OpenDC-STEAM:

Realistic Modeling and Systematic Exploration of Composable Techniques for Sustainable Datacenters

Dante Niewenhuis, Sacheendra Talluri, Alexandru Iosup, Tiziano De Matteis

Dante Niewenhuis

D.niewenhuis@vu.nl

@danteniewenhuis

Learn more: opendc.org

github.com/atlarge-research/opendc-steam

https://github.com/atlarge-research/opendc

@Large Research
Massivizing Computer Systems

Dante Niewenhuis, Sacheendra Talluri, Alexandru Iosup, Tiziano De Matteis: OpenDC-STEAM: Realistic Modeling and Systematic Exploration of Composable Techniques for Sustainable Datacenters. CCGrid (2026)

GRAPH MASSIVIZER

VU VRIJE UNIVERSITEIT AMSTERDAM

CCGRID 2026

Data Center Sustainability

How Los Angeles Wildfires Showcase Climate Change and Community Action

Wildfires rage in Los Angeles, leaving destruction and hard questions in their wake

[1]

The Big Thaw

As the climate warms, how much, and how quickly, will Earth's glaciers melt?

[2]

2024 is on track to be hottest year on record as warming temporarily hits 1.5°C

[3]

Climate change and droughts: What's the connection?

As average temperatures continue to climb, drought has become a permanent part of our vocabulary.

[4]

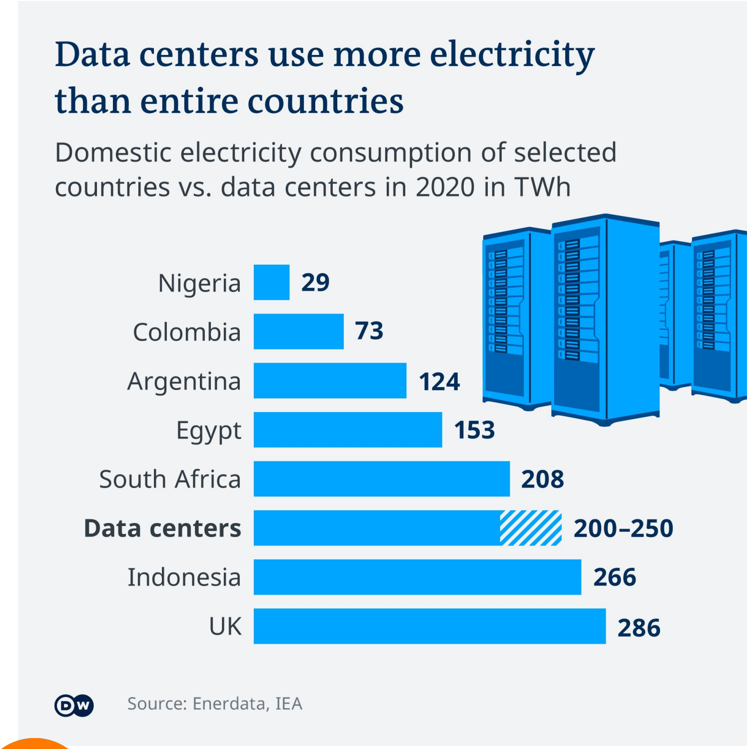

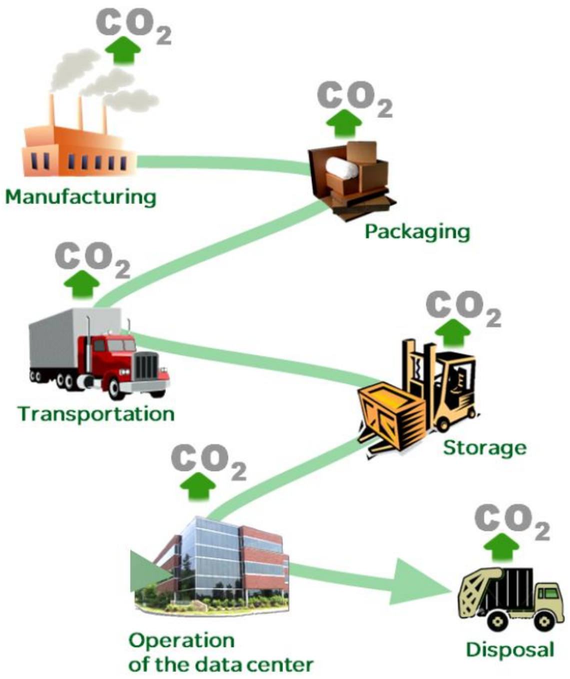

Carbon-producing phases of a data center[5]

! Reducing carbon footprint is challenging!

1. https://www.mcgilldaily.com/2025/01/how-los-angeles-wildfires-showcase-climate-change-and-community-action/
2. https://www.nationalgeographic.com/environment/article/big-thaw
3. https://wmo.int/news/media-centre/2024-track-be-hottest-year-record-warming-temporarily-hits-15degc
4. https://www.mcgilldaily.com/2025/01/how-los-angeles-wildfires-showcase-climate-change-and-community-action/
5. https://www.insight.com/content/dam/insight/en_US/pdfs/apc/apc-estimating-data-centers-carbon-footprint.pdf

Sustainability Techniques

⚠ We identify three problems with past research

- Evaluated using analytical model
- Tested on (a few) specific configurations
- Combining techniques is rarely discussed

OpenDC-STEAM

A **composable** data center simulator used to explore sustainability techniques.

- Trace Based
  - Replay workload traces
  - Determine performance and sustainability
- Composability
  - ✓ Flexible and extensible
  - ✓ Explore **what-if** scenarios

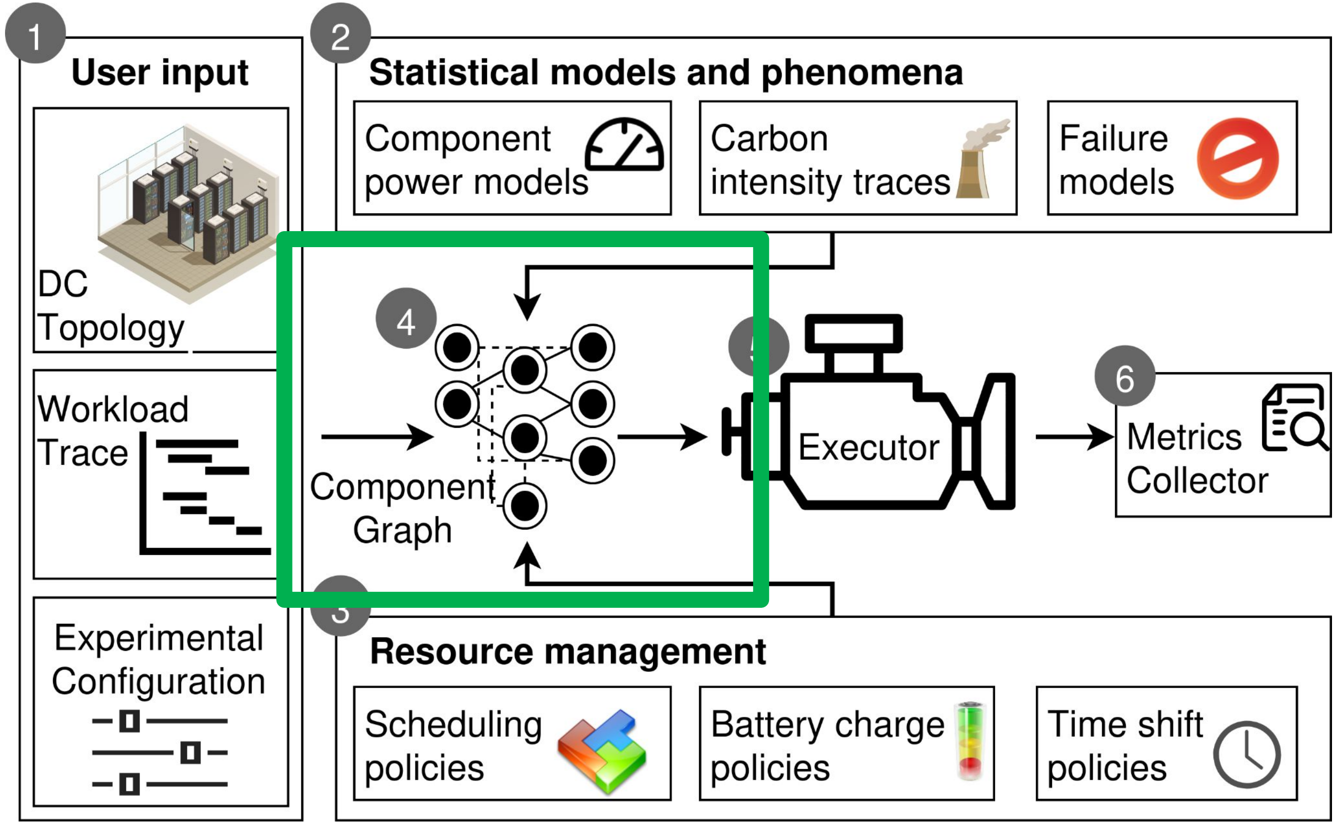

Fig: Diagram of the OpenDC-STEAM architecture

Simulating Sustainability

Operational Carbon

Carbon emitted as a result of the energy used by a data center

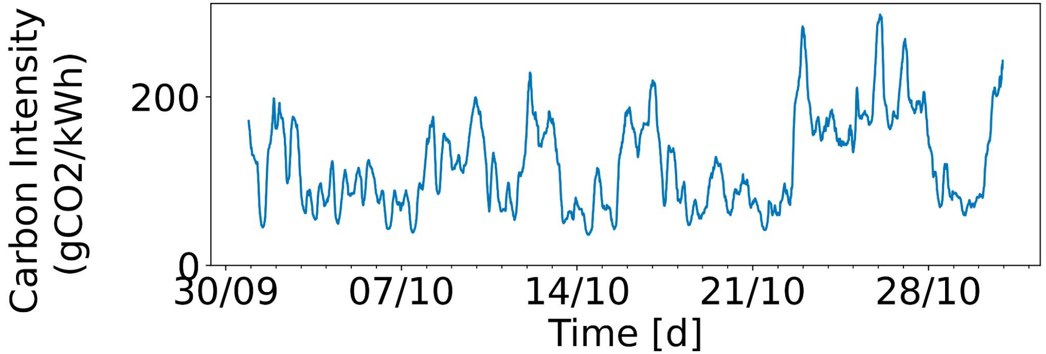

All experiments are run on 158 different carbon traces[1]

Embodied Carbon

Carbon related to the manufacturing of hardware

Increasingly **Important**!

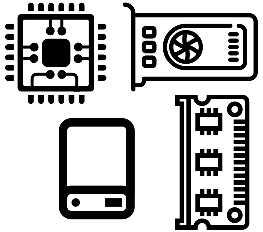

Estimated using the lifecycle method[2]

$$EB_T = \frac{T}{LT} E_{LT}$$

1, Electricity Maps, “All countries 2024 1 hour carbon intensity data,” 2025. [Online]. Available: https://www.electricitymaps.com
2. U. Gupta, M. Elgamal, G. Hills, G.-Y. Wei, H.-H. S. Lee, D. Brooks, and C.-J. Wu, “Act: designing sustainable computer systems with an architectural carbon modeling tool,” in ISCA, 2022.

Workload Traces

All workload traces are based on real-world data centers and replayed using OpenDC

Source	NoH	Host resources	Embodied Carbon
Surf	277	128 GB RAM, 16-core 2.1GHz Intel Xeon Silver 4110	1,022
Marconi	972	196 GB RAM, 2x24-core 2.1GHz Intel Xeon 8160, 4xNVIDIA VoltaV100 16GB	3,542
Borg	1,534	Obfuscated (see Paper)	2,250

Table 1: Data centers used in analysis

Workload	Start	Runtime	Number of Tasks	Average Task Duration
Surf	7-2022	124 days	194,917	01:49:38
Marconi	9-2022	30 days	73,882	06:20:12
Borg	5-2019	31 days	14,867,803	02:01:51

Table 2: Workloads used in analysis

Sustainability Techniques

Horizontal Scaling

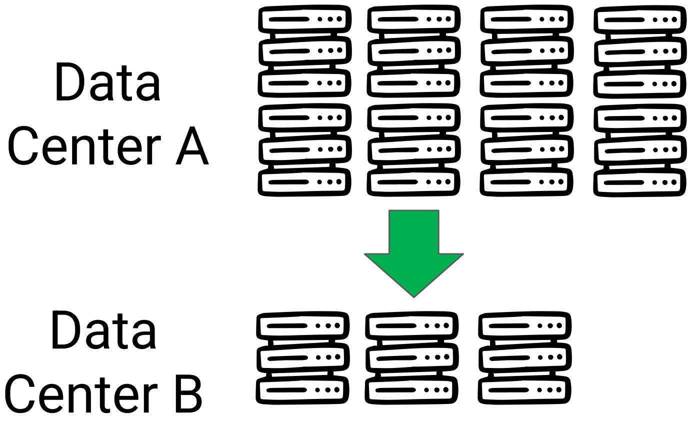

Batteries

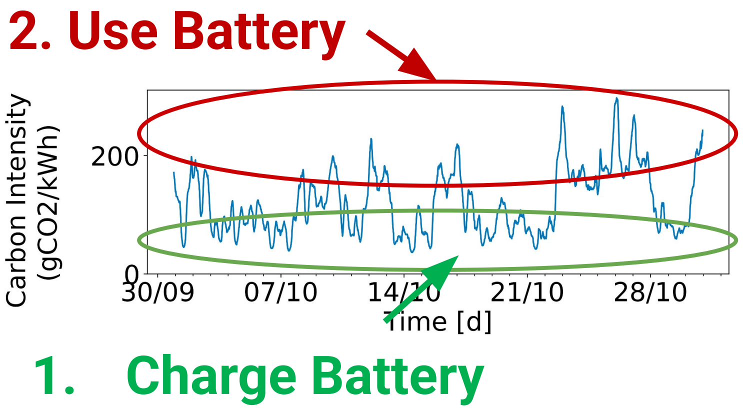

Temporal Shifting

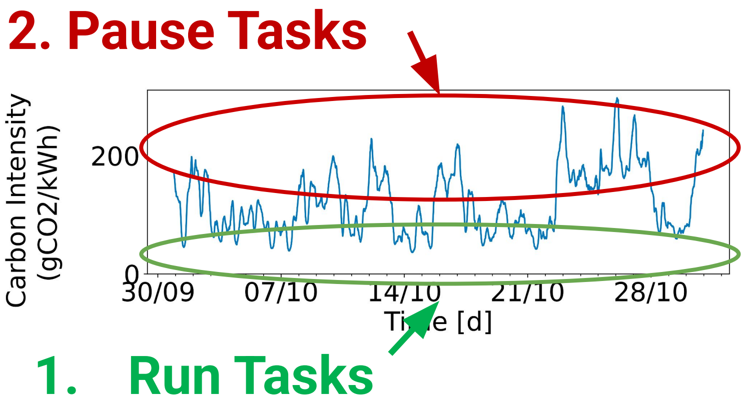

Pros

Reduces both carbon types

Cons

SLO misses

Decreased flexibility

Pros

No impact on performance

Cons

Embodied carbon

High power draw

Pros

No extra costs

Cons

Task have to be delayable

Potential QoS impact

Impact of Batteries

A. On average, using batteries reduces total carbon emissions

B. For >5%, batteries reduce total carbon by >20%

C. For >30% of the regions, the total carbon actually increased!

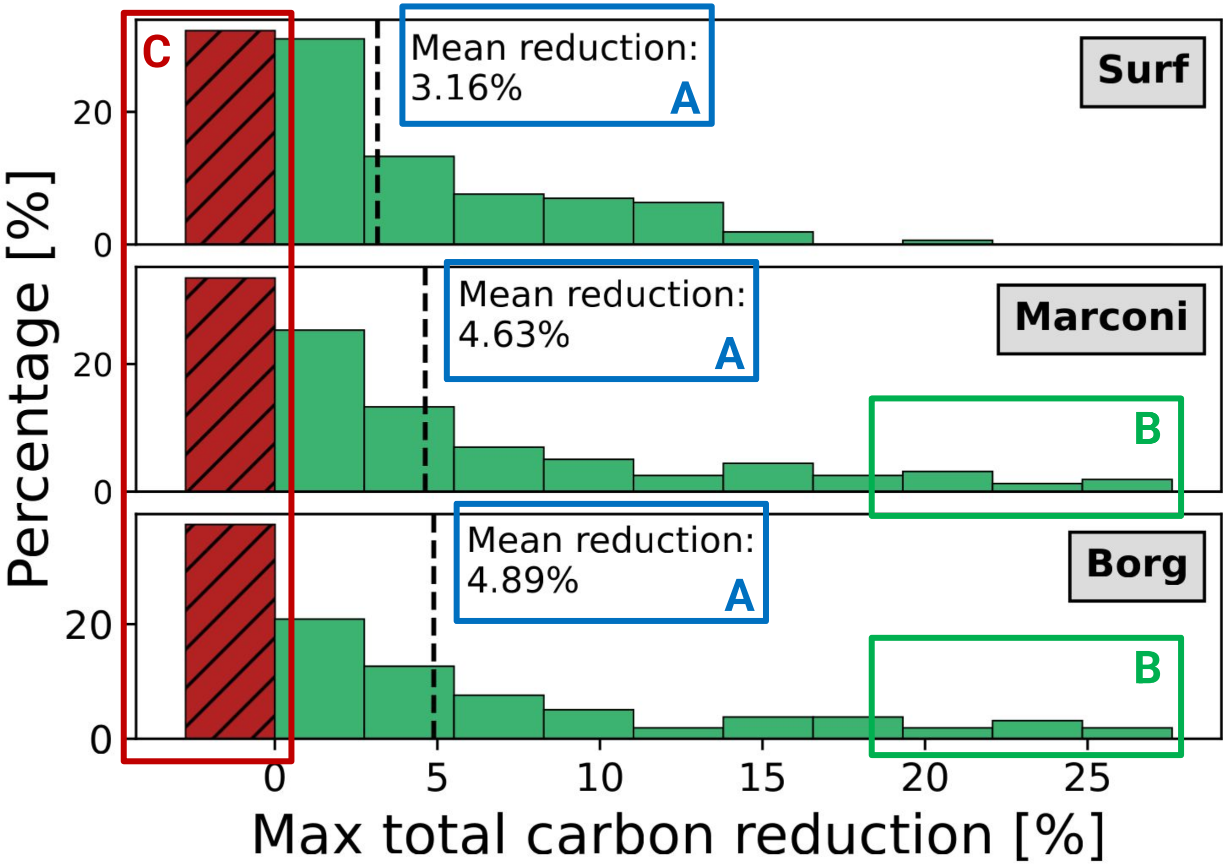

Fig: Total carbon emission reduction using batteries in 158 carbon regions on three workloads. Red bars mark all regions where batteries increased emission.

More results can be found in our paper

Charging Speed

How does the charging speed impact performance?

The same experiment is run with different charging speed.

Higher charging speed
-> Better performance

Large diminishing returns!

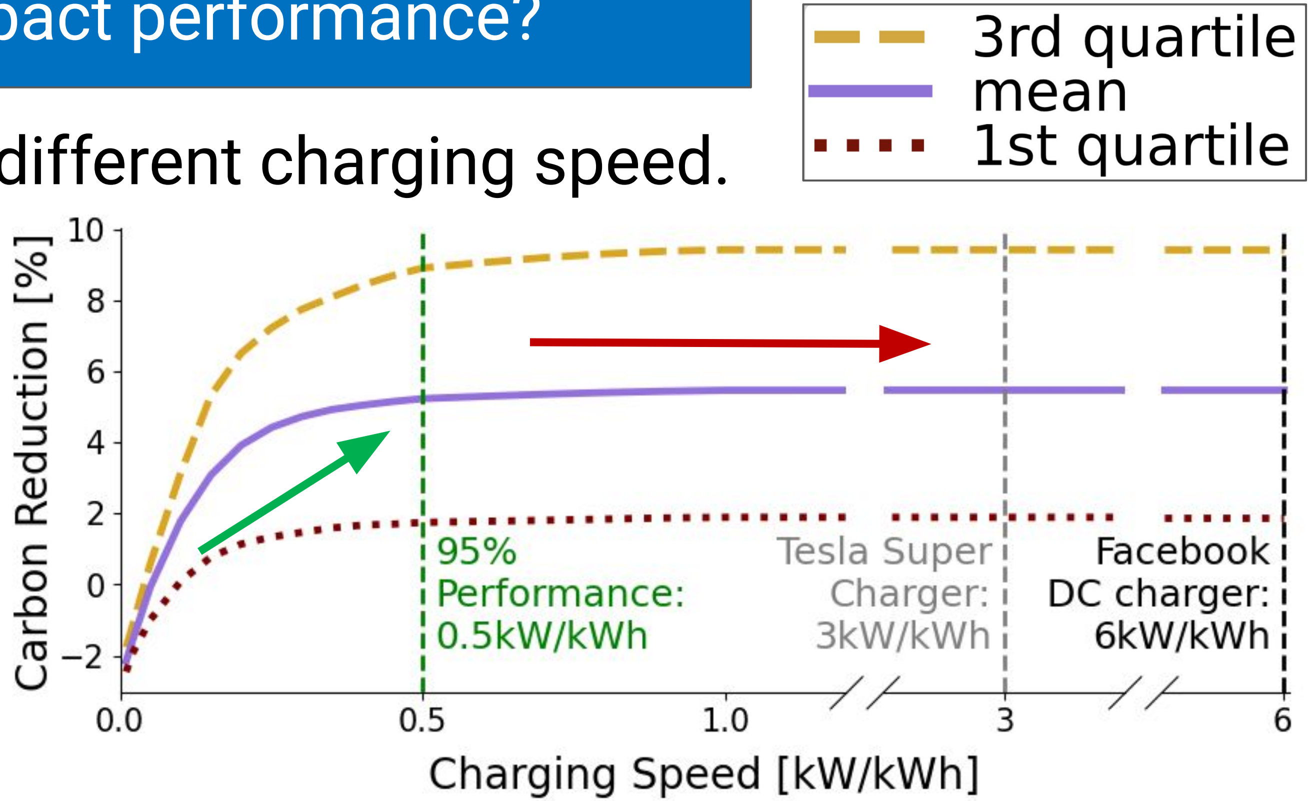

Fig: The impact of battery charging speed on the carbon reduction achieved when using batteries.

More results can be found in our paper

Combining Techniques

A. The impact of techniques varies between workloads

B. Combining techniques can further improve performance

C. Combining some techniques can lead worse performance

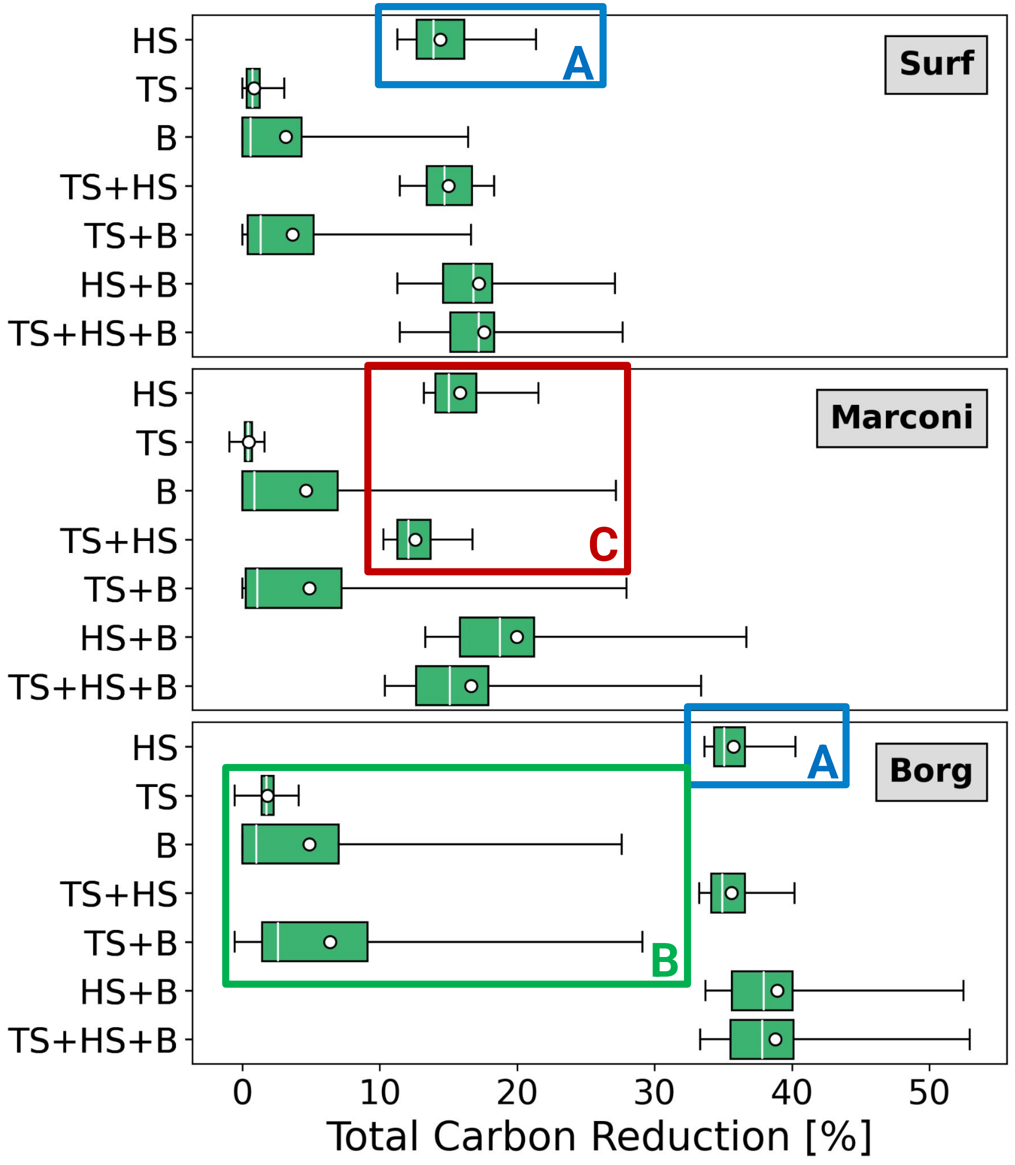

More results can be found in our paper

Combining Techniques

Combining techniques can impact the optimal setup

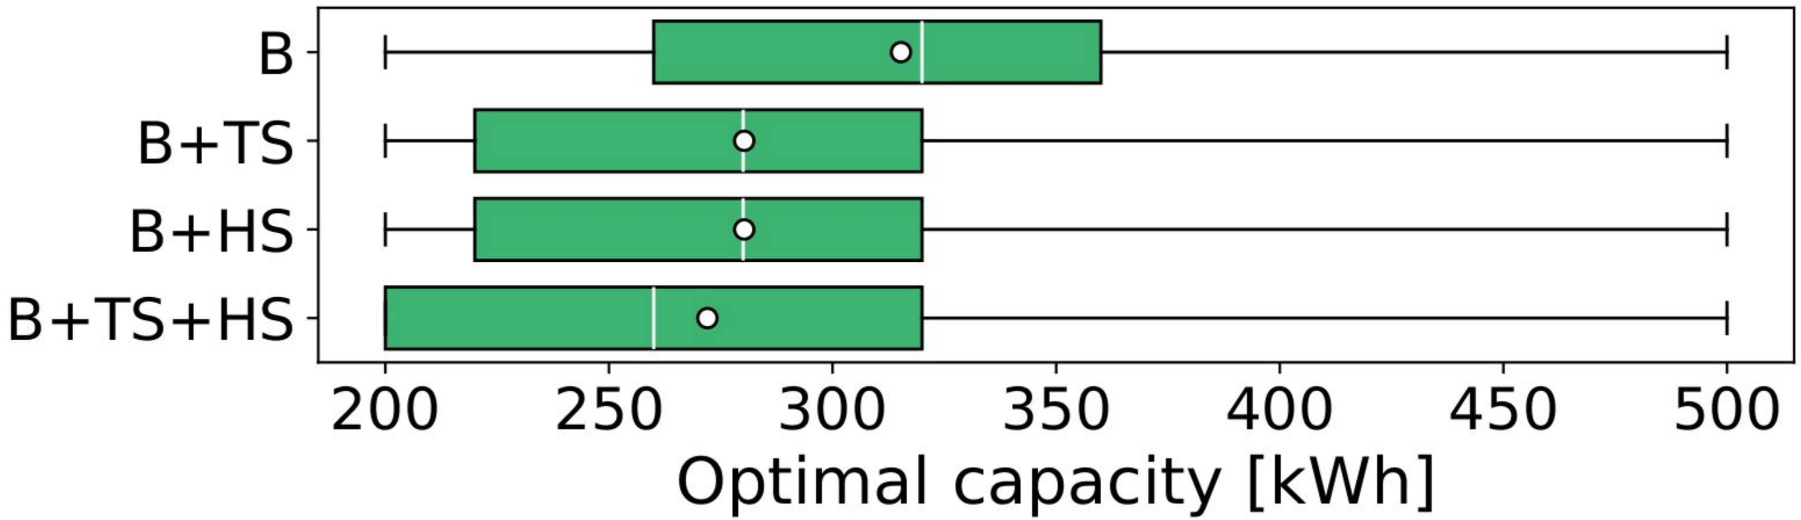

Fig: Optimal battery size when used in combination with other techniques

Take Home Messages

OpenDC-STEAM quantifies the impact of sustainability techniques individually or in combination

Sustainability techniques can be effective, but careful consideration of workload, data center, and carbon region is required

Combining techniques can improve sustainability further, but can also introduce new trade-offs that need to be navigated

Future Work:

1. Usability and Performance
2. Expand features
3. Digital twins

Learn more: opendc.org

github.com/atlarge-research/opendc-steam

https://github.com/atlarge-research/opendc

Composability Example: Host Types

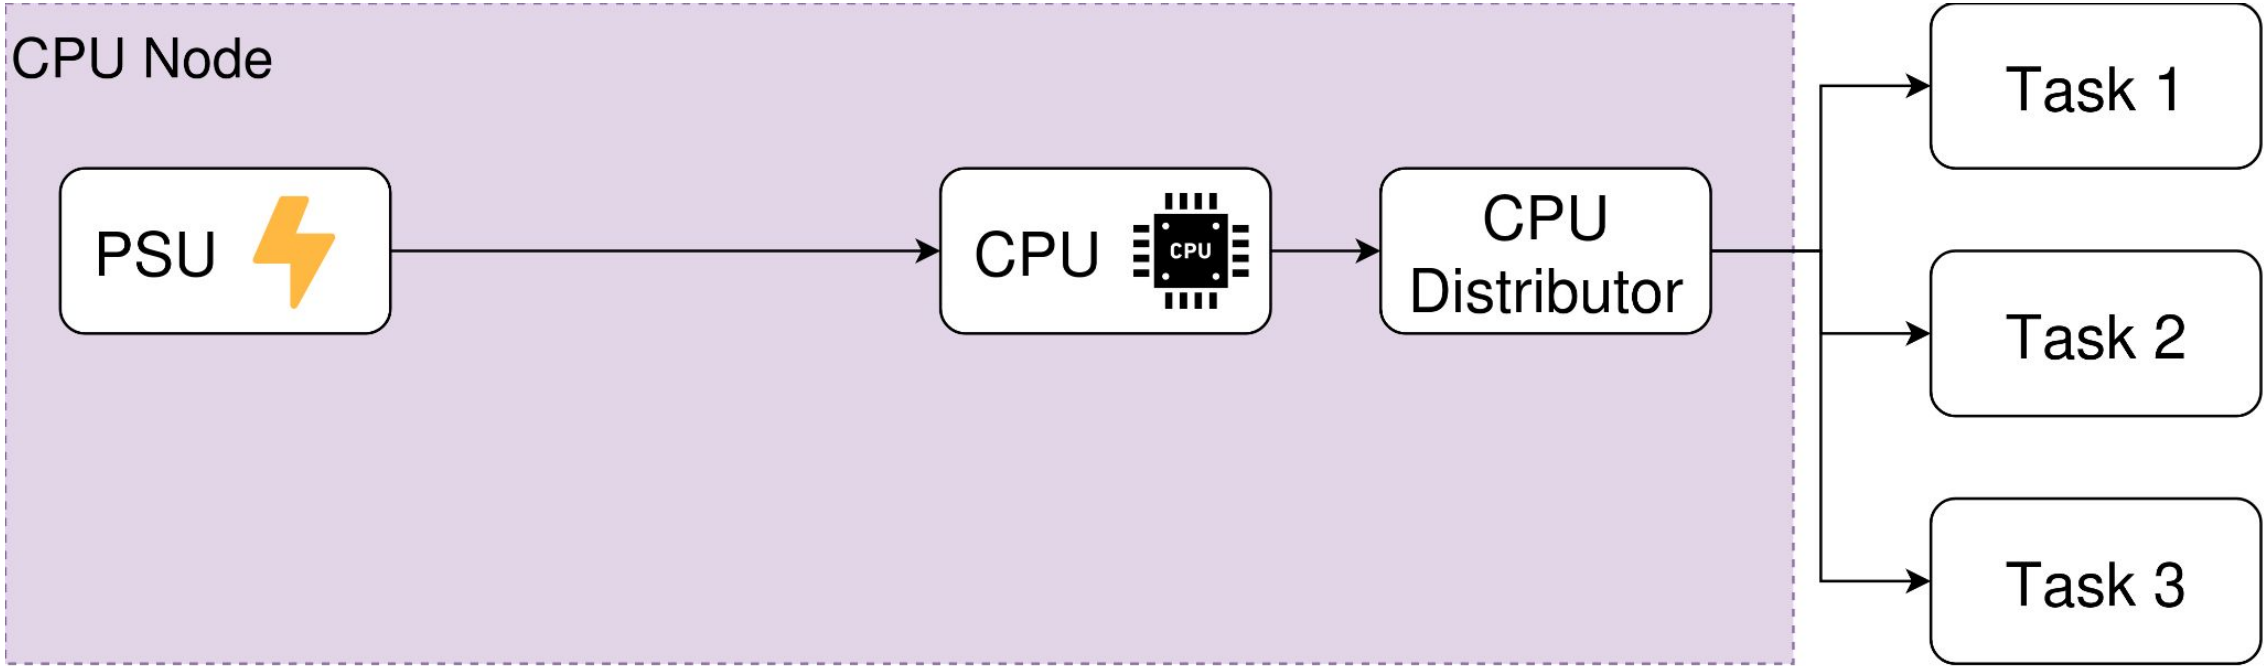

Composability Example: Host Types

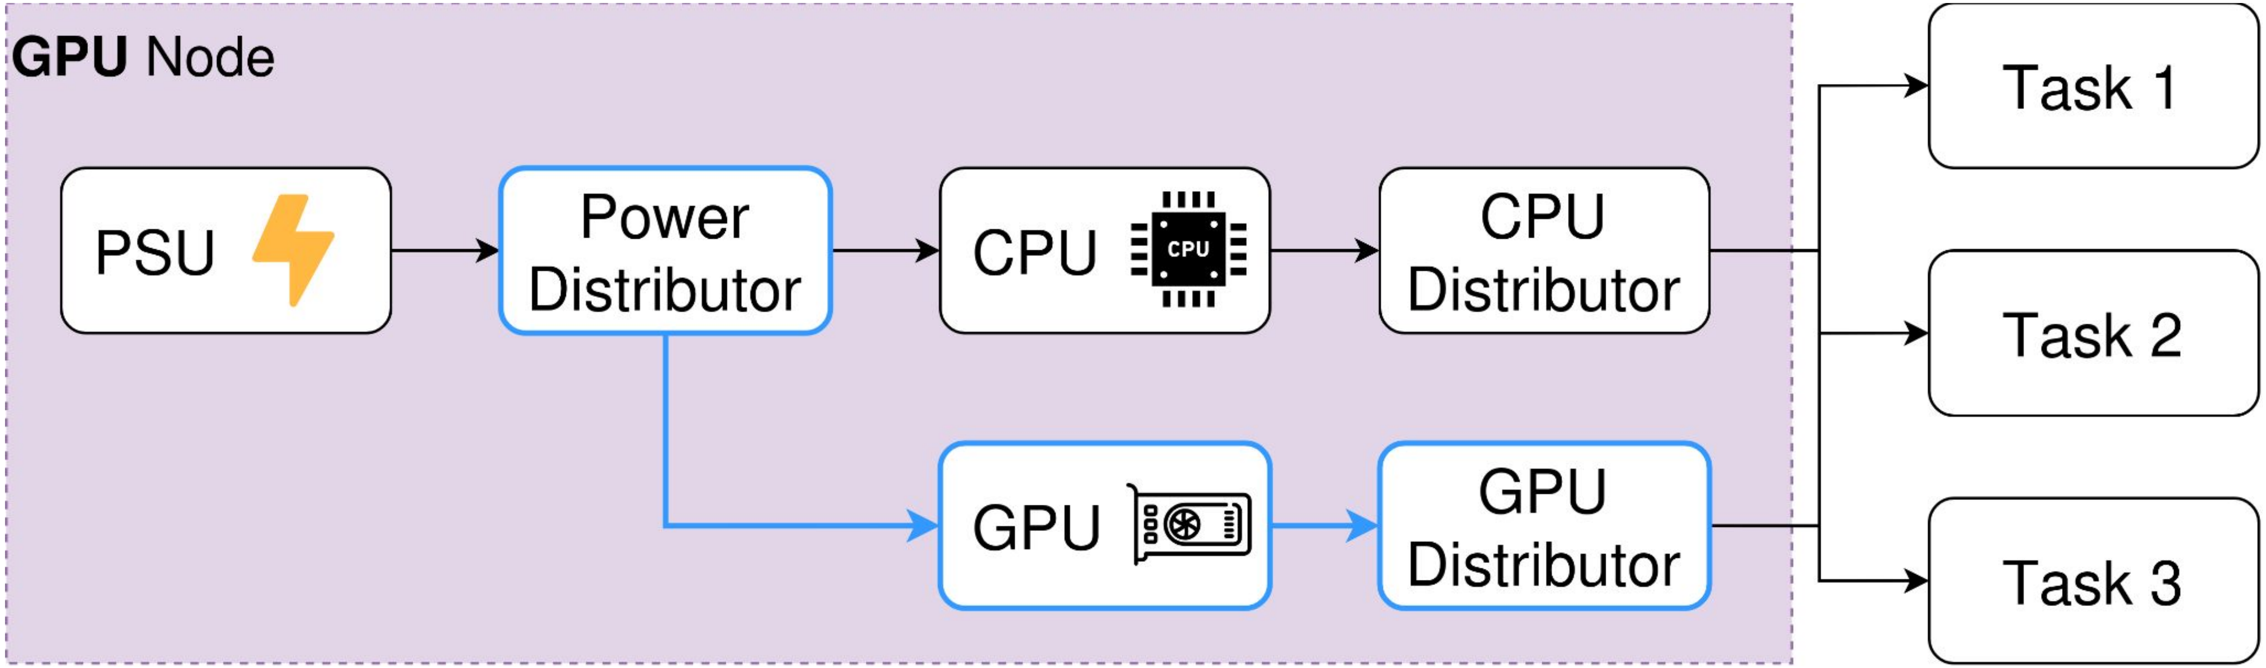

Components only care about their direct neighbours!

Why simulation?

- ✓ Cost and time effective
- ✓ Flexible
- ✓ Low risk
- ✓ Represent complex systems

Delft → Rotterdam | x 55 billion! ...

In Simulation	In Reality
112	**8,600,000,000**
Core-hours	Core-hours
2.1	**116,000,000,000**
kWh	kWh

Exact numbers confidential, depend on topology

Operational Carbon

Carbon emitted as a result of the energy used by a data center

- *Energy * Carbon Intensity*
- The source of energy is **important**[1]
- STEAM uses carbon traces to determine the operational carbon

All experiments are run on 158 different carbon traces[2]

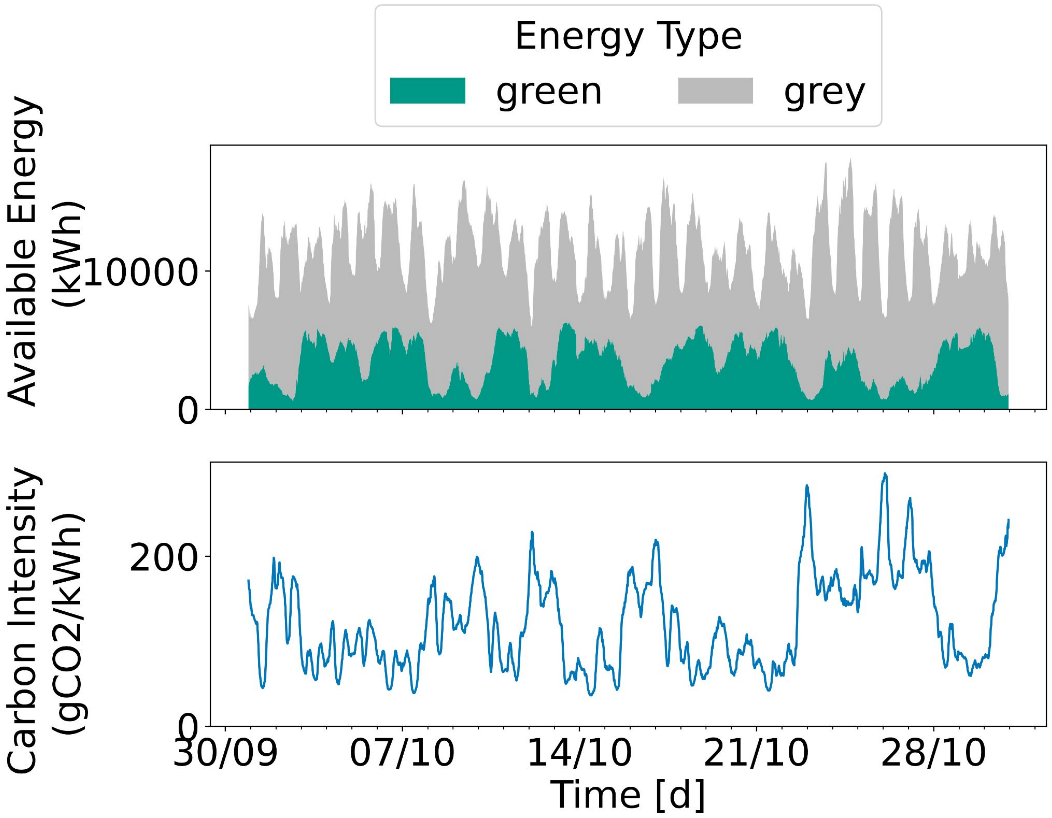

Fig: Energy mix and carbon intensity of the grid in the Netherlands during October 2023[1]

Grid: $$CI_g = \sum_{s \in S} CI_s \frac{E_s}{E_g}$$

1. Dante Niewenhuis, Sacheendra Talluri, Alexandru Iosup, Tiziano De Matteis (2024) FootPrinter: Quantifying Data Center Carbon Footprint. HotCloudPerf.
2, Electricity Maps, "All countries 2024 1 hour carbon intensity data," 2025. [Online]. Available: https://www.electricitymaps.com

Embodied Carbon

Carbon related to the manufacturing of hardware

- Embodied carbon is getting increasingly **important**
- STEAM estimates the embodied carbon of a workload using the lifecycle method[1]

$$EB_T = \frac{T}{LT} E_{LT}$$

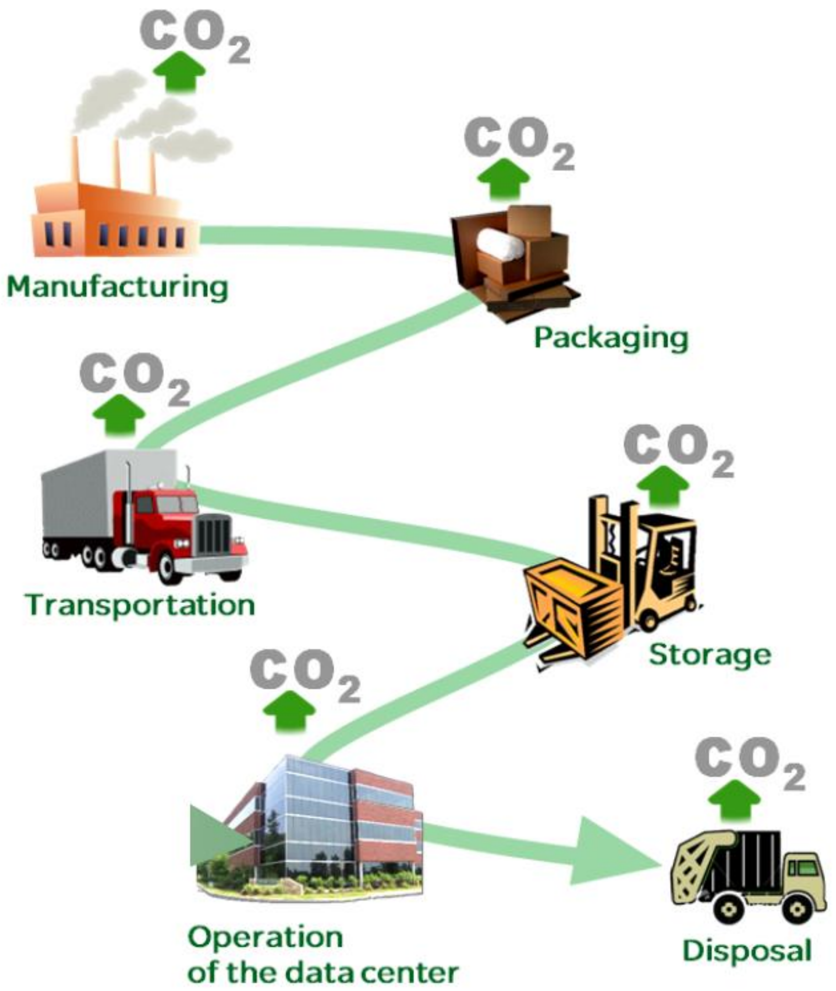

Fig: Carbon-producing phases of a data center[X]

1. U. Gupta, M. Elgamal, G. Hills, G.-Y. Wei, H.-H. S. Lee, D. Brooks, and C.-J. Wu, “Act: designing sustainable computer systems with an architectural carbon modeling tool,” in ISCA, 2022.

Horizontal Scaling

Reducing the number of machine running in a data center

Data Center A

Data Center B

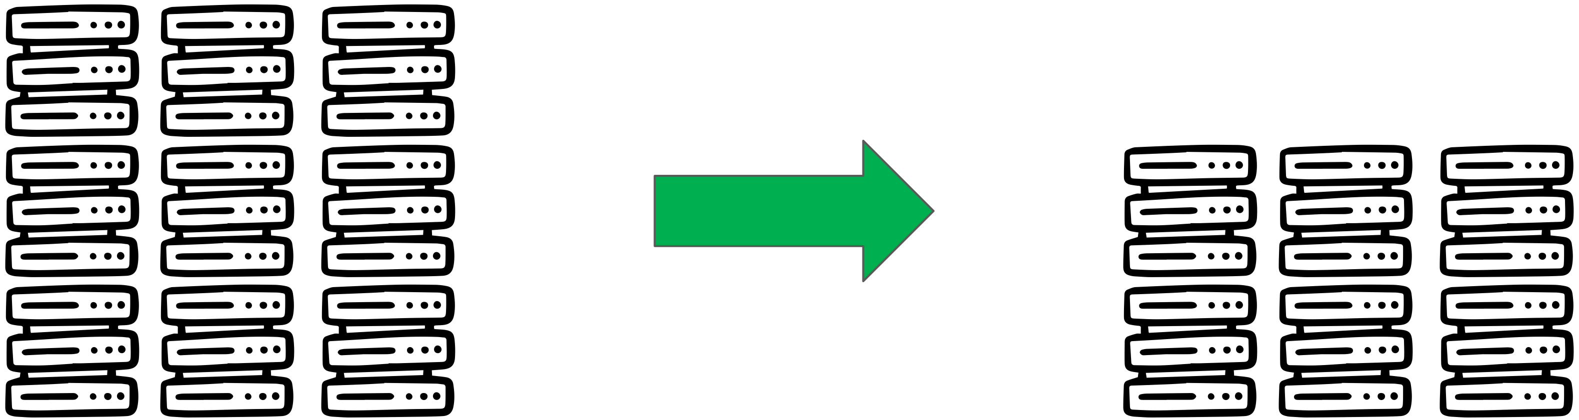

Pros

Reduce both carbon types

Cons

SLO misses

Decreased flexibility

Batteries

Store green energy in batteries to use during high-carbon periods

1. Charge during low-carbon energy
2. Use energy during high-carbon energy

Pros

No impact on performance

Cons

Embodied carbon

High power draw

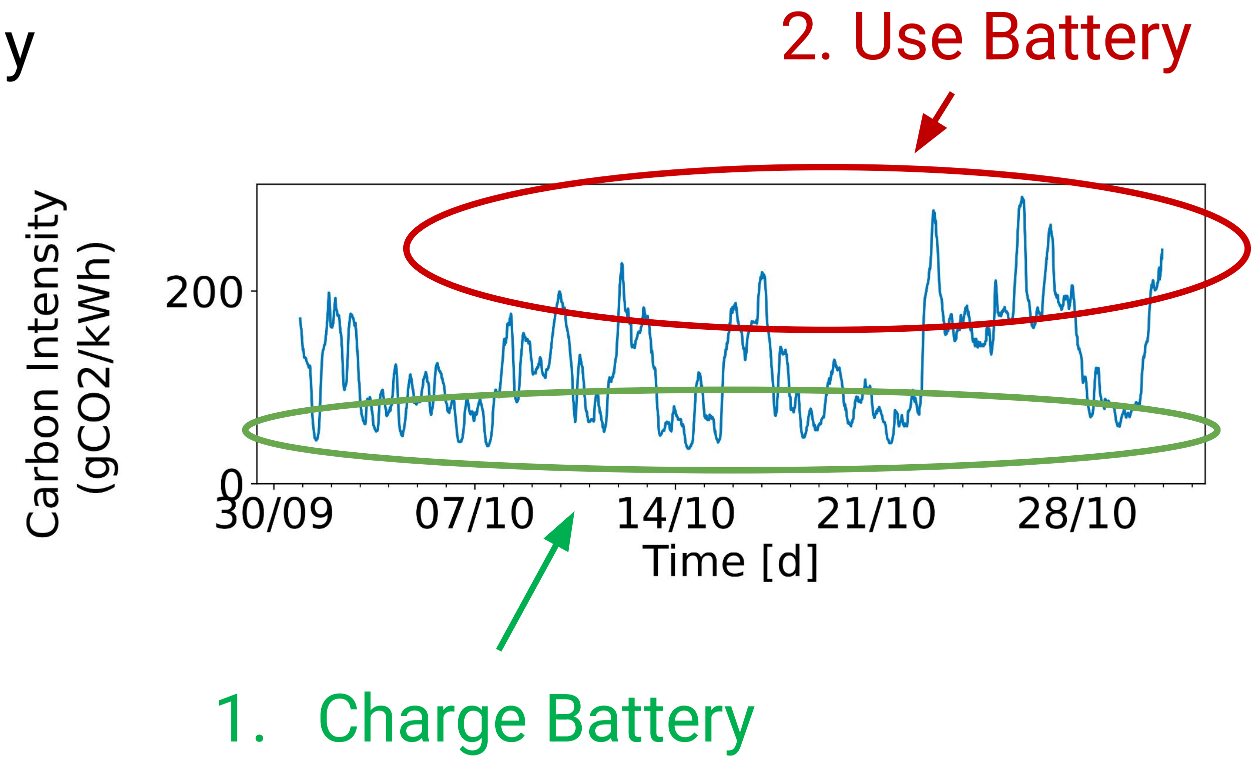

Temporal Shifting

Reducing Carbon Emissions by running tasks when low-carbon energy is available

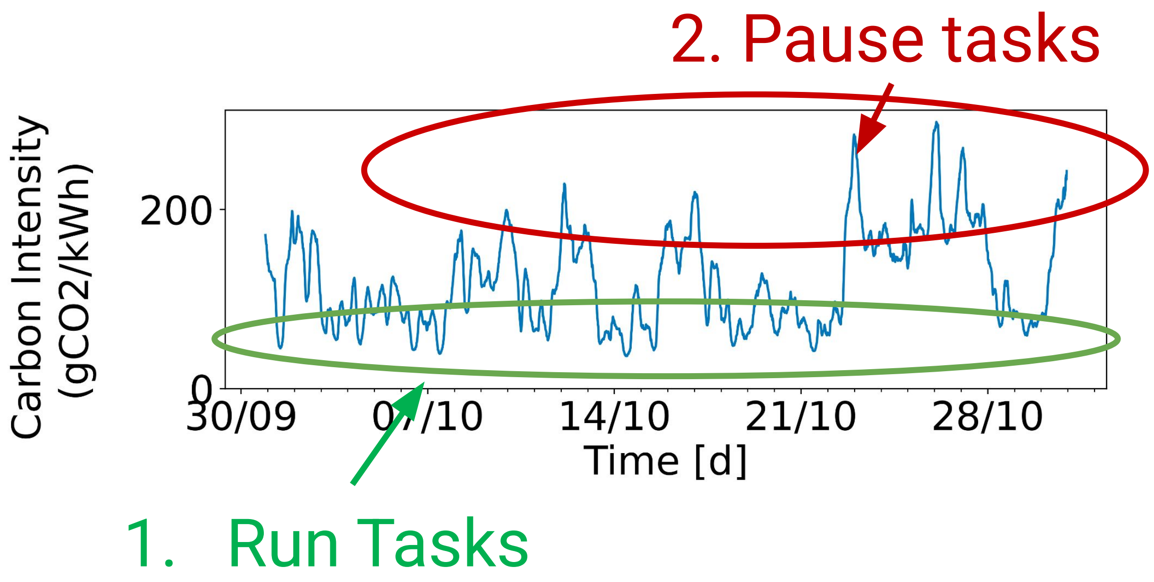

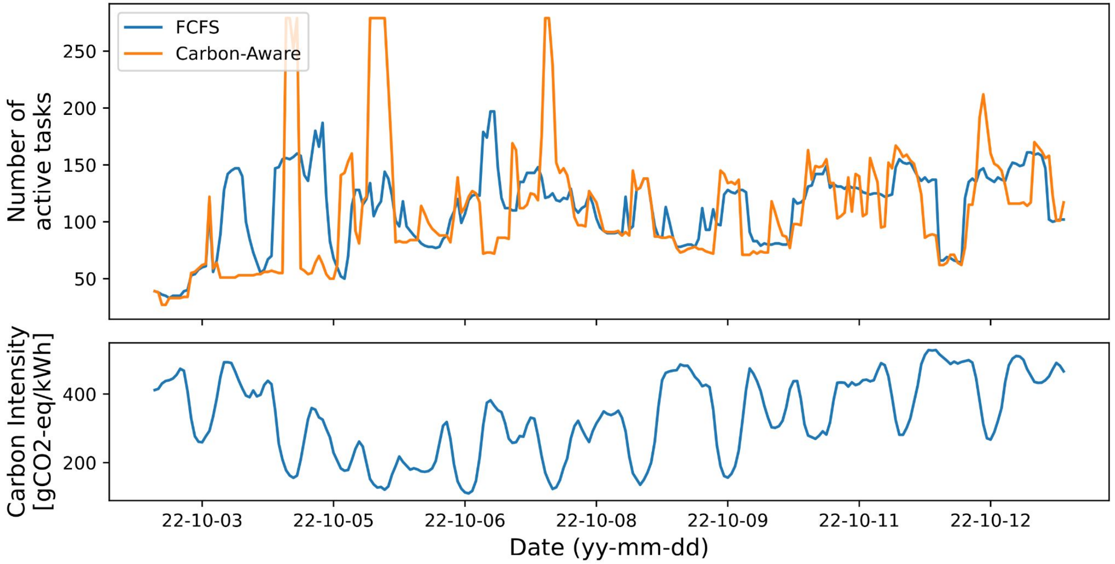

Fig: Number of active task during a workload when using a carbon-aware and a non-carbon aware scheduler.

Pros

No extra costs

Cons

Task have to be delayable

Impact of Horizontal Scaling

We can lower the number of hosts significantly!

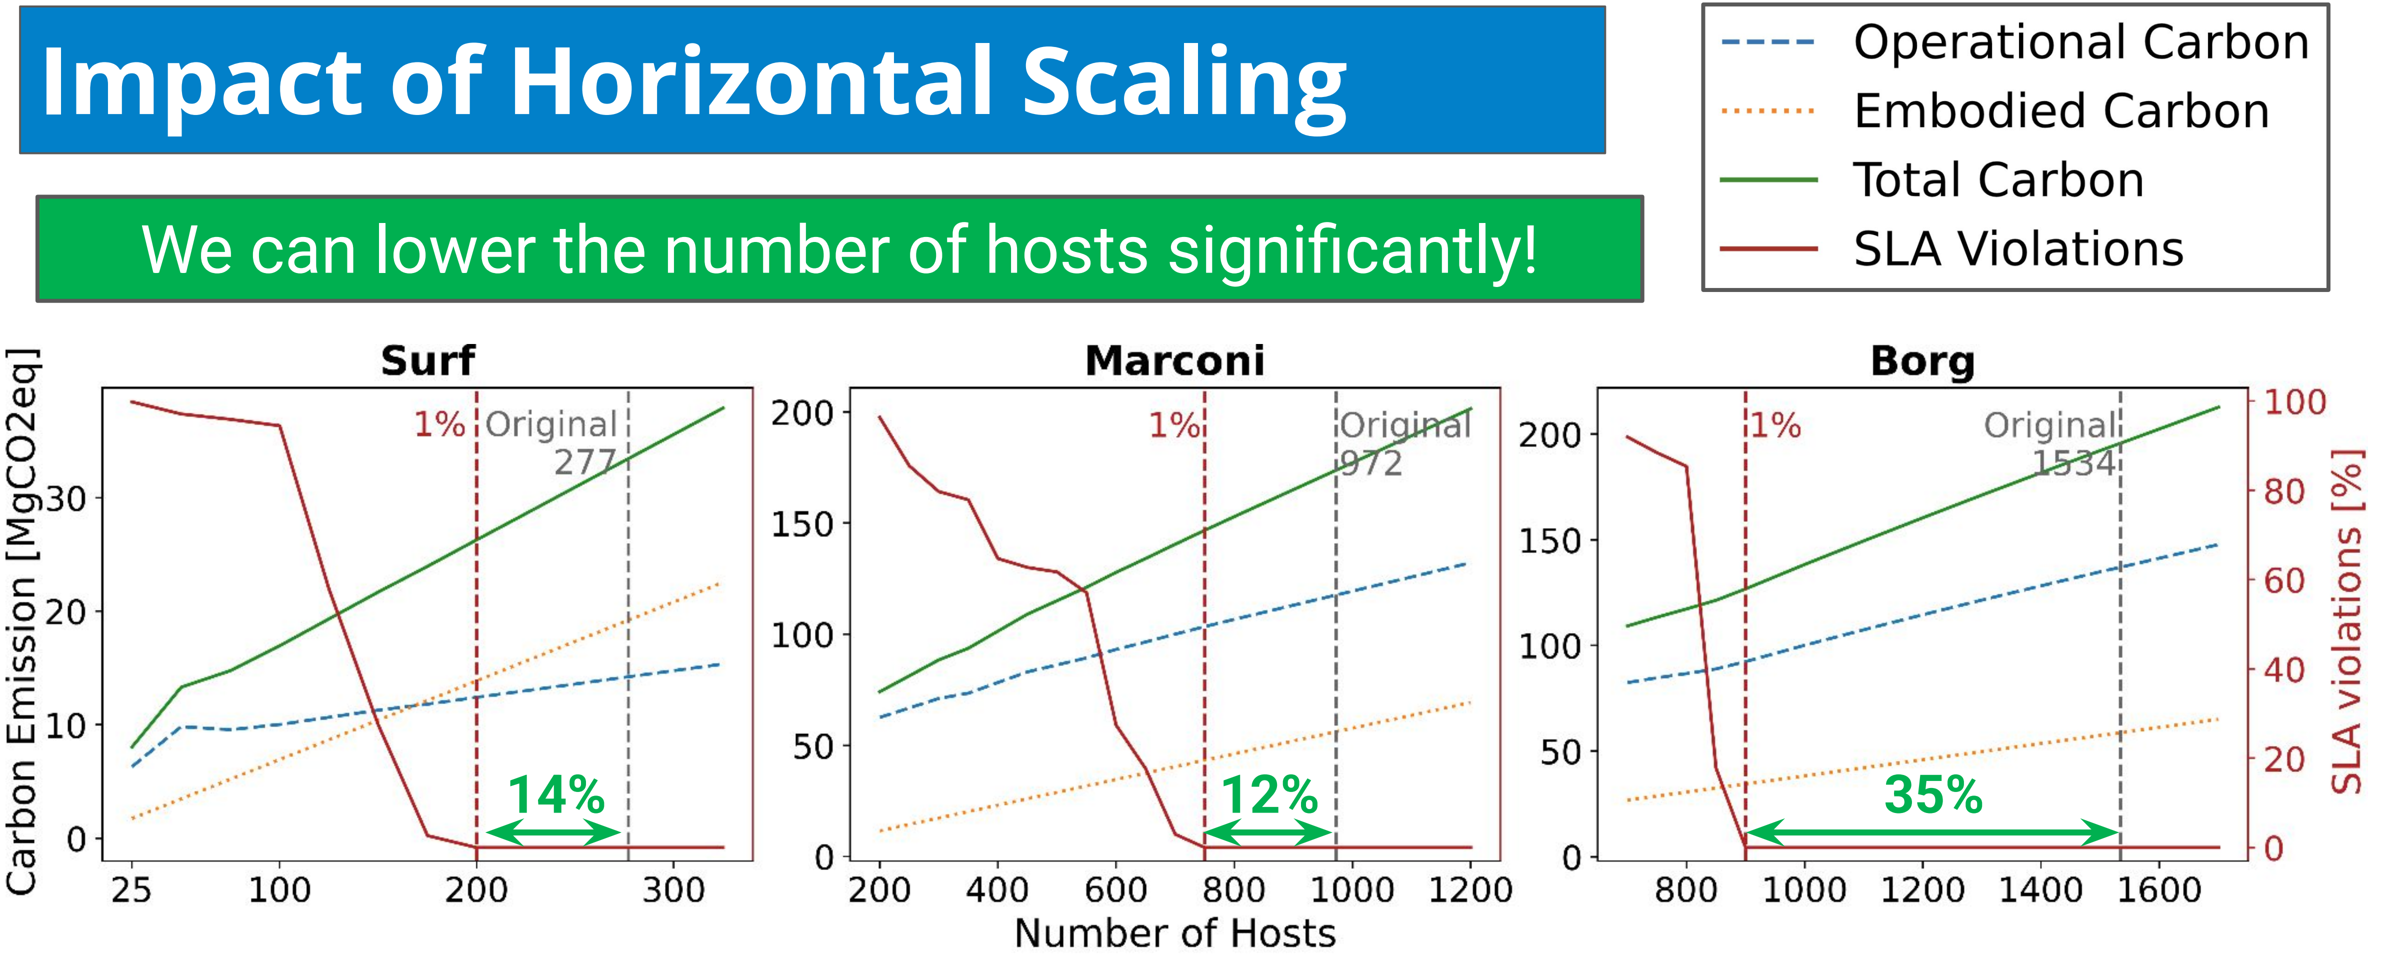

Fig: The impact of horizontal scaling on a data center executing a workload

More results can be found in our paper

Adding Failures

We periodically inject failures based on failure traces[1]

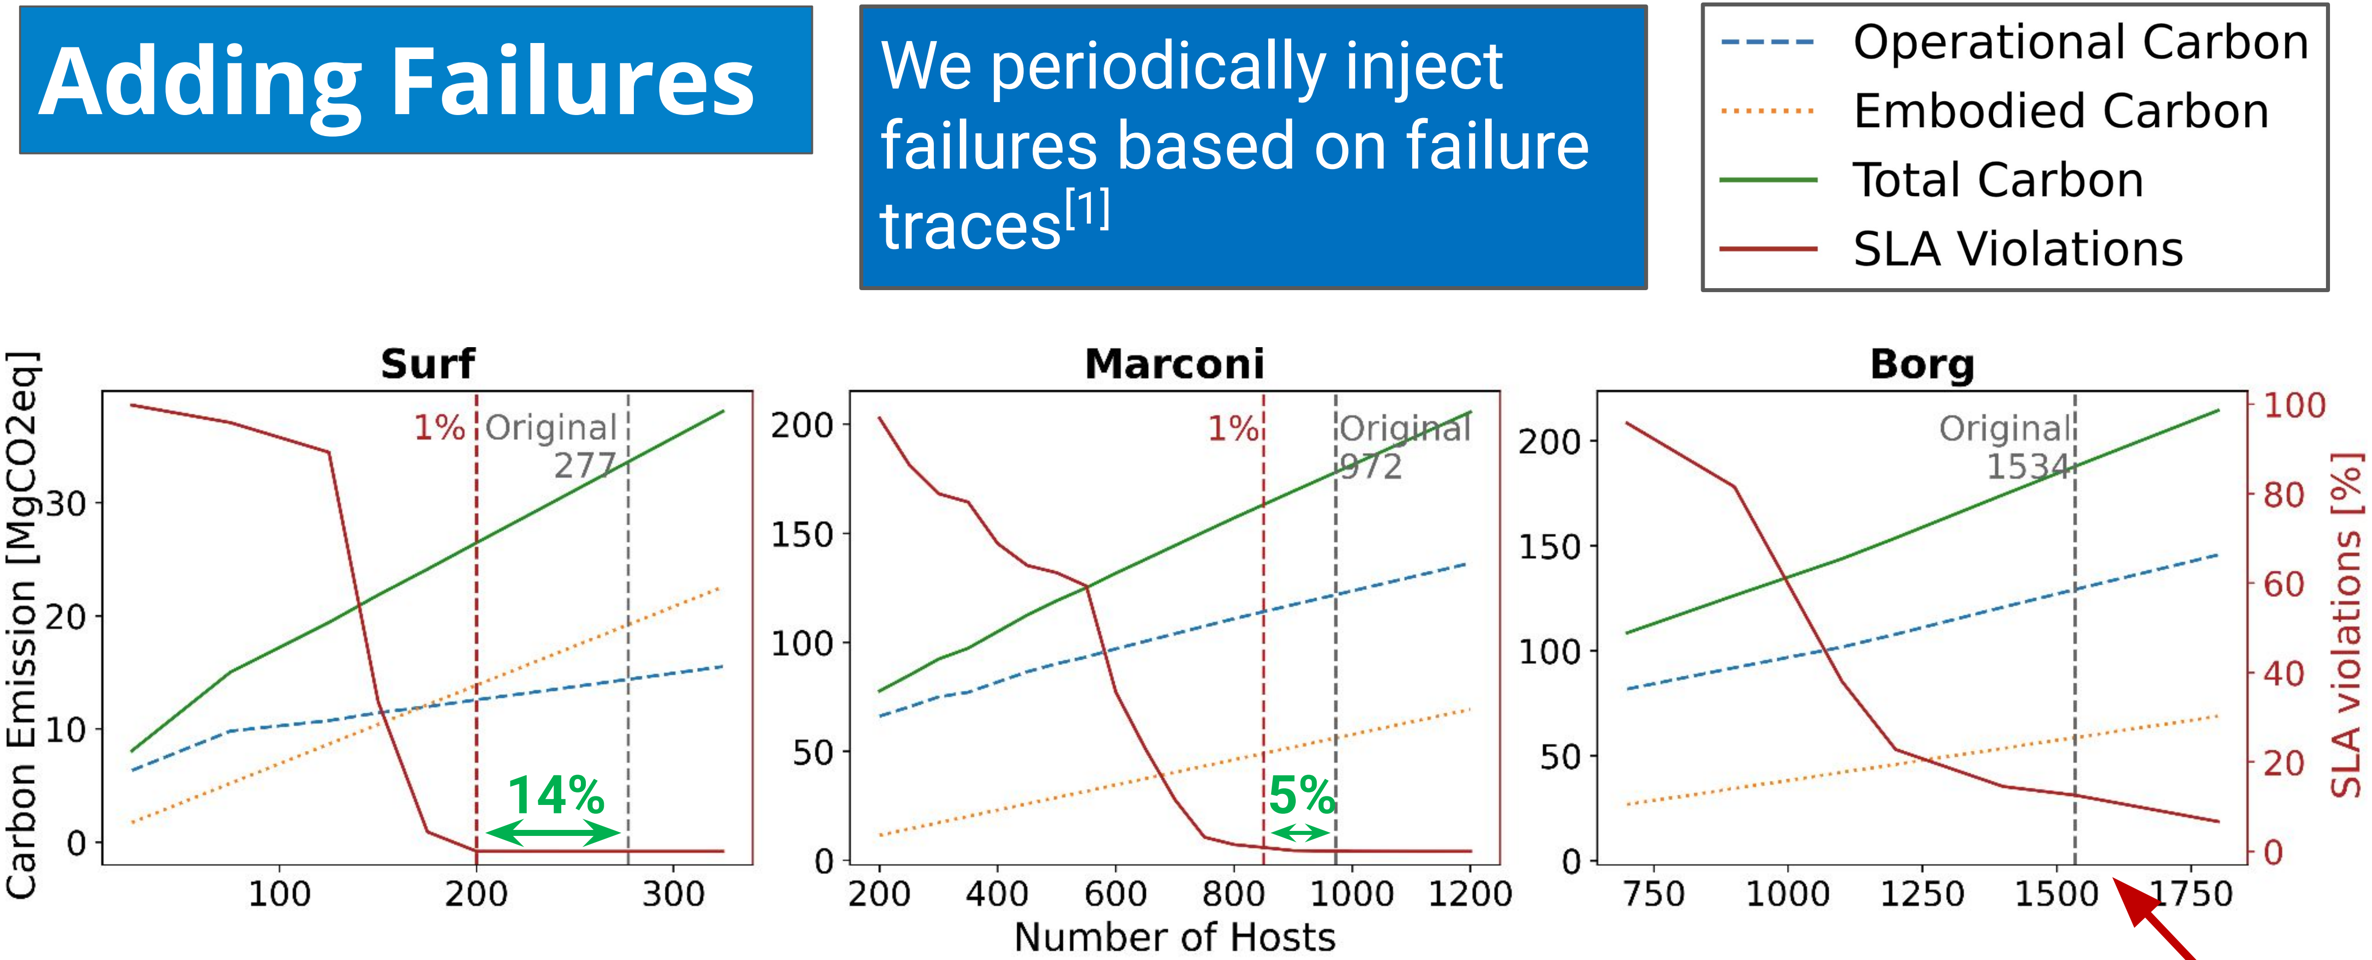

Fig: The impact of horizontal scaling on a data center executing a workload when injecting failures

Still successful in some setups!

We have to be careful!

More results can be found in our paper

@Large Research
Massivizing Computer Systems

1. S. Talluri, D. Niewenhuis, X. Chu, J. Kyselica, M. Cetin, A. Balgavy, and A. Iosup, “Cloud Uptime Archive: Open-access availability data of web, cloud, and gaming services,” 2025.

Impact of Temporal Shifting

1. Using Load Shifting can reduce carbon emissions by up to 7%
2. The average reduction is only 2%

!

Our results are are significantly lower than than reported by prior work!

1. **Shifting technique**
2. **Evaluation method**

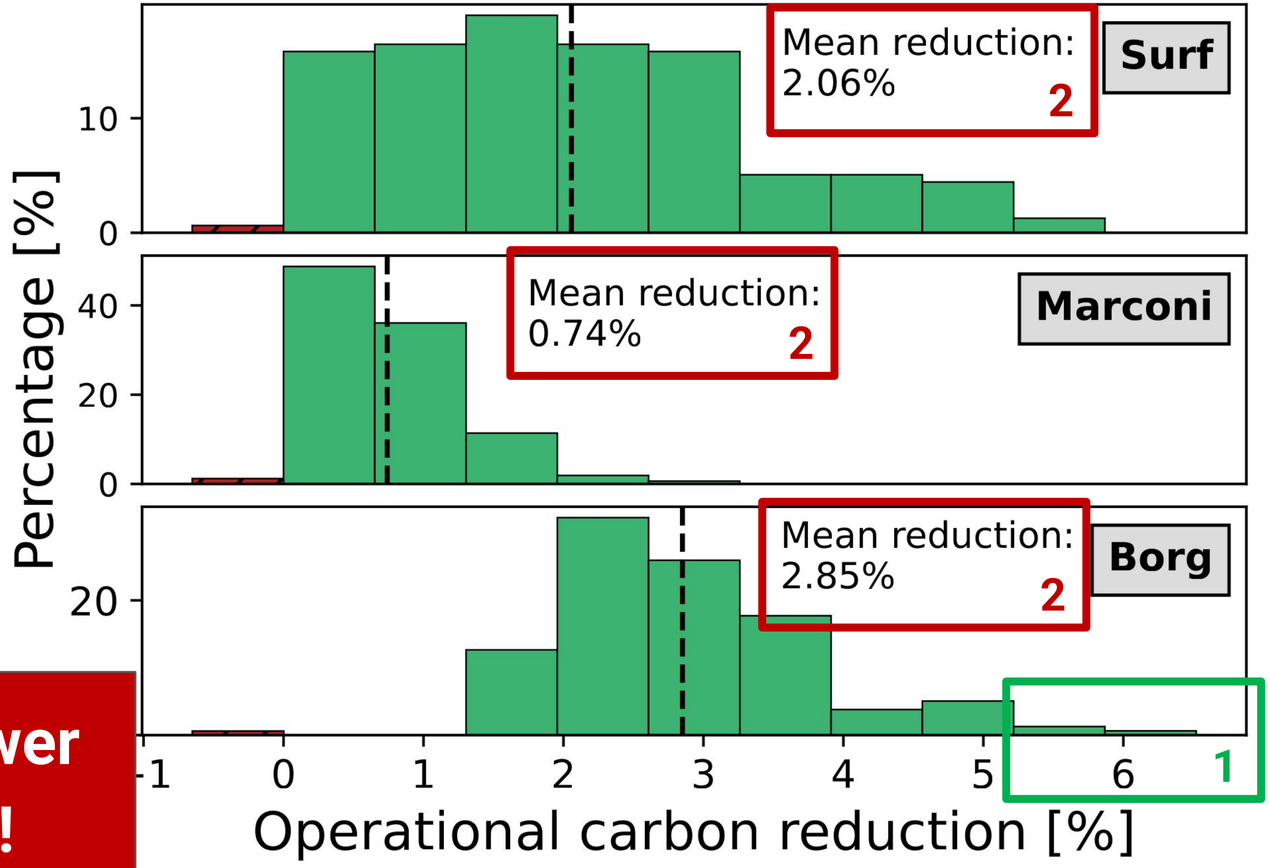

Fig: Total carbon emissions reduction using temporal shifting in 158 carbon regions on three workloads.

More results can be found in our paper

Results 3: Battery Capacity

Does the capacity of the battery impact performance?

The same experiment is run with different battery capacities.

Findings:

1. There is a sweet spot for the optimal battery size.
2. A large range of capacities is within 90% of performance.

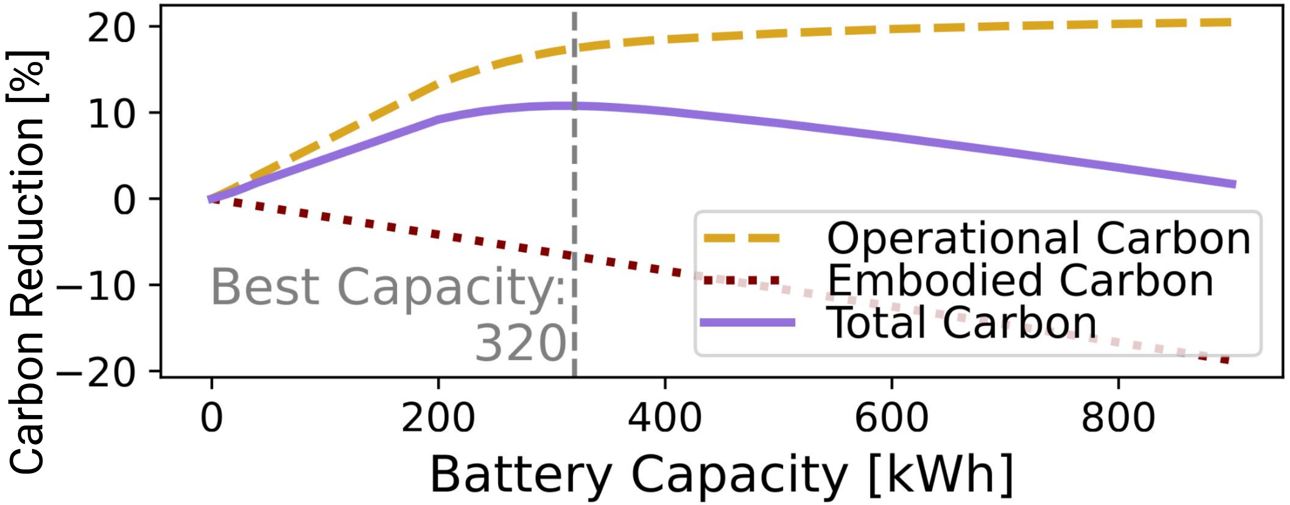

Fig 8: The impact of battery capacity on the carbon reduction achieved when using batteries.

Results 4: Embodied Carbon

Findings:

1. Reducing embodied carbon cost can significantly improve effectiveness
2. In ~20% of the regions, batteries are effective even at high embodied carbon cost
3. In ~10% of the regions, batteries will never be effective

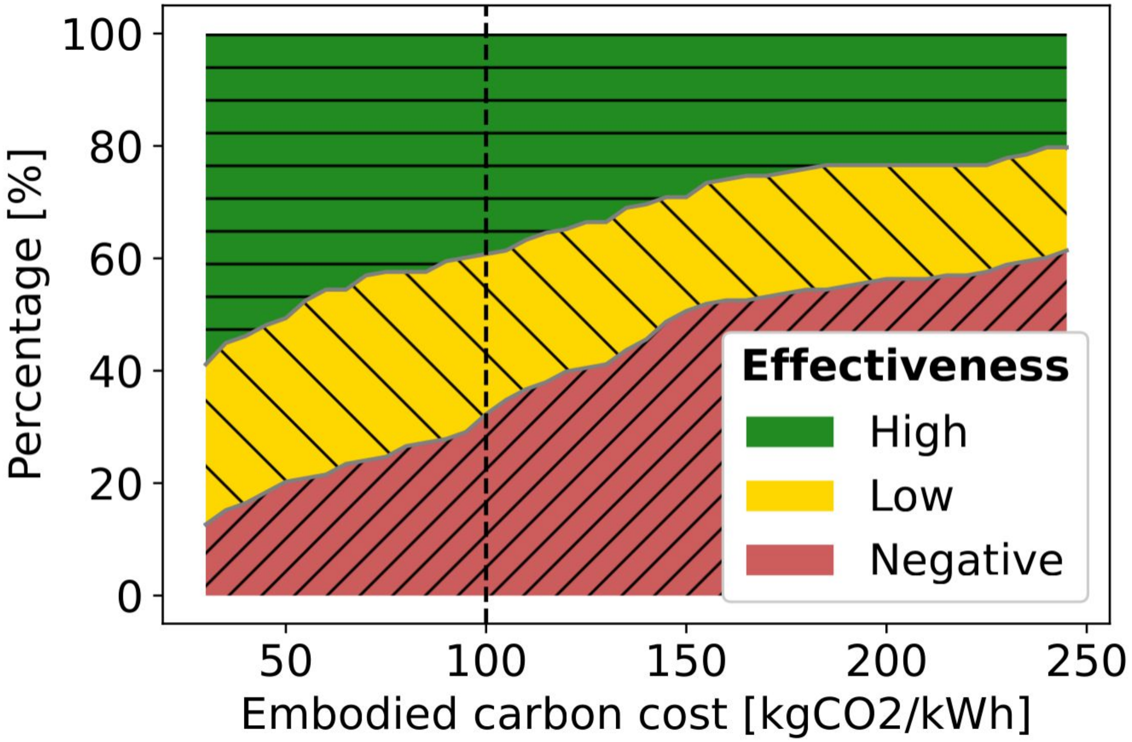

Fig 9: The impact of embodied carbon cost of batteries on total carbon reduction.