Towards Game-Theoretic Approaches to Attributing Carbon in Cloud Data Centers

Leo Han, Jash Kakadia, Benjamin C. Lee, Udit Gupta

Data centers (DCs) account for >0.6% of global GHG emissions

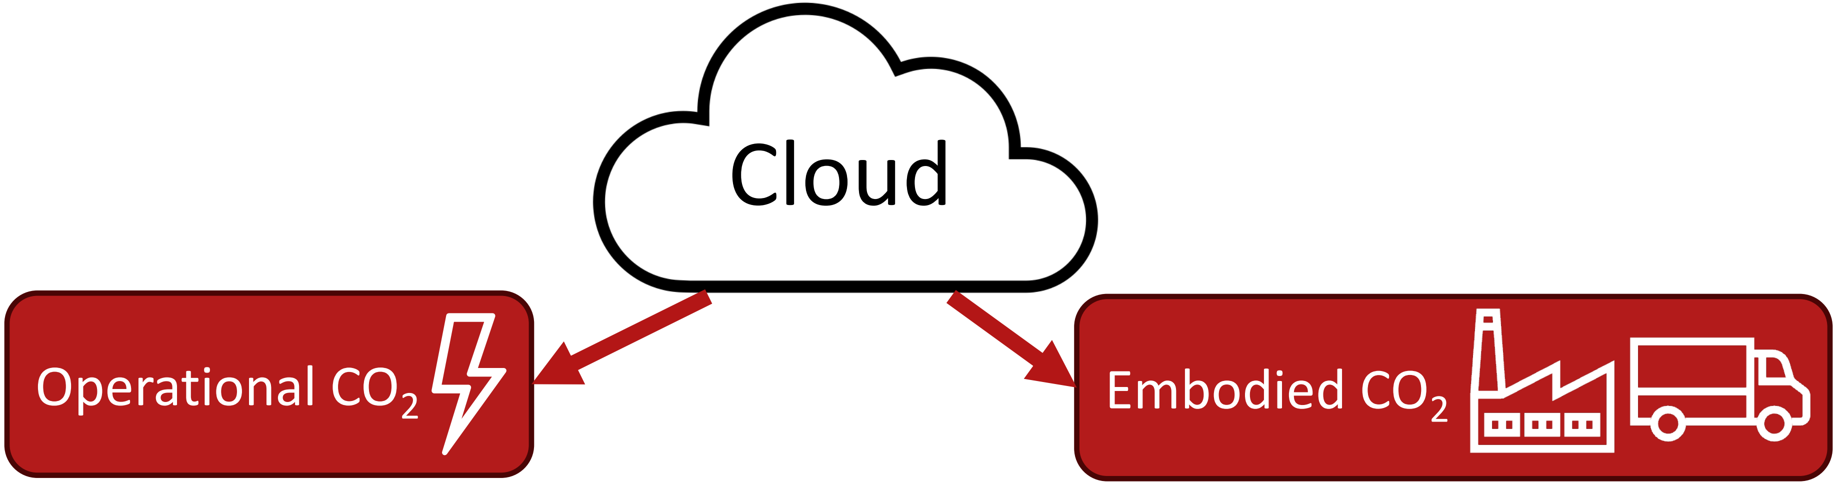

DC operational carbon accounts for ~0.6 % of global GHG emissions (IEA, 2023)

DC embodied carbon accounts for 17.5% to 49.0% of total data center GHG emissions (Lin et al, 2023)

Data center capacity is projected to grow 10% year-on-year (McKinsey, 2023)

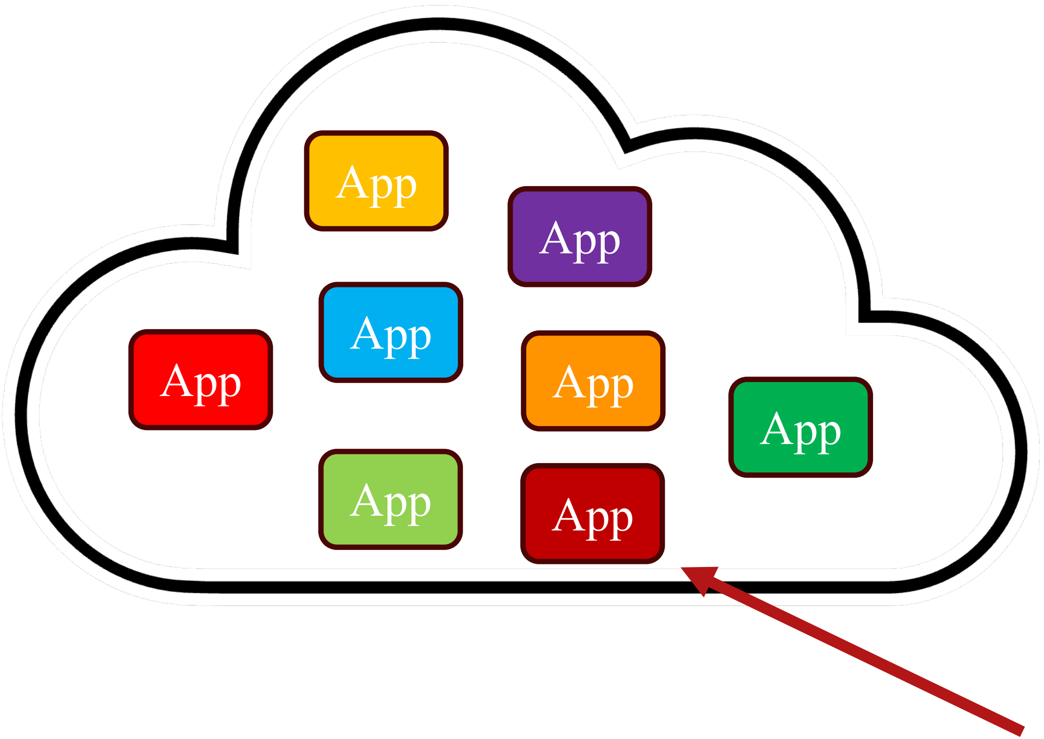

What is the carbon footprint of App ?

Cloud carbon attribution is increasingly important

Cloud Carbon Footprint

Cloud Jewels: Estimating kWh in the Cloud

By Emily Sommer, Mike Adler, John Perkins, Joshua Thiel, Hilary Young, Chelsea Mozen, Dany Daya, Katherine Sundstrom

Apr 23, 2020

Microsoft | Sustainability   Products & solutions   Sustainability guid

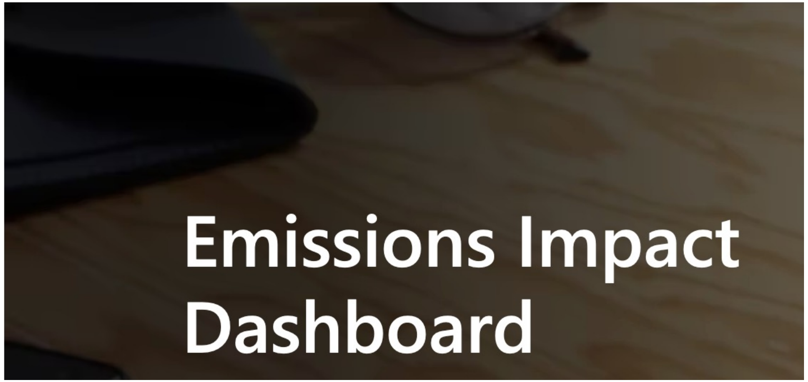

CODE CARBON

aws

AWS Customer Carbon Footprint Tool

Google Cloud   Overview   Solutions   Products   Pricing   Resources

Carbon Footprint

Catch up on the latest product launches, demos, and trainings fror

Carbon accounting in the Cloud: a methodology for allocating emissions across data center users

Ian Schneider*, Taylor Mattia*†

June 2024

Carbon Footprint

Measure, report, and reduce your cloud carbon emissions.

Existing attribution models may be unfair

Colocation effects

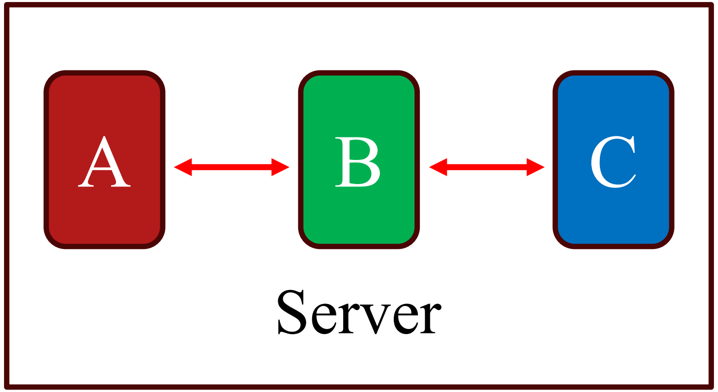

Poor proxies for embodied carbon

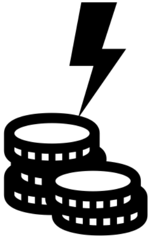

$\not\approx$ embodied CO_2

Intel Xeon CPU: 792 USD/kgCO_2, 37 W/kgCO_2

DDR4 DRAM: 4.29 USD/kgCO_2, 0.4 W/kgCO_2

Shapley value is a game theory solution for fair shared-cost attribution

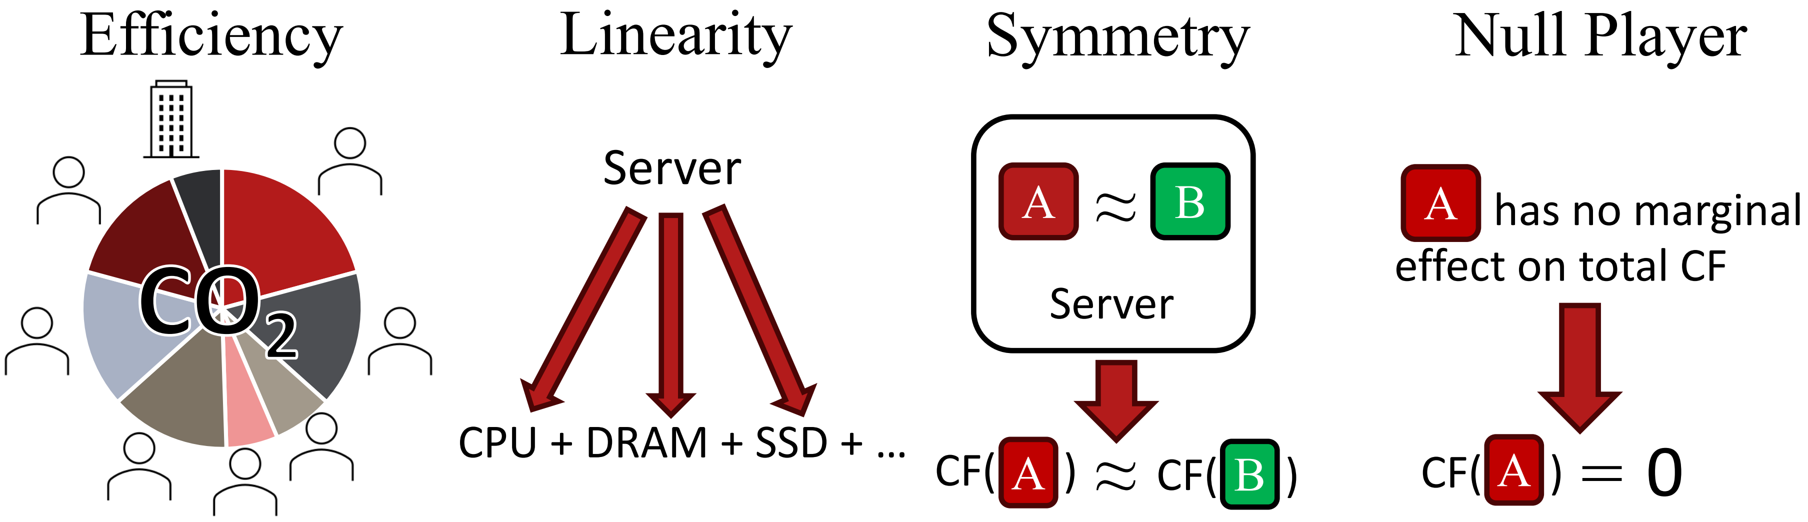

Properties for Fair Attribution

Shapley value averages marginal contributions

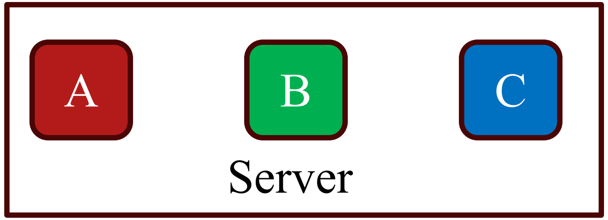

What is A's share of the server carbon footprint (CF)?

CF({A}) - CF({})

CF({B A}) - CF({B})

CF({C A}) - CF({C})

CF({B C A}) - CF({B C})

CF({C B A}) - CF({C B})

Shapley value:
average of A's marginal contribution across all permutations

Case Study: Four CloudSuite 4.0 workloads on CloudLab Dell R650 Node

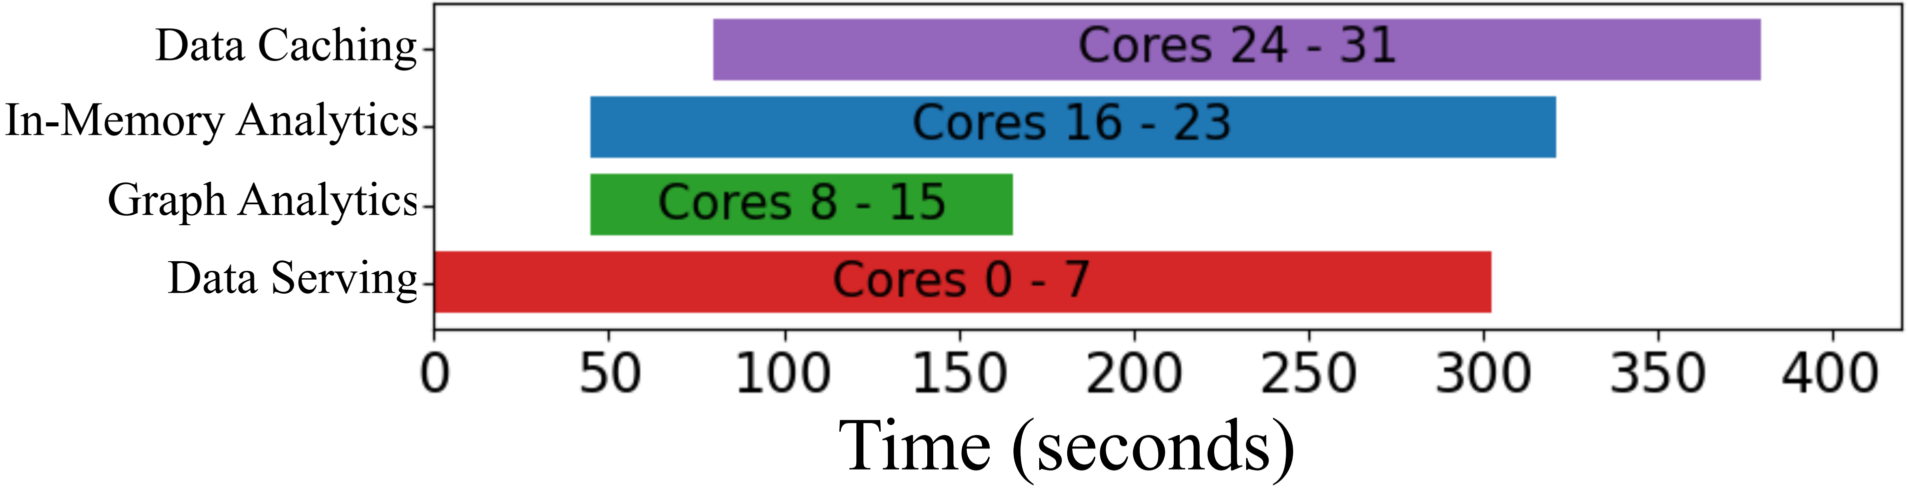

Shapley value power attribution accounts for colocation effects

Offline Power Profiles → Shapley → Power Attribution → CI → Operational Carbon Attribution

DS	GA	IA	DC	Power
				94.6 W
x				100.8 W
	x			137.0 W
. . .				
x	x	x		180.3 W
x	x	x	x	201.2 W

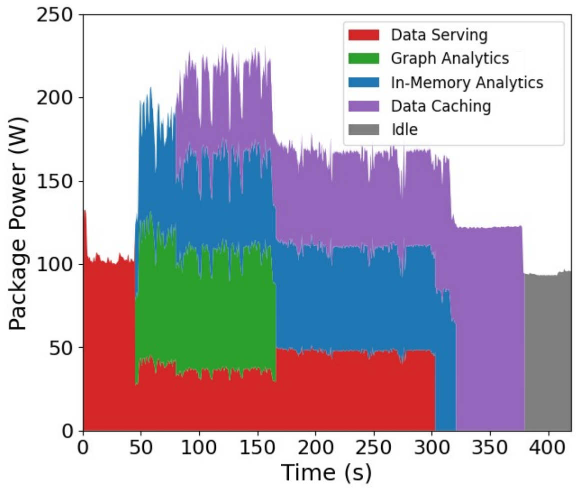

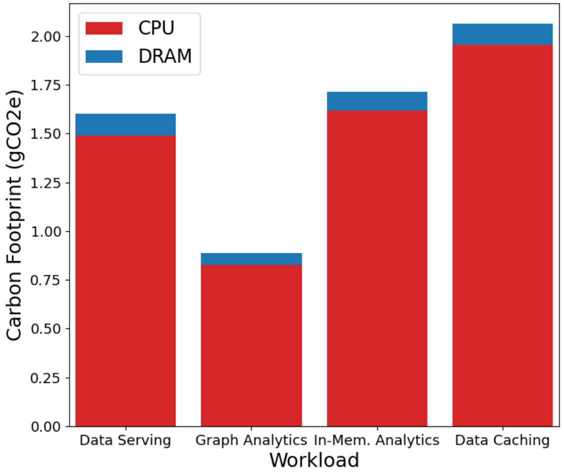

Shapley embodied carbon attribution based on resource utilization

Example: DRAM @ 0.52 $kgCO_2e/GB$

$$CF_{DRAM}(\{DS, GA, IA, DC\}) - CF_{DRAM}(\{DS, GA, IA\})$$

$30.70\ kgCO_2e$ - $24.55\ kgCO_2e$ = $\mathbf{6.15\ kgCO_2e}$

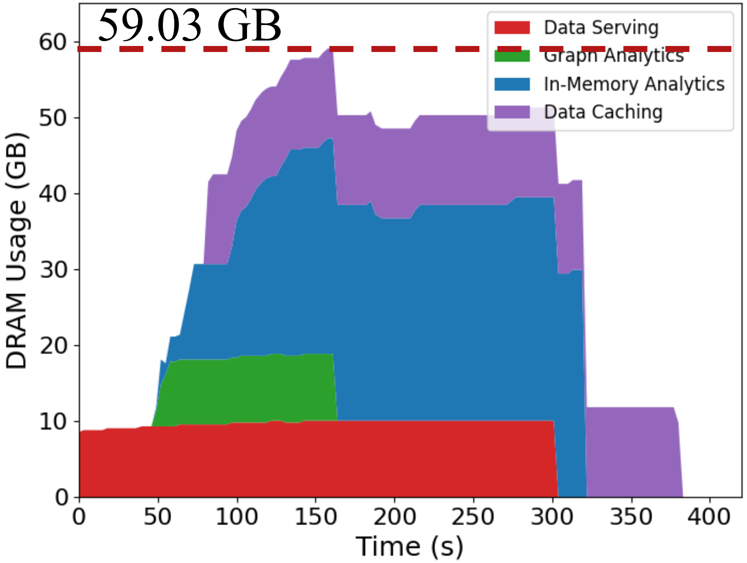

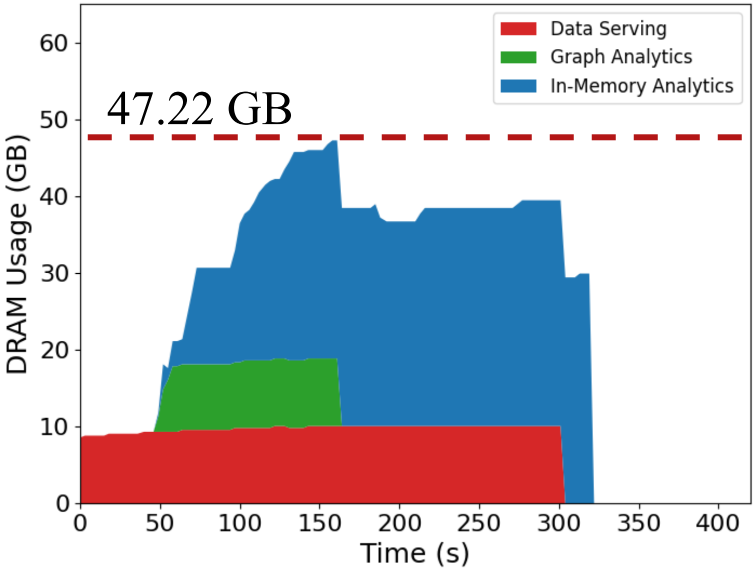

Per-resource Shapley attribution vs. baseline energy-proportional attribution

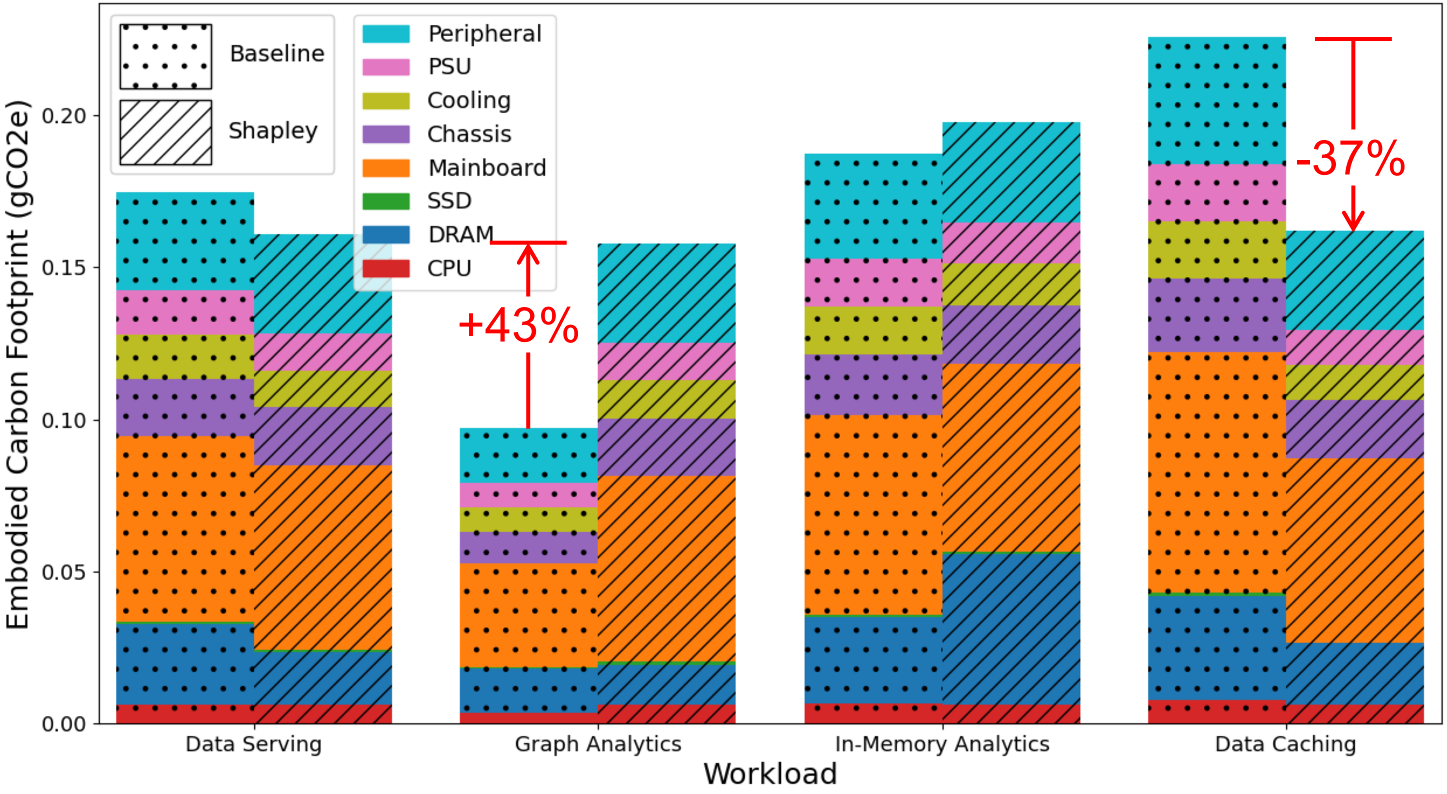

Limitations and Future Work

Limitation: Poor scalability

- Shapley methodology scales at $O(N^2)$ with number of workloads
- Offline profiling not scalable for larger number of workloads

Future work: Scalable Shapley carbon attribution

Future work: Workload optimizations with carbon attribution

Fair Carbon Attribution for Cloud Data Centers Using Shapley Values

Shapley values provide theoretical fairness guarantees for carbon attribution

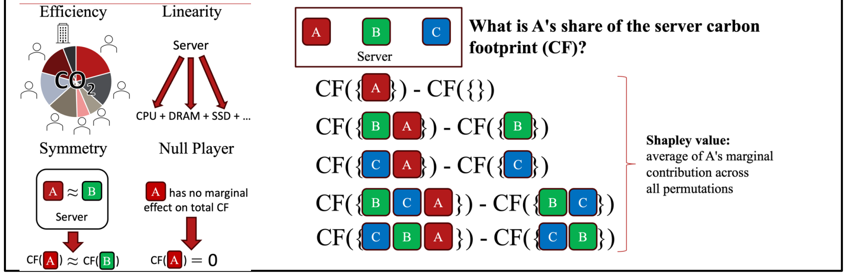

Operational carbon attribution using offline profiling

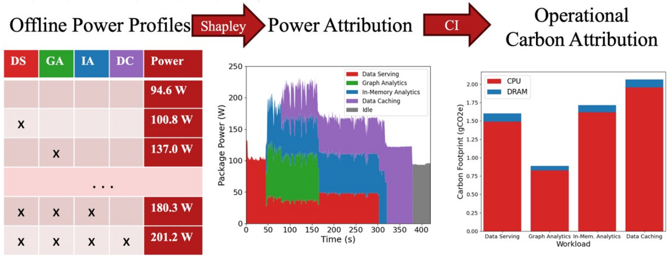

Embodied carbon attribution by tracking per-resource utilization

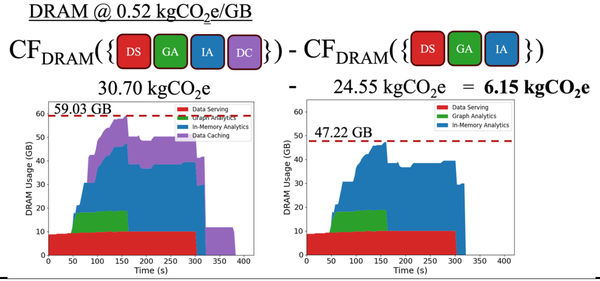

Thank you!

Please connect with us via email at:

lxh4@cornell.edu (Leo)