Published: 19 November 2025

Citation in BibTeX format

SoCC '25: ACM Symposium on Cloud Computing
November 19 - 21, 2025
Online, USA

Conference Sponsors:
SIGOPS
SIGMOD

RESEARCH-ARTICLE

Water Footprint of Datacenter Applications: Methodological Implications of Manufacturing, Operational, and Decommissioning Phases

AMIT SAMANTA, The University of Utah, Salt Lake City, UT, United States

YANKAI JIANG, Northeastern University, Boston, MA, United States

RYAN STUTSMAN, The University of Utah, Salt Lake City, UT, United States

ROHAN BASU ROY, The University of Utah, Salt Lake City, UT, United States

Open Access Support provided by:

The University of Utah

Northeastern University

SoCC '25: Proceedings of the 2025 ACM Symposium on Cloud Computing (November 2025)
https://doi.org/10.1145/3772052.3772216
ISBN: 9798400722769

Water Footprint of Datacenter Applications: Methodological Implications of Manufacturing, Operational, and Decommissioning Phases

Amit Samanta
University of Utah
Salt Lake City, USA

Yankai Jiang
Northeastern University
Boston, USA

Ryan Stutsman
University of Utah
Salt Lake City, USA

Rohan Basu Roy
University of Utah
Salt Lake City, USA

Abstract

Rising computational demands have made cloud datacenters' water footprint a critical concern. We demonstrate how different water footprint accounting methodologies – incorporating operational, manufacturing, and decommissioning water consumption, impact measurements and highlight the need for methodology standardization for water-aware operations. Our analysis reveals opportunities for water-aware scheduling in datacenters by considering regional water variations and lifecycle impacts.

CCS Concepts

• **Computer systems organization** → **Cloud computing**.

Keywords

Cloud Computing, Datacenter Sustainability, Water Footprint.

ACM Reference Format:
Amit Samanta, Yankai Jiang, Ryan Stutsman, and Rohan Basu Roy. 2025. Water Footprint of Datacenter Applications: Methodological Implications of Manufacturing, Operational, and Decommissioning Phases . In *ACM Symposium on Cloud Computing (SoCC '25), November 19–21, 2025, Online, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3772052.3772216

1 Introduction

Why is the water footprint of large-scale cloud computing systems an increasing concern? The rapid growth of compute-intensive workloads including generative AI [13], molecular modeling [49], and climate simulation [7] has increased demand for computational resources. Recent projections indicate that datacenter electricity consumption could reach 1,000 terawatt-hours by 2026 – roughly equivalent to Japan's total electricity demand [50]. While carbon footprint has received substantial attention in sustainable computing research [8, 21, 22, 28], the water consumption of computing infrastructure remains under-explored. Modern datacenters consume vast quantities of water both directly for cooling and indirectly through electricity generation – with facilities like Frontier (Oak Ridge National Laboratory, USA) consuming approximately 3.2 billion gallons annually [9]. Moreover, datacenters are distributed across regions with varying water scarcity levels, where

ACM ISBN 979-8-4007-2276-9/2025/11
https://doi.org/10.1145/3772052.3772216

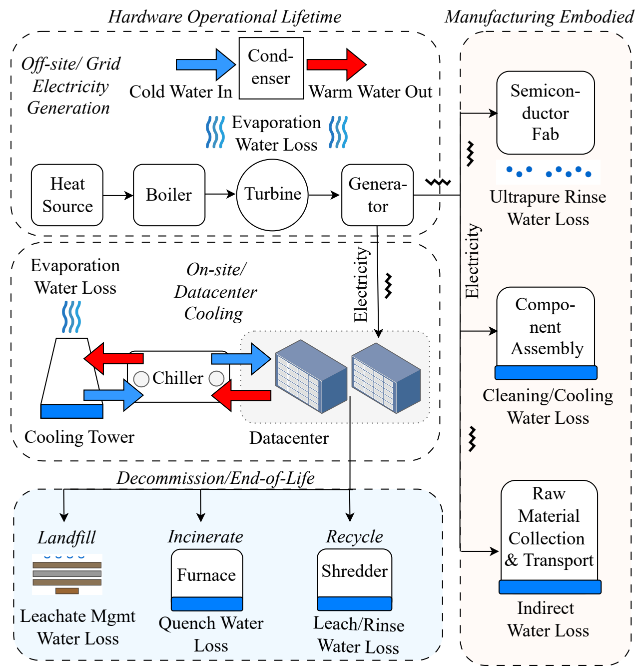

Figure 1: Water footprint sources in datacenters: water loss during operational (on-site cooling and off-site electricity generation), manufacturing, and decommissioning phases.

computational water demands can exacerbate existing resource stress and compete with basic human needs [48].

Sources of datacenter water footprint. Datacenter water consumption occurs across the entire hardware lifecycle through four distinct pathways (Fig. 1). *On-site operational water* is consumed via evaporative cooling in datacenter cooling towers, while *off-site operational water* is used during electricity generation at power plants – varying significantly by energy source [20, 24]. *Manufacturing embodied water* includes ultrapure water for semiconductor fabrication, component assembly processes, and indirect usage from supply chain activities [25]. Hardware *decommissioning water*, previously overlooked in water footprint studies, arises from end-of-life recycling (leaching/rinsing), incineration (quench water), and landfilling (leachate management) processes. While operational water consumption is continuous, manufacturing and decommissioning represent one-time impacts. As we demonstrate (Sec. 4), decommissioning and manufacturing embodied water footprint can be substantial and can significantly influence water-aware scheduling decisions – this adds another dimension to the growing water sustainability challenge in computing infrastructure.

Need for standardized water footprint accounting. Cloud providers have started to report carbon footprints of application execution [3]. There have been carbon accounting standardization efforts [45] to enable fair attribution [23], helping users understand their environmental impact as a step toward carbon-aware execution and scheduling. Yet, to the best of our knowledge, no existing methodology performs similar water footprint accounting, capturing all four components of water consumption: on-site operational, off-site operational, manufacturing embodied, and decommissioning water footprints. The vision of this paper is to bridge this gap.

In summary, we make the following contributions:

1. We present multiple water footprint accounting methodologies for datacenter application execution, capturing on-site operational, off-site operational, manufacturing embodied, and decommission/ end-of-life water footprint (Sec. 2).
2. We demonstrate that different accounting methodological choices can significantly alter measured water consumption (Sec. 4).
3. We demonstrate how to measure application water footprints using open-source tools and public datasets under various accounting methodologies (Sec. 3).
4. We identify opportunities to leverage spatio-temporal variations in water and carbon-related factors for sustainable scheduling and datacenter infrastructure planning (Sec. 4 and Sec. 5)

We present these accounting choices not as definitive solutions but as a foundation for community discourse on water footprint standardization. Next, we show how different methodologies account for various water consumption components in datacenters.

2 Different Methodologies for Accounting Water Footprint

This section presents methodologies for accounting for the water footprint of datacenter application execution. It progressively refines the water model from a simple to a comprehensive water usage model. We discuss six accounting methodologies, covering sources of water consumption and their attribution to applications. We start with an additive approach with on- and off-site water consumption, and then progressively introduce water scarcity factors, regional differentiation, embodied water, infrastructure allocation, and hardware decommissioning impacts. Each methodology builds on the previous, showing how choices impact footprints and the need for standardization.

Next, we provide brief definitions of some of the terms for easier understanding of the methodologies. They are listed below: (a) **Water scarcity** is the lack of sufficient available fresh water resources to meet demand in a specific region, (b) **Wet-bulb temperature** is the lowest temperature air can reach through evaporative cooling, (c) **Water intensity** refers to the amount of water consumed or withdrawn per unit of production.

2.1 Methodology 1: Additive Accounting of On-site and Off-site Water Footprint

Water consumption during application execution occurs through two distinct pathways [31]. *On-site water consumption* refers to the water evaporated and discharged at the datacenter for cooling computing equipment. This is quantified using *Water Usage Effectiveness* (WUE), which measures the liters of water required to dissipate heat per kilowatt-hour of energy consumed (lower values indicate greater water efficiency). WUE varies with the *wet bulb temperature* at the datacenter location: hotter and more humid climates drive higher water use for cooling, while cooler regions enable more efficient water usage. Some facilities employ air cooling to minimize water use, though this often increases energy consumption.

Off-site water consumption arises at power generation facilities supplying electricity to the datacenter. This depends on the *Energy Water Intensity Factor* (EWIF), which measures the liters of water required to produce one kilowatt-hour (L/kWh) of electricity. EWIF varies widely across energy sources: sources like wind and solar generally have low water requirements, while geothermal, hydro, coal, and nuclear energy can be highly water-intensive depending on their cooling technologies. The regional EWIF also fluctuates over time as the grid shifts between energy sources to meet demand. As a result, the off-site water footprint is impacted by energy mix and temporal grid dispatch patterns.

The energy consumption E of an application execution results in on-site and off-site water consumption. This energy is measured via hardware counters (*e.g.*, RAPL). In multi-tenant environments, application-level energy attribution requires resource utilization metrics. The total facility energy exceeds computing equipment energy by a factor, *Power Usage Effectiveness* (PUE), which accounts for cooling, power distribution losses, and other infrastructure overhead. This results in an additive water footprint modeling (W_1), as shown below:

$$W_1 = W_{\text{on-site}} + W_{\text{off-site}} = \overbrace{E \times \text{WUE}}^{\text{On-site}} + \overbrace{E \times \text{PUE} \times \text{EWIF}}^{\text{Off-site}} \tag{1}$$

While simple, this model treats all water consumption uniformly, ignoring variations in regional water availability.

2.2 Methodology 2: Water Scarcity-Weighted Accounting

Water consumption in drought-stricken regions imposes a far greater environmental burden than equivalent consumption in water-abundant areas. The *Water Scarcity Factor* (WSF) provides a coefficient that captures regional water stress, ranging from near zero in water-rich regions to values exceeding 1.0 in severely water-stressed areas [25].

$$W_2 = \overbrace{E \times \text{WUE} \times \text{WSF}}^{\text{On-site w/ Scarcity}} + \overbrace{E \times \text{PUE} \times \text{EWIF} \times \text{WSF}}^{\text{Off-site w/ Scarcity}} \tag{2}$$

An alternative approach to incorporating WSF uses an additive rather factor:

$$W_{2_{\text{Alt}}} = \overbrace{E \times \text{WUE} \times (1 + \text{WSF})}^{\text{Scarcity-Weighted On-site}} + \overbrace{E \times \text{PUE} \times \text{EWIF} \times (1 + \text{WSF})}^{\text{Scarcity-Weighted Off-site}} \tag{3}$$

This formulation ensures that even in regions with low scarcity (WSF approaching zero), the volumetric water consumption remains in the accounting. The choice between multiplicative and additive weighting reflects different approaches to environmental impact assessment—whether water consumption in abundant regions should be considered negligible or merely less impactful.

Note that this formulation applies the same WSF to both on-site and off-site water consumption, implicitly assuming that the datacenter and its electricity sources exist in the same water basin.

The W_2 method emphasizes environmental impact by scaling water use with the Water Scarcity Factor (WSF), making it ideal for high-scarcity regions. $W_{2_{\text{Alt}}}$ ensures that actual water use is counted and is suitable for low-scarcity regions.

2.3 Methodology 3: Accounting Differentiated Water Scarcity

Modern electricity grids aggregate power from geographically distributed generation facilities, each situated in regions with varying water availability. As a result, the WSF at the location of the datacenter (affecting on-site water footprint) can be different from the WSF of the power plant location (affecting off-site water footprint) [20].

$$W_3 = \overbrace{E \times \text{WUE} \times \text{WSF}_{\text{local}}}^{\text{On-site w/ Local Scarcity}} + \overbrace{E \times \text{PUE} \times \text{EWIF} \times \text{WSF}_{\text{grid}}}^{\text{Off-site w/ Grid Scarcity}} \tag{4}$$

Here, $\text{WSF}_{\text{local}}$ adjusts for water scarcity at the datacenter's location, and it reflects the local impact of direct water use for cooling. WSF_{grid} accounts for the weighted-average water stress across all power plants supplying electricity to the grid region, reflecting the indirect upstream water scarcity for energy generation. This differentiation shows that optimizing local (on-site) water use may transfer water stress to grid-level generation regions (off-site). However, these methodologies assume that the entire server's energy consumption can be attributed to a single application, which fails to reflect the reality of modern datacenter operations where multiple applications share computational resources.

2.4 Methodology 4: Attributing Multi-Tenant Resource Allocations

Modern datacenters run multiple applications concurrently on shared servers, with co-located workloads competing for CPU, memory, and accelerators. Since energy is supplied at the server level via power supplies and voltage regulators, direct per-application energy measurement is infeasible. Attribution mechanisms are thus required to divide server-level water consumption among applications. Existing tools like Intel's RAPL (for CPU and DRAM) provide package-level energy counters but cannot separate the energy used by individual processes. Hence, a server's total energy consumption E_{server} needs to be apportioned across all active workloads.

A practical approximation, adopted in previous carbon footprint accounting methodologies, divides energy based on resource utilization [45]. For an application utilizing a fraction $\frac{\text{Compute cycles}_{\text{app}}}{\text{Compute cycles}_{\text{total}}}$ of the server's total compute cycles (or equivalent accelerator compute time), the attributed water is:

$$\begin{aligned} W_4 = {} & \overbrace{E_{\text{server}} \times \frac{\text{Compute cycles}_{\text{app}}}{\text{Compute cycles}_{\text{total}}} \times \text{WUE} \times \text{WSF}_{\text{local}}}^{\text{On-site considering multi-tenant}} \\ & + \underbrace{E_{\text{server}} \times \frac{\text{Compute cycles}_{\text{app}}}{\text{Compute cycles}_{\text{total}}} \times \text{PUE} \times \text{EWIF} \times \text{WSF}_{\text{grid}}}_{\text{Off-site considering multi-tenant}} \end{aligned} \tag{5}$$

This formulation assumes power scales linearly with utilization, which is reasonable for compute-intensive workloads but may overestimate energy for memory- or I/O-bound applications. Despite this limitation, proportional attribution provides a more equitable basis for comparing the water footprint of co-located applications and supports optimization in shared environments.

Observation

The water footprint of an application execution includes on-site water use for cooling and off-site use from grid energy generation. Both are influenced by regional water scarcity, as regional differences in water availability mean the environmental impact of water use is not uniform across regions. In multi-tenant datacenter environments, shared resources make it harder to attribute water use to individual applications, giving rise to various approaches for accounting for the water footprint.

2.5 Methodology 5: Considering Manufacturing Embodied Water Footprint

The previous methodologies account only for *operational water footprint* during datacenter application execution. However, the servers, storage systems, and networking equipment used to run these applications also carry a significant water footprint from their manufacturing processes. Similar to embodied carbon, this footprint arises from both direct and indirect water use during hardware production [25, 31]. Direct water use includes semiconductor fabrication (where ultrapure water is extensively used to rinse silicon wafers), printed circuit board assembly, and metal refining for components such as chassis and heat sinks. Indirect water use stems from the operation of manufacturing facilities themselves, such as the water required for transportation, and for generating electricity and cooling at semiconductor fabs and assembly plants [4]. Although incurred once during production, this manufacturing water footprint should be amortized over the hardware's operational lifetime and attributed proportionally to all applications executed on the system (similar to embodied carbon footprint accounting) [26, 29, 45], as the embodied footprint reflects the hardware's performance capabilities and revenue generation potential.

$$\begin{aligned} W_5 = {} & \overbrace{E_{\text{server}} \times \frac{\text{Compute cycles}_{\text{app}}}{\text{Compute cycles}_{\text{total}}} \times \text{WUE} \times \text{WSF}_{\text{local}}}^{\text{Operational On-site}} \\ & + \underbrace{E_{\text{server}} \times \frac{\text{Compute cycles}_{\text{app}}}{\text{Compute cycles}_{\text{total}}} \times \text{PUE} \times \text{EWIF} \times \text{WSF}_{\text{grid}}}_{\text{Operational Off-site}} \\ & + \underbrace{W_{\text{manufacturing}} \times \frac{t_{\text{execution}}}{t_{\text{lifetime}}}}_{\text{Manufacturing Embodied}} \end{aligned} \tag{6}$$

Here, $W_{\text{manufacturing}}$ represents the total water consumed during production, from raw material extraction to final assembly. For modern servers, this can range from 5,000 to 15,000 liters depending on component complexity and manufacturing location [25]. The temporal allocation factor $\frac{t_{\text{execution}}}{t_{\text{lifetime}}}$ distributes this footprint across applications proportionally to their execution duration ($t_{\text{execution}}$), assuming uniform hardware utilization over its lifetime (t_{lifetime}). Workload co-location complicates fair allocation since resources aren't exclusively dedicated to individual applications.

2.6 Methodology 6: Lifecycle Water Footprint Considering Hardware Decommissioning

Existing studies typically account for the embodied water footprint during manufacturing and the water used during operation but omit impacts arising at the *end-of-life* (EoL) or *decommissioning* stage [24, 25, 31]. This omission occurs because the methodology used in prior work for calculating the embodied water footprint involves first estimating the manufacturing embodied carbon and then converting it to water use by multiplying the associated manufacturing energy with the regional EWIF and applying a local WSF for adjustment [25]. However, hardware decommissioning contributes to additional water consumption for treating water-intensive materials such as precious metals and PCBs [34]. Including EoL/ decommission water use captures trade-offs in hardware refresh cycles and design choices.

Servers, switches, storage shelves, and accelerator cards leave the IT floor in large refresh cycles (typically every 3 - 5 years). Once removed, each device is routed through three mutually–exclusive EoL pathways: (a) *recycling* – materials are shredded, separated, and chemically processed to recover metals and polymers; (b) *incineration* – combustible fractions are burned for energy recovery with flue-gas scrubbing; and (c) *landfilling* – non-recyclable residues are sent to engineered landfills. All three routes consume water: mechanical size-reduction uses spray nozzles for dust control; hydrometallurgical extraction of precious metals relies on rinse and leach solutions; incinerators require quench water and wet scrubbers; landfills need periodic irrigation to suppress particulates and manage leachate [39, 42]. To remain consistent with life-cycle assessment (LCA) practice, we express these flows as *process-specific water intensities* (litres per kg of processed material). Summing over the material bill of a server yields the decommissioning term.

$$W_{\text{decommission}} = \sum_{m \in \text{Materials}} M_m \big(R_m\, W_{\text{recycle},m} + I_m\, W_{\text{incinerate},m} + L_m\, W_{\text{landfill},m} \big) \tag{7}$$

M_m is the mass of material m in a server, R_m, I_m, and L_m are the fractional shares of that material routed to recycling, incineration, and landfill, respectively ($R_m + I_m + L_m = 1$), and $W_{\text{process},m}$ are material-specific water intensities obtained from LCI disclosures.

Complete life-cycle expression. Adding Eq. (7) to W_5 gives a closed-form footprint for an application execution:

$$\begin{aligned} W_6 = {} & \overbrace{E_{\text{server}} \times \frac{\text{Compute cycles}_{\text{app}}}{\text{Compute cycles}_{\text{total}}} \times \text{WUE} \times \text{WSF}_{\text{local}}}^{\text{Operational On-site}} \\ & + E_{\text{server}} \times \frac{\text{Compute cycles}_{\text{app}}}{\text{Compute cycles}_{\text{total}}} \times \text{PUE} \times \text{EWIF} \times \text{WSF}_{\text{grid}} \\ & + \underbrace{W_{\text{manufacturing}} \times \frac{t_{\text{execution}}}{t_{\text{lifetime}}}}_{\text{Manufacturing Embodied}} + \underbrace{W_{\text{decommission}} \times \frac{t_{\text{execution}}}{t_{\text{lifetime}}}}_{\text{Decommission}} \end{aligned} \tag{8}$$

The factor $\frac{t_{\text{execution}}}{t_{\text{lifetime}}}$ amortizes one-time (manufacturing + decommission) impacts over the server's useful life. As shown in Sec. 4, decommission water footprint is substantial due to water-intensive processes like hydrometallurgical recovery of precious metals and PCB recycling [30, 39]. Neglecting $W_{\text{decommission}}$ biases comparisons between hardware generations with different material compositions or datacenter designs with varying refresh cycles. Including it aligns with embodied carbon accounting best practices [21, 29, 45] and enables consistent evaluation of end-of-life strategies.

Observation

Water footprint accounting should extend beyond operational use to include hardware manufacturing and decommissioning impacts. Manufacturing contributes significant water use through semiconductor fabrication and facility operations, while decommissioning adds further consumption via recycling, incineration, and landfilling processes. Allocating these one-time footprints fairly across applications is challenging in shared environments and introduces new complexities, thus requiring the need to develop approaches for capturing the full lifecycle water footprint.

Why is decommissioning water footprint important? Decommissioning at the end of its lifecycle plays a pivotal role in the overall environmental footprint of computing systems, as effective recycling and recovery strategies can dramatically reduce the demand for raw material extraction. However, decommissioning processes themselves, including mechanical shredding to separate components, high-temperature smelting to recover metals, and chemical leaching for material purification, are resource-intensive and consume substantial volumes of water, contributing to freshwater scarcity in water-stressed regions. Furthermore, improper disposal methods, such as landfilling without pretreatment, pose severe risks of long-term water contamination through leachates.

Methods Applicability. The methods (W_1 to W_5) focus on partial lifecycle aspects (e.g., operational water, embodied water, etc.) to allow users with specific goals (e.g., optimizing on-site water footprint with water scarcity) to assess water impacts without needing full lifecycle data. For example, a datacenter manager with only on-site water footprints with water scarcity can use W_2 to estimate operational water without material composition details. Whereas W_6 provides a more accurate and actionable water footprint compared to partial methods, as it considers all aspects of the water footprint. This model can be adapted for servers running multiple applications, and it is needed for a more comprehensive assessment of water footprint.

Next, we discuss measuring water footprint parameters using publicly available data and tools.

3 Water Footprint Measurement

Measuring on-site and off-site water footprint. Both on-site and off-site water footprints are calculated based on the datacenter's energy consumption (E). For CPU and DRAM, Intel RAPL [15] interface reports the energy usage of CPU and DRAM packages by obtaining the number provided by the MSR registers. The energy contribution of the storage device is minimal compared to CPU and DRAM [37]. The Water Usage Effectiveness (WUE), which impacts on-site water footprint, is a function of the outside wet bulb temperature [20, 25, 31], which in turn depends on the dry bulb temperature and ambient humidity. The wet bulb temperature is of various datacenter regions is available in Meteologix [38]. The Power Usage Effectiveness (PUE) reflects a datacenter's architectural design and energy optimization practices, and is defined as the ratio of total facility energy to IT-equipment energy. In our experiments, we use 1.2 for the PUE, which is typical in datacenters [47]. The Energy Water Intensity Factor (EWIF) is determined by the energy mix used in the power grid, collected from Electricity Map [36]. The individual values of EWIF of different energy sources are obtained from a widely-used open-source dataset [32, 33]. Regional EWIF is computed as the weighted sum of the individual EWIFs of each energy source. The water footprint is then scaled by the Water Scarcity Factor (WSF), using county-wise data [5, 27].

Measuring manufacturing embodied water footprint. The manufacturing embodied water footprint represents water consumption during server design, manufacturing, and transportation, amortized over the hardware's lifespan of four or five years [19]. Following the same methodology as prior work in this space [25] we measure it by first collecting the server's embodied carbon footprint [10, 16], multiplying by the manufacturing location's carbon intensity to derive energy consumption, then applying the Energy Water Intensity Factor (EWIF) and water scarcity factor (WSF). This approach yields similar values to industry reports, as vendors like Lenovo [2] and Seagate [6] publish embodied water footprint data for components. Additionally, embodied water footprint can be approximated based on die area and fabrication node size, consistent with prior work's [25] energy-based methodology that we follow. Transportation water usage, which counts toward manufacturing embodied water footprint, is estimated using transport distance, fuel, and truck water usages [41, 44].

Measuring decommission water footprint. To measure the water footprint of decommission, we use a detailed assessment of water consumption associated with end-of-life processes, which includes recycling, landfilling, incineration, transportation, and material preprocessing, as depicted in W_6 (Sec. 2). The recycling water usage is calculated for each material (aluminum, paper, plastic, glass, etc.) using its mass and incorporating water savings from recycling as reported by WasteTrade [52] and other public sources [18, 51]. For landfilling, we incorporate the leachate management and dust suppression water usage [46]. The incineration process includes the ash quenching, which is available through European Commission, JRC Technical Reports [35], and scrubber water usage available through U.S. Environmental Protection Agency (EPA) report [40]. We collect material composition of servers to determine the decommission water footprint from server lifecycle assessment reports [1].

Using the water footprint accounting methodologies and measurement sources, we next analyze how methodological choices impact results and discuss water consumption reduction strategies.

4 Implications of Methodologies and Water Footprint Analysis

In this section, we first show that different methodological choices result in significantly different measured water footprints, and then discuss scheduling opportunities to reduce water consumption.

Differences in measured water footprint using different methodologies. In Sec. 2, we discussed the different methodological choices to determine the water footprint of a datacenter application execution (W_1 to W_6). For the evaluation, we study six different types of applications from the CloudSuite benchmark suite [17]. For example, Media streaming's I/O is really high as it relies on constant data transfer. Graph analytics's CPU usage is really high, due to complex algorithms (e.g., PageRank, shortest path) requiring intensive computation. Fig. 2 shows that these methodologies result in significantly different water footprint measurement for an example application (memory analytics from CloudSuite [17]), which has a 8060.822 J of energy consumption, running on an Intel Xeon Platinum Scalable server with 48 physical cores in a datacenter at Santa Clara, California on June 2025. Since the datacenter is located in a relatively water-stressed region, due to high WSF (> 1), both on-site and off-site water footprint increases using W_2 compared to W_1. Using $W_{2_{\mathrm{Alt}}}$, the water footprint further increases as water scarcity is considered with (1+WSF) as a multiplicative factor instead of just WSF in W_2. Considering different $\mathrm{WSF}_{\mathrm{grid}}$ and $\mathrm{WSF}_{\mathrm{local}}$ causes the water footprint using W_3 to be slightly different than the one using W_2. Since W_4 accounts only for the fraction of the server's energy proportional to the application's compute cycles relative to total compute cycles (excluding idle energy and co-located workloads), both off-site and on-site water footprints of the application execution decrease compared to W_3. Manufacturing embodied water footprint contributes to a significant increase in the total water footprint of the application using W_5. Finally, considering decommission water footprint further increases the total water footprint for the application execution.

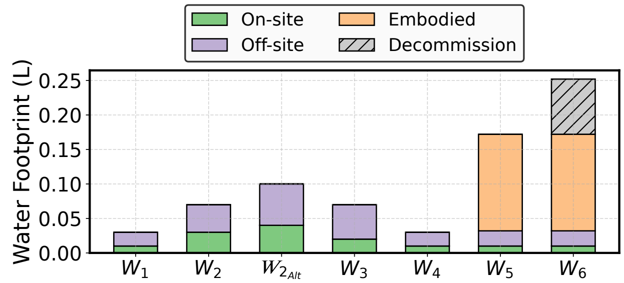

Figure 2: Depending on the chosen methodology, the total water footprint and the individual contributions from on-site, off-site, manufacturing embodied and decommission water footprint vary significantly.

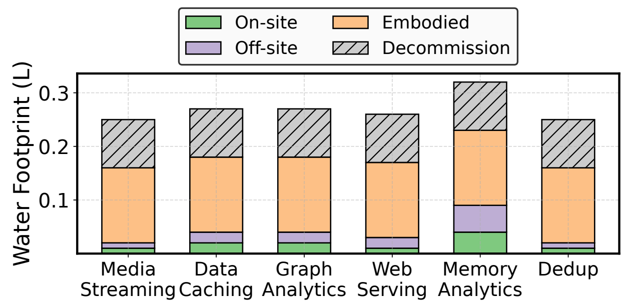

Figure 3: The water footprint of different applications varies depending on the execution time and energy consumption.

Following W_6, in Fig. 3 we show the water footprint of different applications from CloudSuite [17] and Parsec [11] benchmark suites, executed under the same WUE, EWIF, and WSF conditions on the same Intel Xeon Platinum Scalable server. The water footprint varies depending on the energy consumption (affecting on-site and off-site water footprint) of the application and the execution

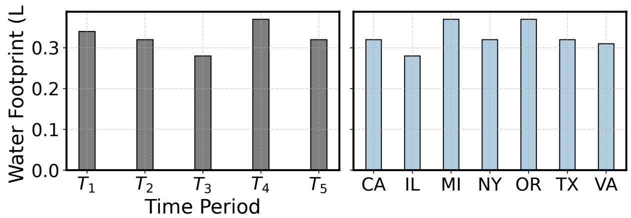

Figure 4: Water footprint of an application execution varies spatio-temporally due to variations in EWIF, WUE, and WSF.

time (affecting decommission and manufacturing embodied carbon footprint). Thus, application-level water footprint depends on hardware, location, and workload execution characteristics, requiring system-level optimizations for water-aware datacenter operations.

Observation

Different accounting methodologies can yield vastly different water footprints for the same application execution, underscoring the need for standardized approaches to enable fair comparisons and guide effective reduction strategies. Manufacturing embodied water footprints can be substantial and can even exceed operational footprints. Unlike carbon, where decommissioning impacts can be relatively small [1], decommissioning water footprint can be significant due to water-intensive recycling and disposal processes.

Next, we discuss how we can leverage spatio-temporal characteristics of water-related parameters to reduce the water footprint of application execution in datacenters.

Spatio-temporal variation of water footprint. EWIF varies across regions with the energy mix and over time as grids shift generation sources. WUE depends on local climate and fluctuates with temperature and humidity, causing regional and seasonal variation in on-site water use [20]. In contrast, WSF remains stable over time but differs significantly across regions (*e.g.,* 0.681 in California, and 0.197 in Michigan), reflecting variations in water availability. These factors make the total water footprint of an application on the same server type vary spatially and temporally, as shown in Fig. 4. This figure shows the water footprint of a CloudSuite media streaming application across different days in June 2025 (T_1–T_5) and datacenter locations on identical Intel Xeon Scalable servers.

Such spatio-temporal dynamics enable water-aware schedulers to optimize application placement. Workloads can shift between datacenters to avoid periods or regions with high EWIF/WUE or be deferred locally to exploit low water-use intervals. Latency-tolerant jobs can align with favorable water conditions, while latency-sensitive ones are prioritized. These variations can guide schedulers to dynamically route or defer tasks based on water footprint profiles and QoS constraints. Note that embodied and decommissioning water footprints remain constant post-deployment and do not vary spatially or temporally. In addition to water footprint, carbon footprint from grid carbon intensity (CI) is a key sustainability metric [21, 22], but minimizing water use may not align with carbon reduction goals. Next, we discuss the reasons for this.

Focusing solely on minimizing water footprint can risk increasing another key sustainability metric, carbon footprint. Carbon footprint, expressed in grams of CO_2, includes operational

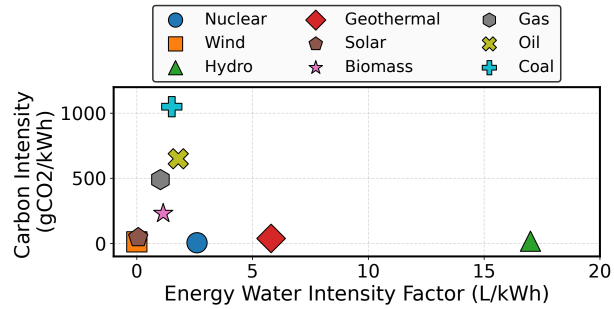

Figure 5: Greener energy sources (low CI) can have higher water footprint (high EWIF) and vice-versa.

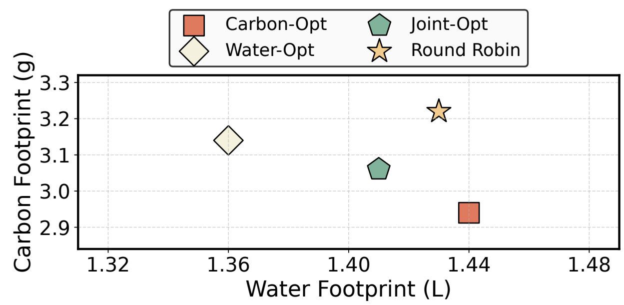

Figure 6: There is a trade-off between minimizing carbon and water footprints that joint optimization can balance.

emissions (energy consumption multiplied by the grid's carbon intensity, CI) and embodied emissions from hardware manufacturing, amortized over its lifetime. Since CI varies across grids and over time with shifting energy sources, operational carbon footprint exhibits spatio-temporal variation – renewables like wind and solar have low CI, while fossil fuels like coal and oil have high CI.

However, low-CI sources are not always water-efficient (*i.e.,* low EWIF). As shown in Fig. 5, geothermal and hydroelectric power have low CI but high EWIF because geothermal relies on water-intensive cooling systems, and hydroelectric dams lose water through evaporation, which cannot be recovered. This mismatch creates tension for sustainability-aware schedulers: minimizing carbon footprint may increase water use and vice versa. Joint optimization requires balancing these trade-offs, especially when deferring workloads or shifting them across regions with different CI and EWIF profiles.

To demonstrate the trade-offs between minimizing carbon and water footprints, we design an experiment where applications (shown in Fig. 3) are invoked periodically for an hour across all regions shown in Fig. 4, with five Intel Xeon Platinum Scalable servers in each location. Fig. 6 shows the results. The carbon-optimized scheduler (Carbon-Opt) assigns workloads to regions with the lowest carbon intensity (CI), filling each datacenter before moving to the next lowest CI region. The water-optimized scheduler (Water-Opt) selects regions with the lowest total water footprint, similarly filling each before moving to the next. Joint optimization (Joint-Opt) balances both metrics by weighting carbon and water footprints equally in the placement decision. It achieves a middle ground: it improves water footprint by 2.08% over Carbon-Opt and reduces carbon footprint by 2.55% compared to Water-Opt. The Round

Robin scheduler, which cycles through datacenter regions without considering sustainability metrics, performs worse on both axes.

Observation

Water footprint varies significantly across regions and time due to variations in EWIF, WUE, and WSF, creating opportunities for spatio-temporal scheduling. However, optimizing solely for water can increase carbon footprint, underscoring the need for joint optimization to balance these different sustainability goals.

Next, we discuss recommendations and best practices for datacenter operators to perform water-aware operations.

5 Discussion on Sustainable Operation

Grid-Aware infrastructure and lifecycle planning. The regional variations in grid composition and climate conditions revealed can enable strategic infrastructure decisions. Datacenters should prioritize locations where the grid relies on truly water-efficient renewables (solar, wind) rather than water-intensive sources like hydroelectric or geothermal, which despite low carbon emissions, can have EWIF values exceeding fossil fuels. This grid-aware site selection, combined with lifecycle planning that schedules hardware decommissioning during wet seasons and routes e-waste to water-abundant recycling facilities, can significantly reduce peak water stress. Deploying efficient processors in water-scarce regions while extending older hardware lifecycles in water-abundant areas can optimize the performance-water trade-off.

Dynamic workload orchestration and market incentives. Dynamic workload orchestration based on real-time water metrics offers sustainability optimization opportunities. Short-duration jobs can opportunistically utilize regions experiencing temporary EWIF spikes without significantly amortizing embodied water footprints, while long-running applications can target consistently water-efficient locations. Like carbon, with standardized water footprint accounting, cloud providers can transparently report water consumption in liters per execution. This enables fair comparisons and creates market incentives – when developers see their applications' true water footprint, they optimize execution time and region selection for water-conscious execution.

Energy Attribution. Developing more sophisticated models that explicitly incorporate factors such as memory bandwidth utilization, diverse I/O patterns (e.g., read/write throughput and latency), and accelerator usage (e.g., GPU/TPU offloading efficiency) would significantly enhance the accuracy of energy consumption estimates of our model. However, this remains an active and challenging area of research, drawing parallels to ongoing advancements in related fields like carbon accounting for data centers, where standardized methodologies for attributing emissions across hardware layers are still evolving to account for dynamic workloads.

Integration of Uncertainty Analysis. Methodologies should incorporate uncertainty analysis to reflect variability inherent in each phase, such as fluctuating water efficiency in manufacturing tools or climate-dependent operational cooling needs. This approach not only enhances methodological validity under real-world scenarios but also makes the models more resilient. More importantly, it transforms static footprint assessments into dynamic and more actionable frameworks for global sustainability.

Observation

Sustainable datacenter operations require both infrastructure-level decisions (grid-aware siting, lifecycle planning, hardware heterogeneity) and operational strategies (dynamic workload orchestration, transparent accounting). Standardized water footprint accounting enables market-driven optimization by making water consumption visible to both providers and customers.

6 Related Works

Due to growing computational demand and resulting water footprint increases, water footprint optimization has emerged as a critical challenge in sustainable datacenter research [20]. Prior work enables water footprint quantification [31, 43, 48], with Li et al. [31] estimating operational water use across spatial and temporal dimensions – though manufacturing and decommissioning impacts remain unaddressed. The Water Scarcity Factor (WSF) impacts the quantification of the water footprint, and several studies provide diverse methodologies for determining WSF [12, 14, 27, 53]. For example, SCARF [53] presents a framework that incorporates both spatial and temporal dimensions to better capture regional water stress. The field of water footprint research is expanding rapidly, with recent datasets [20] enabling new optimization opportunities.

To reduce the water footprint of large-scale systems, various schedulers have been proposed [24, 25] that leverage spatio-temporal variations in water metrics. For example, WaterWise [25] co-optimizes carbon and water footprints using regional scarcity factors. However, these optimization efforts assume standardized water accounting methodologies exist. It does not consider decommission for their modelling, and also does not consider the various possible methodological choices. However, the lack of comprehensive accounting that captures all water sources leads to inconsistent measurements and can misguide optimization decisions.

7 Conclusion

In this paper, we establish multiple methodologies for quantifying datacenter application execution's water footprint. We demonstrate how different accounting choices yield different measurements for identical workloads. Our analysis shows that comprehensive accounting should include all four water consumption sources: on-site operational, off-site operational, manufacturing embodied, and decommissioning water, with decommissioning being previously overlooked despite its significant impact. We demonstrate that methodological choices and spatio-temporal variations in water-related factors (EWIF, WUE, WSF) significantly impact water footprint measurements and optimization opportunities. We hope this work serves as a step toward the standardization of water footprint accounting for sustainable datacenter operations.

Acknowledgments: We thank the anonymous reviewers for their valuable feedback. This work is supported by the University of Utah's Kahlert School of Computing, John and Marcia Price College of Engineering, and Scientific Computing & Imaging Institute.

References

[1] 2019. Life Cycle Assessment of Dell R740). https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/Full_LCA_Dell_R740.pdf.

[2] 2024. Lenovo FY 2024 ESG Report. https://investor.lenovo.com/en/sustainability/reports/FY2024-lenovo-sustainability-report.pdf.

[3] 2024. Viewing your carbon footprint - AWS Billing — docs.aws.amazon.com. https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/what-is-ccft.html.

[4] 2024. The water challenge for semiconductor manufacturing: What needs to be done? uthttps://www.weforum.org/stories/2024/07/the-water-challenge-for-semiconductor-manufacturing-and-big-tech-what-needs-to-be-done/?.

[5] 2024. Water Use and Stress — ourworldindata.org. https://ourworldindata.org/water-use-stress.

[6] 2025. Seagate FY 2025 ESG Performance Report. https://www.seagate.com/content/dam/seagate/assets/esg/resources/files/seagate_fy2024_esg_performance_report_pages_521.pdf.

[7] Sameh Abdulah, Allison H Baker, George Bosilca, Qinglei Cao, Stefano Castruccio, Marc G Genton, David E Keyes, Zubair Khalid, Hatem Ltaief, Yan Song, et al. 2024. Boosting earth system model outputs and saving petabytes in their storage using exascale climate emulators. In *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–12.

[8] 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*. 118–132.

[9] Maxwell Adler. 2023. The World's Fastest Supercomputer Needs 74 Miles of Cable and Tons of Water. *Bloomberg* (2023). https://www.bloomberg.com/news/features/2023-01-05/frontier-supercomputer-world-s-fastest-needs-74-miles-of-cable

[10] Clément Auger, Benoit Hilloulin, Benjamin Boisserie, Maël Thomas, Quentin Guignard, and Emmanuel Rozière. 2021. Open-source carbon footprint estimator: Development and university declination. *Sustainability* 13, 8 (2021), 4315.

[11] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC benchmark suite: Characterization and architectural implications. In *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*. 72–81.

[12] Anne-Marie Boulay, Jane Bare, Lorenzo Benini, Markus Berger, Michael J Lathuillière, Alessandro Manzardo, Manuele Margni, Masaharu Motoshita, Montserrat Núñez, Amandine Valerie Pastor, et al. 2018. The WULCA consensus characterization model for water scarcity footprints: assessing impacts of water consumption based on available water remaining (AWARE). *The International Journal of Life Cycle Assessment* 23, 2 (2018), 368–378.

[13] Franck Cappello. [n.d.]. AuroraGPT: Exploring AI Assistant for Science. ([n. d.]).

[14] Li Chen and Aaron P Wemhoff. 2022. Characterizing Data Center Cooling System Water Stress in the United States. In *ASHRAE Annual Conference papers*.

[15] Howard David, Eugene Gorbatov, Ulf R Hanebutte, Rahul Khanna, and Christian Le. 2010. RAPL: Memory power estimation and capping. In *Proceedings of the 16th ACM/IEEE international symposium on Low power electronics and design*. 189–194.

[16] Benjamin Davy. 2021. Building an AWS EC2 carbon emissions dataset. *Medium, September* (2021).

[17] Michael Ferdman, Almutaz Adileh, Onur Kocberber, Stavros Volos, Mohammad Alisafaee, Djordje Jevdjic, Cansu Kaynak, Adrian Daniel Popescu, Anastasia Ailamaki, and Babak Falsafi. 2012. Clearing the clouds: a study of emerging scale-out workloads on modern hardware. *Acm sigplan notices* 47, 4 (2012), 37–48.

[18] Franklin Associates and The Association of Plastic Recyclers. 2020. *LIFE CYCLE IMPACTS FOR POSTCONSUMER RECYCLED RESINS: PET, HDPE, AND PP*. Technical Report. The Association of Plastic Recyclers (APR). https://plasticsrecycling.org/wp-content/uploads/2024/08/APR-Recycled-vs-Virgin-LCA-May2020.pdf

[19] Sriram Govindan, Anand Sivasubramaniam, and Bhuvan Urgaonkar. 2011. Benefits and limitations of tapping into stored energy for datacenters. In *Proceedings of the 38th annual international symposium on Computer architecture*. 341–352.

[20] 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.

[21] 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 *ACM/IEEE ISCA*. 784–799.

[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 *IEEE HPCA*. 854–867.

[23] Leo Han, Jash Kakadia, Benjamin C Lee, and Udit Gupta. 2025. Fair-CO2: Fair attribution for cloud carbon emissions. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture*. 646–663.

[24] 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.

[25] 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.

[26] Yankai Jiang, Rohan Basu Roy, Baolin Li, and Devesh Tiwari. 2024. EcoLife: Carbon-Aware Serverless Function Scheduling for Sustainable Computing. In *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–15.

[27] 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.

[28] 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 *ACM/IEEE SC*. 1–15.

[29] Baolin Li, Yankai Jiang, Vijay Gadepally, and Devesh Tiwari. 2024. Sprout: Green generative AI with carbon-efficient LLM inference. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 21799–21813.

[30] Huan Li, Jacques Eksteen, and Elsayed Oraby. 2018. Hydrometallurgical recovery of metals from waste printed circuit boards (WPCBs): Current status and perspectives–A review. *Resources, Conservation and Recycling* 139 (2018), 122–139.

[31] Pengfei Li, Jianyi Yang, Mohammad A Islam, and Shaolei Ren. 2023. Making ai less" thirsty": Uncovering and addressing the secret water footprint of ai models. *arXiv preprint arXiv:2304.03271* (2023).

[32] Jordan Macknick, Robin Newmark, Garvin Heath, and KC Hallett. 2011. Review of operational water consumption and withdrawal factors for electricity generating technologies. (2011).

[33] Jordan Macknick, Robin Newmark, Garvin Heath, and Kathleen C Hallett. 2012. Operational water consumption and withdrawal factors for electricity generating technologies: a review of existing literature. *Environmental Research Letters* 7, 4 (2012), 045802.

[34] Hema Madaka, Callie W Babbitt, and Erinn G Ryen. 2022. Opportunities for reducing the supply chain water footprint of metals used in consumer electronics. *Resources, Conservation and Recycling* 176 (2022), 105926.

[35] Magaldi. 2025. *Every Drop Matters: Rethinking Water Use in IBA Handling*. https://www.magaldi.com/en/every-drop-matters-rethinking-water-use-iba-handling

[36] Electricity Maps. 2025. *Electricity Maps Live 24/7*. https://app.electricitymaps.com/map

[37] David Meisner, Christopher M Sadler, Luiz André Barroso, Wolf-Dietrich Weber, and Thomas F Wenisch. 2011. Power management of online data-intensive services. In *Proceedings of the 38th annual international symposium on Computer architecture*. 319–330.

[38] meteologix. 2024. *meteologix Live 24/7*. https://meteologix.com/in

[39] Shaila Mir and Nikhil Dhawan. 2022. A comprehensive review on the recycling of discarded printed circuit boards for resource recovery. *Resources, Conservation and Recycling* 178 (2022), 106027.

[40] Daniel Mussatti and Paula Hemmer. 2002. *Wet Scrubbers for Particulate Matter*. Technical Report. U.S. Environmental Protection Agency. https://www.epa.gov/sites/default/files/2020-07/documents/cs6ch2.pdf

[41] National Transportation Library. [n.d.]. Environmental - Advantages of Inland Barge Transportation. https://rosap.ntl.bts.gov/view/dot/4544/dot_4544_DS1.pdf?

[42] Julio Navarro and Fu Zhao. 2014. Life-cycle assessment of the production of rare-earth elements for energy applications: a review. *Frontiers in Energy Research* 2 (2014), 45.

[43] Bora Ristic, Kaveh Madani, and Zen Makuch. 2015. The water footprint of data centers. *Sustainability* 7, 8 (2015), 11260–11284.

[44] RMI. 2023. *Hydrogen Reality Check: Distilling Green Hydrogen's Water Consumption*. https://rmi.org/hydrogen-reality-check-distilling-green-hydrogens-water-consumption/

[45] Rohan Basu Roy, Raghavendra Kanakagiri, Yankai Jiang, and Devesh Tiwari. 2024. The Hidden Carbon Footprint of Serverless Computing. In *Proceedings of the 2024 ACM Symposium on Cloud Computing*. 570–579.

[46] Tonmoy Sarker. 2017. Water and Energy Costs of Landfilled Food Waste. (2017).

[47] 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).

[48] 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 (2021), 064017.

[49] Ryan Stocks, Jorge L Galvez Vallejo, CY Fiona, Calum Snowdon, Elise Palethorpe, Jakub Kurzak, Dmytro Bykov, and Giuseppe MJ Barca. 2024. Breaking the Million-Electron and 1 EFLOP/s Barriers: Biomolecular-Scale Ab Initio Molecular Dynamics Using MP2 Potentials. In *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–12.

[50] The Guardian. [n.d.]. Water shortage fears as Labour's first AI growth zone sited close to new reservoir. *The Guardian* ([n. d.]). https://www.theguardian.com/technology/2025/jan/13/labour-ai-datacentre-growth-zone-water-shortages-abingdon-reservoir
[51] Utah State University. [n.d.]. *Facts & Figures*. https://www.usu.edu/facilities/recycling/facts-and-figures#:~:text=Recycling%20one%20ton%20of%20paper%20saves%3A&text=7000%20gallons%20(27%2C000%20liters)%20of,pounds%20(27%20kilograms)%20of%20pollutants
[52] WasteTrade. 2025. *How Effective Recycling Conserves Water and Reduces Waste*. https://www.wastetrade.com/resources/environmental/how-effective-recycling-conserves-water-and-reduces-waste/
[53] Yanran Wu, Inez Hua, and Yi Ding. 2025. Not All Water Consumption Is Equal: A Water Stress Weighted Metric for Sustainable Computing. *arXiv preprint arXiv:2506.22773* (2025).