Understanding the Operational Carbon Footprint of Storage Reliability and Management

Varsha Rao[1] and Andrew A. Chien[1,2]
[1]University of Chicago
[2]Argonne National Laboratory

THE UNIVERSITY OF CHICAGO

Argonne NATIONAL LABORATORY

CERES Center for Unstoppable Computing

Contributions

- First study of the carbon footprint of background tasks for storage systems.
- Background tasks footprint can be MOST of the storage system's operational carbon footprint (75% at 90% SSDs).
- Background tasks footprint can be nearly eliminated with temporal shifting (82.8% today and 96.9% in 2035).

Storage Systems are Growing Rapidly

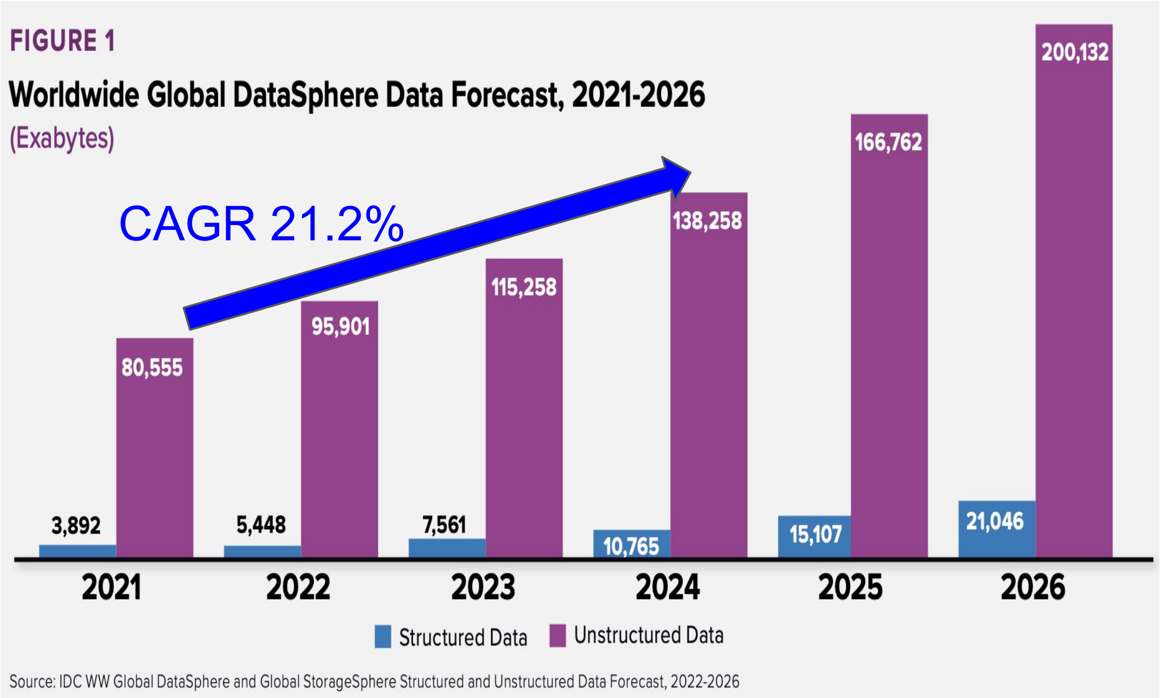

Efficiency improvements in storage technology are not sufficient to compensate.

Storage power and operational carbon footprint continue to increase!

Objective & Approach

Objective: Understand operational carbon footprint of large storage systems and show how to reduce it.

Approach:

- Build background tasks workload model (flexible)
- Estimate potential carbon reduction (temporal shifting)

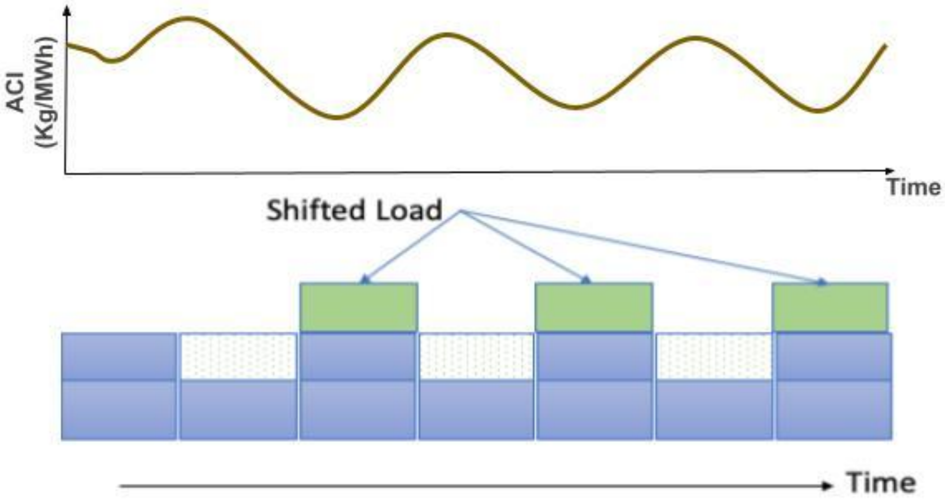

Example: UChicago HEP Storage System

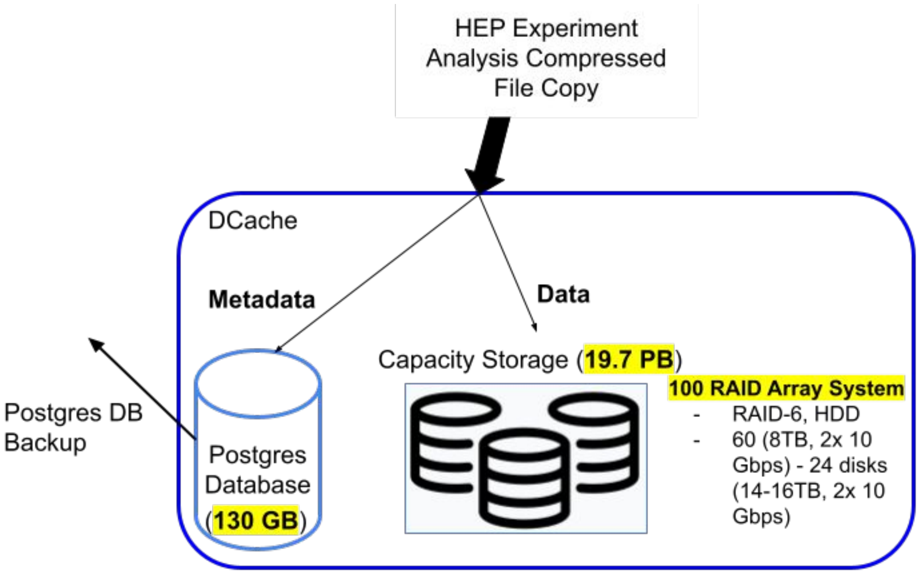

Background Tasks

- **Redundancy & Reliability**: RAID Rebuild, Metadata Backup
- **Reliability**: Data Scrubbing
- **Consistent Data View**: Catalog Summary
- **Efficient Storage Space Utilization**: Storage Space Reclamation, Deduplication

Background tasks handle a large data volume and are deferrable.

Background Tasks Workload Model & Scenarios

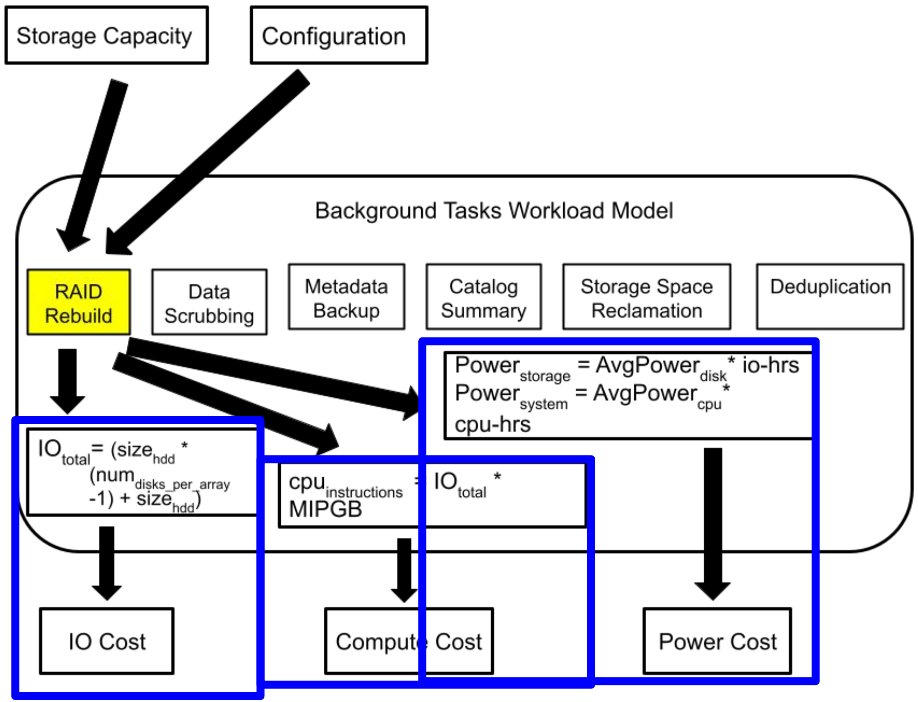

Datacenter Scenarios

Parameter	HEP-L	Cloud	GlobalCloud
Storage Capacity	1 EB [23]	6.5 EB est.	4.8 ZB est.
Annual DC Power	4 MW (2022) [4]	33 MW est. 2022	19.41 GW est. 2022
Storage Power (% of DC power)	21% [6]	18% [8]	

Background task's **power consumption** depends on storage system configuration and management choices

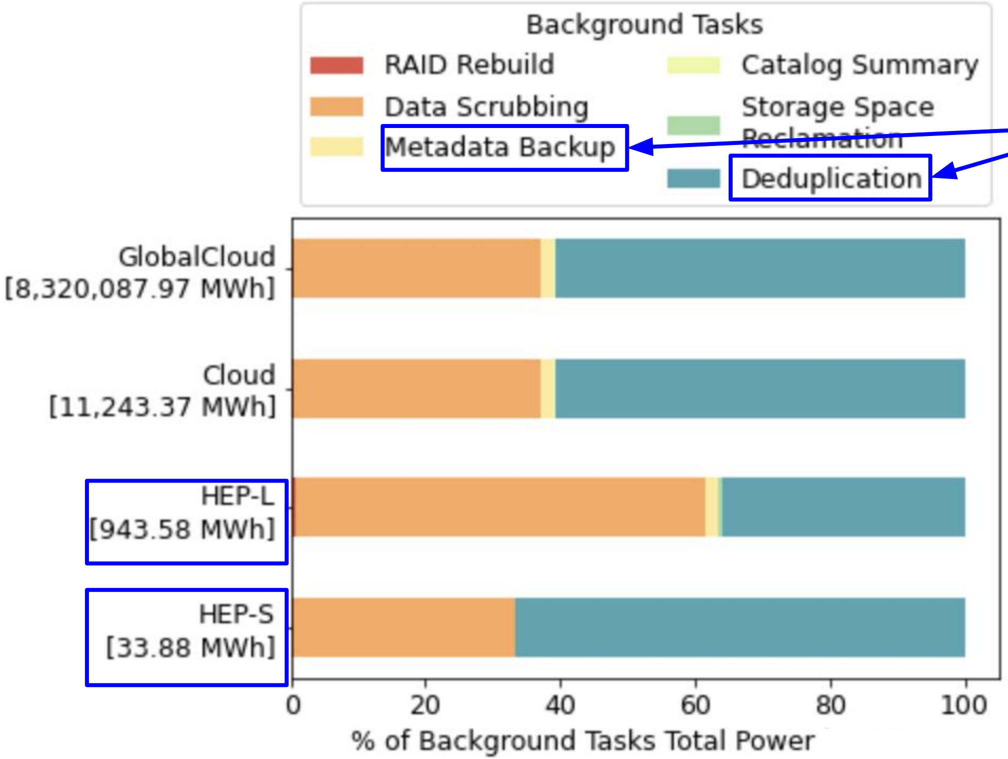

Tasks touch all the data consume the majority of power.

Task frequency dictates the IO, compute, and power costs.

- Deduplication (52 times/year, HEP-S), (12 times/year, HEP-L)

Background Tasks are a Significant Fraction of DC Storage Power

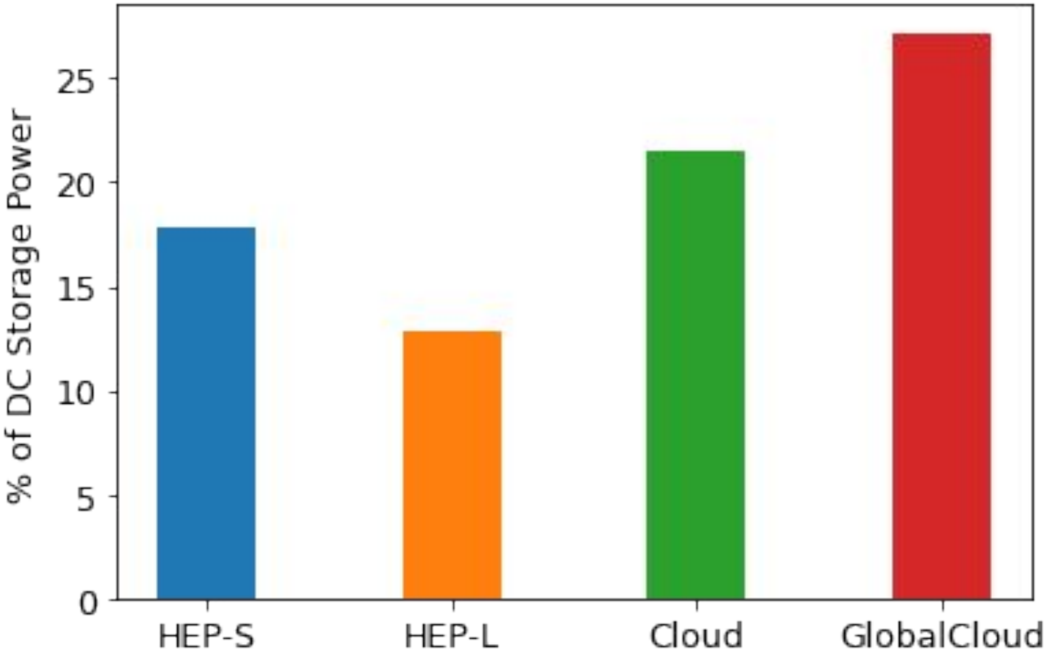

Figure: Background Tasks Power % of DC Storage Power

Data growth drives storage reliability and management power (e.g. Background Tasks)

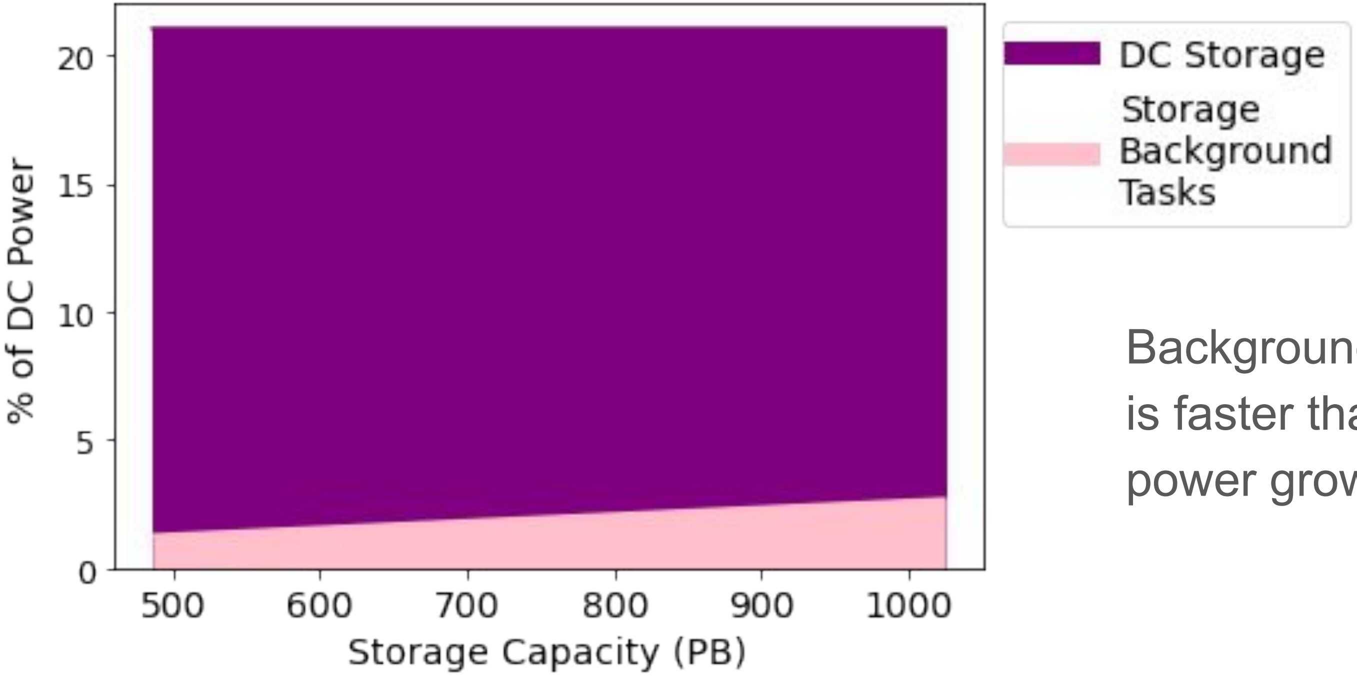

Figure: Background Tasks and Storage Power, CERN data center

Background tasks power growth is faster than storage system power growth.

Background Tasks can be the Dominant Contributor to Storage Power (75% at 90% SSD)

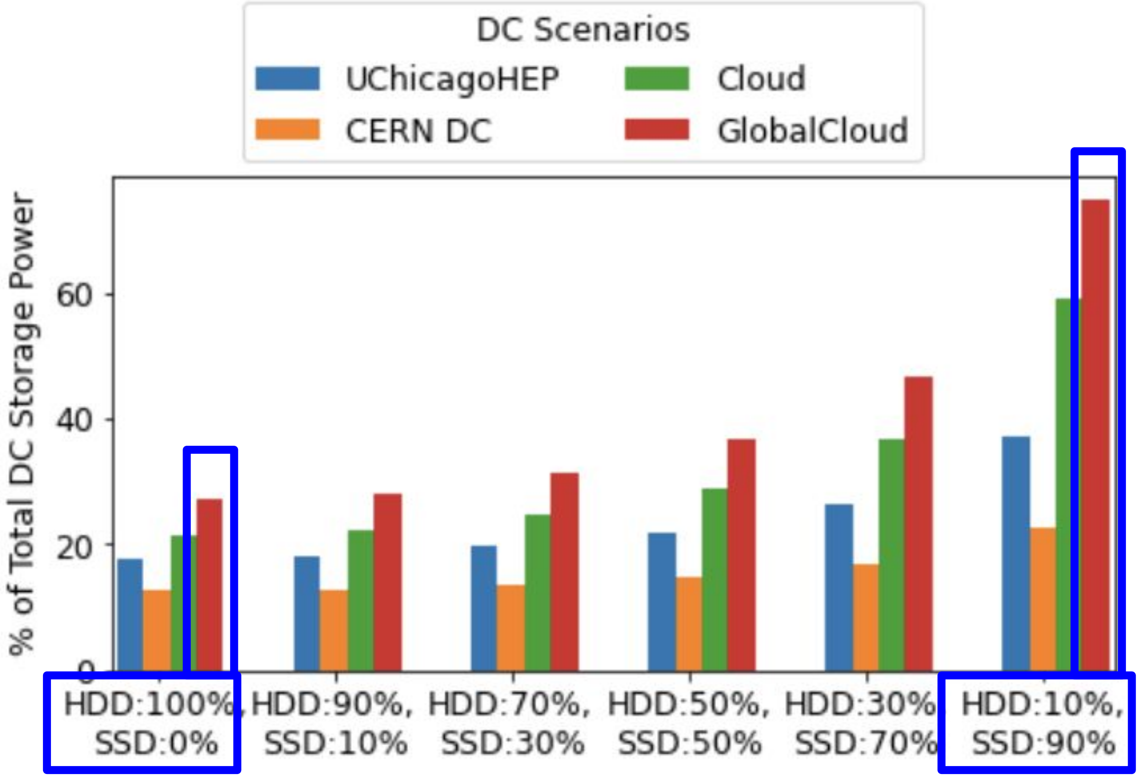

% of SSDs ⇧

Storage Power % DC Power ⇩

Background Tasks Power % Storage Power ⇧

IO is only a part of the power cost for background tasks!

Variation in ACI Provides Opportunity for Shifting Tasks

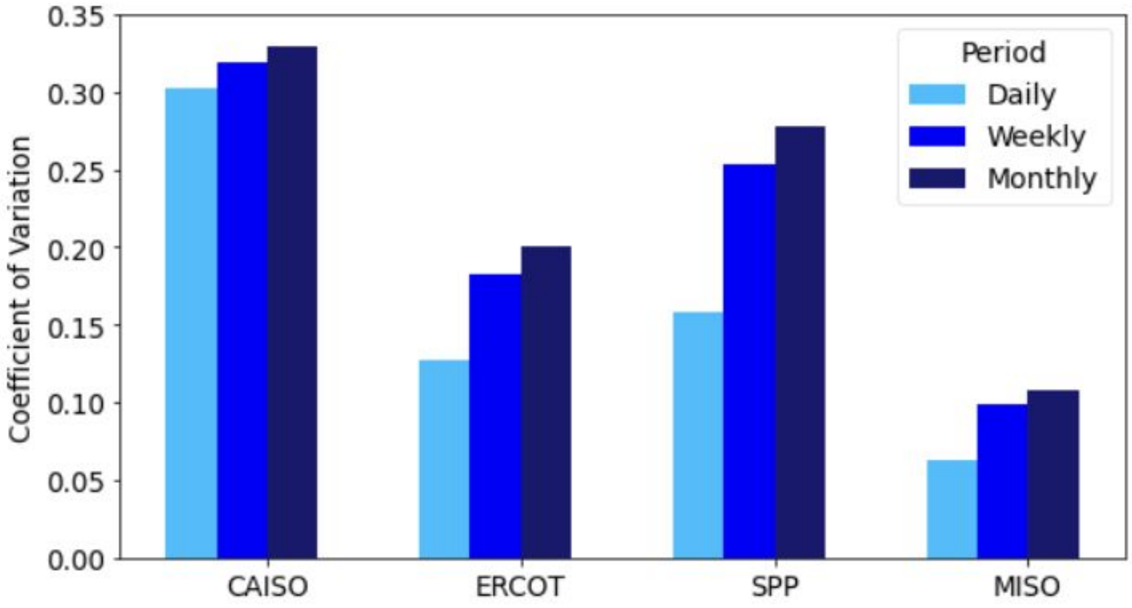

Figure: ACI Coefficient of Variation (CV) for Different Periods by ISO (2023)

$$CV = \frac{\text{Standard deviation}}{\text{Mean}}$$

- High CV indicates larger fluctuations
- Low CV indicates smaller fluctuations

Evaluation of Potential Footprint Savings

Operational carbon footprint equivalent to annual electricity use of US homes: **755K** (MISO) and **323k** (CAISO)

Scheduling Policies:

- ***Naive***: Starts at midnight
- ***NaiveR***: Starts at random time in next 24 hours
- ***GreedyC***: Optimizes for carbon with task period constraints
- ***GreedyNoC***: Optimizes for carbon with no task period constraints

Metric: % Reduction in carbon emissions relative to the baseline.

Shifting Background Tasks can Reduce most Emissions

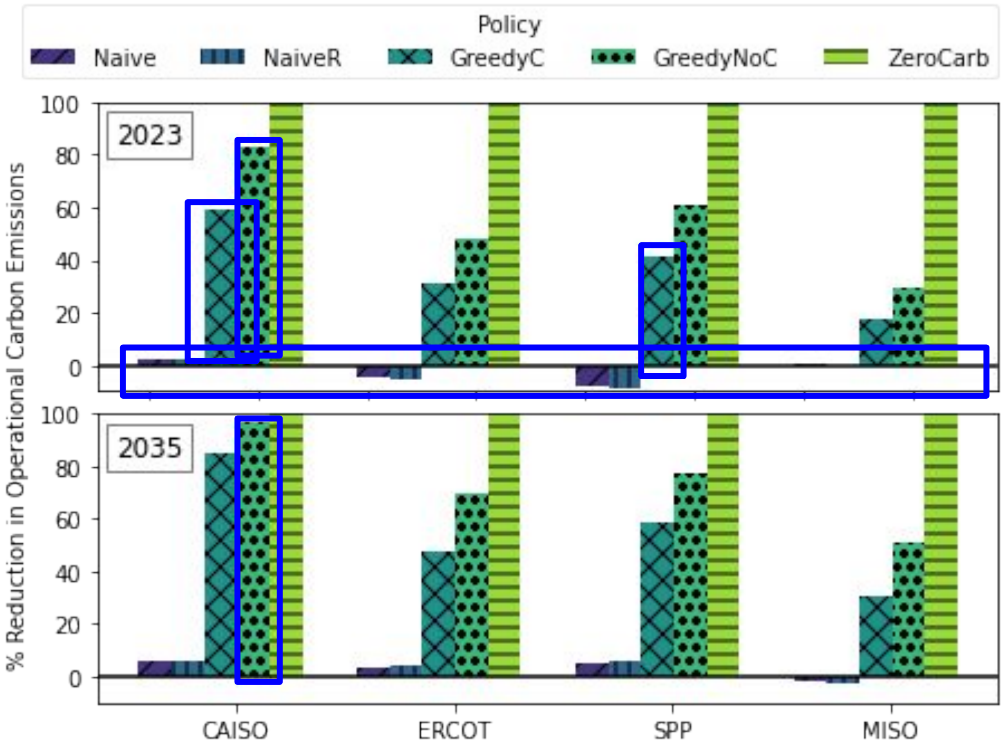

GlobalCloud (Present and Future)

GreedyNoC reduction in emissions: *CAISO* **82.8%**.

GreedyC reduction in emissions: *CAISO* **59.5%** , *SPP* **41%**

In Naive and NaiveR policy results depends on ACI value at task running time.

In 2035, even higher reduction is achieved by GreedyNoC (Up to **96.9%**)

Summary

- Data growth drives storage reliability and management power faster than storage power!
- Background tasks can be MOST of the storage system's operational carbon footprint (**75% at 90% SSDs**).
- Temporal shifting of background tasks is effective in reducing operational carbon footprint of storage system.

Future Work

- How does deferring storage background tasks affect data reliability? (and mitigation)
- How do deferred storage background tasks impact foreground tasks performance?
- What about other datacenter configurations, management policies, storage system designs?
- What about temporal dynamics of other power grids?

Thanks for listening!

This work is part of Zero-Carbon Cloud Project active since 2015 (http://zccloud.cs.uchicago.edu/)

- AI DC Challenges: HotCarbon’23, e-Energy’24
- Grid impacts of datacenter adaptation: e-Energy’21, IEEE Trans. on Renewable Energy’17, e-Energy’23
- Carbon-aware stream processing: CLOUD’23
- Characterization of growing “stranded power”: MISO (AIMS Energy’18), CAISO, ERCOT (UChicago CS TR-2018-07, 2020-06)
- Geographical load shifting: Joule’20
- Challenges for resource management under variable datacenter capacity: JSSPP’21
- Grid generator type inference: e-Energy’20
- Economic viability of Zero-Carbon Cloud datacenters: TPDS’17, IPDPS’16

Thanks to the NSF, Samsung and VMware for support, and thanks to the faculty and students who have worked on ZCCloud!

For any questions, please reach out to varsharao@uchicago.edu