Understanding the Operational Carbon Footprint of Storage Reliability and Management

Varsha Rao
University of Chicago
Chicago, IL, USA
varsharao@uchicago.edu

Andrew A. Chien
University of Chicago & Argonne National Laboratory
Chicago, IL, USA
aachien@uchicago.edu

ABSTRACT

With daily data generation of Zettabytes and exponential growth, our study finds that power needed for data reliability and management is an increasing fraction of storage system power. Thus storage management is an important contributor to data center (DC) energy use and carbon footprint.

We study the University of Chicago's high energy physics storage system (UChicago HEP) background tasks for reliability and management. We build a model for their activity and power costs, and next explore their opportunities for temporal shifting to reduce operational carbon footprint.

We apply our model to varied DC scale: a large HEP facility, large cloud, and global set of large cloud DCs, we find that these storage tasks consume 2.1-4.9% of DC power and 12.8-27.2% of DC storage annually. Translating power use to carbon emissions depends on location and time, so carbon footprint varies. But we show that storage management power can account for as much as the annual electricity consumption of 755,000 US homes. Studying four power grids with varied carbon intensity, we show that aligning background tasks with low-carbon periods can achieve emission reductions of up to 82.8%. Greater reductions (up to 96.9%) are possible, as grids decarbonize further.

CCS CONCEPTS

• **Social and professional topics → Sustainability**; • **Computer systems organization →** *Maintainability and maintenance*; • **Applied computing → Data centers**.

KEYWORDS

Distributed Storage Systems, Storage Reliability and Management, Sustainability, Carbon Emissions

ACM Reference Format:
Varsha Rao and Andrew A. Chien. 2024. Understanding the Operational Carbon Footprint of Storage Reliability and Management. In *Proceedings of 3rd Workshop on Sustainable Computer Systems (HotCarbon'24).* ACM, New York, NY, USA, 8 pages.

HotCarbon'24, July 9, 2024, Santa Cruz, CA

1 INTRODUCTION

In 2022, data centers (DCs) consumed an estimated 1-3% of global power. This load is expected to double by 2026, driven by the increasing adoption of cloud computing, generative AI (large language models), bitcoin mining, and others. This surge will be equivalent to Sweden's annual electricity usage [17, 18], and DC power consumption is projected as high as 9% of US total load by 2030[13].

Data storage systems are the raison d'etre for datacenters, but storage is an often overlooked, but important part of their power consumption. Despite efficiency improvements in storage hardware (cost and power) and capacity optimization (e.g., deduplication) the total power for data center storage continues to increase. And as predicted by Jevon's Paradox, these efficiency improvements drive even faster growth in data creation (21.2% CAGR, which is projected to reach a staggering 221 ZB by 2026)[14, 21]. The result is continued rapid growth of storage power and carbon footprint.

Prior studies of storage carbon footprint have focused on embodied [24, 33], or on foreground storage activity [23, 25, 38]. The latter managed activities such as replication and load distribution to exploit low-carbon power in space or time.

In contrast, we study the background activities of storage systems for management and reliability with a focus on power consumption and carbon footprint. These storage background tasks are critical to ensure storage reliability (e.g., scrubbing, repair, archiving) and efficiency (e.g., deduplication, reclamation). These activities are temporally flexible (deferrable).

In the paper, we first characterize the UChicago HEP storage system, a small high-energy physics DC (HEP-S). Then scale it to a range of DC scenarios (HEP-L: large high-energy physics, Cloud: one hyperscale cloud, and Global Cloud: all of the hyperscale cloud). Next, we consider how exploiting the flexibility in storage background tasks to shift work temporally can reduce carbon emissions. While prior work focused on shifting background tasks to improve foreground tasks performance [1, 22, 39].

Key contributions include:

- Built a workload model for storage background reliability and management tasks (input/output, compute, and power requirements). Scaled the model to three DC scenarios (HEP-L, Cloud, Global Cloud), showed these background tasks consume a significant fraction of DC (2.1-4.9%) and storage power (12.8-27.2%). Data growth increases background task power consumption faster than total storage power.
- Scaled the model to diverse storage types (SSDs and HDDs), the results show that as the fraction of SSD's increases, storage system power decreases, but background tasks as a fraction of that increase (22.7-75% at 90% SSDs)

- Characterized the operational carbon footprint for storage systems in varied power grid (ISO) settings, particularly background tasks. Results show that footprint varies widely by 2.4x depending on ISO, and can be large, as much as that of 755k US homes power use.
- Carbon-aware scheduling can reduce background task emissions up to 82.8% in CAISO. Reductions will increase as future grids decarbonize, and greater average carbon-intensity (ACI) variability. By 2035, reductions of 96.9% may be possible in CAISO.

2 PROBLEM AND APPROACH

Problem: Distributed storage systems, crucial for data reliability and management, are challenged by the exponential growth of data. This data surge strains their ability to meet these demands without significant spikes in energy consumption and carbon footprint. Our research directly addresses this challenge, focusing on reducing the operational carbon footprint of storage systems. To achieve this, we aim to answer the following research questions:

- What are the characteristics of background workload?
- What is its carbon footprint? How does it change based on different DC storage capacity, configurations and management policies?
- How much potential carbon emission reduction could be achieved by temporal shifting of background tasks, both now and in the future?

Approach: The carbon footprint of a computing workload depends on both the time and location of its execution, reflecting the temporal and spatial variations in the ACI of the power grid [6, 8, 10, 26, 28, 35, 37]. We propose temporal shifting of the background workload to low carbon period. However, this approach hinges on the workload's inherent flexibility [7, 9, 11, 27, 40]. Background tasks are ideal candidates due to their lower priority compared to foreground tasks. In our study,

(1) We characterize the data volume and temporal flexibility of different background tasks in UChicago HEP storage system.
(2) We build a workload model for these tasks to understand their IO, compute and power cost. We extend this model to different DC configurations, storage capacity and management policies.
(3) We evaluate carbon aware scheduling policies in diverse ISO settings, today and future (2035).

3 STORAGE SYSTEMS' BACKGROUND WORKLOAD

We study the background workload of the UChicago HEP storage system (configuration and policies), which is part of the Worldwide Large Hadron Collider Computing Grid (WLCG) tiered storage system. The UChicago HEP storage system has a capacity of 19.7 PB for storage of replicated data copies generated from HEP experiment analyses. This data arrives from other WLCG tiers in a compressed file format. The storage system has 100 independent RAID-6 disk array systems, each configured with between 12 and 24 disks. Metadata of 130 GB is stored in a separate PostgreSQL database. Due to access from various geographical locations, the I/O activity exhibits random characteristics. Consistent with the HEP data, the file size distribution, in terms of storage capacity, is dominated by the Gigabyte range. We characterize the background workload based on the data volume, completion time, and period. These metrics were extracted through a comprehensive analysis of system logs spanning six months. The detailed characterization results are presented in the Table 1. We observe all these background tasks handle significant data volume within a year and are deferrable. We conclude background tasks are suitable workload for deferring to low carbon period. Note: While data scrubbing and deduplication are not performed in the UChicago HEP storage system because it stores replicated compressed data. We still consider them for wider storage relevance. For deduplication, we consider 2:1 ratio, 512 KB block size, and use the IPFS merkle tree [20] implementation.

3.1 Model for Storage Background Tasks

We describe the workload model for the background tasks, including their IO, compute, and power costs. Next, we scale the models to various DC scenarios.

IO Workload Model: We build the IO model for the background tasks defined in Table 1 to understand the IO cost of each task. The model is built by empirical modeling of different tasks and evaluated using traces of background tasks performed on HEP experiments analysis data. The model takes in storage capacity ($capacity_{total}$) and few other parameters based on the task as input (see Appendix A for the detailed description). We define the IO workload model for these background tasks as following:

$$IO_{total} = IO_{read} + IO_{write} \tag{1}$$

where IO_{read} and IO_{write} are the total amount of data read and written from/to the storage system respectively.

Compute Workload Model: Next, we build the compute model for background tasks to understand their computation cost. We find on analyzing the performance trace that total IO is a good predictor of instructions for the compute model. We estimate the CPU instructions for each task based ($cpu_{instructions}$) on the average million instructions per gigabyte (*MIPGB*) from it's performance profile.

$$cpu_{instructions} = IO_{total} * MIPGB \tag{2}$$

Note: We do not build compute model for Disk Scrub and RAID Rebuild because they are assigned a low priority and allocated less than 30% of compute resources by default in the hardware RAID controller, resulting in low CPU utilization. This prioritization scheme is recommended for optimal system performance. [2, 19].

Power Workload Model: Finally, we build power consumption model for the background tasks. We define the power consumption of system without storage and storage system as following:

$$Power_{system}(watt\text{-}hours) = AvgPower_{cpu} * cpu\text{-}hrs \tag{3}$$

where *cpu-hrs* is the computation time required to process the task and $AvgPower_{cpu}$ is 33% of CPU's TDP (assuming integer instructions) in Watts.

$$Power_{storage}(watt\text{-}hours) = AvgPower_{disk} * io\text{-}hrs \tag{4}$$

Table 1: Characterization of Background Tasks in UChicago HEP Storage System

Task	Description	Data Volume	Completion Time	Task Period	Deferrable
RAID Rebuild	Data reconstruction process performed on disk failure in a RAID array. Ensures data redundancy and reliability.	96TB-192TB	hours to days	Biweekly est.	Yes
Data Scrubbing	Checks disk for errors and repairs for silent data corruption (SDC).	96TB-192TB	hours to days	-	Yes
Metadata Backup	Metadata in database is backed up in a compressed format.	75GB (9x Compression)	~10 minutes	Every night	Yes
Catalog Summary	Summarize the files/objects stored in the storage system. It ensures all replica copies have consistent data view.	75GB	~25 minutes	Monthly	Yes
Storage Space Reclamation	Deleted file space reclaimed.	~30 M files (MB to GB)	∝ # files deleted	When 1000 or 0.5% of files deleted	Yes
Deduplication	Identify and delete of redundant data.	19.7 PB	Minutes to hours	-	Yes

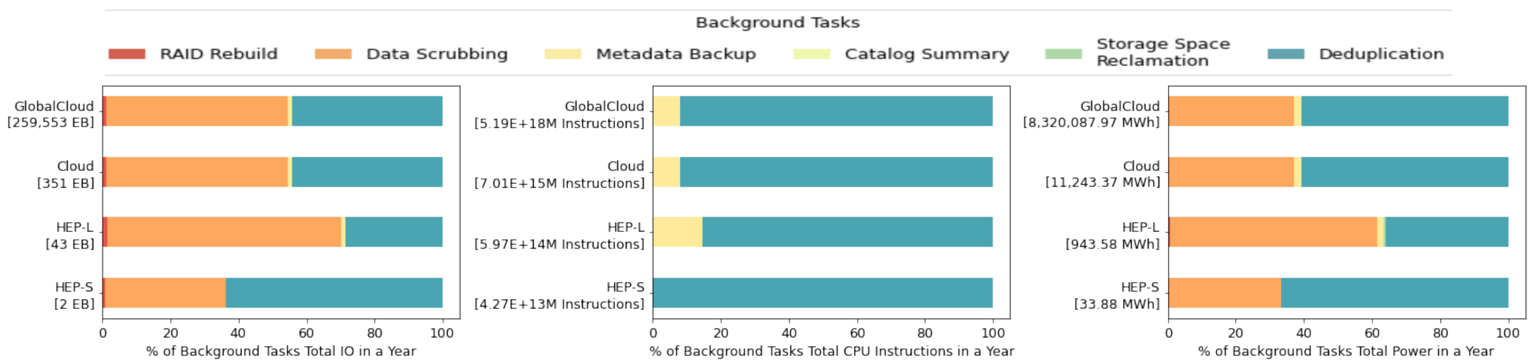

(a) (b) (c)

Figure 1: Background Tasks Workload Model Results for Different DC Scenarios

where *io-hrs* is disk read/write time (considered benchmarked throughput) and $AvgPower_{disk}$ is hard disk's operational power in Watts.

Our workload model captures the impact of DC configuration on IO, compute and power consumption of background tasks (See Figure 1). It incorporates several simplifying assumptions: I/O, compute, and power costs scale with the number of storage array groups. Following industry standards, we assume 10% of disk capacity remains unused. To simplify the model, *io-hrs* and *cpu-hrs* are considered equivalent for each task. Finally, after estimating per-task power consumption, we calculate annual usage based on task period (days) over a year.

3.2 Scaling the Model to Different DC Scenarios

We scale the workload model to three larger DC scenarios: large high-energy physics (HEP-L), single hyperscale cloud (Cloud), and all of the hyperscale cloud (GlobalCloud). Table 2 documents the workload model parameters for each DC scenario. These parameters depend on the DC activity. The HEP-L model parameters are based on CERN DC management policies and configuration. Cloud DC storage capacity is estimated using IDC data for cloud DC [31] and the number of hyperscale DCs in the world. Power consumption is estimated based on IEA data [18, 32]. For managing background tasks in the cloud scenarios, we use best practices [3, 16, 29, 36].

Table 2: Workload Model Parameters, Various DC Scenarios

Parameter	HEP-L	Cloud	GlobalCloud
Storage Capacity	1 EB [34]	6.5 EB est.	4.8 ZB est.
Annual DC Power	4 MW (2022) [5]	33 MW est. 2022	19.41 GW est. 2022
Storage Power (% of DC power)	21% [12]	18% [15]	
Avg Data Written Daily (% storage)	0.96%	0.053% est.	
Avg Data Deleted Daily (% storage)	0.97%	0.046% est.	
Deduplication Ratio	2:1	5:1 [30]	

In our workload model results, we observe that data scrubbing and deduplication dominates the IO because it reads all of the data (see Figure 1a). Metadata backup, catalog summary, and storage space reclamation read only the metadata. Deduplication is the most compute intensive as it analyzes and compares data to eliminate redundancy (see Figure 1b). The tasks with largest fractions of IO and compute (data scrubbing and dedupe), consume the majority of background tasks power (See Figure 1c). In all these results, we observe the total cost depends on the task frequency (1/period).

This depends on DC configuration (hardware type, failure rate, redundancy scheme) and management policies (target reliability). For example, in HEP-S, metadata backup's total IO for a single period (~120 GB) is 1.5 times less than catalog summary's total IO (~172 GB). But metadata backup task have a much higher total IO over all task periods in a year because it has a much shorter task period (21.7x).

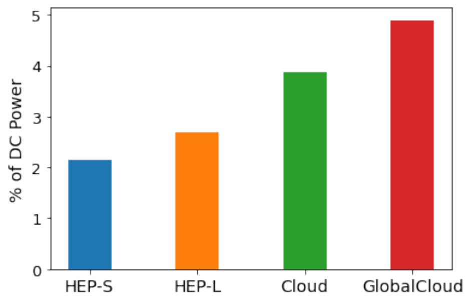

Figure 2: Background Tasks Power % of DC Power

Next, we consider the power consumption of storage system background tasks in the context of overall DC and storage system power. In Figure 2, we find each DC scenario's power consumption for background tasks accounts for 2.1-4.9% of DC annual power. It varies based on storage configuration and management policies. Although the cloud scenarios are alike in their configuration and activity they support, yet the fraction is greater for GlobalCloud compared to Cloud. This is because GlobalCloud's (collection of Cloud DCs) storage capacity increases by 737.5 times compared to Cloud (single cloud DC). However their storage power relative to the DC is same, as shown in Figure 3. The amount of work required to maintain storage reliability and management scales with data, leading to a significantly larger relative increase (26.7%) in their power share. We find similar trends on analyzing three years of CERN DC (HEP-L) storage capacity [34] and power consumption [4, 5] (see Figure 4). While storage capacity increases by 2.1x, efficiency improvements keep the DC power consumption consistent. Thus increasing the background tasks power fraction of overall DC and storage power.

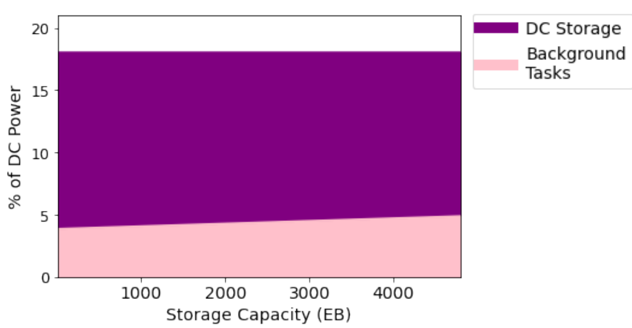

Figure 3: Background Tasks and Storage Power, Cloud Scenarios

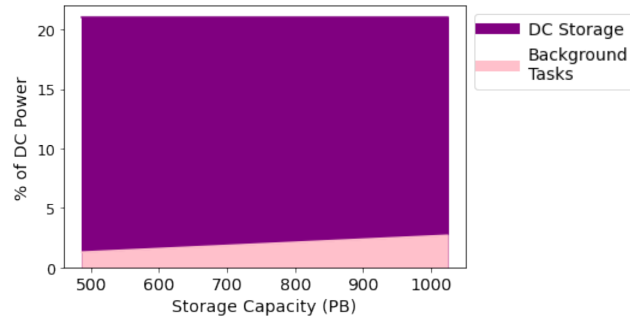

Figure 4: Background Tasks and Storage Power, HEP-L

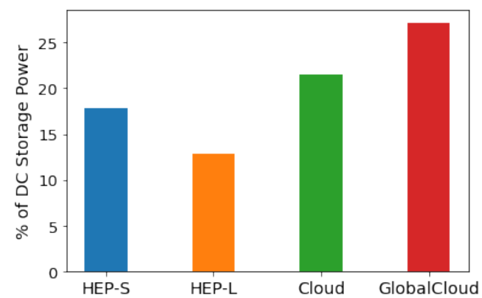

Figure 5: Background Tasks Power % of DC Storage Power

In Figure 5, we find that background tasks consume a significant portion of DC storage power (12.8-27.2%). This is larger than many expect, and a product of the continuous background activity (run independent of application activity) and systematic vigilance for data integrity (all data is touched independent of application data use). The focus of a DC's activity (e.g., compute, data archive, mixed) affects its storage power, causing this percentage to vary.

We conclude that background tasks, influenced by storage size, policies, and hardware configuration, can be major consumers of both overall and storage-specific DC power. Combined with rapid data growth, produces increasing power consumption.

3.3 Modeling Mixes of HDDs and SSDs

We adapt the workload model to encompass a broader hardware configuration with varying storage capacities provided by SSDs and HDDs. No updates are made to the IO and compute model because the total IO is independent of data storage type. However, we make the following changes to the power model:

$$\begin{aligned} Power_{storage}(watt\text{-}hours) = AvgPower_{disk} * io\text{-}hrs_{disk} + \\ AvgPower_{ssd} * io\text{-}hrs_{ssd} \end{aligned} \quad (5)$$

where $AvgPower_{ssd}$ is average power consumption of the SSD in Watts and $io\text{-}hrs_{ssd}$ is the total time required to read or write data to/from the SSD (considered benchmarked throughput).

Figure 6 illustrates the impact of varying storage capacity configurations on background tasks power consumption. The 100% HDD configuration represents the baseline. Results show how the configuration affects background task power consumption depends

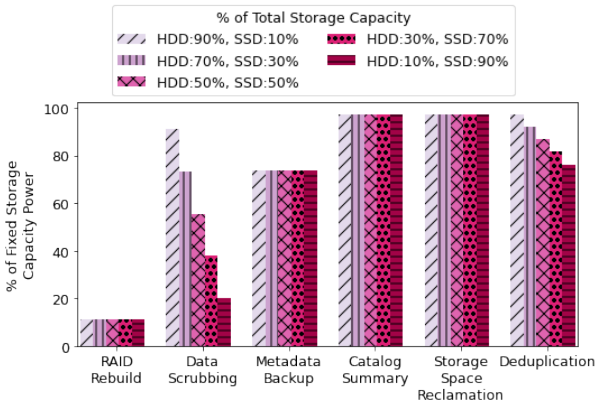

Figure 6: Background Tasks' Power Consumption Relative to Fixed Storage Capacity, Broader Hardware Configuration

on the task activity. For instance, data scrubbing exhibits an inverse relationship to SSD fraction because it reads data from both SSDs and HDDs. In contrast, tasks like catalog summary and those that only access data or metadata on SSDs maintain constant power consumption regardless of the SSD fraction. However, compute-related power consumption remains unaffected across all configurations. This means the overall decrease in power consumption is less pronounced for tasks like deduplication. As a result, even with a higher fraction of SSDs, background tasks continue to be a significant contributor to storage power.

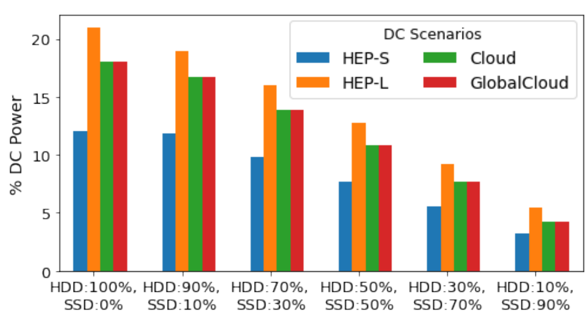

Figure 7: DC Storage Power % of DC Power, Broader Hardware Configuration

In Figure 7, we find that increased SSD fraction decreases DC storage power, but increases background tasks share of DC storage power (see Figure 8). This is because the reduction in overall storage power consumption is less significant for compute-intensive tasks or those that utilize data mostly stored on SSDs. While overall DC storage power consumption falls with increasing storage capacity provided by SSD, background tasks comprise a significant and growing fraction of DC storage power (22.7-75% at 90% SSDs). In this scenario, background tasks are the dominant factor of storage power consumption.

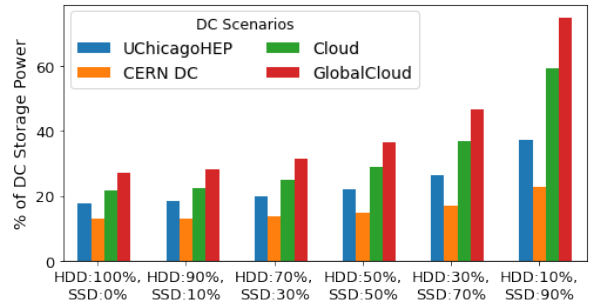

Figure 8: Background Tasks' Power % of DC Storage Power, Broader Hardware Configuration

4 SCHEDULING TO REDUCE OPERATIONAL CARBON

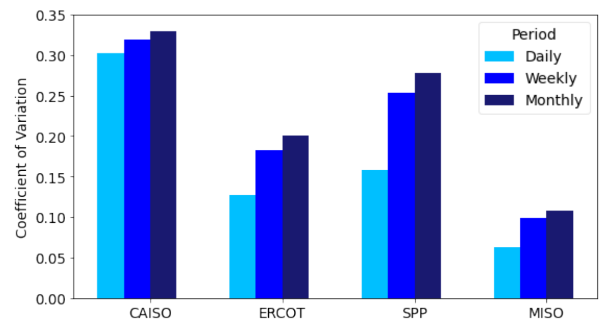

Figure 9: ACI Coefficient of Variation for Different Periods by ISO (2023)

We assess how temporal shifting of storage background tasks can reduce operational carbon emissions. Essential to this is temporal variation of average carbon intensity (ACI) of power. We use coefficient of variation (CV) to understand the significance of this variation. Figure 9 illustrates the CV of hourly ACI for different periods and ISOs in the North America. As the figure shows, a high CV, like in CAISO, indicates significant fluctuations in ACI compared to the average. Conversely, a low CV, like in MISO, suggests a more consistent pattern. Daily and weekly ACI variability within each ISO stems from differing generation mix (wind, solar, gas, etc.). For instance, CAISO experiences daily cycles with lowest ACI during the day because it is dominated by solar generation. In contrast, SPP is wind-dominated, and thus has more irregular cycles, typically several days to a week driven by weather fronts passing through areas. These variations provide opportunities to reduce carbon footprint by moving flexible computing load to low carbon periods.

It's also worth noting that the long-term average of carbon intensity varies by 2.4x across power grids. This geographic variation means that GlobalCloud would produce 3.8 million MT if in MISO (equivalent to annual electricity use of 755K US homes) but only 1.6 million MT if in CAISO (equivalent to annual electricity use of 323k US homes).

4.1 Scheduling Policies and Constraints

We explore several policies that encompass commonly used scheduling configurations and carbon optimal approach: (i) Naive: starts tasks at midnight; (ii) Naive Random (NaiveR): starts tasks at a random time in the next 24 hours; (iii) Greedy Carbon Constraints (GreedyC): schedules tasks during low-carbon periods with task period constraints; (iv) Greedy Carbon No Constraints (GreedyNoC): schedules tasks during low-carbon periods with no task period constraints. We use the metric, reduced carbon emissions, and compare the execution with average ACI values in each respective ISO (data source: RiPiT[35]).

4.2 Results

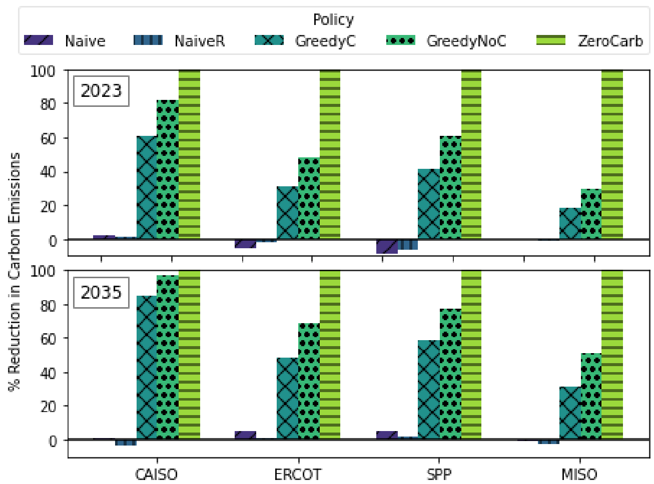

Figure 10: % Reduction of Scheduling Approaches, GlobalCloud (Present and Future)

Figure 10 shows the percentage reduction in carbon emissions achieved by different policies. For brevity, we present the results for GlobalCloud; other DC configuration results were similar.

Because Naive and NaiveR do not use any carbon information, they do not achieve a large change in carbon. In fact, the realized change depends somewhat arbitrarily on the carbon emissions at the time the tasks happen to run.

In 2023, GreedyC achieves large reductions of 59.5% and 41% in CAISO and SPP respectively, due to their high ACI variation. In others (ERCOT, MISO), significant, but lesser reductions of 18-30.8% are achieved. Without the scheduling window constraint, GreedyNoC achieves even greater reductions of 82.8% in CAISO. It is closer to the ideal scenario of Zero Carbon (ZeroCarb) where tasks are scheduled only during periods with zero carbon emissions.

This highlights the effectiveness of temporally shifting background tasks to periods of lower carbon intensity. Also, greater ACI variation leads to larger reductions. This is further supported by the future grid scenario (2035, modeled as described in [9]) with its higher ACI variation. As shown in Figure 10, GreedyC achieves significant reductions between 30.7% and 84.9% across different ISOs. While GreedyNoC achieves substantial reduction of 96.9% in CAISO, remarkably close to ZeroCarb scenario. [1]

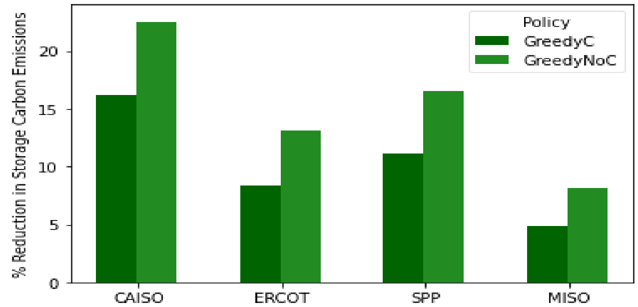

Figure 11: % Reduction of DC Storage Carbon Footprint, 2023, GlobalCloud

On translating the emission reduction achieved by carbon aware policies into the percentage reduction of the DC storage operational carbon footprint, we find significant reductions. GreedyC achieves a 4.9-16.2% reduction, and GreedyNoC achieves an 8.1-22.5% reduction (see Figure 11).

Our study shows that background tasks are significant but overlooked source of data center carbon emissions. We demonstrate that strategically scheduling these flexible tasks during low-carbon periods can substantially reduce the carbon footprint of storage systems.

5 SUMMARY AND FUTURE WORK

Our workload model characterizes the power consumption for storage background tasks, showing their significant power consumption relative to total DC and storage power. As storage capacity increases, the background tasks power consumption grows as a fraction of storage power. By exploiting the tasks' flexibility to align with low-carbon intensity power, we show how storage operators can reduce their operational carbon footprint. Further, projections to 2035 show that the benefit is increasing as grids decarbonize. Several promising future research directions include study of how deferring storage background tasks affects data reliability, and tradeoff with carbon optimization, and study of other datacenter configurations, and of course, more power grids (and future!).

ACKNOWLEDGMENTS

We thank the reviewers for their insights. We are grateful to Lincoln Bryant, Judith Stephens, Ilija Vukotic and Professor Robert Gardner for their support and access to HEP experiment data, test infrastructure. This work is supported in part by NSF Grant CNS-1901466, and the VMware University Research Fund. We also thank the Large-scale Sustainable Systems Group members for their ideas, enthusiasm, and contributions to this work!

[1] Amongst the four ISOs studied, only one exhibits a diurnal pattern in ACI. This correlation between low-carbon periods and high storage system utilization periods could potentially impact foreground task's performance.

REFERENCES

[1] George Amvrosiadis, Angela Demke Brown, and Ashvin Goel. 2015. Opportunistic storage maintenance. In *Proceedings of the 25th Symposium on Operating Systems Principles* (Monterey, California) *(SOSP '15)*. Association for Computing Machinery, New York, NY, USA, 457–473. https://doi.org/10.1145/2815400.2815424

[2] BROADCOM. 2023. RAID Rebuild Rate. https://techdocs.broadcom.com/us/en/storage-and-ethernet-connectivity/enterprise-storage-solutions/megaraid8-tri-mode-software/1-0/v11667295/v11667323/v11667611.html#:~:text=The%20default%20rebuild%20rate%20is%2030%20percent Online; Accessed 1 November 2023.

[3] CEPH. 2023. OSD Config Reference. https://docs.ceph.com/en/latest/rados/configuration/osd-config-ref/#scrubbing Online; Accessed 1 November 2023.

[4] CERN. 2021. Key Facts and Figures – CERN Data Centre. https://information-technology.web.cern.ch/sites/default/files/CERNDataCentre_KeyInformation_Nov2021V1.pdf Online; Accessed 1 November 2023.

[5] CERN. 2022. Key Facts and Figures – CERN Data Centre. https://information-technology.web.cern.ch/sites/default/files/CERNDataCentre_KeyInformation_July2022_V3.pdf Online; Accessed 1 November 2023.

[6] Andrew A Chien. 2018. *Characterizing opportunity power in the California independent system operator (CAISO) in years 2015-2017*. Technical Report. TR-2018-07.

[7] Andrew A Chien. 2023. GenAI: Giga ,TeraWatt-Hours,andGigaTonsofCO2. *Commun. ACM* 66, 8 (2023), 5–5.

[8] Andrew A. Chien and Liuzixuan Lin. 2024. As Grids Reach 100% Renewable at Peak, Growing Curtailment of 8 Gigawatts Looms as a Challenge to Decarbonization. *SIGENERGY Energy Inform. Rev.* 4, 1 (feb 2024), 3–10. https://doi.org/10.1145/3649432.3649434

[9] Andrew A Chien, Liuzixuan Lin, Hai Nguyen, Varsha Rao, Tristan Sharma, and Rajini Wijayawardana. 2023. Reducing the Carbon Impact of Generative AI Inference (today and in 2035). In *Proceedings of the 2nd Workshop on Sustainable Computer Systems* (Boston, MA, USA) *(HotCarbon '23)*. Association for Computing Machinery, New York, NY, USA, Article 11, 7 pages. https://doi.org/10.1145/3604930.3605705

[10] Andrew A Chien, Fan Yang, and Chaojie Zhang. 2016. Characterizing curtailed and uneconomic renewable power in the mid-continent independent system operator. *arXiv preprint arXiv:1702.05403* (2016).

[11] Andrew A Chien, Chaojie Zhang, Liuzixuan Lin, and Varsha Rao. 2022. Beyond PUE: Flexible Datacenters Empowering the Cloud to Decarbonize. *USENIX Hot Carbon Workshop* (2022).

[12] Dave Britton, Simone Campana and Bernd Panzer. 2023. A holistic study of the WLCG energy needs for HL-LHC. https://indico.jlab.org/event/459/contributions/11499/attachments/9236/14205/WLCGEnergyNeedsCHEP2023.pdf Online; Accessed 1 February 2024.

[13] EPRI. 2024. EPRI Study: Data Centers Could Consume up to 9% of U.S. Electricity Generation by 2030. https://www.epri.com/about/media-resources/press-release/q5vU86fr8TKxATfX8IHf1U48Vw4r1DZF#:~:text=EPRI%20Study%3A%20Data%20Centers%20Could,U.S.%20Electricity%20Generation%20by%202030&text=PALO%20ALTO%2C%20Calif.,the%20data%20centers%20processing%20AI. Online; Accessed 1 July 2024.

[14] Eric Burgener and John Rydning,IDC. 2022. High Data Growth and Modern Applications Drive New Storage Requirements in Digitally Transformed Enterprises. https://www.delltechnologies.com/asset/en-my/products/storage/industry-market/h19267-wp-idc-storage-reqs-digital-enterprise.pdf Online; Accessed 21 March 2024.

[15] Gartner. 2022. How to Make the Data Center Eco-Friendly. https://www.gartner.com/en/newsroom/press-releases/2022-12-06-how-to-make-the-data-center-eco-friendly Online; Accessed 1 January 2024.

[16] Google Cloud. 2023. Garbage collection overview. https://cloud.google.com/bigtable/docs/garbage-collection Online; Accessed 1 November 2023.

[17] IEA. 2023. Data Centres and Data Transmission Networks. https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks# Online; Accessed 1 November 2023.

[18] IEA. 2023. Electricity 2024. https://www.iea.org/reports/electricity-2024/executive-summary#abstract Online; Accessed 1 February 2024.

[19] Intel. 2023. Intel RAID Software User Guide for full featured and entry level RAID Controllers. https://www.intel.com/content/dam/support/us/en/documents/server-products/raid-products/SWUG_full_featured_entry_RAID.pdf Online; Accessed 1 November 2023.

[20] IPFS. 2023. Merkle dag github repository. https://github.com/ipfs/boxo/tree/main/ipld/merkledag Online; Accessed 1 July 2023.

[21] Willam Stanley Jevons. 1865. *The Coal Question; An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal Mines*. Macmillan and Company. London.

[22] M. Kachmar and D. Kaeli. 2020. A Smart Background Scheduler for Storage Systems. In *2020 28th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*. IEEE Computer Society, Los Alamitos, CA, USA, 1–8. https://doi.org/10.1109/MASCOTS50786.2020.9285967

[23] William Katsak, Inigo Goiri, Ricardo Bianchini, and Thu D Nguyen. 2015. Greencassandra: Using renewable energy in distributed structured storage systems. In *2015 Sixth International Green and Sustainable Computing Conference (IGSC)*. IEEE, 1–8.

[24] Baolin Li, Rohan Basu Roy, Daniel Wang, Siddharth Samsi, Vijay Gadepally, and Devesh Tiwari. 2023. Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '23)*. Association for Computing Machinery, New York, NY, USA, Article 19, 15 pages. https://doi.org/10.1145/3581784.3607035

[25] Daping Li, Xiaoyang Qu, Jiguang Wan, Jun Wang, Yang Xia, Xiaozhao Zhuang, and Changsheng Xie. 2018. Workload Scheduling for Massive Storage Systems with Arbitrary Renewable Supply. *IEEE Transactions on Parallel and Distributed Systems* 29, 10 (2018), 2373–2387. https://doi.org/10.1109/TPDS.2018.2820070

[26] Liuzixuan Lin and Andrew A Chien. 2020. Characterizing stranded power in the ERCOT in years 2012-2019: A preliminary report. *University of Chicago CS Tech Report* (2020).

[27] Liuzixuan Lin and Andrew A Chien. 2023. Adapting Datacenter Capacity for Greener Datacenters and Grid. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*. 200–213.

[28] Electricity Maps. 2023. Electricity Maps. https://app.electricitymaps.com/map.

[29] Microsoft. 2023. Advanced Data Deduplication settings. https://learn.microsoft.com/en-us/windows-server/storage/data-deduplication/advanced-settings Online; Accessed 1 November 2023.

[30] Steven Scully and W Benjamin. 2010. Improving Storage Efficiencies with Data Deduplication and Compression. *IDC, May* (2010). Online; Accessed 1 November 2023.

[31] Statista. 2020. Total installed based of data storage capacity in the global datasphere from 2020 to 2025. https://www.statista.com/statistics/1185900/worldwide-datasphere-storage-capacity-installed-base/ Online; Accessed 1 November 2023.

[32] Synergy Research Group. 2024. Pipeline of Over 300 New Hyperscale Data Centers Drives Healthy Growth Forecasts. https://www.srgresearch.com/articles/pipeline-of-over-300-new-hyperscale-data-centers-drives-healthy-growth-forecasts Online; Accessed 1 February 2024.

[33] Swamit Tannu and Prashant J. Nair. 2023. The Dirty Secret of SSDs: Embodied Carbon. *SIGENERGY Energy Inform. Rev.* 3, 3 (oct 2023), 4–9. https://doi.org/10.1145/3630614.3630616

[34] Tim Smith, CERN. 2023. An exabyte of disk storage at CERN. https://home.web.cern.ch/news/news/computing/exabyte-disk-storage-cern Online; Accessed 1 November 2023.

[35] University of Chicago. 2021. Right Place, Right Time (RiPiT) Carbon Emissions Service. http://ripit.uchicago.edu Online; Accessed 1 January 2024.

[36] Vivek Menezes, CockroachDB. 2023. Trust, but verify: How CockroachDB checks replication. https://www.cockroachlabs.com/blog/trust-but-verify-cockroachdb-checks-replication/ Online; Accessed 1 November 2023.

[37] WattTime. 2023. WattTime. https://www.watttime.org.

[38] Zichen Xu, Nan Deng, Christopher Stewart, and Xiaorui Wang. 2015. Cadre: Carbon-aware data replication for geo-diverse services. In *2015 IEEE International Conference on Autonomic Computing*. IEEE, 177–186.

[39] Feng Yan, Alma Riska, and Evgenia Smirni. 2012. Busy bee: how to use traffic information for better scheduling of background tasks. In *Proceedings of the 3rd ACM/SPEC International Conference on Performance Engineering* (Boston, Massachusetts, USA) *(ICPE '12)*. Association for Computing Machinery, New York, NY, USA, 145–156. https://doi.org/10.1145/2188286.2188308

[40] Jiajia Zheng, Andrew A. Chien, and Sangwon Suh. 2020. Mitigating Curtailment and Carbon Emissions through Load Migration between Data Centers. *Joule* (October 2020). https://doi.org/10.1016/j.joule.2020.08.001

A BACKGROUND TASKS WORKLOAD MODEL

A.1 Workload Model: IO

Table 3 describes the parameters considered for the workload model in detail. We evaluate the model using the traces of background tasks performed on HEP data. The traces were collected using command line tools like iostat, top. We used two x86_64 bit Intel(R) Xeon(R) Gold 6226 CPU @ 2.70GHz processors for evaluation and the data was stored on the RAID-6 array of 12 hard disk drives, where size of each disk ($size_{hdd}$) is 16TB. We compute the model constant by taking the root mean square error percentages on the model and actual task results.

Table 3: IO model for storage management tasks

Task	Parameter	Total IO Model (GB) (IO_{total}=IO_{read}+IO_{write})	Notes
RAID Rebuild	$capacity_{hdd}$, $size_{hdd}$, $num_{disks\_per\_array}$	($size_{hdd}$ * ($num_{disks\_per\_array}$-1) + $size_{hdd}$)	We consider 12 disks per raid array group ($num_{disks\_per\_array}$).
Data Scrubbing	$capacity_{hdd}$, $size_{hdd}$, $num_{disks\_per\_array}$	($size_{hdd}$* $num_{disks\_per\_array}$ + 2 * $size_{hdd}$)	The number of disks per array group is similar to RAID rebuild assumption.
Data Backup	$capacity_{hdd}$	$capacity_{hdd}$+ $capacity_{hdd}$*$ratio_{compression}$	$ratio_{compression}$ denotes the compression ratio.
Catalog Summary	$capacity_{hdd}$, $size_{file}$	(num_{files} * $size_{index}$) + (num_{files} * $size_{result}$)	The number of files (num_{files}) is computed based on file size distribution in the UChicago HEP storage system and $size_{index}$ denotes the average file entry size. We consider file path output size ($size_{result}$) based on average output file path length.
Storage Space Reclamation	$capacity_{hdd}$, $size_{file}$, num_{files_del}	(num_{files} * $size_{index}$) + (num_{files_del} * $size_{update\_index}$)	Number of files to be deleted is denoted by num_{files_del} and $size_{update\_index}$ denotes the entry table index size. The total number of files (num_{files}) is estimated similar to Catalog Summary.
Deduplication	$capacity_{hdd}$, $ratio_{dedup}$	$capacity_{hdd}$ + ($capacity_{hdd}$ - $capacity_{hdd}$/$ratio_{dedup}$) * $size_{link}$	The IPFS merkle tree [20] implementation is used and dedupe block size of 512 KB is considered. Along with deduplication ratio ($ratio_{dedup}$) to be 2:1 and linking block size ($size_{link}$) to be 4 bytes.