounting for on-site and off-site water stress, WWS avoids these failure modes. Across all workload types, it consistently achieves large reductions in stress-adjusted water while also reducing carbon emissions, avoiding the “optimize-one-harm-the-other” behavior exhibited by single-objective baselines. These results demonstrate that volumetric water savings alone are insufficient to assess sustainability, particularly when workloads are geographically flexible.

Impact of delay tolerance. We vary the delay tolerance from 1 to 24 hours (Fig. 10). For temporally flexible workloads, longer deadlines yield higher savings by expanding the set of feasible low-impact hours. For spatio-temporal workloads, even a small delay tolerance provides strong benefits because the scheduler can combine modest temporal shifts with geographic diversity. Overall, spatio-temporal flexibility provides substantial gains even at low




Figure 12: Impact of data center capacity headroom.

delay tolerance, while purely temporal flexibility benefits most from long deadlines.

Impact of the number of data centers. We vary the number of available markets from 1 to 20. Savings increase rapidly as additional markets introduce greater spatiotemporal heterogeneity, then exhibit diminishing returns as marginal diversity decreases.

Impact of capacity headroom. We scale data center capacity up to 20 $\times$ the baseline. Additional capacity enables workloads to be concentrated into lower-impact times and locations, improving both carbon and stress-adjusted water outcomes. However, benefits saturate beyond moderate scaling, indicating that heterogeneity and flexibility—rather than unlimited capacity—are the primary drivers of sustainability gains.

4.2 Rainwater Harvesting

Rainwater harvesting provides a complementary lever to offset on-site withdrawals by substituting municipal supply with locally captured precipitation. Beyond offsetting freshwater demand, rainwater harvesting can mitigate stormwater runoff pollution and alleviate pressure on sewage infrastructure [7]. Several U.S. states encourage rainwater harvesting through incentives and rebates, particularly in drought-prone regions [6]. In data centers, harvesting can also support sustainability certifications, reduce exposure to water price volatility, and improve resilience under drought restrictions [9]. Rainwater typically requires minimal treatment (e.g., filtration) for cooling use; moreover, its lower mineral content can improve cycles of concentration, increasing the fraction that can be consumed via evaporative cooling [15, 42]. Major operators have deployed harvesting systems in practice [13, 17]. Fig. 13 shows the variability in precipitation across the major U.S. data center markets throughout the year.

System model. A harvesting deployment consists of (i) a collection surface (roofs, parking lots, and potentially nearby buildings) and (ii) storage (tanks or retention ponds). Effectiveness depends on rainfall seasonality, storage sizing, and available collection area. Accordingly, we quantify feasibility across major U.S. markets as a function of harvesting area and storage capacity.

4.2.1 Evaluation setup. We run a year-long simulation for major data center markets using hourly precipitation data from Weather Underground [3]. We use the same 10 MW workload scaling as in scheduling, with hourly IT load ranging from 30% to 100% of peak. We assume tanks start half full and adopt a cycle-of-concentration of 10, meaning approximately 90% of the harvested water can be used for evaporative cooling.




Figure 13: Annual precipitation patterns across major U.S. data center markets, illustrating strong spatial and seasonal variability.

We consider harvesting areas from 300,000 to 1,500,000 sqft, representing (i) roof-only collection and (ii) expanded surfaces including parking and adjacent buildings. We consider tank sizes from 500,000 to 3,000,000 gallons. We report **sustainability**, defined as the fraction of total annual cooling demand supplied by harvested rainwater.

4.2.2 Results. Fig. 14 shows water-offset versus harvesting area for fixed tank sizes. Larger harvesting surfaces generally increase water offset, but benefits saturate once rainfall supply exceeds storage or demand. Increasing tank size improves buffering and reduces overflow loss, particularly in markets with strong seasonality.

Fig. 15 summarizes water offset across the top markets for a representative configuration (1,000,000-gallon tank and 1,000,000 sqft harvesting area). Most markets have high potential for water offset through rainwater harvesting, while arid markets such as Phoenix and Denver have lower potential due to limited precipitation. Importantly, these low water-offset markets often coincide with high water stress, increasing the value of each liter saved and reinforcing the need for stress-aware evaluation.

4.3 Dry Cooling

Dry cooling eliminates **on-site** cooling water by rejecting heat through air-based systems (e.g., air-cooled chillers or dry coolers), as illustrated in Fig. 16. Advanced high-efficiency cooling systems often use a hybrid approach, combining chillers with free-air cooling by directly (after particulate filtering) pushing outside air into server rooms when it is cold enough (e.g., $< 27^{\circ}\text{C}$). Dry cooling is attractive in water-scarce regions, but it typically increases energy consumption and, in turn, **off-site** water consumption and carbon emissions from electricity generation. Therefore, dry cooling is best understood as shifting water impact from Scope 1 to Scope 2 rather than eliminating it.

4.3.1 Evaluation setup. We evaluate dry cooling across major U.S. markets using year-long simulations. We use the same workload traces and hourly location-specific on-site/off-site water and carbon efficiencies as before. We vary PUE from 1.2 (baseline, efficient water-based evaporative cooling) to 2.0 to represent an increasing energy penalty with dry cooling. Under dry cooling, on-site




Figure 14: Water-offset (water saving due to rainwater use) versus harvesting area under different tank sizes across major data center markets.




Figure 15: Rainwater offset for a 1,000,000-gallon tank and 1,000,000 sqft harvesting area (top markets), with corresponding stress context.




Figure 16: Illustration of dry cooling in data centers.

water consumption is set to zero, and off-site impacts scale with increased electricity use. This sweep captures a wide range of possible efficiency penalties observed across climates and system designs.

4.3.2 Results and implications. Fig. 17 shows water savings and stress-adjusted water savings as PUE increases. We also look at the increase in carbon (from a baseline PUE of 1.2) due to higher grid-power consumption resulting from the higher PUE. Carbon emissions increase approximately linearly with PUE, reflecting proportional increases in electricity consumption. At low PUE (e.g., 1.3–1.5), dry cooling can yield substantial reductions in total water volume by eliminating on-site evaporation. However, as PUE increases, off-site water and carbon penalties grow, potentially offsetting or negating the water benefit.

Stress-adjusted savings reveal a critical nuance: markets can differ substantially in whether dry cooling reduces or increases




Figure 17: Trade-offs between water and carbon impacts under dry cooling as PUE increases (baseline water cooling PUE = 1.2).

stress-adjusted water because the off-site stress depends on the supplying eGRID region. Consequently, evaluating dry cooling using water volume alone can be misleading.

To understand broader market impact, Fig. 18 fixes PUE at 1.5 and compares water and stress-adjusted water changes across major markets. While many markets show positive volumetric water savings, stress-adjusted results can flip sign in regions where electricity is supplied by highly stressed generation, as increased electricity demand amplifies off-site impacts.

Feasibility guidance. Dry cooling is most attractive when its efficiency penalty is small (e.g., $PUE \leq 1.5$), in which case eliminating on-site evaporation outweighs off-site penalties. At higher PUE regimes (e.g., ≥ 1.8), savings diminish, and carbon penalties become large, making dry cooling less attractive unless paired with grid decarbonization or other mitigating measures. Overall, dry cooling should be deployed with a stress-aware lens: it can be highly beneficial in some high-stress markets but counterproductive in markets where electricity is sourced from water-stressed regions.

When is dry cooling beneficial? Dry cooling is recommended when the reduction in on-site stress-adjusted water exceeds the




Figure 18: Dry cooling impact across major markets at PUE=1.5: volumetric water savings versus stress-adjusted water change.

increase in off-site stress-adjusted water induced by higher electricity demand. In practice, this occurs when local water stress is high, and the supplying grid’s water stress is relatively low.

5 Related Work

Sustainability-aware data center operation. Prior work has extensively studied sustainability-aware data center operation by adapting workload placement, scaling, and resource allocation to environmental signals. Most systems focus on carbon-aware optimization, exploiting temporal variation in grid carbon intensity [23, 51] and geographic diversity across data centers [34, 35], or combining both dimensions [46]. These approaches are effective for carbon, whose impacts are global, but implicitly assume spatially homogeneous environmental effects—an assumption that does not hold for water.

Water metrics and accounting. WUE is the standard metric for reporting data center water efficiency [5], but it captures water volume rather than environmental impact. Recent efforts recognize this limitation: SCARF introduces an AWI metric that weights water use by regional scarcity [52], demonstrating that water-aware accounting can substantially alter sustainability conclusions. However, SCARF applies regional adjustments at an aggregate level and does not distinguish between on-site and off-site water pathways or model electricity-generation provenance.

Water-aware scheduling and trade-offs. Early systems reduce total water consumption by shifting spatial and temporal workloads [25, 27, 28], but optimize volumetric water use without accounting for stress. More recently, WaterWise incorporates water stress into scheduling decisions and shows that optimizing water and carbon independently can lead to conflicting outcomes [29]. However, it models off-site water impacts using region-level abstractions and does not account for the geographic locations of water-consuming power plants, leading to significant misestimation of off-site water stress.

Infrastructure-based techniques. Complementary work studies cooling-system design and water-saving infrastructure, including flexible cooling [16] and rainwater harvesting [4]. These studies highlight strong location dependence but typically evaluate volumetric water savings. Our work differs by evaluating both software- and infrastructure-level techniques using *stress-adjusted water*, explicitly accounting for spatiotemporal water value and on-site/off-site impacts.

6 Conclusion

As data centers continue to grow in scale and importance, their environmental footprint—particularly water consumption—has become an increasingly critical yet under-addressed sustainability challenge. Unlike carbon emissions, whose impacts are largely global, water impacts are inherently local, seasonal, and spatially heterogeneous. This paper advances the notion of *stress-adjusted water*, arguing that sustainable data center operation must account not only for how much water is consumed, but also for *where* and *when* that consumption occurs, and whether it arises on-site or through electricity generation.

We introduce a stress-adjusted water accounting framework based on the AWARE-US model that jointly captures on-site and off-site water stress and explicitly models the provenance of electricity consumption. Our analysis reveals substantial regional variation in electricity-water intensity, showing that commonly used fixed national factors can significantly misestimate off-site and stress-adjusted water impacts. Using this framework, we demonstrate that software-based workload scheduling can reduce stress-adjusted water while simultaneously lowering carbon emissions, particularly when spatiotemporal flexibility is available.

Beyond scheduling, we evaluate complementary infrastructure-level interventions. We show that rainwater harvesting can provide an effective supplemental water source in precipitation-rich regions, achieving near-complete or even full on-site water offset under realistic storage constraints. In contrast, dry cooling shifts water impact from on-site to off-site and can increase stress-adjusted water and carbon when the associated energy penalty is high or electricity is sourced from water-stressed regions.

Overall, our findings demonstrate that meaningful improvements in data center water sustainability are achievable through stress-aware scheduling, alternative water sourcing, and carefully evaluated cooling technologies. More broadly, this work underscores the need to treat water as a first-class sustainability metric and to move beyond uniform accounting toward context-aware, impact-driven evaluation of digital infrastructure.

Acknowledgments

This work is supported in part by the U.S. National Science Foundation under grant numbers ECCS-2152357, CCF-2324915, CCF-2324916, and CCF-2324941.

References

- [1] 2024. AQUEDUCT Water Risk Atlas. <https://www.wri.org/applications/aqueduct/water-risk-atlas> Accessed: March 7, 2025.
- [2] 2024. U.S. Drought Monitor. <https://droughtmonitor.unl.edu/> Accessed: March 7, 2025.
- [3] 2024. Weather Underground. <https://www.wunderground.com/> Accessed: March 7, 2025.
- [4] Kishwar Ahmed, Mohammad A Islam, Shaolei Ren, and Gang Quan. 2014. Can Data Center Become Water {Self-Sufficient}? In *6th Workshop on Power-Aware Computing and Systems (HotPower 14)*.
- [5] Dan Azevedo, Symantec Christian Belady, and Jack Pouchet. 2011. Water usage effectiveness (WUE): A green grid datacenter sustainability metric. *The Green Grid* 32 (2011).
- [6] Texas Water Development Board. 2020. *Texas Rainwater Harvesting Guide*. Texas Water Development Board.
- [7] A. Campisano and C. Modica. 2012. Optimal sizing of storage tanks for domestic rainwater harvesting in Sicily (Italy). *Water Science and Technology* 66, 1 (2012), 218–227.
- [8] Microsoft Corporation. 2023. Microsoft Environmental, Social, and Governance (ESG) Report 2023. <https://www.microsoft.com/en-us/corporate-responsibility/sustainability>
- [9] U.S. Green Building Council. 2019. *LEED v4.1 Building Design and Construction Guide*. U.S. Green Building Council.
- [10] DataCenterMap. 2026. U.S. Data Center Locations. <https://www.datacentermap.com/usa/>. Accessed: Jan. 2026.
- [11] EIA. 2024. Thermoelectric cooling water data. <https://www.eia.gov/electricity/data/water/> Accessed: March 7, 2025.
- [12] EPA. 2024. Emissions & Generation Resource Integrated Database (eGRID). <https://www.epa.gov/egrid> Accessed: March 7, 2025.
- [13] Equinix. 2021. Sustainability at Equinix: Water Management. Available at: <https://www.equinix.com/>.
- [14] Malin Falkenmark and Johan Rockström. 2004. *Balancing water for humans and nature: the new approach in ecohydrology*. Earthscan.
- [15] E. Ghisi and S.M. Oliveira. 2007. Potential for potable water savings by using rainwater and greywater in a multi-storey residential building in southern Brazil. *Building and Environment* 42, 7 (2007), 2512–2522.
- [16] Moustapha Ghibga and Shaolei Ren. 2024. FlexCoolDC: Flexible Data Center Cooling for Water–Carbon Trade-offs. In *Proceedings of the ACM International Conference on Energy-Efficient Computing and Networking (e-Energy)*. ACM.
- [17] Google. 2022. Google Data Centers: Water Stewardship. Available at: <https://www.google.com/datacenters/>.
- [18] Google. 2023. Google Transparency Report: Traffic Overview. <https://transparencyreport.google.com/traffic/overview>. Accessed: Jan. 2026.
- [19] Greenhouse Gas Protocol. 2024. GHG Protocol Corporate Standards. <https://ghgprotocol.org> Accessed: March 7, 2025.
- [20] Emily Grubert and Kelly T. Sanders. 2018. Water Use in U.S. Power Production: A Multiregional Analysis. *Environmental Research Letters* 13, 1 (2018), 014033.
- [21] Pranjol Sen Gupta, Md Rajib Hossen, Pengfei Li, Shaolei Ren, and Mohammad A Islam. 2024. A dataset for research on water sustainability. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 442–446.
- [22] 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, 854–867.
- [23] 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 (2023), 1–28.
- [24] Arjen Hoekstra, Ashok K Chapagain, Maite M Aldaya, and Mesfin M Mekonnen. 2012. *The water footprint assessment manual: Setting the global standard*. Routledge.
- [25] Mohammad Islam and Shaolei Ren. 2014. Water-Constrained Geographic Load Balancing in Data Centers. In *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*. IEEE.
- [26] Mohammad A Islam, Kishwar Ahmed, Shaolei Ren, and Gang Quan. 2014. Exploiting Temporal Diversity of Water Efficiency to Make Data Center Less "Thirsty". In *11th International Conference on Autonomic Computing (ICAC 14)*. 145–152.
- [27] Mohammad A Islam, Kishwar Ahmed, Hong Xu, Nguyen H Tran, Gang Quan, and Shaolei Ren. 2016. Exploiting spatio-temporal diversity for water saving in geo-distributed data centers. *IEEE Transactions on Cloud Computing* 6, 3 (2016), 734–746.
- [28] Mohammad A Islam, Shaolei Ren, Gang Quan, Muhammad Z Shakir, and Athanasios V Vasilakos. 2015. Water-constrained geographic load balancing in data centers. *IEEE Transactions on Cloud Computing* 5, 2 (2015), 208–220.
- [29] Yankai Jiang, Rohan Basu Roy, Raghavendra Kanakagiri, and Devesh Tiwari. 2025. WaterWise: Co-optimizing Carbon-and Water-Footprint Toward Environmentally Sustainable Cloud Computing. In *Proceedings of the 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*. 297–311.
- [30] Uisung Lee, Joseph Chou, Hui Xu, Derrick Carlson, Aranya Venkatesh, Erik Shuster, Timothy J Skone, and Michael Wang. 2020. Regional and seasonal water stress analysis of United States thermoelectricity. *Journal of Cleaner Production* 270 (2020), 122234.
- [31] Uisung Lee, Hui Xu, Jesse Daystar, Amgad Elgowainy, and Michael Wang. 2019. AWARE-US: Quantifying water stress impacts of energy systems in the United States. *Science of the total environment* 648 (2019), 1313–1322.
- [32] Baolin Li, Rohan Basu Roy, Daniel Wang, Siddharth Samsi, Vijay Gadepally, and Devesh Tiwari. 2023. Toward sustainable hpc: Carbon footprint estimation and environmental implications of hpc systems. In *Proceedings of the international conference for high performance computing, networking, storage and analysis*. 1–15.
- [33] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2025. Making AI Less 'Thirsty'. *Commun. ACM* 68, 7 (June 2025), 54–61. doi:10.1145/3724499
- [34] Zhenhua Liu, Minghong Lin, Adam Wiernan, Steven H Low, and Lachlan LH Andrew. 2011. Greening geographical load balancing. *ACM SIGMETRICS Performance Evaluation Review* 39, 1 (2011), 193–204.
- [35] Diptyarop 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.
- [36] M. M. Mekonnen and A. Y. Hoekstra. 2015. Global Gray Water Footprint and Water Pollution Levels Related to Anthropogenic Nitrogen Loads to Fresh Water. *Environmental Science & Technology* 49, 21 (2015), 12860–12868.
- [37] M. M. Mekonnen and A. Y. Hoekstra. 2016. Four Billion People Facing Severe Water Scarcity. *Science Advances* 2, 2 (2016), e1500323.
- [38] David Mytton. 2021. Data centre water consumption. *npj Clean Water* 4, 1 (2021), 11.
- [39] United Nations. 2020. *The United Nations World Water Development Report 2020: Water and Climate Change*. UNESCO.
- [40] U.S. Department of Energy. 2023. Energy and Water Nexus in U.S. Data Centers: Projections and Challenges. <https://www.energy.gov/>
- [41] Intergovernmental Panel on Climate Change (IPCC). 2021. *Climate Change 2021: The Physical Science Basis*. Cambridge University Press. <https://www.ipcc.ch/report/ar6/wg1/>
- [42] Kh Rahmani. 2017. Reducing water consumption by increasing the cycles of concentration and Considerations of corrosion and scaling in a cooling system. *Applied thermal engineering* 114 (2017), 849–856.
- [43] Paul Reig, Tianyi Luo, Eric Christensen, and Julie Sinistore. 2020. Guidance for calculating water use embedded in purchased electricity. *World Resources Institute* (2020).
- [44] Arman Shehaby, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakkar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, Dale Sartor, et al. 2024. 2024 United States Data Center Energy Usage Report. (2024). <https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf>
- [45] Md Abu Bakar Siddik, Arman Shehaby, and Landon Marston. 2021. The environmental footprint of data centers in the United States. *Environmental Research Letters* 16, 6 (2021), 064017.
- [46] 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*. 924–941.
- [47] Zahidur Talukder, Imtiaz Bin Rahim, Pranjol Sen Gupta, Shaolei Ren, and Mohammad A Islam. . https://osf.io/dv6s3/overview?view_only=69f2e0835e1d45cdbe373f1e54e5bd1f Water-Stress at the U.S. Data Center Locations. *Open Science Framework (OSF)* (). https://osf.io/dv6s3/overview?view_only=69f2e0835e1d45cdbe373f1e54e5bd1f
- [48] Google Sustainability Team. 2023. Google Environmental Report 2023. <https://sustainability.google/reports>
- [49] U.S. Environmental Protection Agency. 2023. eGRID Summary Data. <https://www.epa.gov/egrid/summary-data> Accessed: Feb. 18, 2025.
- [50] Rishab Vardhan and S. Harikrishnan. 2025. Water Energy Nexus in Data Center Design. In *2025 OCP Global Summit*. Open Compute Project. https://www.youtube.com/watch?v=Qvw0OF_GCe0 Presented at the OCP Global Summit 2025, San Jose, CA.
- [51] 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*. 260–272.
- [52] Yanran Wu, Inez Hua, and Yi Ding. 2025. Not all water consumption is equal: A water stress weighted metric for sustainable computing. *ACM SIGENERGY Energy Informatics Review* 5, 2 (2025), 84–90.
- [53] La Zhuo, Mesfin M Mekonnen, Arjen Y Hoekstra, and Yoshihide Wada. 2016. Inter- and intra-annual variation of water footprint of crops and blue water scarcity in

the Yellow River basin (1961–2009). *Advances in water resources* 87 (2016), 29–41.

A Stress-Adjusted Water Modeling and Optimization Formulation

This appendix presents the formal modeling details underlying our analysis. We define stress-adjusted water consumption, explicitly decompose on-site and off-site impacts, and describe the offline optimization formulation used to compute an upper bound on achievable sustainability gains from workload scheduling.

A.1 Stress-Adjusted Water Consumption

Traditional volumetric water accounting implicitly assumes that all water consumption has a uniform environmental impact. In reality, the consequences of water withdrawal depend strongly on regional availability and seasonal scarcity. To capture this heterogeneity, we define *stress-adjusted water consumption* as volumetric water use weighted by a location- and time-specific water-stress factor.

We adopt the Available **W**ater **R**emaining for the United States (AWARE-US) model, which provides county-level monthly water stress characterization. Higher AWARE values correspond to more water-stressed conditions.

For a data center n at time t , the total stress-adjusted water impact is the sum of on-site and off-site components:

$$w_{n,t}^{\text{stress}} = \sigma_{n,t}^{\text{on}} \cdot w_{n,t}^{\text{on}} + \sigma_{n,t}^{\text{off}} \cdot w_{n,t}^{\text{off}}, \quad (7)$$

where $\sigma_{n,t}^{\text{on}}$ denotes the local water stress factor at the data center location, and $\sigma_{n,t}^{\text{off}}$ captures the water stress associated with electricity generation supplying the data center.

A.2 On-Site Water Consumption

On-site water consumption arises primarily from evaporative cooling systems, such as cooling towers. We model on-site water use as proportional to IT power consumption:

$$w_{n,t}^{\text{on}} = \omega_{n,t}^{\text{on}} \cdot p_{n,t}, \quad (8)$$

where $\omega_{n,t}^{\text{on}}$ (L/kWh) represents the on-site water usage effectiveness under local weather and cooling conditions, and $p_{n,t}$ is the IT power consumption at data center n and time t . The on-site stress factor $\sigma_{n,t}^{\text{on}}$ is derived directly from the AWARE-US score of the county in which the data center is located.

A.3 Off-Site Water Consumption from Electricity Generation

In addition to on-site water use, data centers incur indirect (off-site) water consumption through electricity generation. Many electricity generation technologies, including thermal and nuclear plants, consume water for cooling. We model off-site water consumption as:

$$w_{n,t}^{\text{off}} = \omega_{n,t}^{\text{off}} \cdot \eta_n \cdot p_{n,t}, \quad (9)$$

where η_n is the power usage effectiveness (PUE) of data center n , and $\omega_{n,t}^{\text{off}}$ denotes the average water intensity of electricity supplied to the data center at time t .

Crucially, electricity consumed by a data center may originate from multiple power plants located in different regions. To account for this geographic decoupling, we compute $\sigma_{n,t}^{\text{off}}$ as the

generation-weighted average water stress of upstream power plants. This provenance-aware modeling avoids the systematic bias introduced by region-level or national-average abstractions.

A.4 Carbon Emissions

Carbon emissions are treated separately because their environmental impact is global rather than location-specific. We compute carbon emissions as:

$$c_{n,t} = \gamma_{n,t} \cdot \eta_n \cdot p_{n,t}, \quad (10)$$

where $\gamma_{n,t}$ denotes the carbon intensity of electricity (e.g., kg CO₂/kWh) associated with the grid supplying data center n at time t .

A.5 Offline Stress-Adjusted Water Scheduling Formulation

To quantify the maximum achievable benefit of workload shifting, we formulate an offline optimization problem that assumes perfect future knowledge of water stress, water efficiency, and carbon intensity signals. This formulation is used solely for evaluation and benchmarking and is not intended as an online scheduling algorithm.

We consider a time-slotted horizon $t \in \{1, \dots, T\}$, G workload gateways indexed by g , and N data centers indexed by n . At each time t , the workload arriving at gateway g is denoted by $\lambda_{g,t}$. Workload may be executed immediately or deferred by up to K time slots.

Let $x_{g,n,t,k}$ represent the amount of IT power (or power-equivalent workload) originating from gateway g at time t that is scheduled to execute at data center n after a delay of k slots. Performance overheads, such as routing latency or migration costs, are modeled using a multiplicative factor $h_{g,n}$.

The total IT power at data center n and time t is:

$$p_{n,t} = \sum_{g=1}^G \sum_{k=0}^K x_{g,n,t-k,k} \cdot h_{g,n}. \quad (11)$$

The objective minimizes a weighted combination of total carbon emissions and total stress-adjusted water impact:

$$\min_{\{x_{g,n,t,k}\}} \sum_{t=1}^T \sum_{n=1}^N \left(c_{n,t} + \alpha \cdot \left(\sigma_{n,t}^{\text{on}} w_{n,t}^{\text{on}} + \sigma_{n,t}^{\text{off}} w_{n,t}^{\text{off}} \right) \right), \quad (12)$$

where α controls the trade-off between water and carbon.

The optimization is subject to the following constraints:

$$p_{n,t} \leq P_n, \quad \forall n, t, \quad (13)$$

$$\sum_{n=1}^N \sum_{k=0}^K x_{g,n,t,k} = \lambda_{g,t}, \quad \forall g, t, \quad (14)$$

$$x_{g,n,t,k} \geq 0, \quad \forall g, n, t, k. \quad (15)$$

By varying α , the formulation traces the Pareto frontier between carbon- and water-centric operating regimes.