

MIT Open Access Articles

*Toward Sustainable HPC: Carbon Footprint Estimation
and Environmental Implications of HPC Systems*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Li, Baolin, Basu Roy, Rohan, Wang, Daniel, Samsi, Siddharth, Gadepally, Vijay et al. 2023. "Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems."

As Published: <https://doi.org/10.1145/3581784.3607035>

Publisher: ACM|The International Conference for High Performance Computing, Networking, Storage and Analysis

Persistent URL: <https://hdl.handle.net/1721.1/153145>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems

Baolin Li
Northeastern University

Rohan Basu Roy
Northeastern University

Daniel Wang
Northeastern University

Siddharth Samsi
MIT Lincoln Laboratory

Vijay Gadepally
MIT Lincoln Laboratory

Devesh Tiwari
Northeastern University

ABSTRACT

The rapid growth in demand for HPC systems has led to a rise in carbon footprint, which requires urgent intervention. In this work, we present a comprehensive analysis of the carbon footprint of high-performance computing (HPC) systems, considering the carbon footprint during both the hardware manufacturing and system operational stages. Our work employs HPC hardware component carbon footprint modeling, regional carbon intensity analysis, and experimental characterization of the system life cycle to highlight the importance of quantifying the carbon footprint of HPC systems.

CCS CONCEPTS

• **Hardware** → **Impact on the environment.**

KEYWORDS

Carbon Footprint; Sustainability.

ACM Reference Format:

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 *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '23)*, November 12–17, 2023, Denver, CO, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3581784.3607035>

1 INTRODUCTION

High-performance computing (HPC) has become an essential tool in scientific research, engineering, and many other fields [1–3]. The demand for HPC has experienced rapid growth in recent years. According to the US International Trade Commission, in 2010, there was reportedly 1.2 trillion gigabytes of new data created globally. However, this number is estimated to increase to a staggering 175 trillion gigabytes by 2025 [4]. While the need for HPC resources is expanding, there is a downside to this growth. As more and more HPC systems are built and the size of these systems continues to increase [5], this leads to a rise in carbon footprint. For example,

the Summit supercomputer built in 2017 has a peak power consumption of 13 MW, while in 2021, the next-generational Frontier supercomputer has more than doubled the peak power to 29MW [6].

The carbon footprint of an HPC system cannot be accurately characterized by power consumption alone. The energy source used to power the system is also a key contributor to its environmental impact. Renewable energy sources such as hydropower and solar emit more than 20× less CO₂ than traditional energy sources like coal [7]. Furthermore, a significant amount of carbon emission is incurred during the manufacturing and packaging of the HPC system components [8, 9] before the system is deployed into operation. It is estimated that by 2030, datacenters and HPC systems may account for up to 8% of the worldwide emissions if not intervened [10]. As a result, major technology companies have been heavily invested in offsetting carbon emission in their datacenters [11–14], while more and more research efforts have started to focus on the carbon-friendliness of large-scale systems [15–18].

However, despite the current effort, there are still many unexplored and unanswered questions regarding HPC system sustainability and the carbon footprint of our HPC systems – esp. from a procurement and operational point of view. One key question is how to quantify the carbon footprint of an HPC system holistically where the carbon emissions from the hardware manufacturing to the end of the system life cycle have all been accounted for. **In this work, we analyze the carbon footprint from both the production and operational stages of an HPC system to address a series of unexplored investigations (referred to as Research Questions or RQs).** For each RQ, we conduct detailed modeling and characterization related to the question followed by visualization, discussion, and summarized takeaways. The current state of practice and available data, unfortunately, makes it very challenging to collect/build/analyze a standardized and portable model for carbon footprint accounting – this paper aims to raise awareness and calls for a joint effort between vendors and HPC facilities to address this challenge. **The highlights of our analysis include the following:**

We perform modeling and analysis of the carbon emission during the production stage of individual HPC hardware components and compare their contributions. We highlight that performance benchmarking alone is not sufficient to achieve environmental sustainability. The contribution from the embodied carbon footprint of memory and storage devices in HPC systems cannot be ignored – while storage system has been traditionally viewed as a secondary

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SC '23, November 12–17, 2023, Denver, CO, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0109-2/23/11...\$15.00
<https://doi.org/10.1145/3581784.3607035>

optimization goal for performance and top 500 rankings, the carbon embodied in hard drives and solid storage drives presents a serious challenge to sustainability. As the HPC centers prepare to serve more memory-intensive scientific applications, they should carefully consider the hidden carbon cost of these memory modules. Carbon-conscious HPC facilities should explicitly request the embodied carbon specifications for all components from the chip vendor as a part of their request for proposal (RFP).

Based on real-world carbon intensity data, we quantify the importance of regional carbon intensity when evaluating the operational carbon emissions of large-scale HPC systems. We demonstrate, as expected, the average carbon intensity of energy sources varies across different regions, but surprisingly, even the regions with the lowest carbon intensity can have significant hourly temporal variations – highlighting the significance of building cross-regional HPC systems. We identify a strong opportunity for systems researchers to design, develop, and deploy carbon-intensity-aware job schedulers to exploit temporal variations. Similar to core-hour accounting and budgeting, we recommend that HPC system operators and allocation programs should allocate a carbon budget to HPC users and some users could be prioritized to reduce their queue wait time if the carbon footprint of their jobs has been economical.

From a holistic point of view, we integrate the two key aspects of sustainability: the modeling of carbon emission during production and the characterization of carbon emission during operation to analyze the environmental impact throughout the HPC system life cycle. Our analysis reveals that hardware upgrades are often attractive from a performance point of view, but surprisingly, they can introduce significant embodied carbon that may not be offset quickly, especially if the center already runs on renewable energy sources or have low utilization. In the past, carbon-unaware system upgrades have not quantified and considered these factors. We provide a framework to help system practitioners make decisions on system upgrades based on hardware, workload, regional carbon intensity, performance, projected system lifetime, and user usage pattern.

Our end-to-end modeling and characterization of an HPC system can serve as a stepping stone for future research in the carbon footprint perspective of HPC systems. HPC system practitioners can utilize our analysis to gain a better understanding of how sustainable the current system is and the layout of next-generational systems. To promote the idea of sustainable HPC and encourage more research efforts toward carbon neutrality in the community, our framework is available at <https://zenodo.org/badge/latestdoi/659090328>.

We acknowledge that this work is naturally prone to certain expected threats to validity because of the limited (but evolving) availability of accurate and widely-accepted carbon footprint data for various HPC components. As the community’s awareness around carbon footprint increases, more robustness in modeling and characterization of carbon footprint will naturally evolve. Nevertheless, we hope that our study and analysis acts as the first step for raising carbon awareness in HPC practitioners and researchers to evaluate the carbon footprint of their systems, and help us grow as a community.

2 BACKGROUND AND METHODOLOGY

To analyze the carbon footprint of a system, we need to do so from two perspectives: the embodied carbon footprint and the operational carbon footprint. Embodied carbon refers to the carbon emissions associated with one-time expenditures like the production, transportation, and disposal of the materials and equipment used in HPC systems. In this work, we focus on modeling the production phase because the transportation and recycling of the component have been reported to be not dominant [7] and tend to be consistent across different generations of the system. On the other hand, the embodied carbon’s complementary is the operational carbon footprint, which refers to the day-to-day operation of a system. This includes the emissions associated with the electricity used to power the servers and other equipment, as well as the emissions associated with the cooling and ventilation systems used to keep the equipment within safe operating temperatures. In our work, we denote the overall carbon footprint, the embodied carbon footprint, and the operational carbon footprint as C_{total} , C_{em} , C_{op} , respectively. We can calculate the carbon footprint of a system using Eq. 1.

$$C_{total} = C_{em} + C_{op} \quad (1)$$

2.1 Embodied Carbon Footprint Modeling

The modeling of embodied carbon footprint is critical for the sustainability of semiconductor products [7, 8, 19]. We model the embodied carbon footprint of the HPC system components using principles similar to the ACT carbon modeling tool [7] and inspired by early studies in the chip-level carbon modeling including Greenchip and ACT modeling efforts [7, 20–22]; some of the carbon modeling efforts related to manufacturing and fabrication were started as early as 2010s. We provide conceptual coverage of embodied carbon footprint modeling, but more details of the manufacturing carbon modeling are also available in prior studies [7, 22]. Although these prior studies do not systematically model, characterize, derive insights, and identify challenges related to the carbon footprint of HPC systems.

Embodied carbon footprint is categorized into manufacturing carbon and packaging carbon. Manufacturing carbon refers to the emissions created from the creation of electronic components, such as transistors and resistors, from raw materials. Packaging carbon refers to the assembly of these components into functional chips and circuit boards. We summarize their relationships in Eq. 2.

$$C_{em} = \text{Manufacturing Carbon} + \text{Packaging Carbon} \quad (2)$$

Modeling the manufacturing carbon footprint of different types of components requires a different approach. To quantify the manufacturing carbon of processors (i.e., CPUs, GPUs), we follow a vendor-generic approach to collect the part-specific information on die area (A_{die}), fab carbon emission per unit area (FPA , related to fab location and lithography), emissions from chemicals and gases per unit area (GPA , related to lithography), emissions from raw materials (MPA , related to lithography), and fab yield ($Yield$, set to a constant value of 0.875, consistent with [7]) to estimate the amount of CO_2 emitted during the manufacturing process. We obtain such information from public product datasheets and sustainability reports. The manufacturing embodied carbon footprint of a processor

Table 1: Modeled individual components.

Type	Component	Part Name	Release Date
GPU	NVIDIA A100	NVIDIA A100 PCIe 40GB	May 2020
GPU	AMD MI250X	AMD INSTINCT MI250X	November 2021
GPU	NVIDIA V100	NVIDIA V100 SXM2 32GB	March 2018
CPU	AMD EPYC 7763	AMD EPYC 7763 CPU	March 2021
CPU	AMD EPYC 7742	AMD EPYC 7742 CPU	August 2019
CPU	Intel Xeon Gold 6240R	Intel Xeon Gold 6240R CPU	February 2020
DRAM	DRAM 64GB	SK Hynix 64GB DDR4	October 2020
SSD	SSD 3.2TB	Seagate Nytro 3530 3.2TB	October 2018
HDD	HDD 16TB	Seagate Exos x16 16TB	June 2019

unit (M_{proc} , unit: gCO_2) can be calculated using Eq. 3:

$$M_{\text{proc}} = \frac{(FPA + GPA + MPA) \cdot A_{\text{die}}}{\text{Yield}} \quad (3)$$

We estimate the manufacturing carbon of memory and storage devices (DRAM, SSD, HDD) carefully in a vendor-specific way because these components have distinctive internal architectures. We first determine the capacity (e.g., GB) of a memory/storage device, and use publicly available sustainability reports of the vendor to estimate how much carbon is emitted per GB of the memory/storage device manufactured, denoted as emission per capacity (*EPC*). We can calculate the manufacturing footprint of a memory/storage device ($M_{\text{m/s}}$, unit: gCO_2) as:

$$M_{\text{m/s}} = EPC \cdot \text{Capacity} \quad (4)$$

We estimate the packaging carbon by counting the number of integrated circuit (IC) packages on the component and using an average packaging overhead of 150 gCO_2 per IC according to industry reports [7, 23]. The packaging carbon (gCO_2) is:

$$\text{Packaging Carbon} = 150 \cdot \text{Number_of_ICs} \quad (5)$$

Note that Eq. 5 is only applicable to our processor and memory components because this is non-trivial for storage components. To mitigate this issue, we compile data from industry reports on the packaging-to-manufacturing ratio from the vendor website [24].

Modeling of individual components. We perform embodied carbon footprint modeling of individual components of an HPC system. We have listed the hardware components modeled in Table 1. In the table, we have selected the frequently deployed GPU and CPU processors in the Top500 list [25] from three major processor vendors: NVIDIA, AMD, and Intel. For example, the AMD MI250X GPU is available in the Frontier and LUMI supercomputers. We have chosen SK Hynix and Seagate as the DRAM and storage vendors due to the availability of sustainability reports from these vendors [24, 26]. Based on the vendor information, we have set the *EPC* of DRAM, SSD, and HDD to 65 gCO_2/GB , 6.21 gCO_2/GB , and 1.33 gCO_2/GB , respectively.

Modeling of HPC systems. We conducted our study for state-of-the-art high-performance computing (HPC) systems to analyze the embodied carbon contribution of each component. Specifically, we analyzed the Frontier, LUMI, and Perlmutter supercomputers, as listed in Table 2. We selected these systems because they are among the top 10 supercomputers in the Top500 list [25] and were built in recent years. For each component present in the system,

Table 2: Studied HPC Systems

System	Location	CPU & GPU	Cores	Year
Frontier [27]	Oak Ridge, TN United States	AMD EPYC 7763, AMD Instinct MI250X	8,730,112	2021
LUMI [28]	Kajaani, Finland	AMD EPYC 7763, AMD Instinct MI250X	2,220,288	2022
Perlmutter [29]	Berkeley, CA United States	AMD EPYC 7763, NVIDIA A100 SXM4	761,856	2021

Table 3: Independent system operators and regions.

Operator Name	Country of Operation	Region of Operation
Kansai (KN) [31]	Japan	Kansai Region
Tokyo (TK) [32]	Japan	Tokyo Region
Electricity System Operator (ESO) [33]	United Kingdom	Great Britain
California Independent System Operator (CISO) [34]	United States	California
Pennsylvania-New Jersey-Maryland Interconnection (PJM) [35]	United States	Mid-Atlantic US
Midcontinent Independent System Operator (MISO) [36]	United States, Canada	Midwest US, Manitoba
Electric Reliability Council of Texas (ERCOT) [37]	United States	Texas

We calculate the C_{em} according to Eq. 2 and multiply it by the total number of components available.

2.2 Operational Carbon Footprint Characterization

The operational carbon footprint is characterized when workloads are running on the system. It can be calculated using the carbon intensity of the power plant that powers the system (I_{sys} , unit: gCO_2/kWh) and the system's operational energy (E_{op} , unit: kWh).

$$C_{\text{op}} = I_{\text{sys}} \cdot E_{\text{op}} \quad (6)$$

Carbon intensity I_{sys} is a metric of how many grams of CO_2 are released into the atmosphere to produce a unit of energy. It depends on the fuel mix from the power plant. Higher carbon intensity means that the energy source generates more carbon emissions when producing the same amount of energy. Sustainable sources of energy such as wind or solar have a carbon intensity of less than 50 gCO_2/kWh while non-renewable sources like coal have a carbon intensity of more than 800 gCO_2/kWh . The operational energy (C_{op}) is the product of the IC component energy and the HPC system power-usage-effectiveness (PUE), which we set to a constant across all systems we characterize. In this work, we use the carbontracker [30] tool to measure a system's operational carbon footprint C_{op} while running certain benchmark suites.

Geographical carbon intensity. According to Eq. 6, the operational carbon footprint is proportional to the carbon intensity which highly depends on the time and the geographical location. In our work, we study the carbon intensity across different geographical regions to get a better understanding of the operational carbon footprint of a system. We have collected carbon intensity data from multiple power system operators across the globe as listed in Table 3. We obtain the ESO (UK) data from ESO's public

Table 4: Benchmarks performed and their respective models.

Benchmark	Models
Natural Language Processing (NLP)	BERT [40], DistilBERT [41], MPNet [42], RoBERTa [43], BART [44]
Computer Vision (Vision)	ResNet50 [45], ResNext50 [46], ShuffleNetV2 [47], VGG19 [48], ViT [49]
CANDLE [50, 51]	Combo, NT3, P1B1, ST1, TC1

Carbon Intensity API [38] and other regions' carbon intensity from Electricity Maps [39]. For all the regions in Table 3, we perform carbon intensity analysis on hourly data (year 2021).

Benchmarking workloads. To calculate the operational carbon, we also need to characterize the operational energy when the system is running. In this work, we perform a benchmarking study on real systems of different generations. By benchmarking on representative workloads, we are able to compare the operational carbon footprint of different systems and their respective performance. We have listed the details of the benchmark sets in Table 4. These benchmark sets represent the deep learning training workload across different research fields. We choose deep learning training because this is the target workload of today's GPU systems. The NLP benchmarks are provided by Huggingface, where we perform the question-answering task on various language models. The Vision benchmarks are provided by Pytorch, and we select models that have highly varied architecture (e.g., residual network in ResNet50, transformer in ViT) to perform image classification. The CANDLE benchmarks are provided by Argonne National Laboratory (ANL), with the initiative to address cancer research challenges with deep learning. We select five benchmarks from the Pilot1 class which represent problems in predicting drug response based on molecular features of tumor cells and drug descriptors.

Next, we provide a holistic carbon footprint analysis from three perspectives of HPC systems: (i) embodied carbon footprint, (ii) geographical carbon intensity variation, and (iii) total carbon footprint and performance benchmarking of real-world workloads.

3 EMBODIED CARBON ANALYSIS

First, we analyze the relative embodied carbon footprint for different HPC system components. We have listed the details of our modeled components in Table 1. These components appear frequently on the top 500 supercomputer list, represent a wide diversity in terms of vendors and time, and have their carbon-related specifications accessible or derivable.

RQ 1. *How does embodied carbon vary among different types of GPUs, CPUs, and memory/storage? How does the embodied carbon vary after being normalized to performance?*

Result and Analysis. In Fig. 1, we compare the embodied carbon footprint of the GPU components and CPU components we study in Table 1. Fig. 1 (a) shows that each GPU devices have higher embodied carbon than the CPU devices by up to 3.4 \times . On the other hand, when we normalize the embodied carbon to the FP64 operation performance in Fig. 1 (b), the trend is reversed: each CPU device has higher embodied carbon per FLOPS than any of the GPU devices. This is because CPUs provide much lower performance




Figure 1: Embodied carbon footprint of GPU/CPU devices, and the footprint normalized to theoretical double-precision floating point performance (TeraFLOPS).




Figure 2: Embodied carbon footprint of DRAM/SSD/HDD devices and the footprint normalized to bandwidth (GB/s).

despite containing less embodied carbon and hence, they are not able to offset the lower performance with their embodied carbon efficiency. We have also observed similar trends on other floating point precisions such as the FP32.

The AMD MI250X GPU has the highest embodied carbon but has the lowest embodied carbon per FLOPS among all devices. This is because AMD has reported this GPU to have almost 5 \times higher peak FP64 FLOPS than an NVIDIA A100 [52]. Note that we used FLOPS simply because it is the most commonly documented and used. However, our methodology is not specific or limited to FLOPS only; other figures of merit can be used for normalization as well.

Observation 1. We observe that GPUs tend to have significantly more embodied carbon than CPUs – this seems to be true across multiple types of CPUs and GPUs that are used in different top 500 supercomputers. Although GPUs tend to have higher overall embodied carbon, the embodied carbon normalized to raw performance ($\text{gCO}_2/\text{FLOPS}$) is lower than CPUs.

Next, we investigate the embodied carbon of memory and storage devices in a system, as presented in Table 1. From Fig. 2 (a), we can observe that each DRAM/SSD/HDD device has an embodied carbon of 5 to 25 kgCO_2 , which is in a comparable range to the GPU/CPU devices. Similar to the normalized analysis on GPUs/CPUs, we are also interested in the embodied carbon per bandwidth (GB/s), which is considered a key metric of memory/storage devices. The trend we observe in Fig. 2 (b) indicates that the embodied carbon per bandwidth of DRAM devices is significantly smaller than SSDs, and is negligible compared to the HDD devices due to a much higher DRAM bandwidth.

Observation 2. A typical single unit of memory and storage device also tends to have a comparable amount of embodied carbon as compute units (CPU/GPU), but it should be kept in




Figure 3: The manufacturing and packaging carbon footprint of the embodied carbon varies with device types.

mind that the capacity of memory and storage devices can affect the embodied carbon.

Implication. The implication is that carbon-conscious HPC facilities should explicitly request the embodied carbon specifications for CPUs and other computer accelerators from the chip vendor as a part of their request for proposal (RFP), in addition to performance benchmarking numbers. Performance benchmarking alone is not sufficient to achieve environmental sustainability. The embodied carbon footprint of memory and storage devices cannot be ignored either. While storage system has been traditionally viewed as a secondary optimization goal for performance and top 500 rankings, the carbon embodied in hard drives and solid storage drives can present a serious challenge to sustainability.

RQ 2. *What is the breakdown of the embodied carbon for different types of GPUs, CPUs, and memory/storage, in terms of manufacturing and packaging carbon?*

Result and Analysis. In Fig. 3, we quantify the carbon footprint associated with the manufacturing and packaging of various computer hardware components. The manufacturing footprint is incurred during the semiconductor wafer fabrication, assembly, and testing, while the packaging footprint represents the emission during the chip packaging process.

Specifically, we focus on the carbon footprint of GPU, CPU, DRAM, SSD, and HDD. Our analysis is presented through a series of ring charts, where each chart represents the carbon footprint of one specific component. The charts are divided into two parts, representing the manufacturing and packaging carbon footprint.

Interestingly, we found that the composition of the carbon footprint varied significantly between different components. For example, while the embodied carbon footprint of SSD and HDD was dominated by the manufacturing process, the packaging carbon footprint of DRAM was found to be 42% of its overall embodied carbon footprint. This is due to the fact that DRAM chips are typically smaller and require more precise and delicate packaging due to their sensitivity to external factors such as temperature, humidity, and electrostatic discharge whereas SSDs and HDDs are larger and require less intricate packaging. The GPUs and CPUs have about 10% of their embodied carbon from packaging, less than DRAM as potentially their manufacturing is more complex from factors such as lower lithography and larger die areas.




Figure 4: Increasing the number of GPUs in a node can improve node performance, but it also leads to a higher embodied carbon footprint. However, as expected, the performance gains tend to plateau, but the carbon footprint continues to increase.

Observation 3. As expected, the manufacturing carbon is the most dominant part of the embodied carbon for most components including GPUs, CPUs, HDDs, and SSDs. However, for DRAM, packaging carbon contributes over 40% of the embodied carbon – contributing toward a higher packaging-induced embodied carbon.

Limitation of this study. We recognize that a complex HPC system tends to have additional components. Specifically, network interconnects such as HPE Slingshot [53] provide high-bandwidth, low-latency communication between nodes; in a distributed file system, storage devices are connected to storage servers that are in turn connected to compute nodes [54]. In this work, these components could not be modeled and characterized due to the unavailability of open-access production carbon emission reports.

Implication. There is a critical shortage of carbon footprint data related to different components including networking equipments in HPC systems. HPC practitioners and vendors should work together to build standardized models for collecting and sharing embodied carbon of different components.

RQ 3. *How do the embodied carbon and performance vary for different workloads as the number of GPUs increases in a compute node?*

Result and Analysis. In Fig 4, show how the embodied carbon footprint and the system performance vary when we increase the number of GPUs in a node. In our node, we have two Intel Xeon Gold 6240R CPUs, and we vary the number of NVIDIA V100 GPUs between 1, 2, and 4. We characterize the 1-GPU, 2-GPU, and 4-GPU system performance using the three sets of benchmarks in Sec. 2 (Table 4). We compare the performance against the embodied carbon footprint of the node. We have kept the batch size per GPU in these benchmarks consistent as we increase the number of GPUs.

As expected, the embodied carbon footprint increase is proportional to the number of GPUs added. For all benchmark sets, when we increase the number of GPUs to 2, both the embodied carbon and the node performance are increased by approximately 30% to 40% of the normalized carbon footprint and the corresponding performance, meaning the performance-to-embodied-carbon ratio is approximately 1. However, as we further increase the number of GPUs to 4, the performance increase cannot keep up with the




Figure 5: Carbon footprint contribution from different components in three leadership supercomputers: Frontier, LUMI, and Perlmutter.

embodied carbon footprint due to heavier communication overhead between the GPUs, and the performance-to-embodied-carbon ratio has dropped to approximately 0.88 for the NLP and CANDLE benchmarks, and 0.79 for the Vision benchmarks.

Observation 4. Our experimental results demonstrate that the carbon footprint per unit of achieved performance of a workload may get worse as we increase the number of GPUs. As expected, with increasing the number of GPUs, the performance may not increase linearly, but the embodied carbon increases linearly. Therefore, the overall carbon footprint per unit of achieved performance may increase.

RQ 4. For leading supercomputers, does the contribution from different components toward the overall embodied carbon change? Which is the most dominant embodied carbon in today’s supercomputers: GPU, CPU, memory, SSD, or HDD?

Result and Analysis. To better understand the sources and distribution of embodied carbon in supercomputing, we present a comparative analysis of three of the world’s powerful supercomputers: Frontier, LUMI, and Perlmutter that are ranked 1st, 3rd, and 8th respectively in the latest Top-500 list as of November 2022 [25]. We have listed a more detailed description of these systems in Table 2.

We note that the magnitude of the absolute carbon footprint of each supercomputer is not listed because it is not our intent to showcase that one is better than the other, or even compare them. Instead, we want to highlight that the composition of a system greatly affects the embodied carbon footprint breakdown.

In Fig. 5, each ring chart shows the proportion of carbon footprint contributed by different components of the system, including the CPU, GPU, DRAM, SSD, and HDD. By examining the relative contributions of these components, we can gain insights into the amount of carbon emission embodied in the components when building different supercomputer architectures. We recognize that even the current systems can potentially undergo upgrades, the carbon footprint estimates of their components are subject to inaccuracy and limited to publicly available information, and these can change over time/generations. Therefore, instead of absolute numbers, we are more interested in projected estimates about how system composition can shift the carbon bottlenecks - this also guided our choice of selected supercomputers.

First, our analysis reveals interesting carbon footprint differences between the three selected supercomputers in terms of their composition. For example, while Frontier and LUMI have significantly more sizeable embodied carbon footprints due to their GPUs,

Perlmutter has a more balanced embodied carbon distribution between CPUs and GPUs. This is because Perlmutter has a large CPU partition whereas LUMI has a relatively small CPU partition and Frontier is also relatively GPU-heavy.

Second, the proportion of carbon emissions from storage devices varies between the three systems, reflecting differences in their storage architectures. Frontier has 695 PB of HDD storage that makes up large embodied carbon footprint, while Perlmutter deploys an all-flash file system. This is particularly important and highlights that the carbon contributions from storage can be significant and should be treated as a first-class citizen, even though they often take the backseat in performance rankings.

Interestingly, when we compare the compute components (CPU and GPU) against the memory and storage devices (DRAM, SSD, HDD) in terms of their embodied carbon, the memory and storage have made up approximately 60% of the carbon in Frontier and Perlmutter, and almost 50% in LUMI. This indicates that although these memory and storage devices do not consume as much power as compute devices, they inherently have a higher embodied carbon footprint, which cannot be neglected when building a supercomputer. This is in alignment with multiple concurrent and recent studies which have highlighted the importance of the carbon footprint of storage devices [55–57].

Finally, in Fig. 5, the GPUs have consistently higher embodied carbon footprint than CPUs in all three supercomputers, especially in Frontier, where the embodied carbon in GPUs is more than 7× that of the CPUs. This shows a trend that the GPUs are becoming important in today’s supercomputers, not just from the performance point-of-view, but also for their carbon footprint.

Observation 5. Our analysis reveals that the breakdown of their embodied carbon differs significantly among different supercomputers – even though, these supercomputers appear to be comparable in peak raw performance (among the top ten in the top 500 supercomputer list). Depending upon the supercomputer architecture and organization, GPU, memory or even SSD can be the most dominating factor. Surprisingly, DRAM contributes significantly to overall embodied carbon for all evaluated supercomputers. This is in contrast with an earlier result, where we showed that a DRAM card’s embodied carbon was lower compared to a single CPU/GPU. The reason for an overall relatively higher contribution of DRAM is because of the high number of DRAM cards in the system – which adds up.

Implication. The implication is that HPC facilities should explicitly document and understand which components are contributing to their overall embodied carbon footprint. Currently, there is limited awareness of these factors in supercomputing center design. As energy sources powering the supercomputers become “greener”, this aspect will become the most dominant factor in the overall carbon footprint of a supercomputing center. The memory footprint of computational science applications has been on rapid rise [58–60]. Many supercomputing centers provision a large amount of

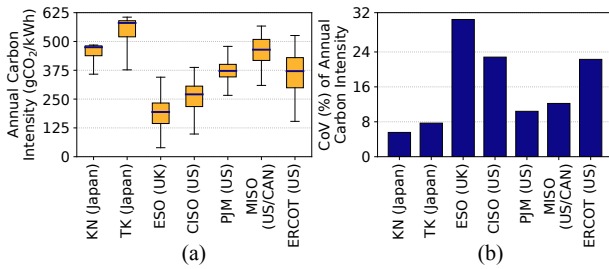


Figure 6: The average carbon intensity of the energy sources differs across geographical regions in Table 3, and a significant temporal variation exists for each geographical region.

memory to serve such applications. As the HPC centers prepare to serve more memory-intensive scientific applications, they should carefully consider the hidden carbon cost of these memory modules. Memory often has the largest failure rate and gets replaced [61, 62], therefore, lack of attention around minimizing or mitigating embodied carbon cost for DRAM can be undesirable.

4 GEOGRAPHICAL CARBON INTENSITY

In Sec. 3, we delved into the analysis of embodied carbon, which refers to the carbon emissions associated with the manufacturing and packaging of individual devices. Now, in this section, we shift our focus to operational carbon, which relates to the emissions resulting from the day-to-day operation of a large-scale system. It is worth noting that operational carbon strongly depends on regional carbon intensity, which reflects the amount of carbon emissions associated with electricity production in a given region. This means that datacenters and supercomputers located in regions with high carbon intensity will have a higher operational carbon footprint than those located in regions with lower carbon intensity. Therefore, it’s crucial to analyze the regional carbon intensity when evaluating the operational carbon emissions of a system.

RQ 5. *How does the carbon intensity vary across geographical regions?*

Result and Analysis. In Fig. 6, we illustrate the annual carbon intensity in the year 2021 of seven different power system operators distributed across different countries and regions. We have provided more details of the system operators in Table 3. In Fig. 6 (a), we use a box plot to compare the annual carbon intensity for different regions, which display distinctive carbon intensity trends.

Overall, the ESO (Great Britain, UK) region has the lowest carbon intensity among all regions, with a median carbon intensity of less than 200 gCO₂/kWh. The TK (Tokyo, Japan) region has the highest carbon intensity among all regions, whose medium annual carbon intensity is three times ESO’s. In recent years, there has been a growing interest in assessing the environmental impact of supercomputers, for example, the Green500 list [63]. However, it is important to note that the carbon intensity of electricity generation can vary significantly among different regions. Therefore, when comparing the "greenness" of supercomputers, it is crucial to take into account the geographical location of the facility. This is because the electricity grid mix in one region may be heavily reliant on fossil




Figure 7: Hourly variation in carbon intensity across three most carbon-friendly regions. Although not shown explicitly, season variations also naturally exist.

fuels, while in another region, it may be predominantly powered by renewable energy sources. As such, a supercomputer located in a region with a higher proportion of renewable energy in its electricity grid mix would have a lower operational carbon footprint than a supercomputer located in a region with a higher proportion of fossil fuel-based electricity. Thus, in order to accurately evaluate the environmental impact of supercomputers, it is important to consider the carbon intensity difference among different regions.

Remarkably, Fig. 6 (b) shows a different side of the story when we show the coefficient of variation (CoV) in %, which represents the standard deviation as a percentage of the average carbon intensity in the region. The two regions with the lowest medium carbon intensity – ESO (Great Britain, UK) and CISO (California, US), also have the most variations in their carbon intensity. On the other hand, the regions with the highest medium carbon intensity – TK (Tokyo, Japan) and KN (Kansai, Japan) have the least carbon intensity variation among all regions. This means that while the average carbon intensity of a region may be relatively low, there may be significant fluctuations in carbon emissions over time. One factor that contributes to variation in carbon intensity is the use of intermittent renewable energy sources such as wind or solar power, which may result in fluctuations in carbon emissions.

Insight 6. As expected, the average carbon intensity of the energy sources differs across geographical regions. On average, the ESO (Great Britain, UK) region overall and the California region within the USA has the lowest carbon intensity. But, interestingly, the temporal variance in those regions is among the highest. Therefore, simply building a data center in the least-carbon intensity region is not an optimal solution at all times – due to significant temporal variation.

Implication. The implication is that the federal agencies within a country and across countries should continue to strongly consider deploying similar architecture supercomputers across multiple geographical regions. Incidentally, this model has been followed in the past in the USA for other reasons (programmability, staffing, science missions). Our analysis reinforces that model from a carbon-aware computing perspective – esp. in the collaborative international context, too. When ranking supercomputers based on their "greenness" (Green 500 ranking), we should also consider the geographical location of the facility and energy-mix, and its temporal variations – which is not currently practiced.

RQ 6. Can the temporal variation in carbon intensity be exploited at finer timescale (e.g., hours)?

Result and Analysis. In Fig. 7, we pick the three operator regions with the lowest medium carbon intensity and compare their carbon intensities during the same hour of the day. Since they are distributed in different geographical locations, we account for the difference between time zones (GMT, PST, CST) and convert them to JST (UTC+9) time during the analysis. On the y-axis, we show the number of days that the region experiences the lowest carbon intensity among all regions during that hour. For example, for the 1st hour of the day, ESO (Great Britain) is the region with the lowest carbon intensity in about 150 days out of a year while CISO (California) is the lowest carbon intensity region in about 215 days out of a year.

Fig. 7 reveals that the number of days that each region has the lowest carbon intensity during a given hour varies significantly throughout the year, with no region consistently having the lowest carbon intensity for any given hour. This is due to the different energy generation mixes and demands in different time zones in different regions. The hours during which ESO (Great Britain) is the region with the lowest carbon intensity, hour 8 to hour 20, are the midnight to noon time in the UK when electricity demand is expected to be low. Therefore, it would be beneficial to have more jobs running in the HPC centers in the ESO to exploit the availability of renewable energy. However, during other hours, running the job in the ESO region is more likely to yield more carbon emissions as CISO (California) is a “greener” region during most of the days. Overall, our analysis highlights the temporal variability in the carbon intensity of different regions and underscores the need for flexible and dynamic approaches to exploit the opportunity in distributing jobs across regions.

While visually not depicted, we verified that even when two regions have very similar carbon intensity (e.g. Mid-Atlantic US and Texas), it is possible to optimize for carbon footprint further by distributing jobs between data centers in these regions. This is because the regions exhibit temporal variations and these variations are aligned due to geographical characteristics – for example, when wind power is more easily available in Texas, it may be not available in New Jersey at the same time.

Insight 7. Our analysis reveals that even among the greenest region, there is a significant hourly variation in carbon intensity. This variation is strong enough that no single region is a consistent winner for all hours of the day for all days in a year, and the number of days they are a winner in a year also varies – motivating a case for geographically distributed data centers where jobs can be distributed. However, exploiting this opportunity is not trivial since the temporal variation on different days of the year varies. There are additional challenges related to latency and the energy consumption associated with data transfers when distributing the workload across geographically distributed HPC centers. Therefore, workload distribution policies should consider such a tradeoff.

Table 5: Different generations of nodes analyzed.

Name	GPU	CPU
P100	4× NVIDIA Tesla P100 PCIe	2× Intel Xeon CPU E5-2680
V100	4× NVIDIA Tesla V100 SXM2	2× Intel Xeon Gold 6240R
A100	4× NVIDIA A100 PCIe 40GB	4× AMD EPYC 7542

Table 6: Performance improvement from the node upgrade.

Upgrade Option	NLP Improv.	Vision Improv.	CANDLE Improv.	Average Improv.
P100 to V100	44.4%	41.2%	45.5%	43.4%
P100 to A100	59.0%	60.2%	68.3%	62.5%
V100 to A100	25.6%	35.8%	44.4%	35.9%

Implication. An important implication around these observations is an incentive structure for end users to exploit these fine-grained carbon intensity patterns. In particular, an incentive structure and accounting methods to encourage users to submit/run their jobs during low-carbon intensity would be useful. Similar to core-hour accounting and budgeting, HPC users should also be provided a carbon budget as a part of their allocation, and they could be prioritized to reduce their queue wait time if the carbon footprint of their jobs have been economical. There is a strong need to design, develop, and deploy carbon-intensity-aware job schedulers to exploit these opportunities across geographically distributed HPC centers. Currently, we already have some CloudBank-related programs which allow users to submit their jobs to different centers. However, robust system software support for real-time and automatic distribution of jobs is needed.

5 OPERATIONAL AND EMBODIED CARBON

In previous sections, we have addressed two key aspects of sustainability in HPC centers: embodied carbon in supercomputer components and operational carbon, which is influenced by regional carbon intensity. *In this section, we aim to integrate these factors and examine the total carbon footprint of supercomputer upgrades. By taking a holistic view of embodied and operational carbon, we can gain a more comprehensive understanding of the environmental impact of these upgrades.*

RQ 7. What are carbon footprint and performance trade-offs as supercomputing facilities consider hardware upgrades, esp. multi-generation GPU upgrades? Will the introduction of embodied carbon due to upgrade be offset by savings in operational carbon footprint due to more energy-efficient newer generation hardware? Do these trade-offs depend on the “greenness” of the energy mix?

Result and Analysis. To investigate the impact of system upgrades on performance and carbon footprint, we perform benchmarking on three nodes of different generations, denoted as P100, V100, and A100 nodes shown in Table 5. These three generations represent NVIDIA’s three major datacenter GPU architectures released in the past: Pascal, Volta, and Ampere.

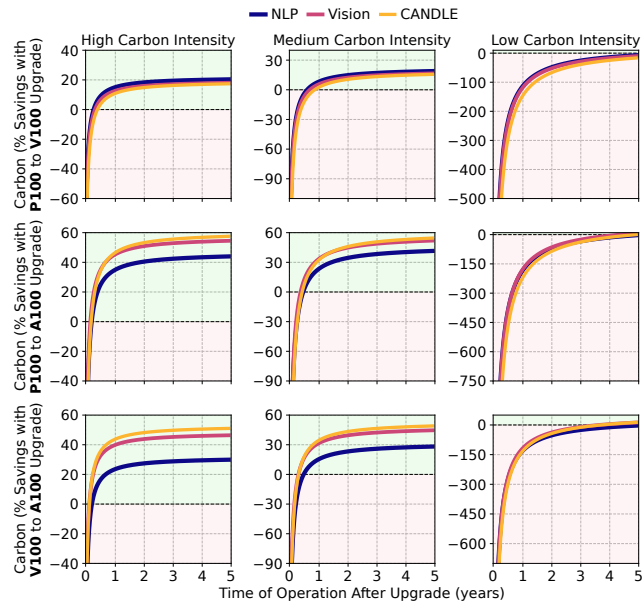


Figure 8: By upgrading a GPU system, one can reduce carbon emission over time despite a large carbon footprint increase initially. The amortization rate depends on the region’s carbon intensity.

We clarify that these experiments and analyses are primarily based on GPUs for simplicity, as they are likely to be among the most dominant contributing factors toward the overall carbon footprint of the data center (embodied and operational).

Our evaluation considers three upgrade options between node generations: P100 to V100, P100 to A100, and V100 to A100. To assess performance improvements, we conducted benchmarks listed in Table 4 and recorded the results for each benchmark set in Table 6. The data indicates that all upgrade options delivered notable performance improvements, as expected, with the largest gains when upgrading from P100 to A100 due to a longer gap in time between those generations. The performance improvements ranged from 25% to almost 70%. Notably, the CANDLE benchmark demonstrated greater performance improvements than the other two benchmarks across all three upgrade options. Overall, the performance gains from the upgrades were significant. In the following section, we explore the environmental impact of the upgrades in terms of their carbon footprint and make a case that performance improvements alone may not be sufficient for considering the upgrade decision.

The figure shown in Fig. 8 presents the results of our analysis of three upgrade scenarios, each represented by a row (P100 to V100, P100 to A100, and V100 to A100), and different levels of average carbon intensity, represented by columns (high, medium, and low). Three lines on each subplot represent three different workloads (NLP, Vision, and CANDLE).

We evaluated the carbon footprint savings achieved by upgrading the system over a five-year period following the node upgrade. The red region indicates when the upgraded option resulted in a higher carbon footprint than not upgrading, while the green region indicates the upgrade has resulted in carbon footprint savings.

First, we analyze the first column (that is, the carbon intensity is held constant) of Fig. 8. As expected, all curves start from a

negative point because an upgrade immediately incurs embodied carbon cost, and it takes some time before this “tax” can be paid. Along the way, it is offset by saving operational energy over time (newer hardware is typically more energy efficient and hence, results in lower energy consumption). This is why almost all curves go toward the “green” region, albeit at a different rate.

The rate/steepness of the curve depends on the workload (NLP vs. CANDLE) and upgrade tier (P100 to A100 vs V100 to A100). In general, the energy efficiency improvements are the highest when upgrading from P100 to A100, and hence, the embodied carbon “tax” is paid quickly. NLP curve is typically below other Vision and CANDLE workloads because NLP receives the least performance improvement, and hence, the least energy improvement.

These findings align with the conventional wisdom that when newer, faster, and more energy-efficient hardware is available, we should upgrade the system. Our results partially support that from a carbon-consciousness aspect too, albeit with a caveat that the upgrades cannot be too fast and the window before the tax is offset can vary depending upon the workload being run in the system.

Next, we analyze the effect of carbon intensity on the same decisions. Across columns in Fig. 8, we evaluate the carbon footprint reduction in three different carbon intensities: high intensity with an average of 400 gCO₂/kWh, medium intensity with an average of 200 gCO₂/kWh, and low intensity with an average of 20 gCO₂/kWh which is the carbon intensity of hydropower [7].

We make an interesting observation that hardware upgrade benefits can heavily depend on the energy source of the HPC center. At high carbon intensity, it takes less than half a year to amortize the embodied carbon incurred at system upgrade; at medium carbon intensity, it takes less than a year to amortize the embodied carbon; but at low carbon intensity when the energy source is highly renewable, the amortization time is about five years or more.

Overall, upgrading is beneficial in terms of performance, but the carbon footprint perspective needs more consideration in the carbon intensity and expected system service life. In regions with high carbon intensity, upgrades can happen when the new generation is released since the new system will quickly amortize its embodied carbon. In regions with an abundant amount of green energy, upgrading would be carbon-friendly only if the system is expected to serve for at least five years approximately. We note that these periods are dependent on the embodied carbon and energy consumption of these devices, and hence, they should be interpreted in the context of the modeled situation.

Insight 8. Hardware upgrades are always attractive due to significant performance improvements, but the upgrade introduces significant embodied carbon which may not be offset quickly – esp. if the center already runs primarily on renewable energy sources, as could be the case in the future for many centers. In such cases, extending the hardware lifetime could be a worthy option. If the energy source is less green, a quicker upgrade may be desirable. Although we note that there are many other indirect side effects associated with the upgrades, including monetary investment, potential redesign of the data center, etc. which may further increase the carbon footprint.




Figure 9: Carbon saving curve after the GPU system upgrade under different usage patterns.

RQ 8. Does carbon footprint trade-offs at supercomputing facilities consider hardware upgrades depending upon the average load or utilization of the current system, and the workload being run?

Result and Analysis. HPC centers have high utilization levels, but not all GPUs are utilized at all times. In fact, multiple HPC and data centers have reported low GPU utilization levels [64–66]. Therefore, we investigate this aspect too.

In Fig. 9, we keep the average carbon intensity constant at 200 gCO₂/kWh and vary the average GPU usage rate of the nodes, which represents the percentage of time the GPU is being used. The three carbon-saving curves in each subplot represent three different GPU utilization levels. We first assume that all the GPU nodes are allocated by users 100% of the time, and choose 40% GPU usage as the medium usage to align with a production trace [64–66]. The high and low usage numbers are selected based on 1.5× more and less of the medium usage, respectively. Our previous analysis in Fig. 8 is conducted with medium usage.

In Fig. 9, we show that depending on the user usage pattern, the time it takes to amortize the embodied carbon varies. The difference is not as significant as the carbon intensity, where it can be multiple years of difference. Nevertheless, it would still substantially impact the saved carbon footprint at a certain point in time after the upgrade. Using the NLP benchmark set as an example, when we upgrade from the V100 system to the A100 system, after one year, a high/medium usage pattern would result in approximately 20% carbon footprint reduction, whereas the low usage pattern has just paid off the initial embodied carbon of the upgrade.

Insight 9. If the center has limited GPU utilization, extending the hardware lifetime could be a worthy option, but it also depends upon the mix of workloads being run and their

energy consumption characteristics. If the GPU utilization is high, a quicker upgrade may be desirable, since the savings in operational carbon footprint could quickly offset the introduction of embodied carbon. Nevertheless, this decision is still heavily influenced by the “greenness” of the energy source.

Implication. The HPC centers should consider carbon footprint in their hardware upgrade decisions. We recognize that carbon awareness alone is not the determining factor – traditionally, cost and performance improvements have dictated the timing of such decisions. However, as carbon net-zero aims become more commonplace, the centers should have methods, such as those introduced in this paper, to evaluate the lifetime of a hardware generation and if extending it would be useful. The HPC centers need to continuously monitor and evaluate the carbon-intensity and GPU utilization load to determine when to upgrade because carbon-intensity and GPU utilization are among the most dominant factor in the hardware-upgrade decision-making which attempts to minimize carbon footprint. Assessing the carbon-friendliness of HPC centers involves considering several factors. These factors include the initial embodied carbon footprint, the operational carbon that is linked to the regional carbon intensity and user usage pattern, and the expected operating lifetime of the HPC system. These factors are not standard yet. Large-scale HPC applications would have a large operational carbon footprint due to the heavy computation carried out across multiple nodes.

6 THREATS TO VALIDITY AND DISCUSSION

We recognize and acknowledge that our analysis is not immune to threats to validity. For instance, while our current analysis models the embodied carbon footprint of the hardware and operational carbon footprint, it does not account for carbon emissions related to other side effects of upgrading the system. Such emissions arise not only during manufacturing and packaging but also during transportation, installation, and recycling. System upgrades often require increasing building capacity, changes in cooling infrastructure, additional staffing, and acquisition of other compute/storage components to match the increase in the compute capacity – these aspects lead to additional carbon emission. Hardware cost is also an important aspect, but we do not explicitly include it because hardware cost may depend on many factors and vary across geographical regions, the scale of the system, and may change over time. Unfortunately, as a community, we are at a very early stage to have sufficiently detailed and accurate modeling of these by-product effects. We are hoping that this work will spur a new research direction in the HPC community to standardize and document these aspects – which have not been central for performance benchmarking, but will be critical for HPC carbon benchmarking.

We acknowledge that fab yield differs across manufacturers and lithography. Due to the lack of a standardized public database providing accurate yield data, we configured it to be based on prior works. Similarly, the PUE metric, while challenging to estimate with seasonal variation, can be approximated well with IT and

cooling energy monitors available in most HPC datacenters. The operational carbon can be calculated using Eq. 6 (e.g., energy from power measurement tools (e.g., NVML, RAPL) multiplied by the carbon intensity). Note that the embodied carbon relative to the operational carbon varies across P100, V100, and A100 GPUs.

Our choice to select the three workloads from a broad range of HPC-relevant benchmarks was because they cover a diverse range of resource consumption characteristics, resulting in different trends that capture the spectrum. Adding more benchmarks makes the interpretation and presentation complex, but they should be considered in the future as they can lead to more insights, esp. given that HPC workloads are rapidly evolving with new AI and quantum use cases. The operational carbon footprint of large-scale applications is also an important consideration and potential threat to validity, but these can be performed in-house by different HPC systems as they may not always make their representative workloads publicly available. The key actions include but are not limited to, measuring the carbon footprint of large-scale applications on different computing hardware and making carbon-aware HPC scheduling and hardware upgrade decisions based on geographical carbon intensity and application characteristics. The HPC site may have a contract with a local power provider which needs to be considered when accounting for the carbon footprint. Additionally, migrating users to new architectures (e.g., GPUs) also incurs an indirect carbon footprint and has trade-offs in terms of performance, esp. if the system is not being procured for a high FLOP count but instead for running a specific workload. These aspects are not explicitly accounted in this study.

The popular energy efficiency metric (FLOPS/Watt) alone does not fully capture the environmental impact of an HPC system because it does not capture embodied carbon and carbon intensity of the energy source. A system with higher energy efficiency does not necessarily mean it has lower operational carbon footprint – because operational carbon also depends on the carbon intensity of the energy being fed (e.g., operation of system A (20GFLOPS/Watts) may be “greener” than B (50GFLOPS/Watts) if A uses hydropower while B uses gas).

We recognize that cost is an important consideration during hardware upgrades. For example, some low-cost systems can indeed have a carbon lower footprint as well, depending upon the load and carbon intensity. However, hardware cost may depend on many factors and vary (across geographical regions, scale of the system, and may change over time, etc.), hence integrating it into the study requires careful consideration with restrictive assumptions.

7 RELATED WORK

Study of HPC system components. Various works have addressed the energy efficiency of CPU [67–69], DRAM [70–72], and storage components [73–75] in a HPC system. GPUs have been widely integrated into today’s HPC systems to accelerate deep learning applications, hence a variety of efforts have focused on improving the GPU device energy efficiency [76–81]. These works have only focused on the energy efficiency of individual components, not on the carbon emission during manufacturing and operation.

Green computing in large-scale GPU systems. The environmental impact of HPC systems has been an important research

topic in the past. GPU-NEST [82] has examined the power and performance behaviors of multi-GPU inference systems. Helios [83] proposed a cluster energy-saving service for GPU datacenter running deep learning workloads. Patki et. al. [84] applied GPU power and frequency capping in a scientific workflow to improve power efficiency while preserving performance. Our work goes beyond the investigation of energy consumption: we analyze the regional carbon intensity to relate the energy consumption to the actual carbon emission. Various works have investigated the operational carbon footprint when running workloads such as bioinformatics and astrophysics in large-scale systems [1, 85–87]. Our work is distinct from the previous ones as we combine the embodied carbon footprint, operational carbon footprint, regional carbon intensity, performance, system service life, and usage pattern to conduct a holistic analysis. This study is complementary to various emerging energy-efficient HPC network topologies [88, 89].

Carbon footprint modeling. Wang [19] and Totally Green [8] proposed to take the production-operation-recycling product life cycle into sustainability analysis. In regard to the production perspective, ACT [7] released an embodied carbon footprint modeling tool that is useful for conducting embodied carbon analysis for mobile devices. Our work is not limited to embodied carbon modeling and we also model the operational carbon on systems of different generations under different regional carbon intensities and usage patterns for HPC systems. This is the first work to provide carbon modeling methods and tools for HPC practitioners and identify various areas of opportunities for HPC system facilities.

8 CONCLUSION

We present a carbon footprint analysis framework that addresses a series of research questions related to HPC system sustainability. We have conducted modeling of the embodied carbon footprint of HPC system components, investigations into how regional carbon intensity affects the system, and end-to-end characterizations of the carbon footprint throughout the system life cycle. We hope that our methodology, framework, and study encourage HPC researchers to promote sustainable HPC and promote more research efforts toward carbon neutrality in the community.

ACKNOWLEDGMENTS

This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702-15-D-0001, and United States Air Force Research Laboratory Cooperative Agreement Number FA8750-19-2-1000. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering, or the United States Air Force. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

We thank the anonymous reviewers for their constructive feedback. We used ChatGPT, an AI language model developed by OpenAI, for partial assistance in writing, and all such texts were verified and edited for correctness.

REFERENCES

- [1] Simon Portegies Zwart. The ecological impact of high-performance computing in astrophysics. *Nature Astronomy*, 4(9):819–822, 2020.
- [2] Dhableswar Kumar Panda, Hari Subramoni, Ching-Hsiang Chu, and Mohanmadreza Bayatpour. The mvapich project: Transforming research into high-performance mpi library for hpc community. *Journal of Computational Science*, 52:101208, 2021.
- [3] Christian Feichtinger, Stefan Donath, Harald Köstler, Jan Götz, and Ulrich Rüde. Walberla: Hpc software design for computational engineering simulations. *Journal of Computational Science*, 2(2):105–112, 2011.
- [4] United States International Trade Commission. Data centers around the world: A quick look, 2021. URL https://www.usitc.gov/publications/332/executive_briefings/ebot_data_centers_around_the_world.pdf.
- [5] Statista. Data centers - statistics & facts, 2023. URL <https://www.statista.com/topics/6165/data-centers/>.
- [6] Oak Ridge National Laboratory. Frontier's architecture, 2023. URL <https://olcf.ornl.gov/wp-content/uploads/Frontiers-Architecture-Frontier-Training-Series-final.pdf>.
- [7] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S Lee, David Brooks, and Carole-Jean Wu. Act: Designing sustainable computer systems with an architectural carbon modeling tool. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, pages 784–799, 2022.
- [8] Jichuan Chang, Justin Meza, Parthasarathy Ranganathan, Amp Shah, Rocky Shih, and Cullen Bash. Totally green: evaluating and designing servers for lifecycle environmental impact. *ACM SIGPLAN Notices*, 47(4):25–36, 2012.
- [9] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. Chasing carbon: The elusive environmental footprint of computing. *IEEE Micro*, 42(4):37–47, 2022.
- [10] Anders SG Andrae and Tomas Edler. On global electricity usage of communication technology: trends to 2030. *Challenges*, 6(1):117–157, 2015.
- [11] Amazon. Amazon sustainability report, 2021. URL <https://sustainability.aboutamazon.com/2021-sustainability-report.pdf>.
- [12] Google. Google environmental report, 2022. URL <https://www.gstatic.com/gumdrop/sustainability/google-2022-environmental-report.pdf>.
- [13] Meta. Meta sustainability report, 2021. URL <https://sustainability.fb.com/wp-content/uploads/2022/06/Meta-2021-Sustainability-Report.pdf>.
- [14] Apple. Environmental progress report, 2022. URL https://www.apple.com/environment/pdf/Apple_Environmental_Progress_Report_2022.pdf.
- [15] Zhiwei Cao, Xin Zhou, Han Hu, Zhi Wang, and Yonggang Wen. Towards a systematic survey for carbon neutral data centers. *IEEE Communications Surveys & Tutorials*, 2022.
- [16] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. 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*, pages 118–132, 2023.
- [17] Yi Gu and Chandu Budati. Energy-aware workflow scheduling and optimization in clouds using bat algorithm. *Future Generation Computer Systems*, 113:106–112, 2020.
- [18] Baolin Li, Siddharth Samsi, Vijay Gadepally, and Devesh Tiwari. Green carbon footprint for model inference serving via exploiting mixed-quality models and gpu partitioning. *arXiv preprint arXiv:2304.09781*, 2023.
- [19] David Wang. Meeting green computing challenges. In *2008 10th Electronics Packaging Technology Conference*, pages 121–126. IEEE, 2008.
- [20] Alex K Jones, Yiran Chen, William O Collinge, Haifeng Xu, Laura A Schaefer, Amy E Landis, and Melissa M Bilec. Considering fabrication in sustainable computing. In *2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 206–210. IEEE, 2013.
- [21] Donald Kline, Nikolas Parshook, Xiaoyu Ge, Erik Brunvand, Rami Melhem, Panos K Chrysanthis, and Alex K Jones. Holistically evaluating the environmental impacts in modern computing systems. In *2016 Seventh International Green and Sustainable Computing Conference (IGSC)*, pages 1–8. IEEE, 2016.
- [22] Donald Kline Jr, Nikolas Parshook, Xiaoyu Ge, Erik Brunvand, Rami Melhem, Panos K Chrysanthis, and Alex K Jones. Greenchip: A tool for evaluating holistic sustainability of modern computing systems. *Sustainable Computing: Informatics and Systems*, 22:322–332, 2019.
- [23] SPIL. Corporate social responsibility report, 2020. URL <https://www.spil.com.tw/Files/pdf-en/2019-en.pdf>.
- [24] Seagate. Seagate product sustainability, 2023. URL <https://www.seagate.com/global-citizenship/product-sustainability/>.
- [25] Top-500 List. November 2022, 2022. URL <https://www.top500.org/lists/top500/2022/11/>.
- [26] SK Hynix Sustainability Report, 2022. URL <https://www.skhynix.com/sustainability/UI-FR-SA1601/>.
- [27] Oak Ridge National Laboratory. Frontier, 2023. URL <https://www.olcf.ornl.gov/frontier/>.
- [28] LUMI Consortium. LUMI Supercomputer, 2023. URL <https://lumi-supercomputer.eu/>.
- [29] NERSC. Perlmutter: High Performance Computing Optimized for Science, 2023. URL <https://perlmutter.carrd.co/>.
- [30] Lasse F Wolff Anthony, Benjamin Kanding, and Raghavendra Selvan. Carbon-tracker: Tracking and predicting the carbon footprint of training deep learning models. *arXiv preprint arXiv:2007.03051*, 2020.
- [31] Kansai Electric Power Company. Kansai Electric Power, 2023. URL <https://www.kepco.co.jp/english/>.
- [32] Tokyo Electric Power Company. TEPCO Power Grid, 2023. URL <https://www.tepco.co.jp/en/pg/index-e.html>.
- [33] National Grid ESO. Welcome to ESO, 2023. URL <https://www.nationalgrideso.com/>.
- [34] California ISO. California Independent System Operator, 2023. URL <https://www.caiso.com/Pages/default.aspx>.
- [35] PJM. PJM Interconnection, 2023. URL <https://www.pjm.com/>.
- [36] MISO. Midcontinent Independent System Operator, 2023. URL <https://www.misoenergy.org/about/>.
- [37] ERCOT. Electric Reliability Council of Texas, 2023. URL <https://www.ercot.com/>.
- [38] National Grid ESO. Carbon intensity api, 2023. URL <https://carbonintensity.org.uk/>.
- [39] Electricity Maps. Reduce carbon emissions with actionable electricity data, 2023. URL <https://www.electricitymaps.com/>.
- [40] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [41] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [42] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33:16857–16867, 2020.
- [43] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [44] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- [45] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [46] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [47] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*, pages 116–131, 2018.
- [48] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [49] Xiaofeng Mao, Gege Qi, Yuefeng Chen, Xiaodan Li, Ranjie Duan, Shaokai Ye, Yuan He, and Hui Xue. Towards robust vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12042–12051, 2022.
- [50] Ecp-candle. <https://github.com/ECP-CANDLE/Benchmarks>, 2022.
- [51] Xingfu Wu, Valerie Taylor, Justin M Wozniak, Rick Stevens, Thomas Brettin, and Fangfang Xia. Performance, energy, and scalability analysis and improvement of parallel cancer deep learning candle benchmarks. In *Proceedings of the 48th International Conference on Parallel Processing*, pages 1–11, 2019.
- [52] AMD. AMD INSTINCT MI200 SERIES ACCELERATOR, 2023. URL <https://www.amd.com/system/files/documents/amd-instinct-mi200-datasheet.pdf>.
- [53] HPE. HPC Slingshot Interconnect, 2023. URL <https://www.hpe.com/us/en/compute/hpc/slingshot-interconnect.html>.
- [54] Peter Braam. The lustre storage architecture. *arXiv preprint arXiv:1903.01955*, 2019.
- [55] Swamit Tannu and Prashant J Nair. The dirty secret of ssds: Embodied carbon. *arXiv preprint arXiv:2207.10793*, 2022.
- [56] Aviad Zuck, Donald Porter, and Dan Tsafir. Degrading data to save the planet. In *Proceedings of the 19th Workshop on Hot Topics in Operating Systems*, pages 61–69, 2023.
- [57] Gabriel Mersy and Sanjay Krishnan. Toward a life cycle assessment for the carbon footprint of data. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*, pages 1–9, 2023.
- [58] Samyag Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE, 2020.

- [59] Esha Choukse, Michael B Sullivan, Mike O'Connor, Mattan Erez, Jeff Pool, David Nellans, and Stephen W Keckler. Buddy compression: Enabling larger memory for deep learning and hpc workloads on gpus. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pages 926–939. IEEE, 2020.
- [60] Ivy Peng, Kai Wu, Jie Ren, Dong Li, and Maya Gokhale. Demystifying the performance of hpc scientific applications on nvm-based memory systems. In *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 916–925. IEEE, 2020.
- [61] Jie Ren, Kai Wu, and Dong Li. Exploring non-volatility of non-volatile memory for high performance computing under failures. In *2020 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 237–247. IEEE, 2020.
- [62] Zhao Zhang, Lei Huang, Ruizhu Huang, Weijia Xu, and Daniel S Katz. Quantifying the impact of memory errors in deep learning. In *2019 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 1–12. IEEE, 2019.
- [63] Green-500 List. November 2022, 2022. URL <https://www.top500.org/lists/green500/2022/11/>.
- [64] Qizhen Weng, Wencong Xiao, Yinghao Yu, Wei Wang, Cheng Wang, Jian He, Yong Li, Liping Zhang, Wei Lin, and Yu Ding. MLaaS in the wild: Workload analysis and scheduling in large-scale heterogeneous GPU clusters. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 945–960. USENIX Association, 2022.
- [65] Baolin Li, Rohin Arora, Siddharth Samsi, Tirthak Patel, William Arcand, David Bestor, Chansup Byun, Rohan Basu Roy, Bill Bergeron, John Holodnak, et al. Ai-enabling workloads on large-scale gpu-accelerated system: Characterization, opportunities, and implications. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 1224–1237. IEEE, 2022.
- [66] Myeongjae Jeon, Shivaram Venkataraman, Amar Phanishayee, Junjie Qian, Wencong Xiao, and Fan Yang. Analysis of large-scale multi-tenant gpu clusters for dnn training workloads. In *USENIX Annual Technical Conference*, pages 947–960, 2019.
- [67] George Papadimitriou, Manolis Kaliorakis, Athanasios Chatzidimitriou, Dimitris Gizopoulos, Peter Lawthers, and Shidhartha Das. Harnessing voltage margins for energy efficiency in multicore cpus. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 503–516, 2017.
- [68] Wanghong Yuan and Klara Nahrstedt. Energy-efficient cpu scheduling for multimedia applications. *ACM Transactions on Computer Systems (TOCS)*, 24(3): 292–331, 2006.
- [69] Andrej Podzimek, Lubomir Bulej, Lydia Y Chen, Walter Binder, and Petr Tuma. Analyzing the impact of cpu pinning and partial cpu loads on performance and energy efficiency. In *2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, pages 1–10. IEEE, 2015.
- [70] Mike O'Connor, Niladrish Chatterjee, Donghyuk Lee, John Wilson, Aditya Agrawal, Stephen W Keckler, and William J Dally. Fine-grained dram: Energy-efficient dram for extreme bandwidth systems. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 41–54, 2017.
- [71] Bharan Giridhar, Michael Cieslak, Deepankar Duggal, Ronald Dreslinski, Hsing Min Chen, Robert Patti, Betina Hold, Chaitali Chakrabarti, Trevor Mudge, and David Blaauw. Exploring dram organizations for energy-efficient and resilient exascale memories. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pages 1–12, 2013.
- [72] Hasan Hassan, Minesh Patel, Jeremie S Kim, A Giray Yaglikci, Nandita Vijaykumar, Nika Mansouri Ghiasi, Saugata Ghose, and Onur Mutlu. Crow: A low-cost substrate for improving dram performance, energy efficiency, and reliability. In *Proceedings of the 46th International Symposium on Computer Architecture*, pages 129–142, 2019.
- [73] Erica Tomes and Nihat Altiparmak. A comparative study of hdd and ssd raids' impact on server energy consumption. In *2017 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 625–626. IEEE, 2017.
- [74] Seonyeong Park, Youngjae Kim, Bhuvan Urgaonkar, Joonwon Lee, and Euseong Seo. A comprehensive study of energy efficiency and performance of flash-based ssd. *Journal of Systems Architecture*, 57(4):354–365, 2011.
- [75] Bryan Harris and Nihat Altiparmak. Ultra-low latency ssds' impact on overall energy efficiency. *USENIX HotStorage'20*, 2020.
- [76] Akhil Arunkumar, Evgeny Bolotin, David Nellans, and Carole-Jean Wu. Understanding the future of energy efficiency in multi-module gpus. In *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 519–532. IEEE, 2019.
- [77] Abhinandan Majumdar, Leonardo Piga, Indrani Paul, Joseph L Greathouse, Wei Huang, and David H Albonese. Dynamic gpgpu power management using adaptive model predictive control. In *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 613–624. IEEE, 2017.
- [78] Vijay Kandiah, Scott Peverelle, Mahmoud Khairy, Junrui Pan, Amogh Manjunath, Timothy G Rogers, Tor M Aamodt, and Nikos Hardavellas. Accelwattch: A power modeling framework for modern gpus. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 738–753, 2021.
- [79] Jingwen Leng, Tayler Hetherington, Ahmed ElTantawy, Syed Gilani, Nam Sung Kim, Tor M Aamodt, and Vijay Janapa Reddi. Gpuwattch: Enabling energy optimizations in gpgpus. *ACM SIGARCH Computer Architecture News*, 41(3): 487–498, 2013.
- [80] Ankit Sethia and Scott Mahlke. Equalizer: Dynamic tuning of gpu resources for efficient execution. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 647–658. IEEE, 2014.
- [81] Mahmoud Khairy, Vadim Nikiforov, David Nellans, and Timothy G Rogers. Locality-centric data and threadblock management for massive gpus. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 1022–1036. IEEE, 2020.
- [82] Ali Jahanshahi, Hadi Zamani Sabzi, Chester Lau, and Daniel Wong. Gpu-nest: Characterizing energy efficiency of multi-gpu inference servers. *IEEE Computer Architecture Letters*, 19(2):139–142, 2020.
- [83] Qinghao Hu, Peng Sun, Shengen Yan, Yonggang Wen, and Tianwei Zhang. Characterization and prediction of deep learning workloads in large-scale gpu datacenters. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15, 2021.
- [84] Tapasya Patki, Zachary Frye, Harsh Bhatia, Francesco Di Natale, James Glosli, Helgi Ingólfsson, and Barry Rountree. Comparing gpu power and frequency capping: A case study with the mummy workflow. In *2019 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS)*, pages 31–39. IEEE, 2019.
- [85] Jason Grealey, Loic Lannelongue, Woei-Yuh Saw, Jonathan Marten, Guillaume Méric, Sergio Ruiz-Carmona, and Michael Inouye. The carbon footprint of bioinformatics. *Molecular biology and evolution*, 39(3):msac034, 2022.
- [86] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, et al. Sustainable ai: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems*, 4:795–813, 2022.
- [87] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021.
- [88] Francisco J Andújar, Salvador Coll, Marina Alonso, Juan-Miguel Martínez, Pedro López, José L Sánchez, and Francisco J Alfaro. Energy efficient hpc network topologies with on/off links. *Future Generation Computer Systems*, 139:126–138, 2023.
- [89] Giorgis Georgakoudis, Nikhil Jain, Takatsugu Ono, Koji Inoue, Shinobu Miwa, and Abhinav Bhatele. Evaluating the impact of energy efficient networks on hpc workloads. In *2019 IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC)*, pages 301–310. IEEE, 2019.

Appendix: Artifact Description/Artifact Evaluation

ARTIFACT DOI

<https://zenodo.org/badge/latestdoi/659090328>

ARTIFACT IDENTIFICATION

The paper addresses the challenges and unexplored questions regarding the sustainability and carbon footprint of high-performance computing (HPC) systems. While the demand for HPC resources is growing, it also leads to a rise in energy consumption and carbon emissions. The paper builds a carbon footprint analysis framework to address a series of unexplored investigations from both the production and operational stages of an HPC system. The analysis highlights the need to quantify the carbon footprint of an HPC system holistically, where the carbon emissions from the hardware manufacturing to the end of the system life cycle have all been accounted for. The paper provides detailed modeling and characterization of carbon emissions during the production and operational stages of an HPC system, emphasizing the importance of considering the carbon footprint of memory and storage devices in HPC systems, the significance of regional carbon intensity when evaluating the operational carbon emissions, and the embodied carbon introduced by hardware upgrades.

The artifact we provide aligns with the structure of this paper. Our artifact includes both the source data we collect to perform embodied carbon modeling, the benchmarks we created to conduct carbon benchmarking, and the data visualization code written in Jupyter Notebook using the Python `matplotlib` library. More details are listed below following the three major sections of the paper.

- Sec. 3 Embodied Carbon Analysis

In this section, we conducted embodied carbon analysis on various components of an HPC system and performed embodied carbon breakdown of three different supercomputing facilities. In our artifact, we include all the source data that we have collected for the embodied carbon modeling in “.csv” formats and provide the source code to visualize the data. The data includes the embodied carbon footprint we modeled for each component listed below:

- (1) GPU: NVIDIA A100, AMD MI250X, NVIDIA V100
- (2) CPU: AMD EPYC 7763, AMD EPYC 7742, Intel Xeon Gold 6240R
- (3) DRAM: SK Hynix 64GB DDR4
- (4) SSD: Seagate Nytro 3530 3.2TB
- (5) HDD: Seagate Exos x16 16TB

In the artifact, we also include the data we used for the HPC system breakdown for three supercomputers listed below.

- (1) Frontier (<https://www.olcf.ornl.gov/frontier/>)
- (2) LUMI (<https://lumi-supercomputer.eu/>)
- (3) Perlmutter (<https://perlmutter.carrd.co/>)

- Sec. 4 Geographical Carbon Intensity.

In this section, we analyzed the impact of geographical carbon intensity on the carbon footprint of a system. Similar to the artifact for Sec. 3, we include the annual carbon intensity data across different regions in “.csv” formats and provide the visualization script. The carbon intensity data is collected hourly throughout the year 2021, we have listed the regions that we analyze below.

- (1) KN: Kansai, Japan
 - (2) TK: Tokyo, Japan
 - (3) ESO: Great Britain Electricity System Operator, United Kingdom
 - (4) CISO: California Independent System Operator, United States
 - (5) PJM: Pennsylvania-New Jersey-Maryland Interconnection, United States
 - (6) MISO: Midcontinent Independent System Operator, United States/Canada
 - (7) ERCOT: Electric Reliability Council of Texas, United States
- Sec. 5 Operational and Embodied Carbon

In this section, we study the impact of hardware generation upgrades on the overall carbon footprint of the system. We studied three tiers of hardware generations, as listed below:

- (1) P100: 4 NVIDIA Tesla P100 GPUs + 2 Intel Xeon CPU E5-2680
- (2) V100: 4 NVIDIA Tesla V100 SXM2 GPUs + 2 Intel Xeon Gold 6240R CPUs
- (3) A100: 4 NVIDIA A100 PCIe 40GB GPUs + 4 AMD EPYC 7542 CPUs

Our analysis includes both embodied carbon and operational carbon. For embodied carbon, we again provide the modeling details in “.csv” format. For operational carbon, we have created three benchmark suites to model the performance and operational carbon footprint of the three generations of systems. The benchmark suites are listed below:

- (1) Natural Language Processing (NLP) benchmarks: BERT, DistilBERT, MPNet, RoBERTa, BART
 - (2) Computer Vision (Vision) benchmarks: ResNet50, ResNext50, ShuffleNetV2, VGG19, ViT
 - (3) CANDLE benchmarks: Combo, NT3, P1B1, ST1, TC1
- The NLP benchmark is implemented with Huggingface, the Vision benchmark is implemented using Pytorch, and the CANDLE benchmark is implemented using TensorFlow2. We have provided the source code for these benchmarks as well as the data we have collected from running them.

The artifact is accessible via Zenodo: <https://zenodo.org/badge/latestdoi/659090328>

REPRODUCIBILITY OF EXPERIMENTS

We provide the following information for experimental reproducibility:

A complete description of the experiment workflow. The experiment workflow is straightforward for Sec. 3 and Sec. 4 of the paper. We have included all the data needed to visualize, and the readers only need to execute the Jupyter Notebook scripts on any Linux machine. All the scripts are packaged into a Docker image.

For Sec. 5 of the paper, the reader has the option to execute the benchmarks we provide by themselves. We have listed the hardware to run these benchmarks on and provided the corresponding benchmark suites in the Operational and Embodied Carbon artifact description. Simply pull the image from docker and run it on these three different systems. Once the benchmark runs have been completed, we provide the visualization scripts to generate the plots as shown in the paper. The reader may also skip the benchmarking and use the data that we collect from our own experiments to reproduce the paper figures.

An estimation of the execution time to execute the experiment workflow. We estimate the time to execute the visualization scripts to be less than 1 minute. If the reader chooses to run all three benchmark suites (NLP, Vision, and CANDLE) on all three generations of systems (P100, V100, and A100) sequentially, we estimate the execution time to be less than 10 hours. If the benchmarks are conducted in parallel, each benchmark suite would take 1 hour on one system.

A complete description of the expected results and an evaluation of them. The expected results will be figures that are identical to Fig. 1 to Fig. 9 in the original paper.

How the expected results from the experiment workflow relate to the results found in the article. The results for Sec. 3 and 4 of the paper will be completely identical between the artifact-generated figures and the figures presented in the paper. For Sec. 5, if the reader chooses to run benchmarks on their own machines, we would expect there to be a maximum of 5% difference observed in Fig. 8 and 9 of the paper. If skipping benchmarking and use the data we collected, the results for Fig. 8 and Fig. 9 will again be completely identical.

ARTIFACT DEPENDENCIES REQUIREMENTS

Specified in repository.

ARTIFACT INSTALLATION DEPLOYMENT PROCESS

Specified in repository.