

OpenDC-STEAM: Realistic Modeling and Systematic Exploration of Composable Techniques for Sustainable Datacenters

Dante Niewenhuis

Vrije Universiteit Amsterdam
d.niewenhuis@vu.nl

Sacheendra Talluri

Vrije Universiteit Amsterdam
s.talluri@vu.nl

Alexandru Iosup

Vrije Universiteit Amsterdam
a.iosup@vu.nl

Tiziano De Matteis

Vrije Universiteit Amsterdam
t.de.matteis@vu.nl

Abstract—The need to reduce datacenter carbon-footprint is urgent. While many sustainability techniques have been proposed, they are often evaluated in isolation, using limited setups or analytical models that overlook real-world dynamics and interactions between methods. This makes it challenging for researchers and operators to understand the effectiveness and trade-offs of combining such techniques. We design OpenDC-STEAM, an open-source customizable datacenter simulator, to investigate the individual and combined impact of sustainability techniques on datacenter operational and embodied carbon emissions, and their trade-off with performance. Using STEAM, we systematically explore three representative techniques—horizontal scaling, leveraging batteries, and temporal shifting—with diverse representative workloads, datacenter configurations, and carbon-intensity traces. Our analysis highlights that datacenter dynamics can influence their effectiveness and that combining strategies can significantly lower emissions, but introduces complex cost-emissions-performance trade-offs that STEAM can help navigate. STEAM supports the integration of new models and techniques, making it a foundation framework for holistic, quantitative, and reproducible research in sustainable computing. Following open-science principles, STEAM is available as FOSS: <https://github.com/atlarge-research/OpenDC-STEAM>.

Index Terms—OpenDC, STEAM, Sustainability, Data Center, Simulation, Carbon Emission, Load Shifting, Batteries, Horizontal Scaling

I. INTRODUCTION

The Information and Communication Technology (ICT) sector accounts for an estimated 1.5%–4% of global carbon emissions [1], [2], with datacenters contributing for a large and growing share [3], [4]. Policy mandates, user expectations, and financial considerations are pushing the computing community to control and reduce datacenter carbon footprint [5], [6]. In response, many techniques have been proposed to reduce the carbon emissions of a datacenter, including load shifting [7], [8], hardware scaling [9], [10], or batteries [7], [11], [12]. However, these strategies are often evaluated using analytical models or limited setups that do not account for interactions between sustainability techniques and with real-world datacenter dynamics, such as task failures. Their effectiveness also depends on multiple factors, such as the scale of the datacenter, workload characteristics, the energy source *carbon intensity*, and the *embodied carbon* emitted during hardware manufacturing [8], [13]. In such a broad design space, non-trivial trade-offs in costs, carbon emissions, and performance

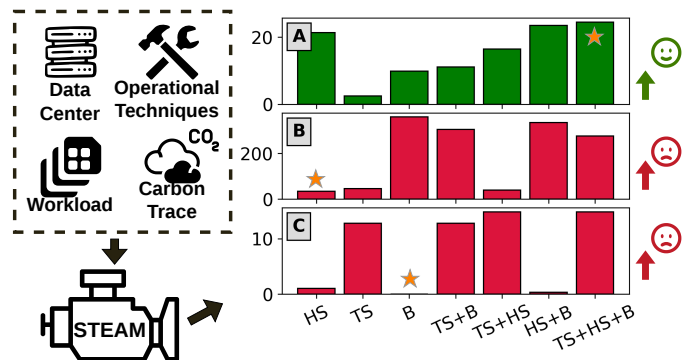


Fig. 1: STEAM quantifies the impact and trade-offs of sustainability techniques. In the figure, (A) results for the Surf workload are shown as total carbon reduction [%], (B) peak power draw [kW], and (C) average task delay [h]. Stars indicate the best performing technique for each metric. HS: Horizontal Scaling, TS: Temporal Shifting, B: Batteries.

frequently arise. To navigate this complexity and make informed decisions on how to design datacenters and configure their operations, *datacenter designers, operators, and sustainability researchers need quantitative insights*. In this work, we introduce **OpenDC-STEAM** (STEAM in the following), an open-source simulation-based tool for quantifying the effects of sustainability techniques on datacenter performance and carbon footprint.

STEAM is a new capability-driven feature in the OpenDC space [14]. STEAM’s composable architecture allows users to analyze trade-offs arising from combining techniques, evaluate new techniques, and tune existing ones. Its flexibility and performance enable datacenter designers, operators, and researchers to rapidly explore a wide range of “what-if” scenarios and gain actionable insights into sustainability decisions.

We use STEAM to quantitatively evaluate three representative sustainability techniques—horizontal scaling, temporal shifting, and batteries—on datacenter performance (e.g., service quality) and carbon footprint, considering diverse and realistic workloads, carbon regions, datacenters, and their dynamics. We show that each technique can reduce carbon emissions with minimal performance impact, but the actual extent of

this reduction depends on the workload, available hardware, location, and operational phenomena such as task failures. We highlight how disregarding datacenter dynamics has led to optimistic carbon emissions reductions in prior work. Figure 1 illustrates how STEAM can be used to understand the impact of sustainability techniques. Combining techniques can further improve their effectiveness (see Figure 1A), but also introduces trade-offs, such as power-draw spikes (Figure 1B) and service-quality degradation (Figure 1C), which STEAM can quantify.

With STEAM, we contribute to the efforts towards better understanding modern computing ecosystems [15]. STEAM provides a framework for the systematic evaluation of sustainability techniques, helping datacenter providers and researchers to make evidence-based decisions on datacenter design and operation.

This work makes the following contributions:

- 1) We motivate and design STEAM¹ (§III and §IV). STEAM is an open-source, flexible framework for quantifying the effect of sustainability techniques on the datacenter carbon footprint and performance.
- 2) Using STEAM through a relevant experimental setup (§V), we conduct an extensive analysis of three representative sustainability techniques, showing that:
 - a) Horizontal scaling can reduce carbon footprint (up to 35%), but ignoring operational phenomena, such as failure, leads to optimistic results (§VI-A).
 - b) Batteries can significantly improve the sustainability of datacenters (up to 28%) when used in the right context, but can also reduce it when misused (§VI-B).
 - c) Temporal shifting is less effective (7%) than shown by previous work when operational effects, such as stacking tasks, are taken into account (§VI-C).
- 3) We leveraged STEAM’s composable design to conduct a first-of-its-kind analysis of using multiple sustainability techniques in conjunction (§VII). Combining techniques can increase effectiveness but add new, complex trade-offs that STEAM helps quantify and explain.

II. BACKGROUND

This section introduces the concepts and techniques that enable realistic modeling of sustainability techniques in STEAM.

A. Workload

A workload contains all jobs executed by a datacenter in a specific period. We use a bag-of-tasks workload model where users submit jobs to a datacenter, each containing one or more tasks. Workloads are defined using traces containing the arrival times of all tasks, resource requirements, durations, deadlines, and sampled resource utilization over time. Workloads are often created by collecting operation logs of existing datacenters.

B. Carbon Emissions

Carbon emissions are all the CO₂ and CO₂-equivalent of other greenhouse gases emitted by a datacenter. We distinguish between *operational* and *embodied* carbon emissions.

Operational Carbon refers to all emissions generated from producing energy used to power servers, cooling systems, and other infrastructure in a datacenter. The type of energy a datacenter uses significantly impacts its operational carbon footprint, with coal-based electricity emitting up to 20x more carbon than renewable sources [16]. Since green energy availability fluctuates [8], accurate estimation requires information about the carbon intensity over time, often provided as traces from services such as ElectricityMaps² or Watttime³.

Embodied Carbon considers emissions from manufacturing datacenter components such as CPUs and GPUs. As the energy efficiency of datacenters improves, embodied carbon has become a major contributor to the overall carbon footprint [17]. A common method to calculate the embodied carbon of a component during a workload is to multiply the total embodied carbon cost of the component by the portion of the component’s lifetime the workload used [16].

C. Sustainability Techniques

Many techniques have been proposed to improve datacenter sustainability. In this work, we focus on techniques aimed at reducing the overall carbon emissions associated with the datacenter deployment and use. We consider three representative techniques: horizontal scaling, the use of batteries, temporal shifting, and their combination.

Horizontal scaling reduces the number of machines in a datacenter to the minimum needed to meet computational requirements, decreasing energy usage and carbon emissions. Datacenters often run at very low utilization [18] as fears of service disruption have led to conservative practices and significant over-provisioning of resources [19]. Although this might protect datacenters from sudden resource requirement spikes or failures, the idle machines still use energy and contribute to operational and embodied carbon emissions.

Batteries can be used to reduce carbon emissions by storing green energy when available, and providing it to the datacenter in periods of high-carbon energy. Their advantage is that they reduce emissions without impacting task execution time. However, they also introduce downsides: manufacturing batteries has a high embodied carbon cost, and charging them can cause spikes in the datacenter’s power draw (see Section VI-B2).

Temporal shifting delays or pauses a task during high-carbon periods, resuming it when greener energy is available. By moving tasks and, in turn, energy usage, to periods with lower carbon intensity, the overall operational carbon emitted is reduced. However, this requires tasks to be interruptible and delay-tolerant [8], and shifting many tasks to the same period might introduce workload spikes, requiring larger-scale datacenters to handle (see Section VII).

III. RATIONALE FOR A SIMULATION-BASED APPROACH

Datacenters are large and complex ecosystems composed of many interacting components. Evaluating the impact of

¹<https://github.com/atlarge-research/OpenDC-STEAM>

²<https://www.electricitymaps.com/>

³<https://watttime.org/>

sustainability techniques requires datacenter designers, operators, and researchers to account for all these components and their interactions. Prior work often used single-purpose analytical models to estimate carbon savings. While useful for theoretical insights, such models struggle to capture datacenter dynamics and cross-technique interactions. To address this gap, we advocate a simulation-based approach. Simulation can enable a more comprehensive analysis by capturing datacenter operational phenomena and considering how different strategies interact with each other and the rest of the system. In this section, we illustrate this using a motivating example concerning the analysis of the temporal shifting technique.

Previous works [8], [20] evaluated temporal shifting using analytical models similar to the following: for each task, emissions are calculated at its original execution time, and when delayed by a carbon-aware scheduling policy. The difference between those two represents the task’s emission savings. The overall impact of temporal shifting is calculated by averaging the savings across all tasks. This approach has two potential downsides. First, it assumes each task can be moved to a low-carbon period independently, ignoring the datacenter’s computational capacity constraints. Although similar considerations can be integrated into the scheduling strategy, their impact is difficult to capture analytically. Moreover, the approach disregards the energy used by idle machines when a task is delayed. Second, the method overlooks operational phenomena such as failures. As a result, the analytical model might overestimate the carbon savings of temporal shifting.

We illustrate the downsides with a simple example depicted in Figure 2. The example concerns the operations of a datacenter containing three hosts (H1-3) that uses temporal shifting to reduce carbon emissions. The top of the figure shows the tasks submitted and scheduled on the different hosts. Dotted boxes indicate delayed tasks, with arrows pointing to the new scheduled time. A task marked with a red cross indicates it has been interrupted due to machine failure. The bottom of the figure depicts the carbon intensity over time. We zoom into four timestamps, marked by $t1$ - $t4$. At timestamp $t1$, three tasks (1-3) are submitted. Due to the high carbon intensity, the scheduling policy delays all tasks to $t3$. At timestamp $t2$, two new tasks (4,5) are submitted. Again, the scheduling policy delays the tasks to $t3$. At timestamp $t3$, no new tasks are submitted, but all five tasks are scheduled to start. Because only three tasks can run in parallel, two tasks (4,5) are delayed further to $t4$. Furthermore, during execution, host H1 fails, interrupting Task 1, causing it to be rescheduled at $t4$.

In this example, all five tasks are delayed by the scheduling policy. Using the analytical model discussed previously, it would be concluded that three tasks (1-3) are successfully shifted from high to low carbon periods, and two tasks (4,5) from medium to low carbon periods. In reality, however, only two tasks (2, 3) are shifted successfully, with the other tasks (1, 4, 5) being executed in a high-carbon period. Although analytical models can be extended to model also failures, they still struggle to capture the actual behavior of a specific workload on a given datacenter, especially when multiple




Fig. 2: Example of a scenario challenging for analytical models to evaluate correctly.

operational phenomena and sustainability techniques interact. The composable simulator approach we propose addresses these challenges. In STEAM, each sustainability technique can be modeled by the user independently. They will then interact automatically with other techniques and datacenter dynamics during simulation, providing an accurate and quantitative view of their combined impact.

IV. STEAM DESIGN

STEAM is a discrete event simulator, based on the OpenDC framework [14], designed to evaluate the efficacy of sustainability techniques. It significantly extends OpenDC by adding support for operational and embodied carbon modeling, sustainability techniques, and GPUs. Notably, STEAM proposes a major redesign of OpenDC’s core architecture to enable a composable and extensible simulator in which techniques can be independently developed and freely combined based on users’ needs (see Section IV-B). This flexibility makes STEAM a powerful and versatile tool for holistic and quantitative research in sustainable datacenter design and operations.

STEAM’s architecture is depicted in Figure 3. STEAM takes a datacenter topology, workload trace, and other experiment configurations as input ①. STEAM comes with ready-to-use statistical models ② the user can attach to components in the datacenter topology for use in simulation, such as power models for datacenter components (e.g., CPUs and GPUs), carbon intensity traces, and failure models. It also provides resource management primitives ③ the user can configure to manage scheduling, battery management, and temporal shifting in the datacenter. Based on the user configuration, the datacenter components, respective statistical models, and resource managers are assembled into a component graph ④. The component graph is then executed by the event-based executor ⑤. The metrics collector ⑥ samples the simulation at periodic intervals and collects data for analysis. The user can select which metrics to export and at which granularity.

A. Overview of STEAM Components

The user specifies the workload trace, the datacenter topology, carbon intensity traces, failure models, and configuration parameters such as a random seed for reproducibility. STEAM supports many common distributed systems workload formats.

The *topology* (part of ①) describes the available hardware in a datacenter. The user defines host types by providing CPUs, GPUs, and memory specifications. A user can provide

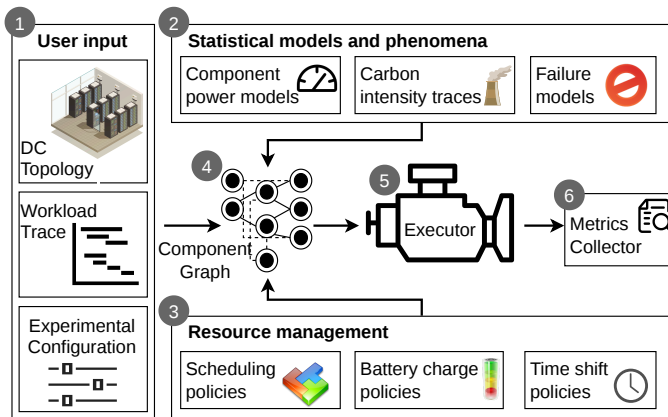


Fig. 3: STEAM system architecture.

different host types and their number. The user defines the power sources the hosts are connected to. Carbon traces can be connected to power sources to measure operational carbon. Users can add batteries to the datacenters. Batteries can be charged and discharged based on battery policies. Particularly relevant to this work are the CPU and GPU statistical power models (part of ②), which convert component utilization (in %) into power demand (in Watts). The user can calibrate and use built-in power models, including linear, square-root, square, and cubic. The user can also implement custom statistical power models.

Resource managers ③ and operational techniques are needed to execute a workload on a datacenter. STEAM provides resource managers and operational techniques that support different scheduling algorithms and policies. The user can associate a resource manager with each resource, such as power, battery, compute, and memory. The user can specify scheduling policies in the experiment configuration (part of ①) using custom filters and weighers.

In STEAM, we implement various sustainability-focused resource management and operational techniques. These include threshold-based battery charge and discharge policies inspired by similar policies for batteries [11], [12], as well as load shifting strategies [20]. We implement task-delaying policies to shift tasks into lower-carbon-intensity periods, and a task-stopping mechanism that stops execution during high-carbon periods, resuming once greener energy becomes available. Users can easily configure and customize these policies with minimal configuration changes.

B. Component Graph, Events, and Composability

The goal of STEAM is to support any operational technique (including sustainability strategies) with minimal implementation effort, regardless of which other techniques are enabled. To achieve this, we use a component graph, inspired by reactive programming [21] at the core of our simulator. This architecture is inherently composable and loosely coupled, enabling the easy integration of new techniques and models.

We define a *component* as any object in the simulator. Components can represent physical datacenter entities, such

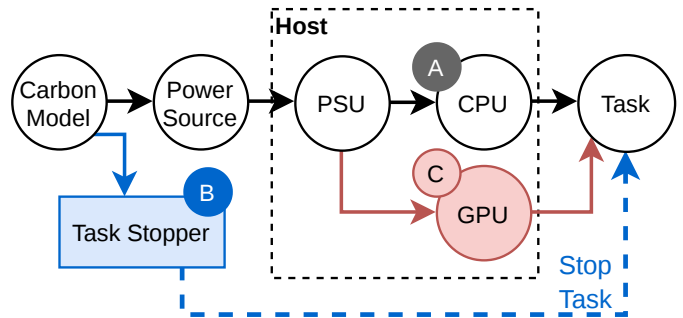


Fig. 4: Example component graph depicting events connecting two new components: task stopper and GPU.

as CPUs and GPUs, or abstract objects, like resource managers and fault injectors. Directed edges connect components, defining supplier-consumer relationships. For example, a CPU consumes the power supplied by a Power Supply Unit (PSU).

Simulation is driven by *events* that change the state of components. An example of an event is a task changing its CPU requirements. Depending on the event, connected nodes might also be affected. When a task updates its CPU requirements, the connected CPU will change its power requirements according to its power model. This creates a cascading effect that could potentially traverse the whole component graph.

The component graph approach of STEAM ensures that adding new features only requires changes to directly connected components. This promotes loose-coupling, ensuring that unrelated components remain independent, and allows individual components to be activated or disabled.

Figure 4 illustrates this composability. We start with a simple component graph ① connecting a power source to the PSU, the PSU to the CPU, and the CPU to the workload task. The carbon model provides the power source with the current carbon intensity. We extend the component graph by adding a task stopper ②. The task stopper monitors the carbon intensity of the power source and stops tasks when the carbon intensity exceeds a threshold to reduce the datacenter’s operational carbon. Next, we add a GPU ③ to the host. A GPU is connected to only two nodes; it demands power from the PSU and provides computation for the task. The GPU is inserted into the component graph and connected to the PSU and the task. Besides the directly connected nodes, no other nodes or edges have to be updated.

We use this composability to implement three sustainability techniques: hardware scaling, batteries, and temporal shifting. This required us to implement components to start and stop tasks based on carbon intensity, batteries, and a battery resource manager. All these techniques are independent and can be activated by the user with a simple configuration change. This flexibility allows STEAM to analyze techniques independently and in conjunction with each other or other simulator features, including fault injection, checkpointing, restarts, and advanced scheduling mechanisms.

V. EXPERIMENTAL SETUP

The following sections discuss the setup used for all experiments and evaluations. To ensure generalizability, we evaluate sustainability techniques across diverse workloads, hardware configurations, and carbon-intensity traces. All techniques, costs, and parameters are configurable in STEAM.

A. Workload and Datacenter setup

1) *Workload Traces*: The characteristics of a workload can significantly influence the effectiveness of sustainability techniques. To capture this, we experiment with three publicly available workloads from different datacenter types: Surf LISA [22], CINECA Marconi M100 [23], and Google Borg [24]. The workloads differ in duration, task frequency, and task execution time (see Table I). The Lisa workload is 5 months long and was collected in 2019 by SURF, an IT cooperative of Dutch education and research institutions⁴. The Marconi M100 workload is collected by CINECA using their EXAMON monitoring framework [23]. Over 90% of the tasks in the Marconi trace utilize GPUs. We use a subset of the dataset, recorded in October 2022. Finally, we use a large-scale industrial workload based on the Google Borg infrastructure [24]. We use a subset of the dataset containing all tasks executed on cell-a located in New York⁵. Tasks can not run on more than one server, but multiple tasks can share the same server. Tasks can be CPU and GPU-based.

2) *Datacenter setup*: The hardware in a datacenter can significantly impact its performance and sustainability. Each experiment has a corresponding datacenter setup, reflecting the hardware (CPU, GPU, RAM) used when collecting the traces (see Table II). Unless stated otherwise, workloads are run on these datacenters.

B. Sustainability techniques

1) *Batteries*: Batteries are used to reduce a datacenter’s carbon emissions (see Section II-C). The policy we use to determine when to charge and discharge is similar to previous works [11], [12] and is based on a carbon threshold: if the current carbon intensity is below the threshold, the battery is charged; if it is above the threshold, the battery is discharged. The threshold is equal to the rolling mean of the past week’s carbon intensity. As an optimization, batteries wait until the carbon intensity stops decreasing before charging.

The *battery capacity* can significantly impact its effectiveness: larger battery capacities allow the datacenter to run longer on battery power, but will also increase the embodied carbon cost. We evaluate multiple capacities and report the configuration achieving the highest total emissions reduction.

We assume that the *charging speed* of batteries increases linearly with their capacity due to parallel charging. The charging speed of batteries scales with a factor of 3 kW/kWh, similar to the charging speed of a Tesla Model 3 using a DC

TABLE I: Workload Traces used in the evaluation. NoT: number of tasks. ATD: Average task duration in HH:MM:ss.

Workload	Start	Runtime	NoT	ATD
Surf [22]	7-2022	124 days	194,917	01:49:38
Marconi [23]	9-2022	30 days	73,882	06:20:12
Borg [24]	5-2019	31 days	14,867,803	02:01:51

TABLE II: Default hardware setup per workload. NoH: Number of Hosts, EC: Host Embodied carbon [kgCO₂-eq].

Source	NoH	Host resources	EC
Surf [22]	277	128 GB RAM, 16-core 2.1GHz Intel Xeon Silver 4110	1,022
Marconi [23]	972	196 GB RAM, 2x24-core 2.1GHz Intel Xeon 8160, 4xNVIDIA VoltaV100 16GB	3,542
Borg [24]	1,534	Obfuscated (see Section VIII)	2,250

charger⁶. In Section VI-B2, we further discuss the impact of battery capacity and charging speed.

2) *Temporal Shifting*: Temporal shifting is the process of delaying tasks to better align with the availability of low-carbon energy. Many past works have used “oracle” models to determine when to schedule a task optimally. In this work, we use a task-scheduling policy similar to [20]. A task is scheduled when the current carbon intensity is below the carbon threshold. Otherwise, they are delayed. The carbon threshold is dynamic and is set to the 35th percentile of next week’s carbon forecast. All tasks have a maximum delay time of 24 hours and are interruptable. When the maximum delay is reached, a regular FIFO scheduler is used.

C. Carbon Emissions

1) *Operational Carbon*: Operational carbon, emitted when using energy, is determined by multiplying energy usage by the power source’s carbon intensity (amount of carbon emissions per unit of energy). The power draw of CPUs and GPUs is modeled using a square root and linear model, respectively, following approaches adopted in prior state-of-the-art work [25]. Historical carbon intensity traces collected from the ElectricityMaps portal [26] are used to determine the carbon intensity during a workload. To capture a broad range of scenarios, we use 158 carbon traces from different regions collected in 2021-2024, which display diverse carbon intensity profiles and coefficients of variation. Some experiments are only run in a single region. In those cases, we use the datacenter’s region. These are the Netherlands, Italy-North, and New York ISO for Surf, Marconi, and Borg, respectively.

2) *Embodied Carbon*: The embodied carbon of a workload is estimated by multiplying the fraction of the hardware’s total lifetime used by the workload’s runtime with the total embodied cost [16]. In this work, the embodied cost of a datacenter is based on the host and battery.

For *hosts*, we assumed a 5-years lifespan. A host’s embodied carbon cost is based on its components. Unfortunately,

⁴<https://www.surf.nl/en>

⁵https://console.cloud.google.com/storage/browser/clusterdata_2019_a;tab=objects?authuser=0&inv=AbuSoA

⁶<https://evbox.com/en/electric-cars/tesla/tesla-model-3>




Fig. 5: Impact of horizontal scaling on total carbon emissions (operational + embodied) and performance, quantified by SLA violations (tasks not scheduled within 24 hours of submission). For all metrics, lower is better. Gray and red lines indicate the original and required scale for <1% SLA violations. Bottom row shows the impact when the datacenter is exposed to failures.

manufacturers rarely disclose their components’ embodied carbon. Therefore, we estimated hosts’ embodied carbon using EC2 instance equivalents from a dataset collected by Teads Engineering [27]. We match Surf and Marconi hardware to the *a1.4xlarge* and *p3.8xlarge* instances, resulting in embodied carbon costs of 1,022 and 3,542 kgCO₂-eq, respectively. For Borg, we do not have hardware specifics (see Section VIII). Instead, we estimate the embodied carbon to contribute 40% of the total carbon emissions, a common distribution for modern datacenters [16]. For Borg, this is 2,250 kg CO₂-eq.

For *batteries*, we assumed a 10-years lifespan. The embodied carbon of a battery depends on its manufacturing process and location, typically ranging from 30 to 500 kg CO₂-eq/kWh [28]. For our experiments, we use an embodied carbon cost of 100 kg CO₂-eq/kWh.

VI. ANALYSIS OF SUSTAINABILITY TECHNIQUES

In the following sections, we demonstrate STEAM capabilities by investigating the impact of horizontal scaling, batteries, and temporal shifting, on datacenter carbon footprint.

A. Horizontal Scaling

We use STEAM to investigate the effects of horizontally down-scaling datacenters on their performance and sustainability. We measure performance using *Service Level Agreement* (SLA) violations: a task meets the SLA if it is completed within 24 hours of its expected completion time. We define a datacenter service as acceptable if it has fewer than 1% SLA violations. We aim to find the smallest datacenter configuration that still meets this threshold. Since horizontal scaling is independent of carbon intensity, we run experiments in a single region determined by the datacenter’s location.

1) *Performance*: Figure 5 shows the impact of horizontal scaling on system performance. Plots A, B, and C show the effects on systems without failures. We define a datacenter as *underprovisioned* if it has more than 1% SLA violations and *overprovisioned* otherwise.

Under-provisioning: When a datacenter does not have sufficient hosts to run incoming tasks, they are delayed until a host becomes available. Since tasks arrive continuously, delays

accumulate, leading to large SLA violations that approach 100%. Adding more machines does not always reduce SLA violations significantly. For example, scaling the Surf datacenter from 25 to 100 machines reduces SLA violations by only 5% (see Figure 5A). Even with 100 machines, the tasks still accumulate. Interestingly, the 4x increase in scale does not significantly increase operational carbon emissions, as the increased energy usage is offset by a shorter runtime.

Over-provisioning: Datacenters must have sufficient hosts to absorb workload peaks and avoid SLA violations. We find that 200, 750, and 900 machines are sufficient for the Surf, Marconi, and Borg workloads, respectively. Adding machines yields minimal SLA improvement but increases carbon emissions due to higher idle power draw and embodied carbon. Reducing the number of machines results in an average carbon reduction of 14%, 12%, and 35%, respectively.

2) *Failures*: Failures (e.g., hardware or task failures) can severely disrupt service quality. To mitigate these risks, operators typically over-provision resources. We use STEAM to simulate the impact of failures of the operations of a datacenter. For this analysis, we configured STEAM to inject machine failures using a trace from the Cloud Uptime Archive, derived from incident reports for Facebook Messenger [29]. To lessen the impact of failures, we use checkpointing, which periodically saves task snapshots. In case of a failure, the task resumes from its latest snapshot rather than starting over. We use a checkpointing frequency of 1 hour similar to [29].

Plots D, E, and F in Figure 5 show the impact of horizontal scaling on a datacenter subjected to failures for the Surf, Marconi, and Borg workloads. The Surf workload is minimally affected, with the scale required for <1% SLA violations staying at 200. Marconi is moderately impacted: the number of hosts needed to get below 1% SLA violations is increased from 750 to 850. This reduces the carbon reduction for Marconi from 12% to 5%. The Borg workload is the most affected: with failures, it is not able to achieve <1%, with the default scale having 12.5% SLA violations. This could mean that Borg has a lower failure rate than we used or uses a different SLA.

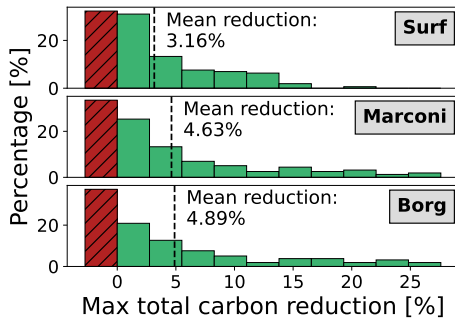


Fig. 6: Total carbon emission reduction using batteries in 158 carbon regions on three workloads. Red bars mark all regions where batteries *increased* emission.




Fig. 7: The impact of battery capacity on the effectiveness of using batteries to reduce carbon emissions.




Fig. 8: The impact of battery charging speed (on the effectiveness of using batteries to reduce carbon emissions).

Key Takeaway 1: Down-scaling over-provisioned datacenters can reduce total carbon-emissions by up to 35% without impacting service quality. *However*, when accounting for operational phenomena such as failures, the achievable reduction falls to 14%. STEAM can accurately quantify these trade-offs to help optimize designs and operations.

B. Batteries

In this section, we use STEAM to assess the impact of using batteries on a datacenter’s carbon footprint.

1) *Battery Performance:* Figure 6 shows the total carbon reduction achieved using batteries in 158 carbon regions while running the Surf, Marconi, and Borg workloads on their respective datacenters. On average, batteries reduced total carbon by 3.16%, 4.63%, and 4.89% for the three workloads, respectively. However, there is a significant difference in battery effectiveness across the carbon regions. Using batteries significantly decreases carbon emissions ($\geq 5\%$) in 24%, 29%, and 33% of the regions for the Surf, Marconi, and Borg workloads, respectively, with the highest reduction being 21%, 27%, 28%. However, in 32%, 34%, and 37% of the regions, carbon emissions actually increased, caused by the additional embodied carbon outweighing the reduced operational carbon.

Key Takeaway 2: Batteries’ effectiveness depends on the datacenter, workload, and energy source. STEAM shows that batteries reduce carbon emissions significantly ($> 5\%$) in 136 out of the 474 scenarios we evaluated. *However*, when the variation in regional carbon intensity is low, batteries can actually increase total carbon emissions.

2) *Battery Capacity And Charging Speed:* Using large fast-charging batteries can be expensive and introduce large power draw spikes, which can stress the power delivery infrastructure. Figure 9A shows that using batteries can increase the peak power draw by 8x when running Surf, with similar results observed for the other workloads.

Capacity: Figure 7 shows the impact of battery capacity on carbon emission reductions for the Surf workload. Increasing




Fig. 9: Peak power draw [kW] (A), mean task delay [h] (B), and total energy usage [MWh] (C) for the Surf workload.

battery capacity reduces the operational carbon, but with diminishing returns. Batteries reduce operational carbon by providing energy when carbon intensity is high. If a battery can handle most periods of high carbon, increasing its size has little to no impact. In contrast, the embodied carbon of a battery increases linearly with its capacity. This creates a sweet spot in which operational carbon reduction exceeds the added embodied carbon, which can be identified using STEAM.

Charging Speed: Figure 8 shows the impact of battery charging speed on carbon emission reductions for the Surf workload. Increasing charging speed increases effectiveness, but again, with diminishing returns. We observed that at a charging rate of 0.5 kW/kWh, 95% of the mean performance has already been reached. This is 6 times smaller than the charging rate of a Tesla Model 3 DC charger.

Key Takeaway 3: Optimally using batteries to reduce carbon emissions requires fast-charging batteries with high capacity, which are expensive and can increase the max power-draw up to 8x. *However*, charging speed can be reduced by 50% while retaining 95% of the expected carbon reduction.

3) *Battery Embodied Carbon:* A major downside of batteries is the carbon cost associated with manufacturing. The embodied carbon of a battery depends on how and where it is manufactured [28]. Figure 10 shows the effectiveness

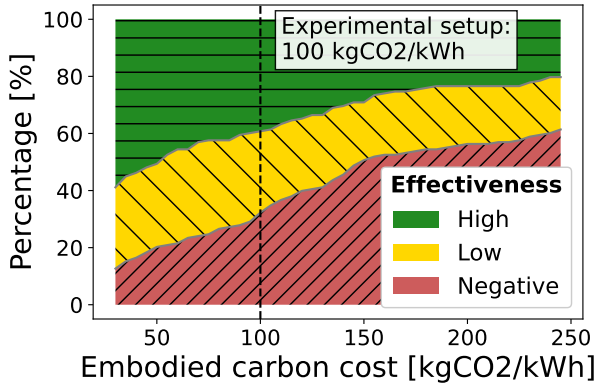


Fig. 10: Percentage of regions where the effectiveness of batteries to reduce total carbon emissions is high ($>5\%$), low ($0\%-5\%$), or negative ($<0\%$). The black line indicates the embodied cost ($100 \text{ kgCO}_2/\text{kWh}$) used in other experiments.

of batteries for embodied carbon costs between 30 and 250 $\text{kgCO}_2\text{-eq}/\text{kWh}$ when running Surf. As expected, reducing the carbon cost improves the impact of using batteries. At $30 \text{ kgCO}_2/\text{kWh}$, the number of regions in which batteries can reduce overall carbon by $>5\%$ increases from 39% to 59%. However, for 13% of the regions, batteries still increase overall carbon emissions. In these regions, operational carbon reduction is minimal, likely due to a low coefficient of variation in carbon intensity. Batteries are not viable for these regions with the current technology.

C. Temporal Shifting

Temporal shifting reschedules tasks to periods with low-carbon intensity (see Section V-B2). Figure 11 depicts the distribution of operational carbon reduction achieved using temporal shifting in 158 carbon traces while running the Surf, Marconi, and Borg workloads. When using temporal shifting, operational carbon is reduced on average by 2.06%, 0.74%, and 2.85% for the three workloads. The effectiveness of temporal shifting varies significantly across workloads and carbon regions, achieving a $>5\%$ reduction in the best scenarios, whereas it has little impact in others. Using temporal shifting introduces an average delay of 14 hours on the Surf trace (see Figure 9B), which is considerably less than the allowed 24 hours. The carbon reduction achieved by temporal shifting in this work is significantly lower than that reported in previous work, which reported reductions up to 40% [8]. We indicate the following reasons for this significant difference: *i) Policy*: Most previous work used oracles that could determine the optimal time to execute a task [8]. However, these models are difficult (or even impossible) to use in real systems. On the contrary, the policy used in this work can easily be adapted to most datacenters. More advanced scheduling policies could improve the effectiveness, but would still not match the oracle models. *ii) Idle Hosts*: Previous work evaluated the impact of temporal shifting by only looking at the tasks [8], [20] (see Section III). This ignores the idle power draw of hosts.

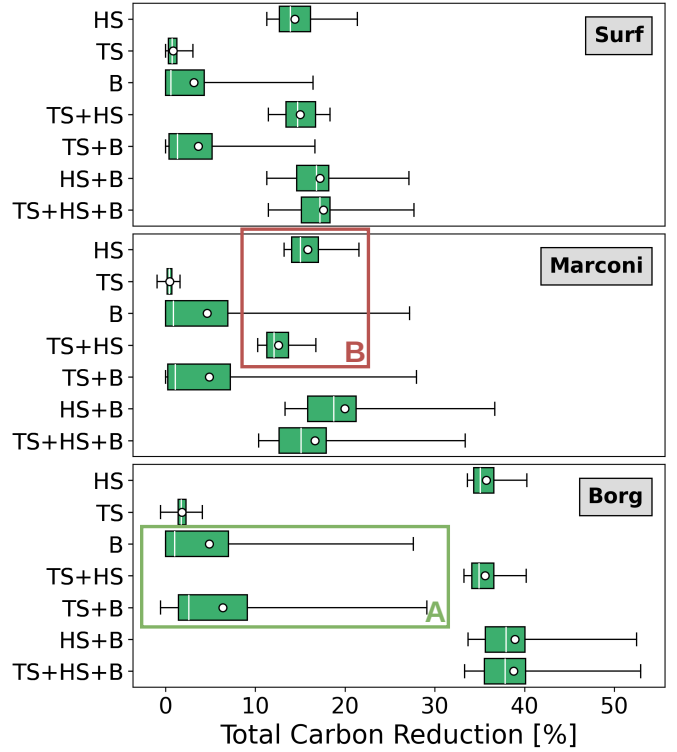


Fig. 11: Carbon emissions reduction of temporal shifting (TS), horizontal scaling (HS), and batteries (B) individually and in combination for three workloads in 158 regions. The white dot and line indicate the mean and median. Whiskers represent outliers that are at least 3 times the mean.

iii) Compounding Tasks: Previous work did not account for tasks stacking when multiple tasks are delayed to the same timestamp (see Section III). STEAM automatically accounts for such operational phenomena, quantifying their impact on performance and carbon emissions.

Key Takeaway 4: Temporal shifting reduces carbon emissions by over 2% in 210 of the 474 considered scenarios. These results are significantly lower than previously reported at the state of the art; STEAM considers the impact of complex dynamics such as compounding tasks.

VII. ANALYSIS OF COMBINED SUSTAINABILITY TECHNIQUES

Prior works have primarily focused on using single sustainability techniques. In contrast, the composable design of STEAM (Section IV) enables assessing the impact of combining multiple techniques. In the following, we explore the effectiveness of compositing the three sustainability techniques discussed so far. An extended version of this paper, including additional results, can be found at <https://arxiv.org/abs/2603.12381>.

Figure 11 shows the distribution of total carbon reduction achieved by individual and combined sustainability techniques.




Fig. 12: The optimal battery size in the range [200, 500] kWh for the combinations of sustainability techniques across 158 regions. The white dot and line indicate the mean and median. Whiskers represent outliers that are at least 3 times the mean.

As expected, in some cases, combining sustainability techniques can reduce carbon emissions more than when used individually. For example, combining batteries with temporal shifting when running Borg increases the average reduction from 4.9% (batteries alone) to 6.4% (Figure 11, box A). This is slightly less than the sum of the individual reductions. However, there are also combinations of techniques, workloads, and locations that have no impact on effectiveness or even reduce it. For example, on Marconi, combining horizontal scaling and temporal shifting lowers the reduction from 15% to 12.5% compared to when using only horizontal scaling (Figure 11, box B).

Key Takeaway 5: STEAM helps systematic analysis of how sustainability techniques interact. Some combinations, such as batteries and horizontal scaling, are compatible, achieving an average carbon reduction close to the sum of their individual effects. *However*, combining other techniques, such as temporal shifting and horizontal scaling, can lead to lower average carbon-reduction when combined. The effectiveness of any combination also depends on the workload characteristics, the datacenter, and its location.

Although combining sustainability may seem straightforward, understanding its implications is not. Further investigation shows that combining techniques optimally could require different setups than when used individually. For example, the optimal battery size decreases when combined with other techniques. Figure 12 shows the optimal battery size when running the Surf workload across different regions. The average optimal battery size decreases from 315 kWh to 272 kWh when all techniques are combined.

Key Takeaway 6: Combining techniques can impact their optimal setup. For instance, the optimal battery size is reduced when combined with temporal shifting.

The findings in this work show that the simulation-based approach of STEAM and its composable design are necessary to accurately assess the impact of sustainability techniques, both individually and in combination, while capturing their interactions with datacenter operational phenomena.

VIII. CALIBRATION, VALIDATION, AND PERFORMANCE

STEAM uses power models to estimate the power draw of hardware components such as the CPU and GPU, based on parameters such as max and idle power draw. These models are calibrated through a two-step process. 1) We *select initial values* for Surf and Marconi based on information provided by manufacturers, indicating expected power draw. 2) We further *calibrate* the Surf and Marconi datacenters on a subset of the existing data (1 month and 1 day, respectively) The accuracy of STEAM is evaluated by comparing the estimated power draw against real-world measurements. After calibration, STEAM achieves a Mean Average Percentage Error of 5.5% and 3.4% for Surf and Marconi, respectively. These results highlight STEAM ability to realistically estimate energy usage, making it a reliable tool for sustainability analysis for both CPU-(Surf) and GPU-based (Marconi) workloads.

Calibrating Borg was less straightforward, as the trace is highly obfuscated to protect confidential or proprietary information [30]. Critical information, such as hardware specification, resource utilization, or power usage, is hashed or converted to a normalized metric. However, not everything is unknown. First, we know the number of nodes is 1534. Second, resources have an average utilization of 60% [24]. Finally, the trace we use is based on a single PDU, with a max power draw of 2 MW [31]. We convert the normalized utilization number to match the average resource utilization of 60% and the max power draw of 2MW. The power models are calibrated to match these derived values. Given this calibration approach, the simulation is consistent with the trace under the assumed conditions but cannot be independently validated.

The simulation time of STEAM depends on the workload size, the datacenter hardware, and the techniques used. As a reference, a single simulation of Surf, Marconi, and Borg took 1, 2, and 10 minutes on average. Simulations are independent and can be executed in parallel. All experiments were executed on a 20-node cluster, each equipped with a single 24-core 2.8 GHz AMD EPYC-2 (Rome) 7402P CPU and 128 GB of RAM. In this paper, we simulated over 2,787 years of datacenter operations using just 60 hours of actual compute time. During these simulations, 83 billion tasks were executed, corresponding to 170 billion simulated compute hours.

IX. RELATED WORK

In past years, many techniques have been proposed to improve the carbon footprint of a datacenter. This work quantifies the impact of three such techniques: horizontal scaling [9], [10], batteries [7], [11], [12], and load shifting [7], [8].

Lin et al. propose dynamic "right-sizing" in which machines are turned off during low-load periods, showing that this can reduce energy usage by at least 15% [9]. Uddin et al. show that dynamic server consolidation can reduce energy usage by 50% [10], but do not quantify the impact on carbon emission. Sukprasert et al. use analytical models to determine the limitations of temporal shifting [8]. SHEME et al. use a finite-state machine to show that batteries can improve green

energy coverage [11]. However, their work is limited to three countries and a single datacenter with constant power.

Carbon Explorer uses analytical modeling to assess the impact of sustainability techniques on achieving NetZero carbon emissions [7]. However, it does not model hosts or tasks, but only the total datacenter energy usage. This makes it hard to understand task-level effects such as task stacking or SLA violations (see Section III). Additionally, the tool does not specify how a datacenter should implement temporal shifting. Instead, it simply shifts the total power draw to highlight the potential carbon reductions. Due to this coarse granularity, only the analysis in Figure 6 can be replicated by Carbon Explorer without requiring significant modifications to the tool design and implementation.

We are not the first to use simulation to explore datacenters and their carbon footprint [32]. STEAM builds on and significantly extends OpenDC [14], an open-source datacenter simulator used in prior works on sustainability [13]. Although inspired by OpenDC’s core concepts, STEAM required a major architectural redesign to: *i*) enable modularity and composability, crucial for modeling and analyzing multiple sustainability techniques and their interactions; *ii*) introduce sustainability techniques and carbon footprint modeling; and *iii*) support GPU-enabled datacenters. This flexibility and novel features are core contributions of this work, allowing service providers and sustainability researchers to explore how different techniques and their interactions impact modern datacenters’ design and operations. Popular alternatives to OpenDC are CloudSim [33], and SimGrid [34], two datacenter simulators that have spawned many more specialized simulators. Using these specialized simulators in combination is difficult, limiting their use. In contrast, STEAM allows users to mix and match all implemented components and techniques.

X. THREADS TO VALIDITY

We recognize and acknowledge that our analysis is not immune to threats to validity, which we outline below to ensure transparency and guide future work.

The *policies used* for batteries and temporal shifting are practical and easy for most datacenter operators to implement. Although this improves applicability, we acknowledge that more advanced strategies may yield better results. OpenDC STEAM enables convenient exploration of different policies, offering a convenient framework for future research. When using horizontal scaling, we used SLA violations to determine the optimal scale. STEAM allows users to define different SLAs (e.g., with stricter completion-time constraints), which could lead to different optimal scales and carbon reductions.

We computed *operational carbon* using historical carbon traces. However, large datacenters can have special contracts with energy providers, resulting in a different energy mix. STEAM can be used to analyze the impact of such contracts.

In this work, energy usage is determined by the computational load of CPUs and GPUs. We acknowledge that other components of datacenters might significantly contribute

to carbon costs, such as storage, networking, and cooling costs. Future work can easily incorporate models for these components, further expanding STEAM capabilities.

We estimate *embodied carbon* based on prior industry reports [16], [27]. Accurate embodied carbon estimation is an active area of research, and different models can yield widely different carbon cost estimates [35]. As more accurate cost models or transparent manufacturer reports become available, these can be easily integrated into STEAM.

We used CO2 emissions as a proxy to estimate the sustainability impact of different techniques and design choices. Although CO2 is the most common metric of sustainability, a more comprehensive picture of environmental impact could be obtained by extending STEAM to model other aspects, such as water wastage.

XI. CONCLUSION

We introduced OpenDC STEAM, a trace-based, composable simulation tool to quantify the impact of sustainability techniques on datacenters’ design and operations. STEAM can provide insights into sustainability aspects, such as carbon emissions, as well as performance, such as scheduling delays or peak power draw. We used STEAM to evaluate three representative sustainability techniques: horizontal scaling, temporal shifting, and batteries. We showed that all three techniques can reduce carbon emissions when combined, but their effectiveness depends on the workload, datacenter configuration, and energy sources. They can also introduce trade-offs: for example, our analysis shows that temporal shifting is less effective than previously reported, and that combining techniques creates new operational challenges, such as power spikes, that must be carefully navigated. This is the first study to systematically quantify both the individual and combined effects of sustainability strategies across diverse datacenters, workloads, and carbon regions. These results highlight the importance of holistic, quantitative tools such as STEAM in guiding sustainable datacenter design and operations.

We are continuing the development of OpenDC-STEAM, adding more sustainability techniques and additional metrics, such as water consumption and monetary costs. Furthermore, we are exploring the use of OpenDC-STEAM’s simulation capabilities to support large-scale digital twin projects.

ACKNOWLEDGEMENT

This work was partially supported by EU MSCA CloudStars (g.a. 101086248) and EU Horizon Graph Massivizer (g.a. 101093202). This research is partly supported by a National Growth Fund through the Dutch 6G flagship project Future Network Services.

REFERENCES

- [1] L. Belkhir and A. Elmeligi, “Assessing ICT global emissions footprint: Trends to 2040 & recommendations,” *J. Cleaner Production*, 2018.
- [2] C. Freitag, M. Berners-Lee, K. Widdicks, B. Knowles, G. S. Blair, and A. Friday, “The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations,” *Patterns*, vol. 2, no. 9, 2021.
- [3] ACM, “Computing and climate change,” *ACM TechBriefs*, 2021.

- [4] J. D. Wilson and Z. Zimmerman, "The era of flat power demand is over," *Grid Strategies*, 2023.
- [5] B. Knowles, K. Widdicks, G. Blair, M. Berners-Lee, and A. Friday, "Our house is on fire: The climate emergency and computing's responsibility," *Commun. ACM*, vol. 65, no. 6, May 2022.
- [6] A. A. Chien, "Computing's grand challenge for sustainability," *Commun. ACM*, vol. 65, no. 10, 2022.
- [7] B. Acun, B. Lee, F. Kazhamiaka, K. Maeng, U. Gupta, M. Chakkaravarthy, D. Brooks, and C.-J. Wu, "Carbon explorer: A holistic framework for designing carbon aware datacenters," in *ASPLOS*, vol. 2, 2023.
- [8] T. Sukprasert, A. Souza, N. Bashir, D. Irwin, and P. Shenoy, "On the limitations of carbon-aware temporal and spatial workload shifting in the cloud," in *EuroSys*, 2024.
- [9] M. Lin, A. Wierman, L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," *Trans. Netw.*, vol. 21, 2012.
- [10] M. Uddin and A. A. Rahman, "Server consolidation: An approach to make data centers energy efficient and green," *IJSER*, 2010.
- [11] E. Sheme, S. Lafond, D. Minarolli, E. K. Meçe, and S. Holmbacka, "Battery size impact in green coverage of datacenters powered by renewable energy: A latitude comparison," in *EIDWT*, 2018.
- [12] Y. Li, A.-C. Orgerie, and J.-M. Menaud, "Balancing the use of batteries and opportunistic scheduling policies for maximizing renewable energy consumption in a cloud data center," in *PDP*, 2017.
- [13] D. Niewenhuis, S. Talluri, A. Iosup, and T. De Matteis, "Footprinter: Quantifying data center carbon footprint," in *ICPE*, 2024.
- [14] F. Mastenbroek *et al.*, "OpenDC 2.0: Convenient modeling and simulation of emerging technologies in cloud datacenters," in *CCGRID*, 2021.
- [15] A. Iosup, A. Uta, L. Versluis, G. Andreadis, E. V. Eyk, T. Hegeman, S. Talluri, V. van Beek, and L. Toader, "Massivizing computer systems: A vision to understand, design, and engineer computer ecosystems through and beyond modern distributed systems," *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*, pp. 1224–1237, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:3319705>
- [16] U. Gupta, M. Elgamal, G. Hills, G.-Y. Wei, H.-H. S. Lee, D. Brooks, and C.-J. Wu, "Act: designing sustainable computer systems with an architectural carbon modeling tool," in *ISCA*, 2022.
- [17] J. Lyu, J. Wang, K. Frost, C. Zhang, C. Irvine, E. Choukse, R. Fonseca, R. Bianchini, F. Kazhamiaka, and D. S. Berger, "Myths and misconceptions around reducing carbon embedded in cloud platforms," in *HotCarbon*, 2023.
- [18] A. Shehabi, S. J. Smith, E. Masanet, and J. Koomey, "Data center growth in the united states: decoupling the demand for services from electricity use," *Environmental Research Letters*, vol. 13, no. 12, 2018.
- [19] J. Glanz. (2012) Power, pollution and the internet. URL: www.nytimes.com/2012/09/23/technology/data-centers-waste-vast-amounts-of-energy-belying-industry-image.html.
- [20] R. Bostandoost, W. A. Hanafy, A. Lechowicz, N. Bashir, P. Shenoy, and M. Hajiesmaili, "Data-driven algorithm selection for carbon-aware scheduling," in *HotCarbon*, 2024.
- [21] E. Bainomugisha, A. L. Carreton, T. V. Cutsem, S. Mostinckx, and W. D. Meuter, "A survey on reactive programming," *CSUR*, vol. 45, 2013.
- [22] L. Versluis, M. Çetin, C. Greeven, K. Laursen, D. Podareanu, V. Co-dreanu, A. Uta, and A. Iosup, "Less is not more: We need rich datasets to explore," *FGCS*, vol. 142, 2023.
- [23] A. Borghesi, C. Di Santi, M. Molan, M. S. Ardebili, A. Mauri, M. Guarrasi, D. Galetti, M. Cestari, F. Barchi, L. Benini *et al.*, "M100 exadata: a data collection campaign on the cineca's marconi100 tier-0 supercomputer," *Scientific Data*, vol. 10, no. 1, 2023.
- [24] M. Tirmazi, A. Barker, N. Deng, M. E. Haque, Z. G. Qin, S. Hand, M. Harchol-Balter, and J. Wilkes, "Borg: the next generation," in *EuroSys*, 2020.
- [25] W. Brewer, M. Maiterth, V. Kumar, R. Wojda, S. Bouknight, J. Hines, W. Shin, S. Greenwood, D. Grant, W. Williams, and F. Wang, "A digital twin framework for liquid-cooled supercomputers as demonstrated at exascale," in *SC*, 2024.
- [26] Electricity Maps, "All countries 2024 1 hour carbon intensity data," 2025. [Online]. Available: <https://www.electricitymaps.com>
- [27] T. Engineering, "Building an AWS EC2 carbon emissions dataset," 2021. [Online]. Available: <https://medium.com/teads-engineering/building-an-aws-ec2-carbon-emissions-dataset-3f0fd76c98ac>
- [28] N. Lutsey and D. Hall. (2018) Effects of battery manufacturing on electric vehicle life-cycle greenhouse gas emissions.
- [29] S. Talluri, D. Niewenhuis, X. Chu, J. Kyselica, M. Cetin, A. Balgavy, and A. Iosup, "Cloud Uptime Archive: Open-access availability data of web, cloud, and gaming services," 2025.
- [30] C. Reiss, J. Wilkes, and J. L. Hellerstein, "Obfuscatory obscurantism: making workload traces of commercially-sensitive systems safe to release," in *NOMS*, 2012.
- [31] V. Sakalkar, V. Kontorinis, D. Landhuis, S. Li, D. De Ronde, T. Blooming, A. Ramesh, J. Kennedy, C. Malone, J. Clidaras *et al.*, "Data center power oversubscription with a medium voltage power plane and priority-aware capping," in *ASPLOS*, 2020.
- [32] J. Song, P. Zhu, Y. Zhang, and G. Yu, "Versatility or validity: A comprehensive review on simulation of datacenters powered by renewable energy mix," *FGCS*, vol. 136, 2022.
- [33] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. De Rose, and R. Buyya, "Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms," *Software: Practice and experience*, vol. 41, no. 1, 2011.
- [34] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter, "Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms," *JPDC*, vol. 74, no. 10, 2014.
- [35] A. Bhagavathula, L. Han, and U. Gupta, "Understanding the implications of uncertainty in embodied carbon models for sustainable computing," in *ICPE*, 2024.