SCARIF: Towards Carbon Modeling of Cloud Servers with Accelerators

IEEE Computer Society Annual Symposium on VLSI, ISVLSI 2024

Shixin Ji*, Zhuoping Yang*, Xingzhen Chen*, Stephen Cahoon*, Jingtong Hu*, Yiyu Shi§, Alex K. Jones*, and Peipei Zhou*

University of Pittsburgh*; University of Notre Dame§

peipei.zhou@pitt.edu
shixin.ji@pitt.edu

https://peipeizhou-eecs.github.io/

https://github.com/arc-research-lab/SCARIF

Carbon Cost for a Server? Dell, Lenovo, HPE

PowerEdge R840

Report produced February, 2019

From design to end-of-life and everything in between, we work to improve the environmental impact of the products you purchase. As part of that process, we estimate the specific impacts throughout the lifecycle. This includes the contributions from materials, manufacturing, distribution, use and end-of-life management.

This product's estimated carbon footprint:

15600 kgCO2e *

Estimated impact by lifecycle stage:

Dell uses PAIA (Product Attribute to Impact Algorithm) to perform product carbon footprints. PAIA is a streamlined LCA tool developed by MIT's Materials System Laboratory. It takes into consideration important attributes of the product which can be correlated to activities in order to calculate the product carbon footprint.

Due to high configurability of servers, the information provided here was calculated based on the products highest selling configuration (see assumptions on page 2).

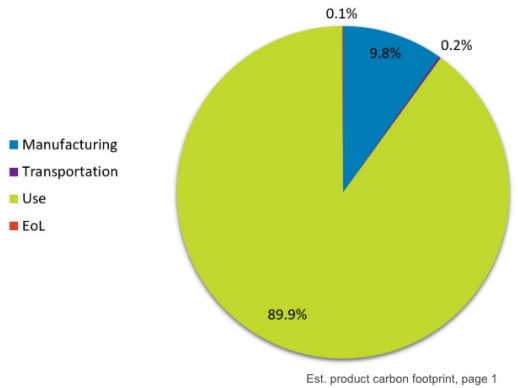

Est. product carbon footprint, page 1

Lenovo Product Carbon Footprint (PCF) Information Sheet

Storage

Commercial Name	Lenovo ThinkSystem DE120S
Model Number	7Y63
Issue Date	Dec. 03, 2021

Product Environmental Attributes	
(a) Product Carbon Footprint Value:	**8270 kg of CO_2e (see Note 1 below)**
(b) Product Picture:	**(c) Life Cycle Detail by Component & Life Stage (Pie Chart):**

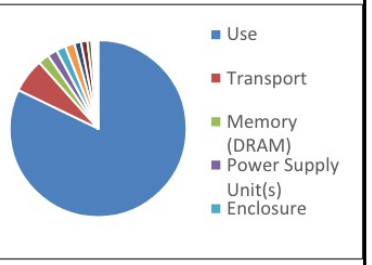

Note 1:
All estimates of carbon footprint are uncertain. Lenovo reports the 95th percentile of the carbon footprint estimate to reflect that uncertainty. For this product, that estimate has a mean of 6770 kg of CO_2e and standard deviation of 342 kg of CO_2e. For a quantity that follows a normal distribution, the 95th percentile value is equal to the mean plus the standard deviation multiplied by 1.64. Other organizations might report this value as 6770 +/- 342 kg of CO_2e. Results are for the server only. Datacenter values are not reported.

This PCF was generated using the Product Attribute to Impact Algorithm model, Version December 13, 2021, Date: December 13, 2021 (Product Type: 7Y63), © Massachusetts Institute of Technology's Materials Systems Laboratory, August 2012. Please refer to the Intended Uses and Limitations of the PAIA Model, © Massachusetts Institute of Technology's Materials Systems Laboratory, March 2017 for further details. Link to Document

This calculation was based upon a Lenovo ThinkSystem DE120S with the assumptions and configuration described in the calculation assumptions in the next page.

This pie chart provides the percent contribution of the mean value for each element of the analysis for the full life cycle CO_2e impacts of the product. Individual elements displaying 0% are less than 0.5%.

HPE Compute

Data sheet

HPE product carbon footprint

HPE ProLiant DL360 Gen10 Server

Hewlett Packard Enterprise recognizes the imperative to help minimize our environmental footprint and it is a core part of our strategy. This product carbon footprint (PCF) sums up the total greenhouse gas (GHG) emissions generated over this product's lifecycle. The product lifecycle includes manufacturing, transportation, use, and end of life. Read more about our approach to the circular economy.

This PCF uses Product Attribute to Impact Algorithm (PAIA) version 1.3.2 for manufacturing and end of life GHG emissions. PAIA is a streamlined lifecycle assessment (LCA) tool developed by the Massachusetts Institute of Technology's Materials Systems Laboratory. Read more about the intended uses and limitations of the PAIA model from "PAIA Intended Use" document.

Product input information

HPE ProLiant DL360 Gen10 Server

Table 1. Configuration inputs used to estimate the PCF

Lifecycle stage	Component	Base	Mainstream	Performance
Manufacturing	CPU	1x Intel® Xeon® Silver 8-core	2x Intel Xeon Silver 16-core	2x Intel® Xeon® Gold 24-core
	DRAM	64 GB	256 GB	768 GB
	Network adapter	1	2	3
	Storage controller	1x P408i-a	1x P408i-a	1x P816i-a
	SSD	2x 480 GB SATA RI	6x 800 GB SAS MU	10x 1.6 TB SAS MU
	Power supply	2x 800W Platinum	2x 800W Platinum	2x 800W Platinum
	Product weight	13 kg	15 kg	16 kg
Use	Product lifetime	4 years	4 years	4 years
	PUE	1.55	1.55	1.55
	Yearly energy (TEC)	841 kWh	1671 kWh	2036 kWh

HPE GreenLake

Carbon Cost for a Server? Dell Server R840

* This product has an estimated standard deviation of +/- 17500 kgCO2e

As part of our commitment to transparency, the chart to the right demonstrates the degree of uncertainty that exists within the PAIA model for product carbon footprinting, based on assumptions we have made for select variables.

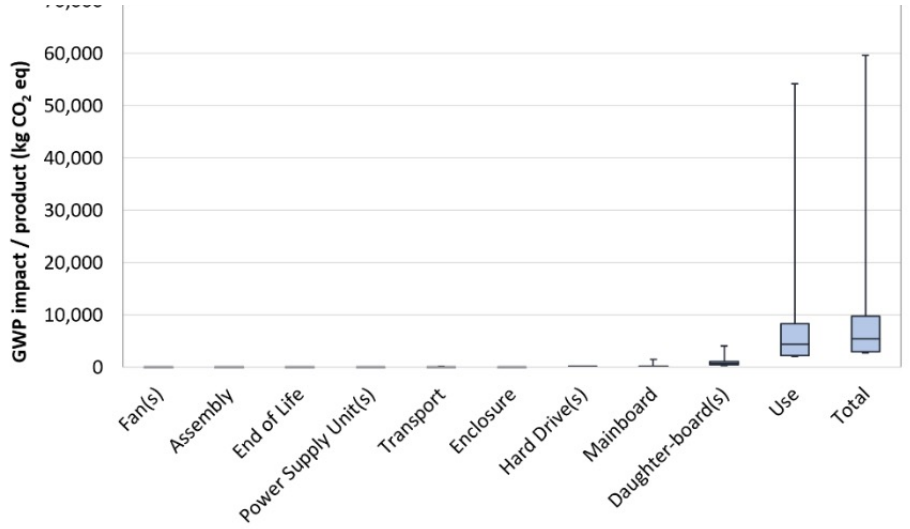

Assumptions for calculating product carbon footprint:

Product Weight	36.6 kg	**Server Type**	Rack	**Assembly Location**	EU
Product Lifetime	4 years	**Use Location**	EU	**Energy Demand (Yearly TEC)**	3325.7kWh
HDD/SSD Quantity	x4 1TB 2.5" HDD x2 300GB 2.5" HDD	**DRAM Capacity**	128GB	**CPU Quantity**	4

15600

kgCO2e

1 of these products… has a footprint approx. equivalent to **driving 38,220 miles** in a passenger car.

10 of these products… have a footprint approx. equal to what **184 acres of US forests** can absorb in a year.

To help our customers and other stakeholders contextualize product carbon footprint values, we provide

Carbon Cost for a Server? Dell Server R840

1 Dell Server R840
HWs: 4 CPUs, 128 GB RAM, 4 TB HDD
Product lifetime: 4 year
Total environmental impact / carbon cost?

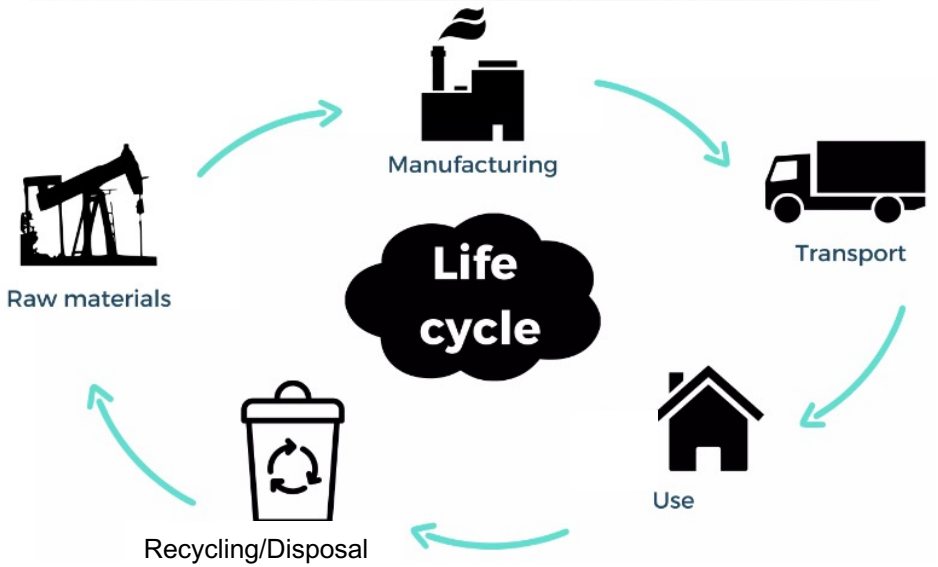

15600 kgCO2e

18 acres of US forests can absorb in a year

How large?

https://i.dell.com/sites/csdocuments/corpcomm_docs/en/carbon-footprint-poweredge-r840.pdf

Crowne Plaza Knoxville Downtown Ur

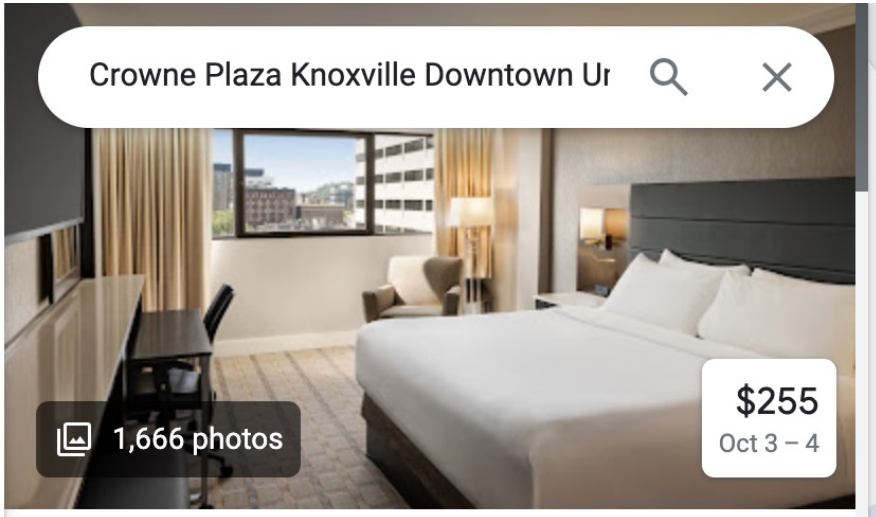

Crowne Plaza Knoxville Downtown University, an IHG Hotel

4.4 ★★★★½ (1,869) · 3-star hotel

Overview | Prices | Reviews | About

Directions | Save | Nearby | Share

CHECK AVAILABILITY

Compare prices

Free cancellation only ☐

Check in / Check out

Thu, Oct 3 | Fri, Oct 4 | 2

Book for Wed, Oct 2 - Thu, Oct 3 to save $70

Nearby hotels | Restaurants | Things to do | Bars | Coffee | Ph

W Vine Ave
S Gay St
W Summit Hill Dr SW
200.00 ft
100.00 ft
400.00 ft
920.06 ft
800.00 ft
700.00 ft
600.00 ft
Crowne Plaza Knoxville Downtown...
4.4 ★ (1869)
3-star hotel
LibertyX Bitcoin ATM
Blink Charging Station
Knox County Public Building Auth
Knox County Schools - Central Office
UT Tower
Layers

Measure distance
Click on the map to add to your path
Total area: 41,083.21 ft² (3,816.75 m²)
Total distance: 920.06 ft (280.43 m)

Map data ©2024 Google | United States | Terms | Privacy | Send Product Feedback | 50 ft

Carbon Cost for a Server?

1 Dell Server R840
HWs: 4 CPUs, 128 GB RAM, 4 TB HDD
Product lifetime: 4 year
Total environmental impact / carbon cost?

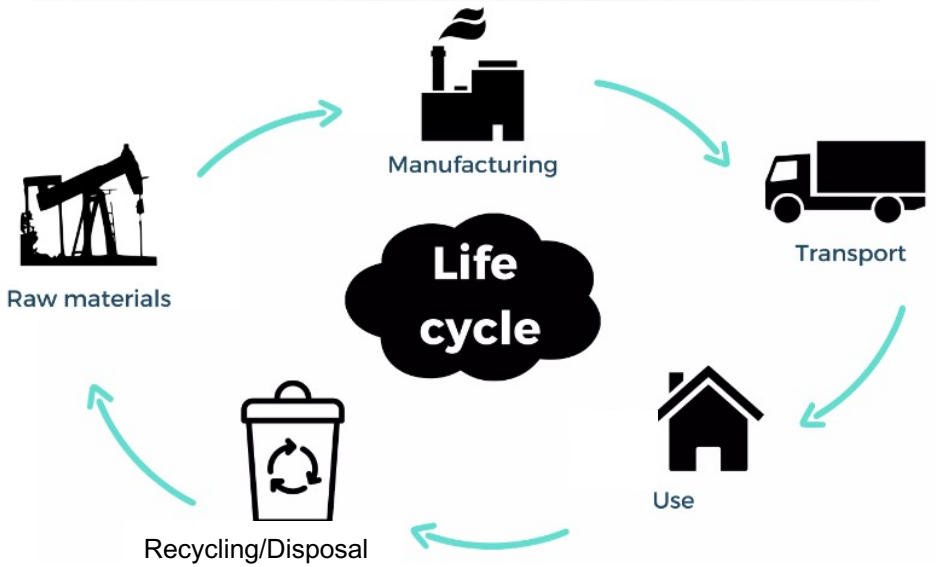

https://i.dell.com/sites/csdocuments/corpcomm_docs/en/carbon-footprint-poweredge-r840.pdf

15600

kgCO2e

18 acres of US forests can absorb in a year

How large? 18 Crowne Plaza Hotel

Carbon Cost for a Server?

More than 100 million data centers servers in US, estimated* in 2021

Total carbon impact?

1.8 billion acres of forests absorb in a year

⇔0.74 total U.S. land

⇔1.28 Amazon rainforest

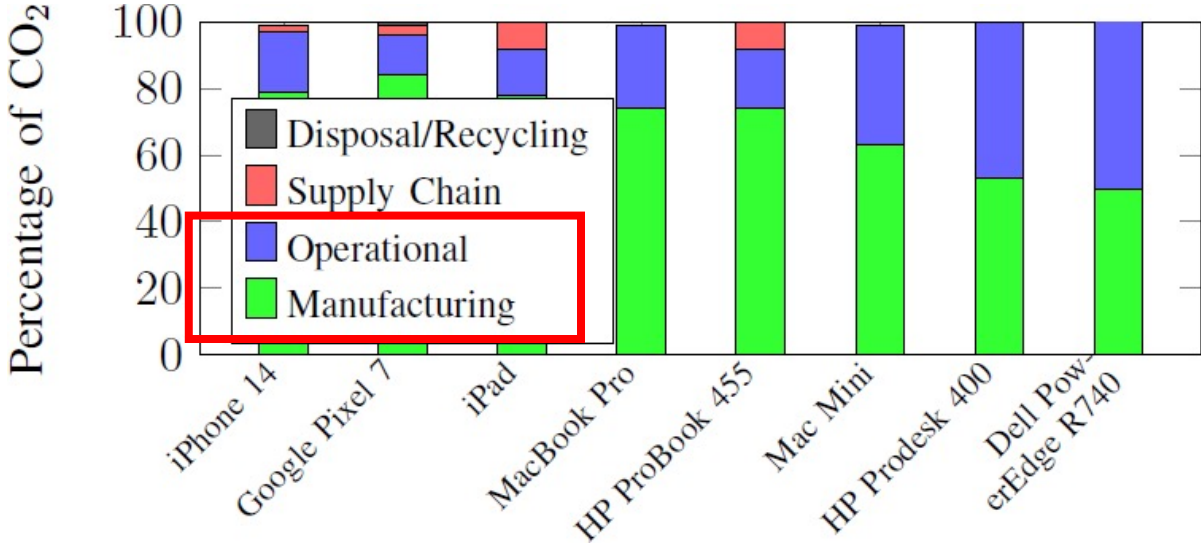

MORE: data center & edge servers, laptop/PC, cellphone, IoT, wearables, analog/mixed-signal, emerging devices,…

Manufacturing + Operational
> 95% Total Carbon

Improve Carbon Cost Efficiency is the Key:

- Improve Carbon Cost Efficiency in Operational Phase
- Improve Carbon Cost Efficiency in Manufacturing Phase

https://www.usitc.gov/publications/332/executive_briefings/ebot_data_centers_around_the_world.pdf
https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator

Carbon Cost for a GPU? => Missing Report???

impacts, and responsible sourcing of materials in the supply chain

- Participation in industry organizations and engagement with suppliers
- Supplier audits to ensure compliance with standards and requirements
- Protection of human rights throughout the supply chain

NVIDIA expects suppliers to comply with all applicable industry compliance and legal requirements, including:

- Conflict Minerals
- IEC 62474
- REACH
- RoHS
- WEEE
- UFLPA

We closely manage our supply chain to deliver innovative products that satisfy our customers' expectations in a socially and environmentally conscious manner.

Social and Environmental Performance

We utilize suppliers, such as Taiwan Semiconductor Manufacturing Company Limited (TSMC) and Samsung Electronics Co. Ltd, to produce our semiconductor wafers. We then utilize independent subcontractors and contract manufacturers, such as Amkor Technology, BYD Auto Co. Ltd., or BYD Auto, Hon Hai Precision Industry Co., or Hon Hai, King Yuan Electronics Co., Ltd., Omni Logistics, LLC, Siliconware Precision Industries Company Ltd., and Wistron Corporation to perform assembly, testing, and packaging of most of our products and platforms. We use contract manufacturers such as Flex Ltd., Jabil Inc., and Universal Scientific Industrial Co., Ltd., to manufacture our standard and custom

from many other suppliers and consign key components to contract manufacturers.

Supplier Environmental Impact

Emissions are generated at every stage of our product lifecycle, including manufacturing within our supply chain. Since 2014, we've expected our key silicon manufacturing and systems contract manufacturing suppliers to report their annual energy and water usage, waste, greenhouse gas (GHG) emissions, and reduction goals and objectives through the RBA Environmental Survey or CDP. We also expect suppliers to have their GHG emissions verified by a third party. We use this supplier data to better understand our product manufacturing impact and allocate carbon emissions to our customers.

We regularly survey key suppliers to better understand the renewable energy performance and capability of our manufacturing supply chain. In FY23, over 60% of these suppliers reported renewable energy use.

Conduct and associated NVIDIA policies, including our Agreement for Manufacture Environmental Compliance. Since 2016, all NVIDIA Master Service Agreements (MSA) executed with suppliers require compliance with the RBA Code of Conduc

Our assessment process involves using th RBA-Online system to evaluate existing a potential new suppliers against product- compliance industry standards, social and environmental criteria, use of conflict min the RBA Code of Conduct, and NVIDIA's code of conduct. We use the results of th assessment, which includes a spending analysis, to determine their overall risk.

Every year, we perform an RBA risk assessment on all strategic suppliers bas on geography and type of industry. All suppliers conduct an annual self-assessm questionnaire (SAQ), and we expect bienn Validated Assessment Program (VAP) on- audits to validate the SAQ. We work with moderate- and high-risk suppliers to revie

Answer: SCARIF, open-source tool to report embodied carbon cost for servers with accelerator hardware (GPUs, FPGAs, etc.)

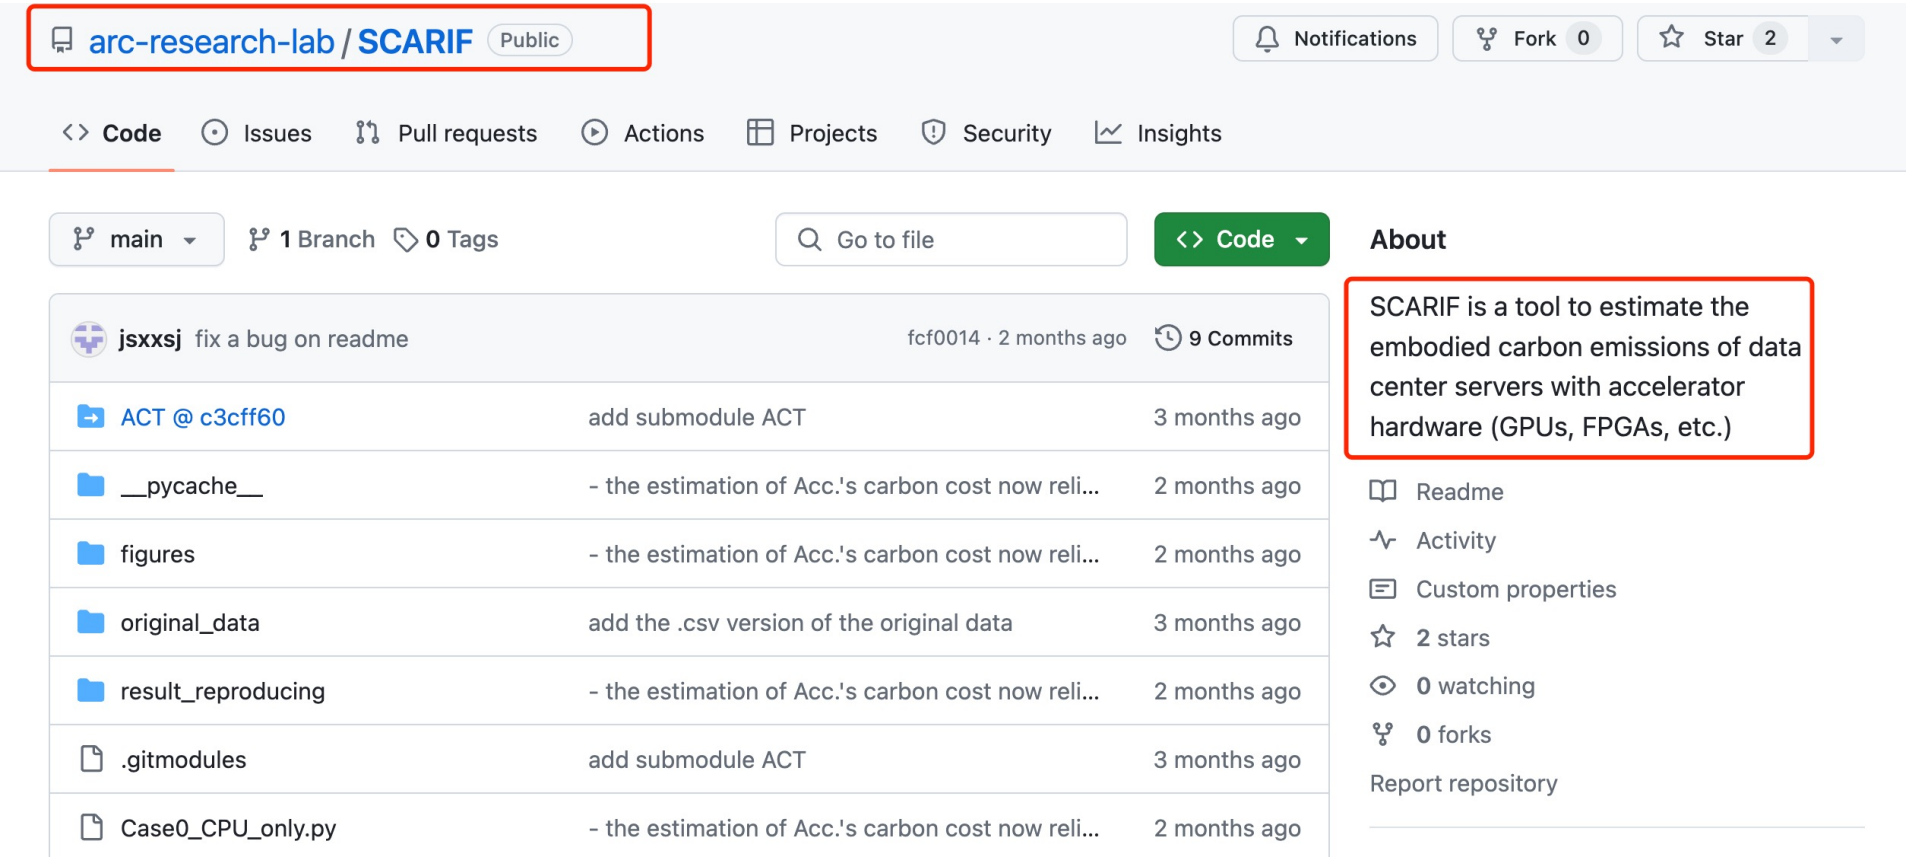

SCAN ME

SCARIF: overview

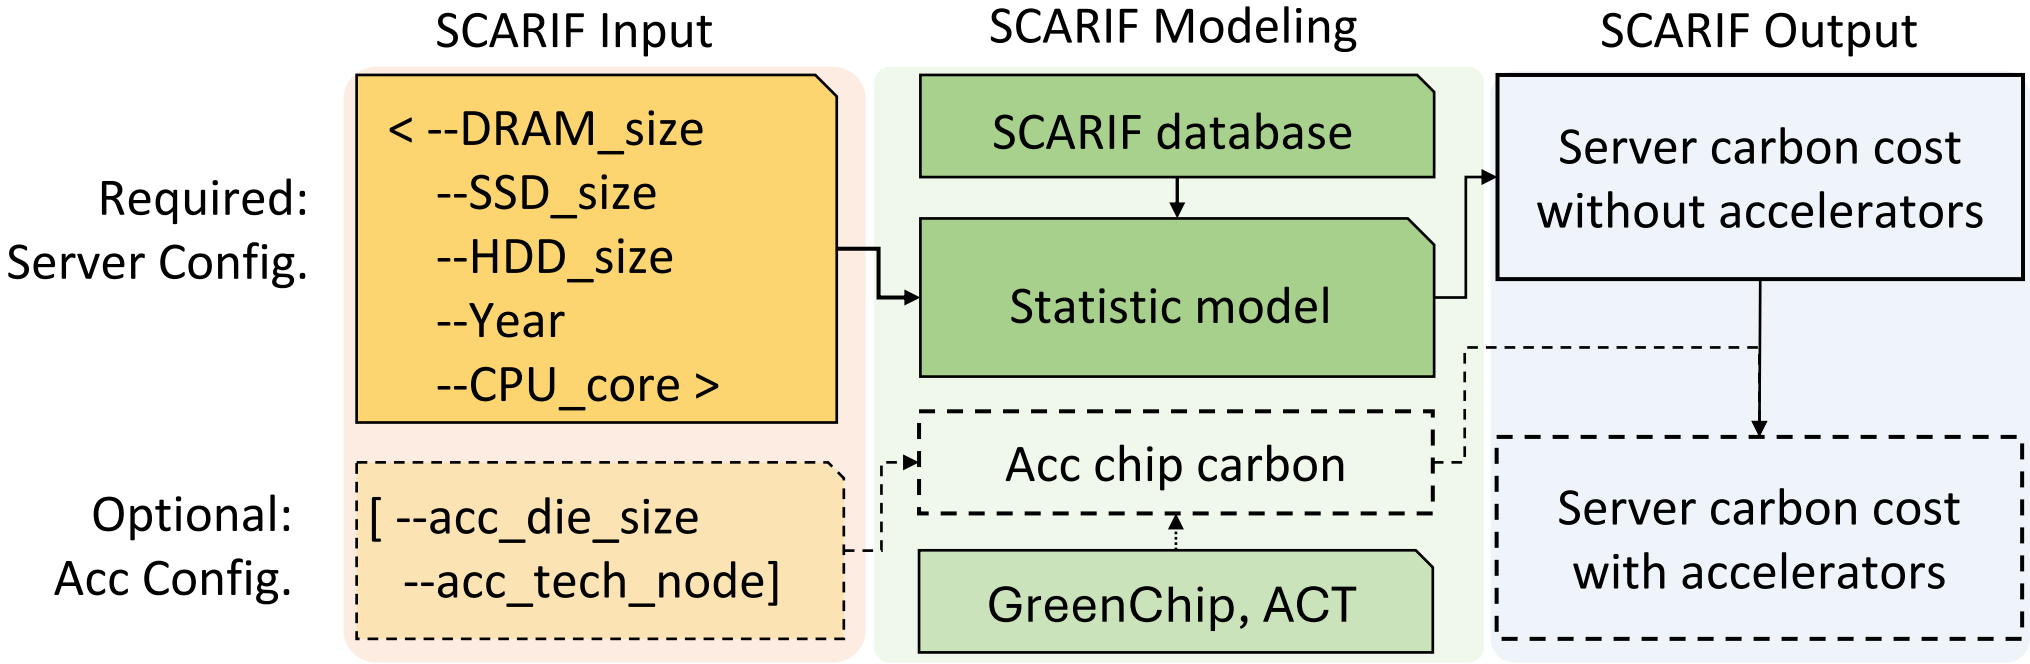

Existing tool, ACT, vs Vendor Report

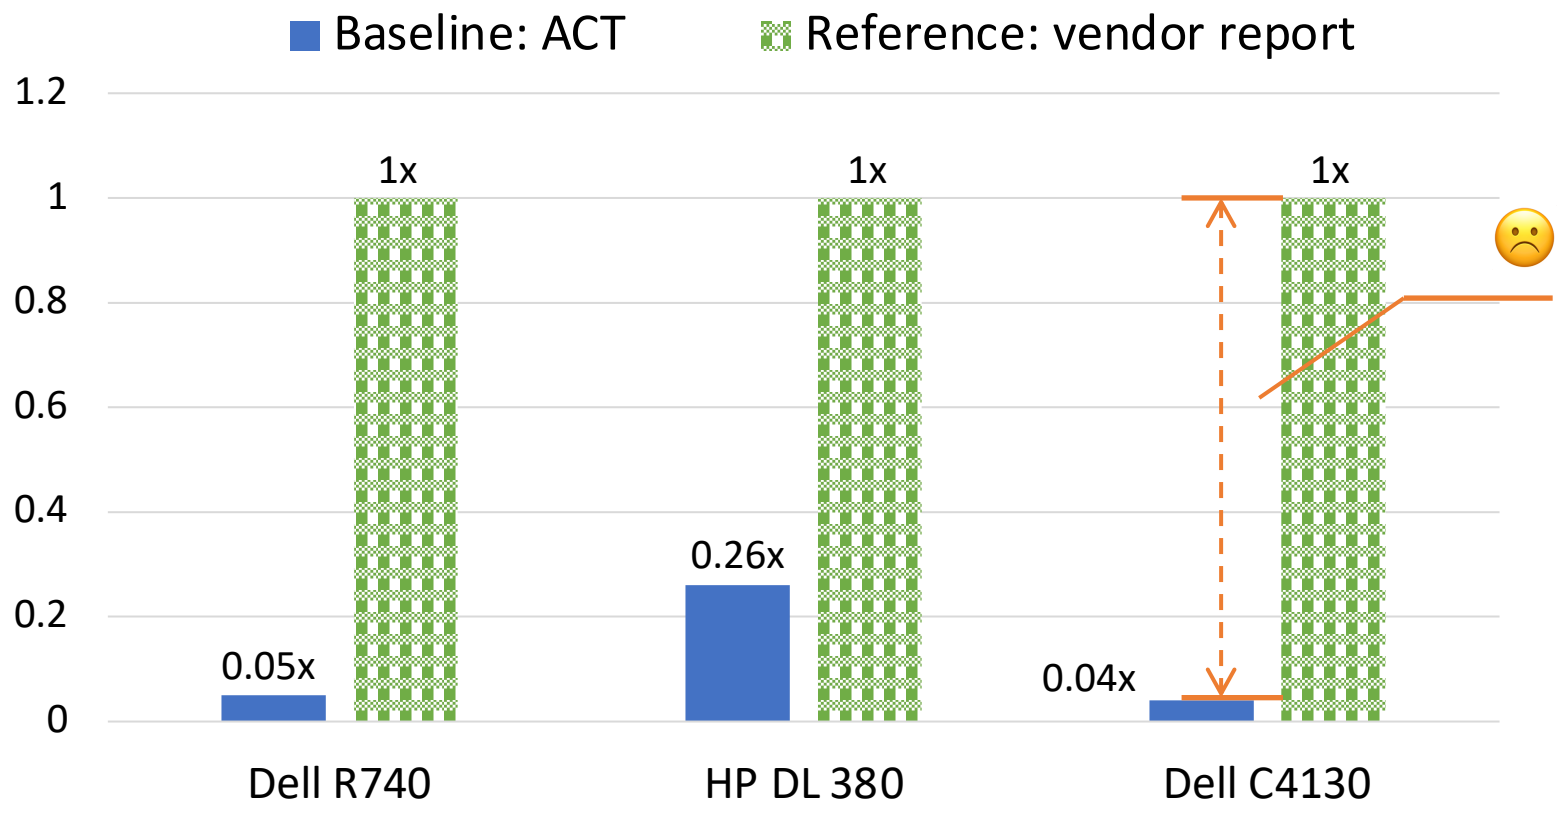

In the domain of servers, existing tool gives incompatible results with reports from the hardware vendors

Existing tool, ACT, vs Vendor Report

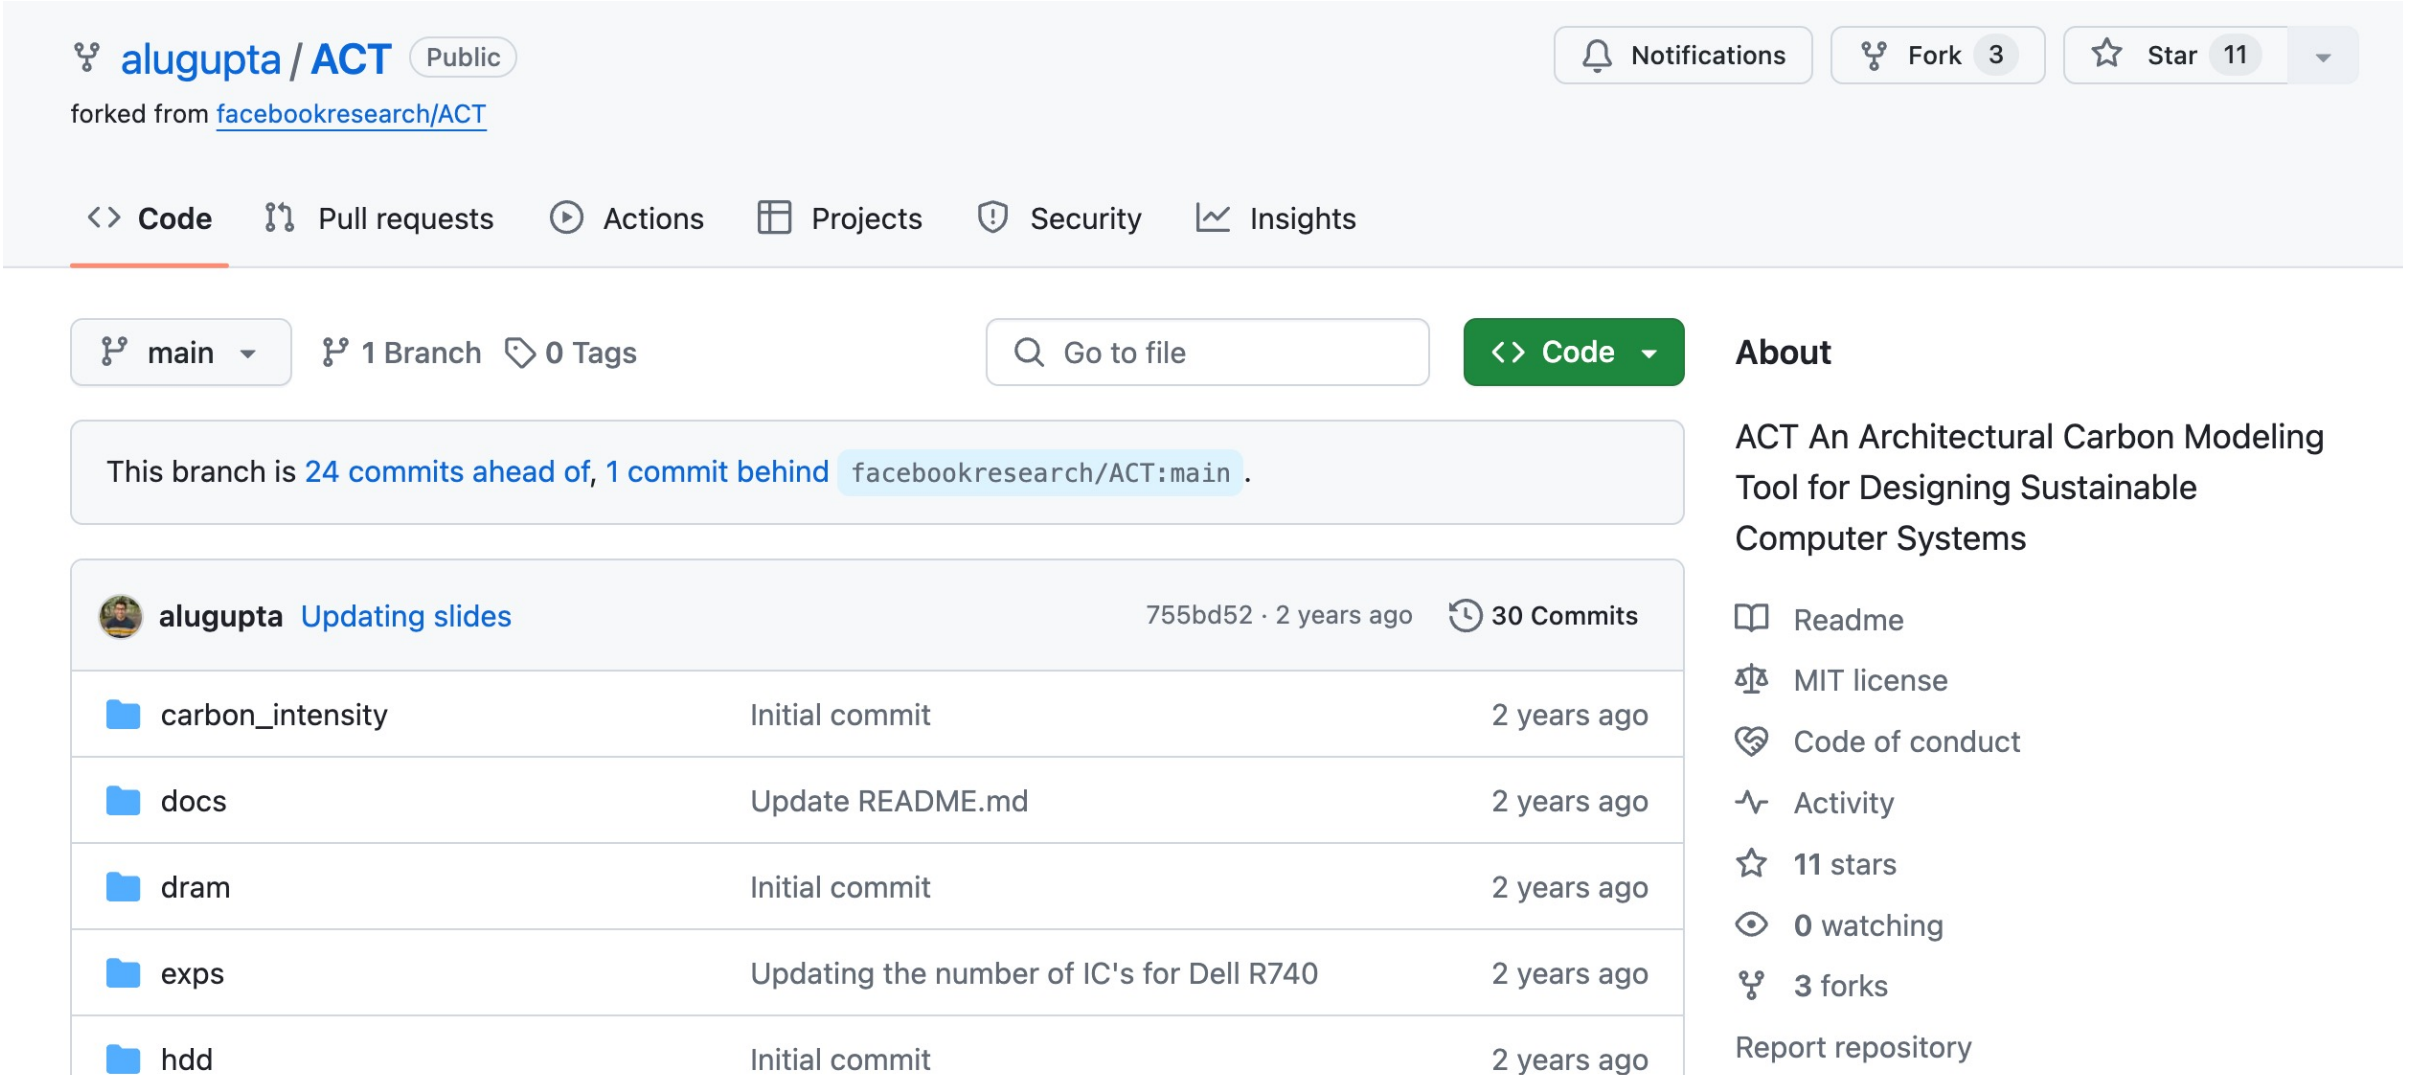

https://github.com/alugupta/ACT

ACT vs. SCARIF vs. Vendor Report

Challenges:

- Existing, bottom-up methods are hard to consider the ‘peripheral’ components
- Lack of reports about Accelerators

Solutions:

- Collect and analysis the reports from vendors
- Build a statistic model for analysis
- Scaling up the chip-level cost to server-level

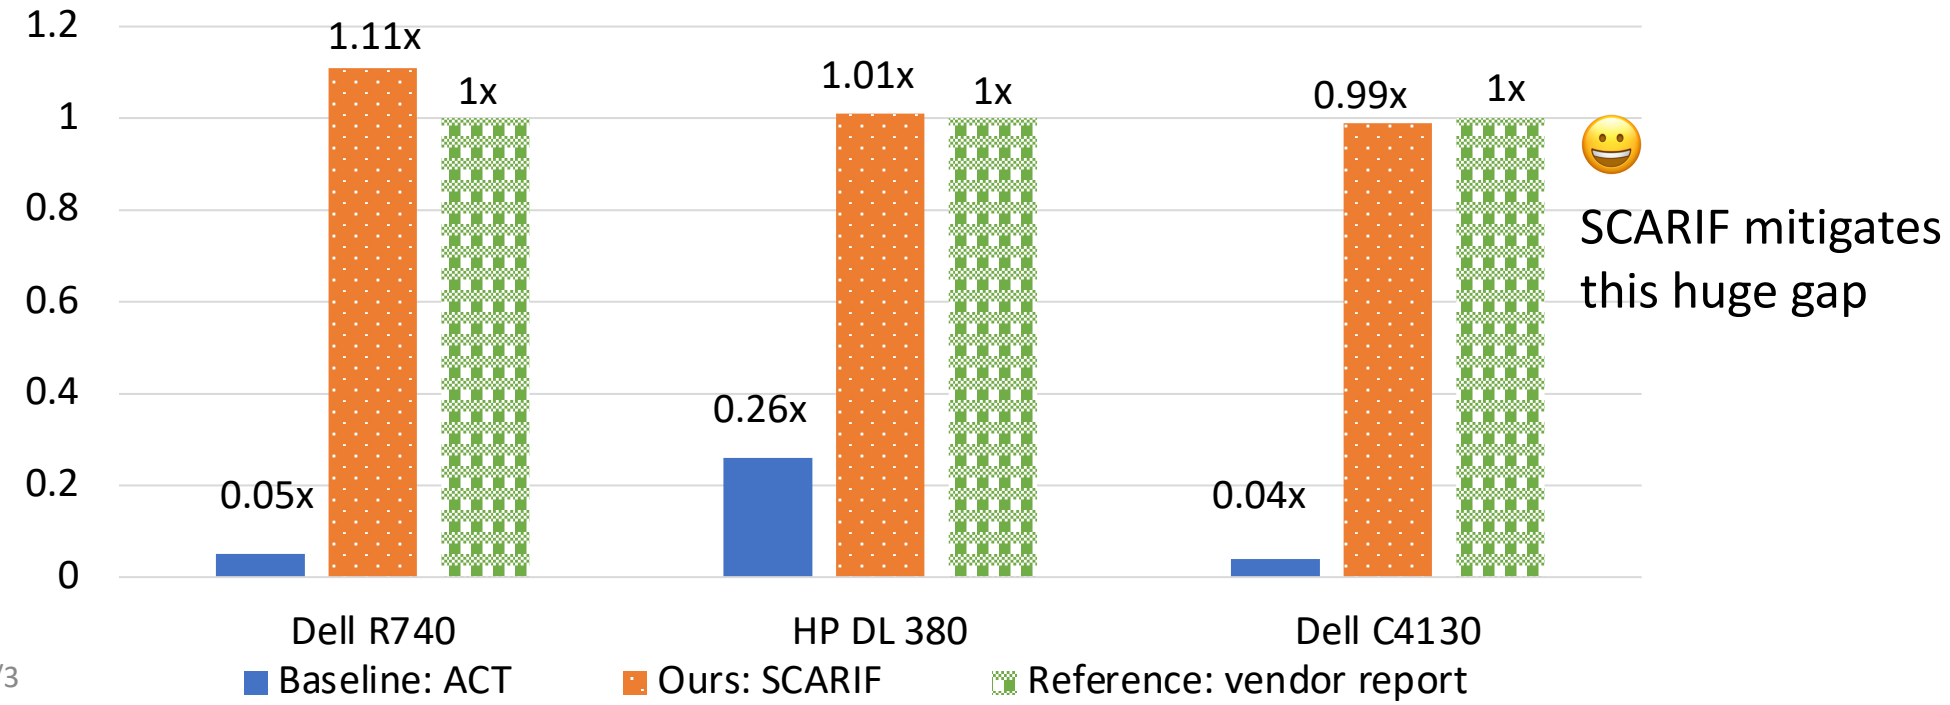

SCARIF mitigates this huge gap

In-depth Analysis for Dell R740

ACT estimation tool [1]

whole server

Dell's Report [2]

whole server

😕? Roughly 2x gap, looks good…?

[1] https://github.com/facebookresearch/ACT
[2] https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/Full_LCA_Dell_R740.pdf

In-depth Analysis for Dell R740

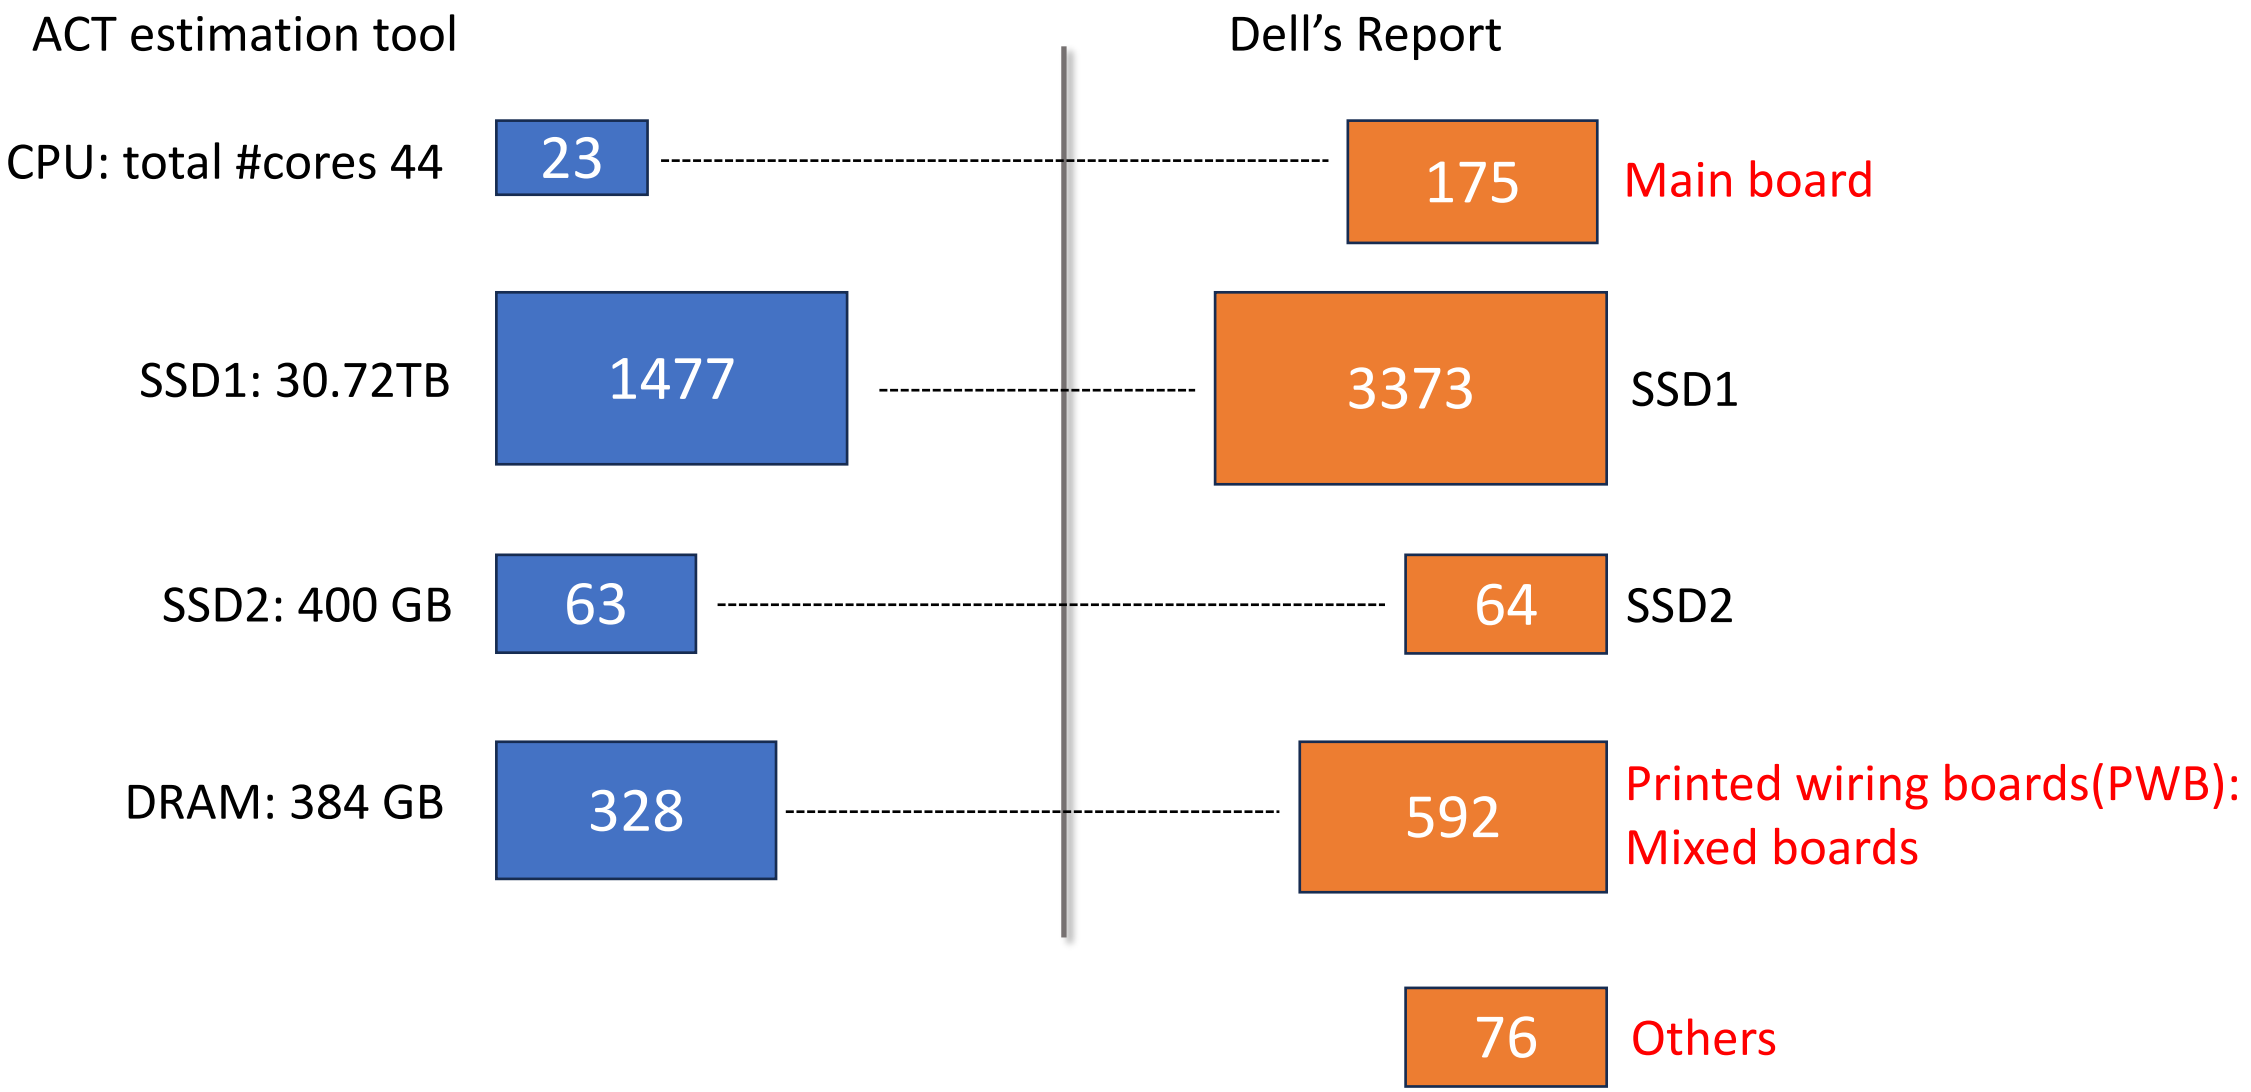

In-depth Analysis for Dell R740

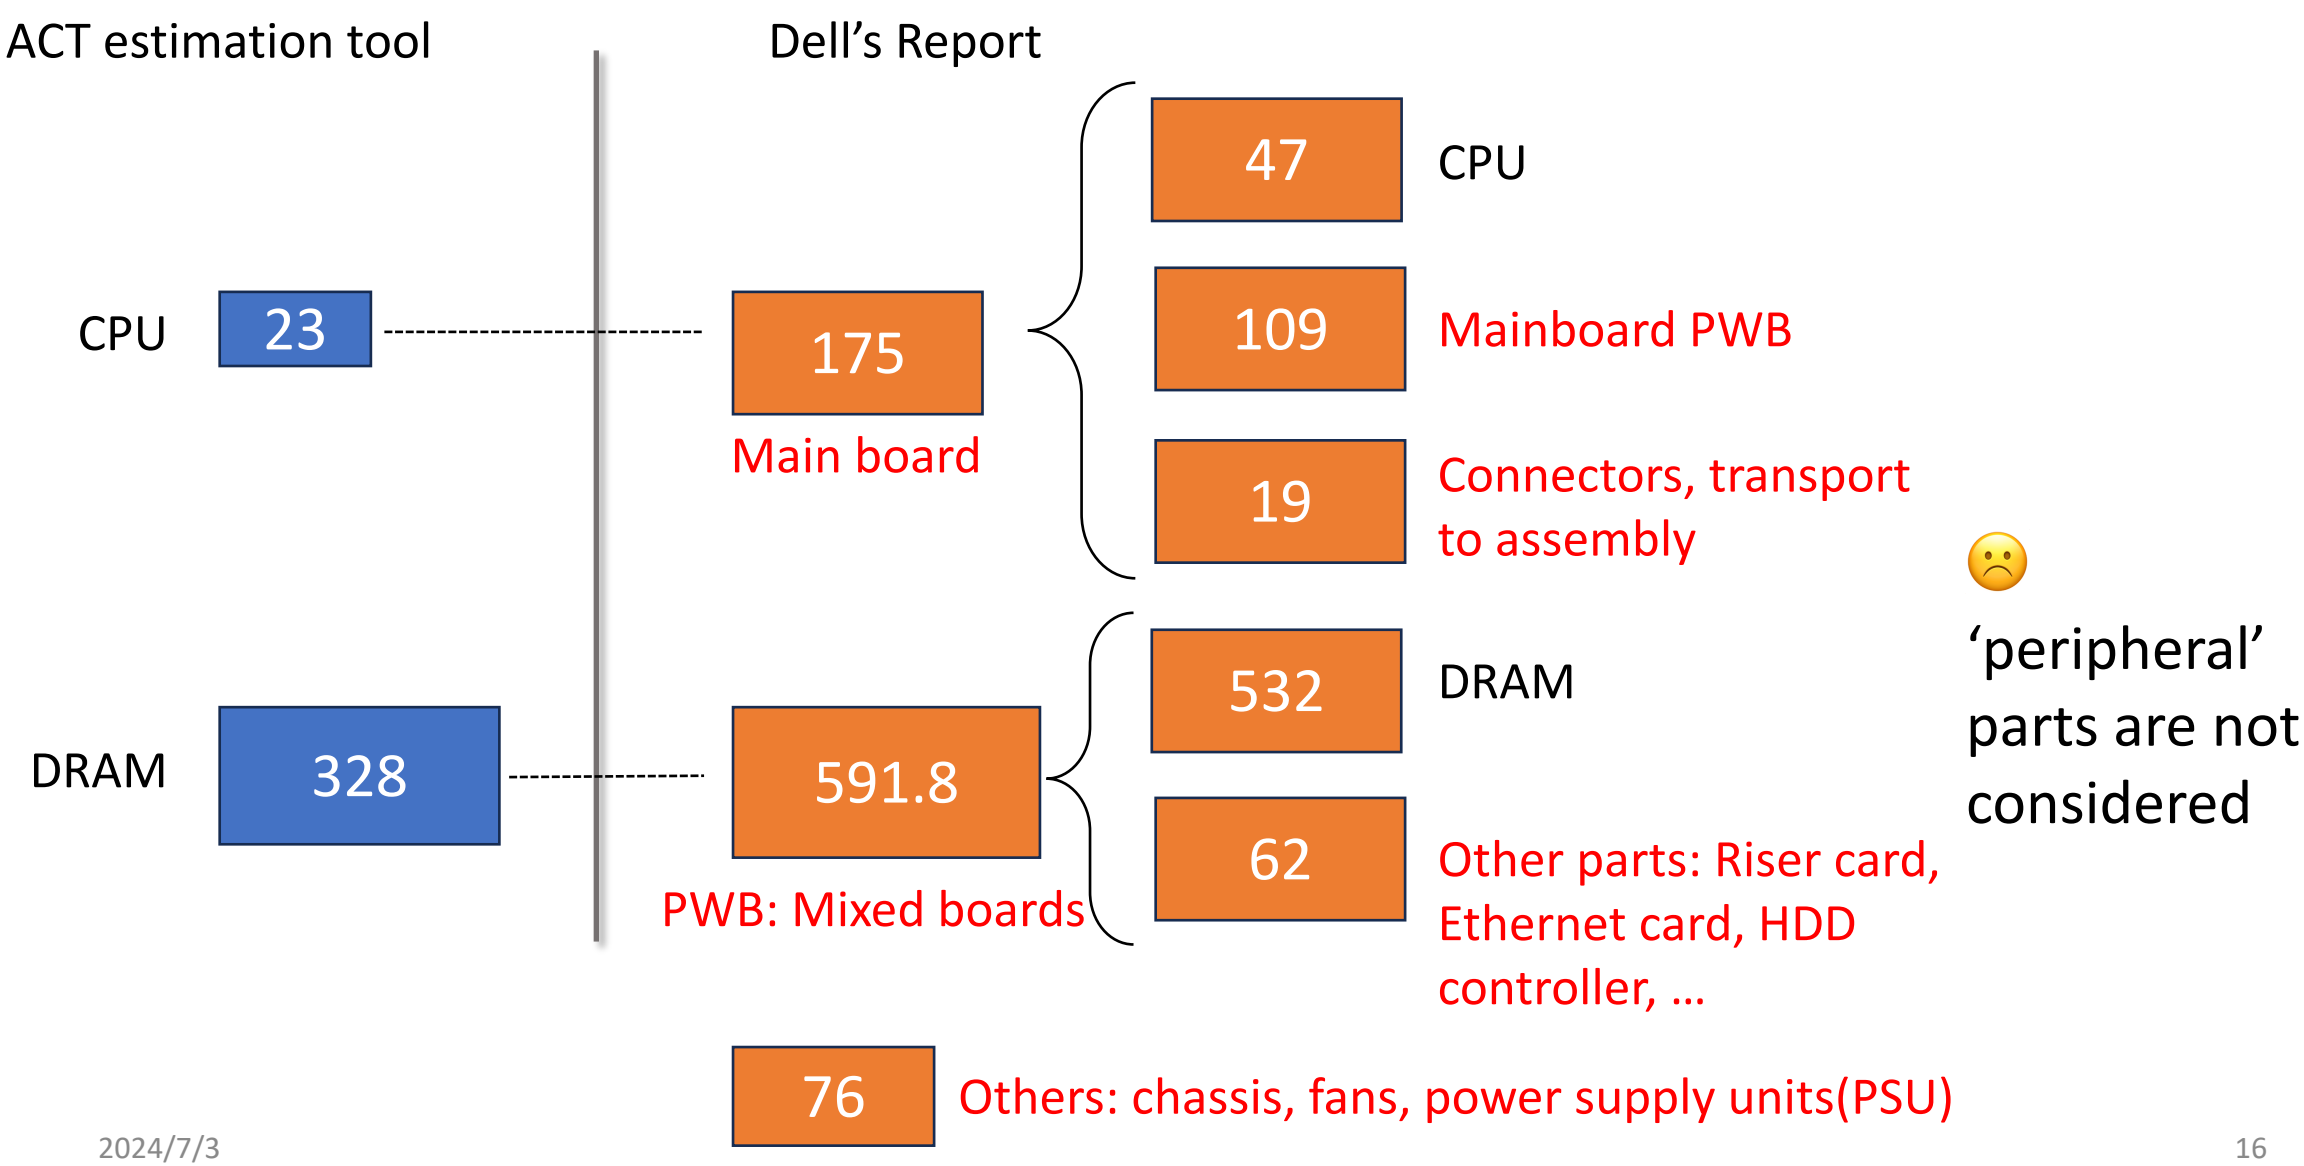

In-depth Analysis for Dell R740

Another Setup with smaller DRAM and less SSD disks

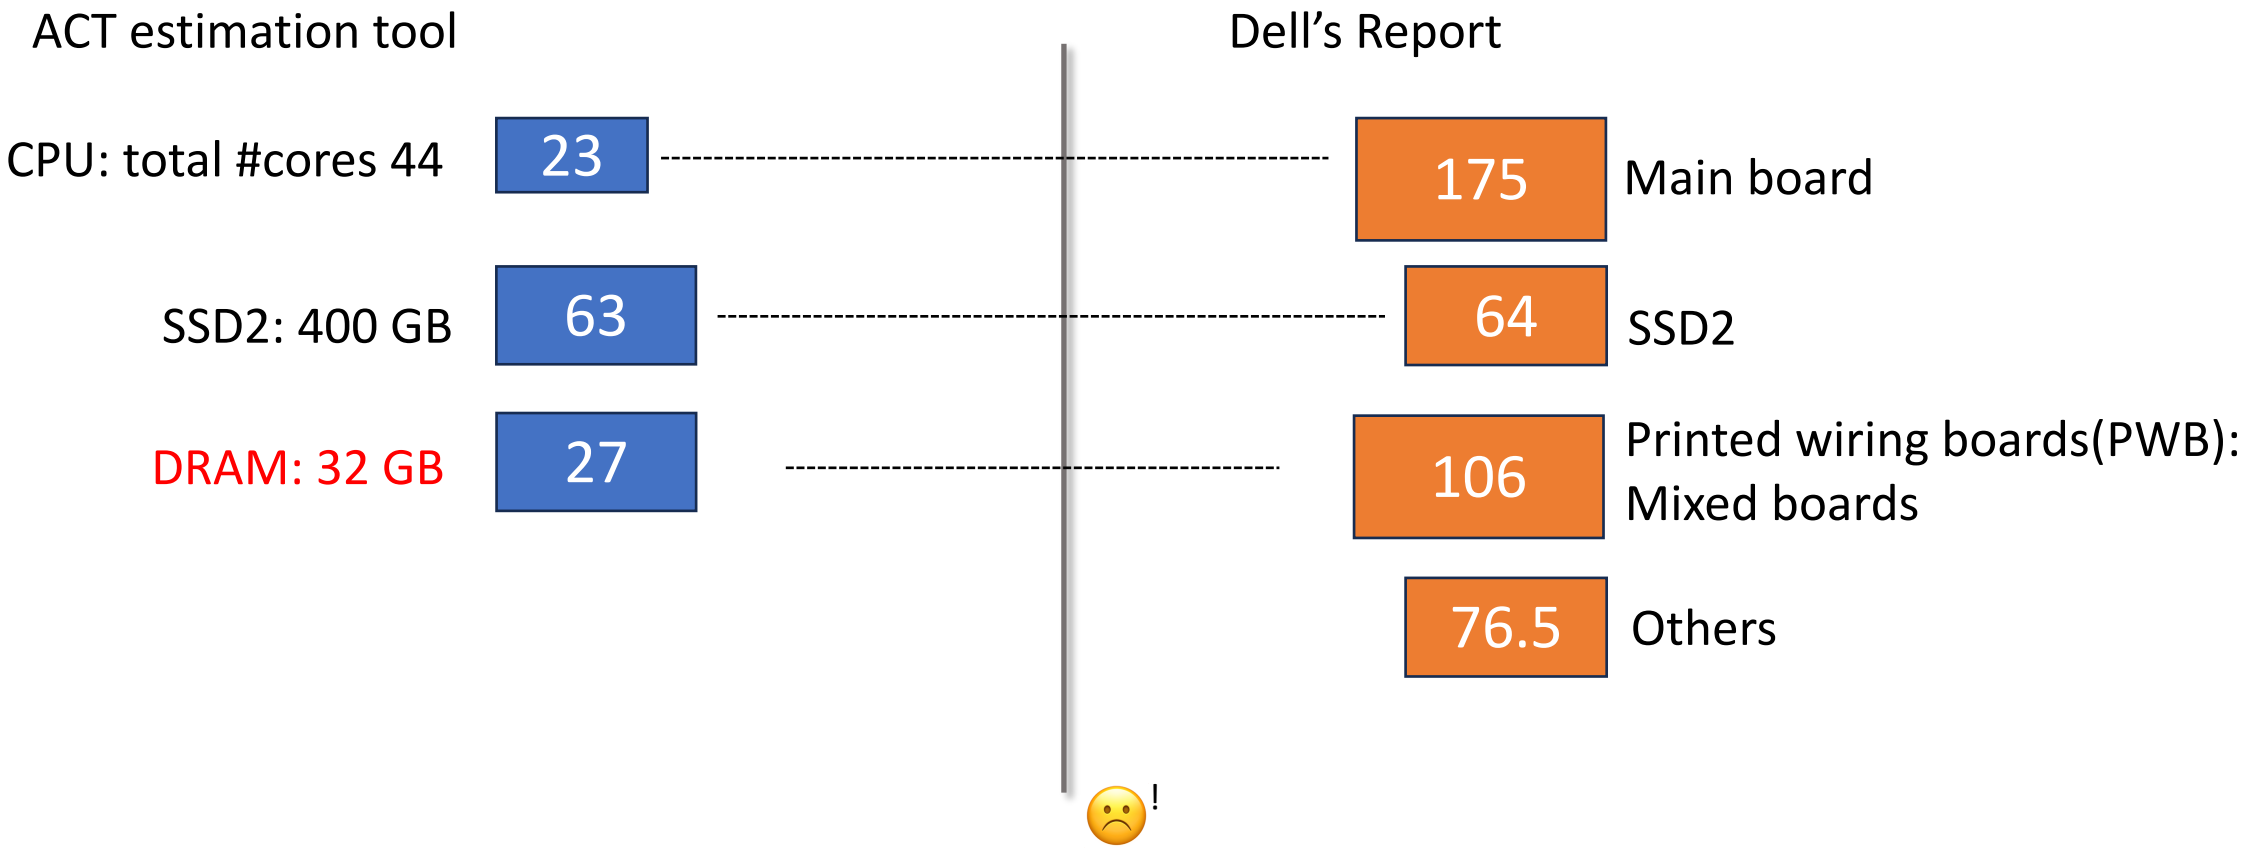

In-depth Analysis for Dell R740

A 3rd setup with GPU

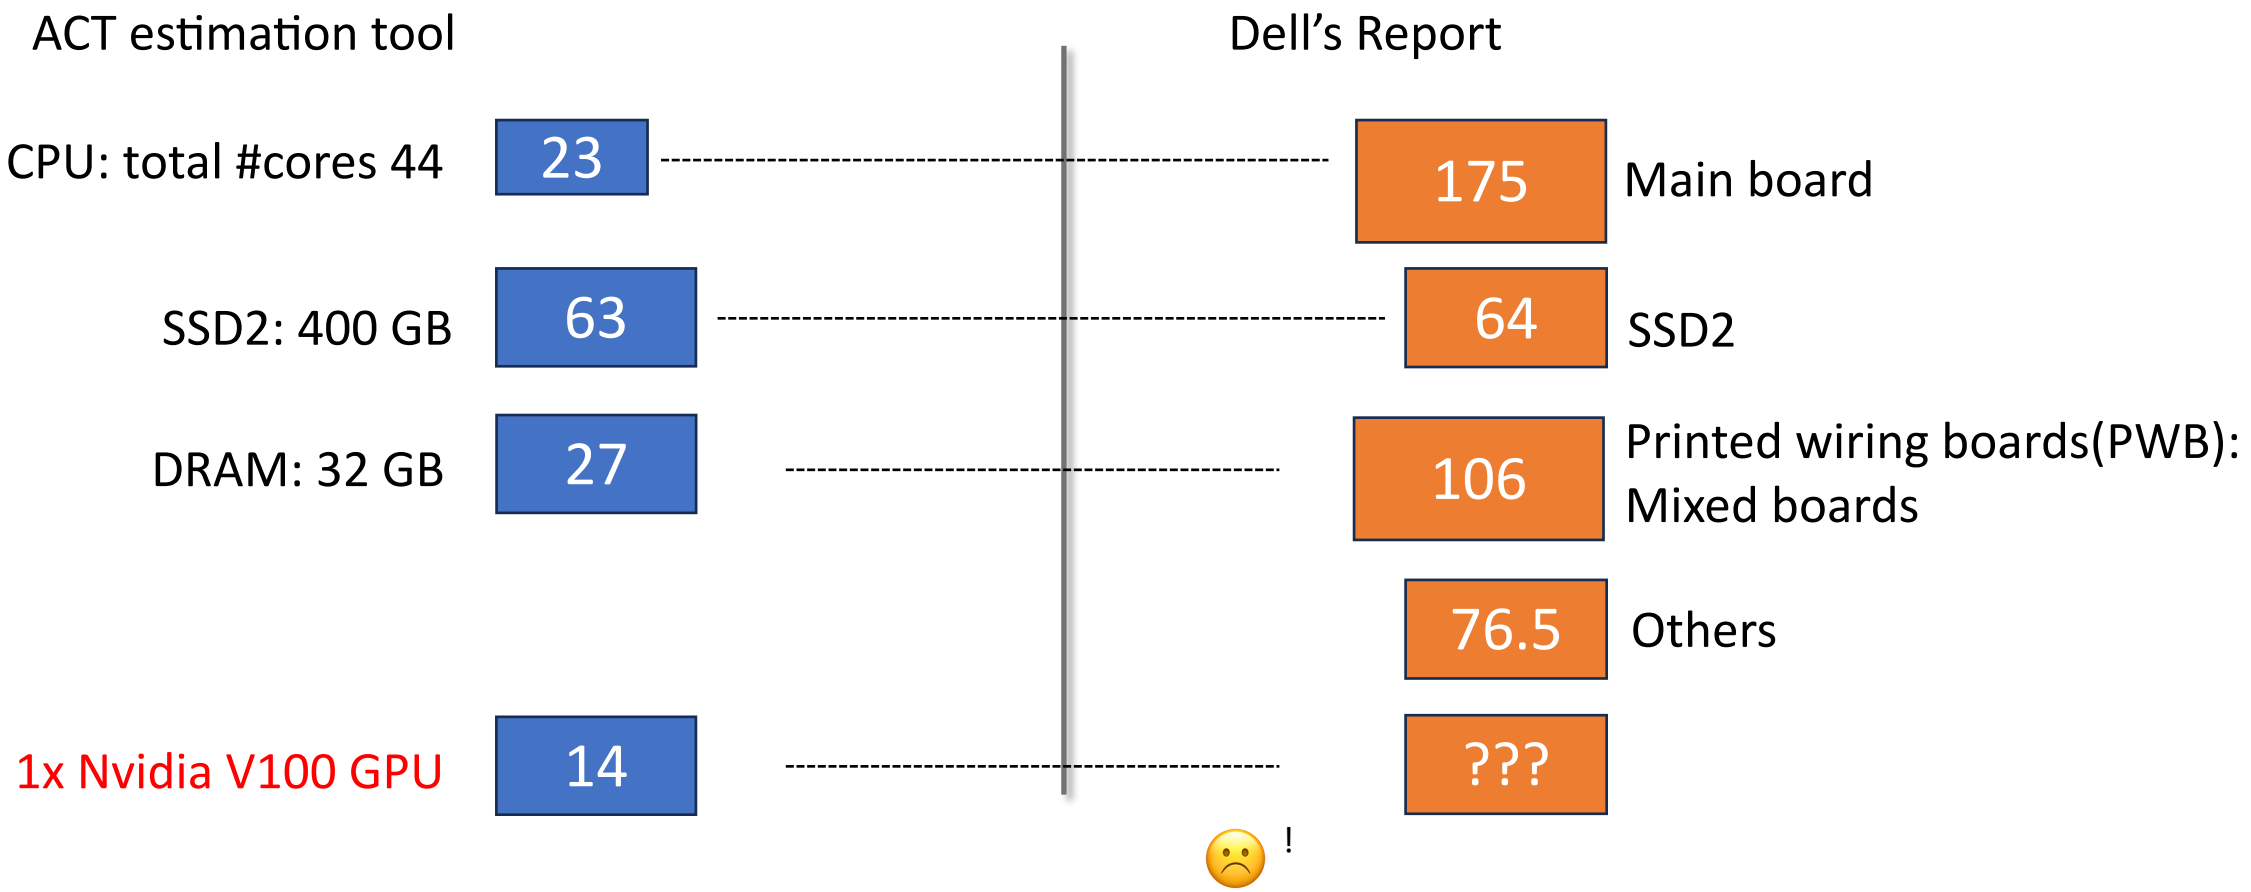

SCARIF Solutions: A Data-Driven Approach

Challenges:

- Existing, bottom-up methods are hard to consider the ‘peripheral’ components
- Lack of reports about Accelerators

Solutions:

- Collect and analysis the reports from vendors
- Build a statistic model for analysis
- Scaling up the chip-level cost to server-level

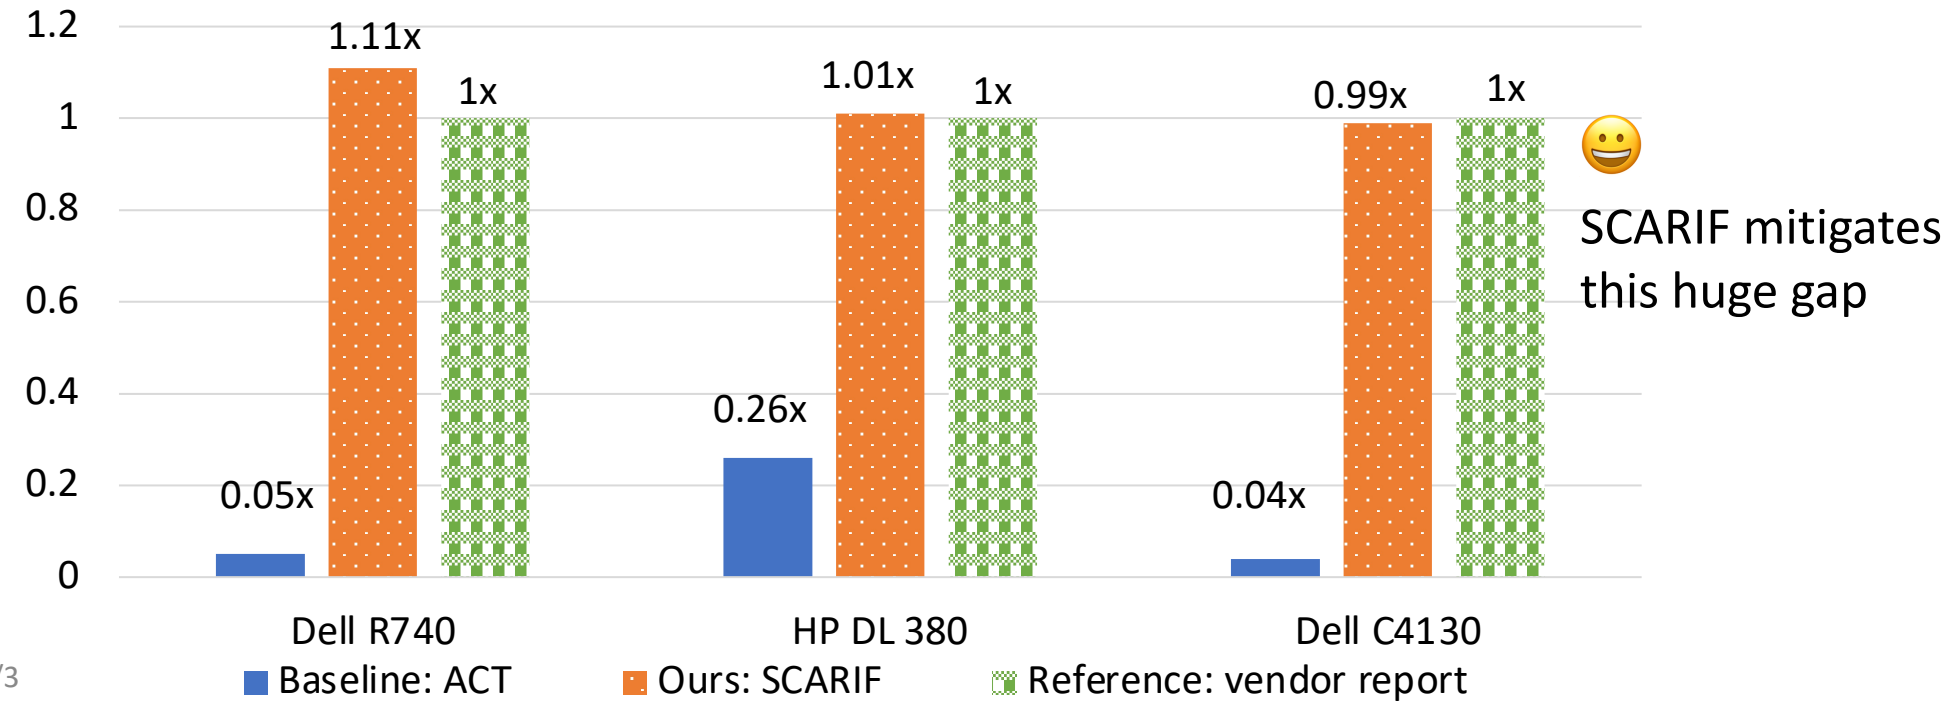

SCARIF mitigates this huge gap

SCARIF: Modeling on Servers

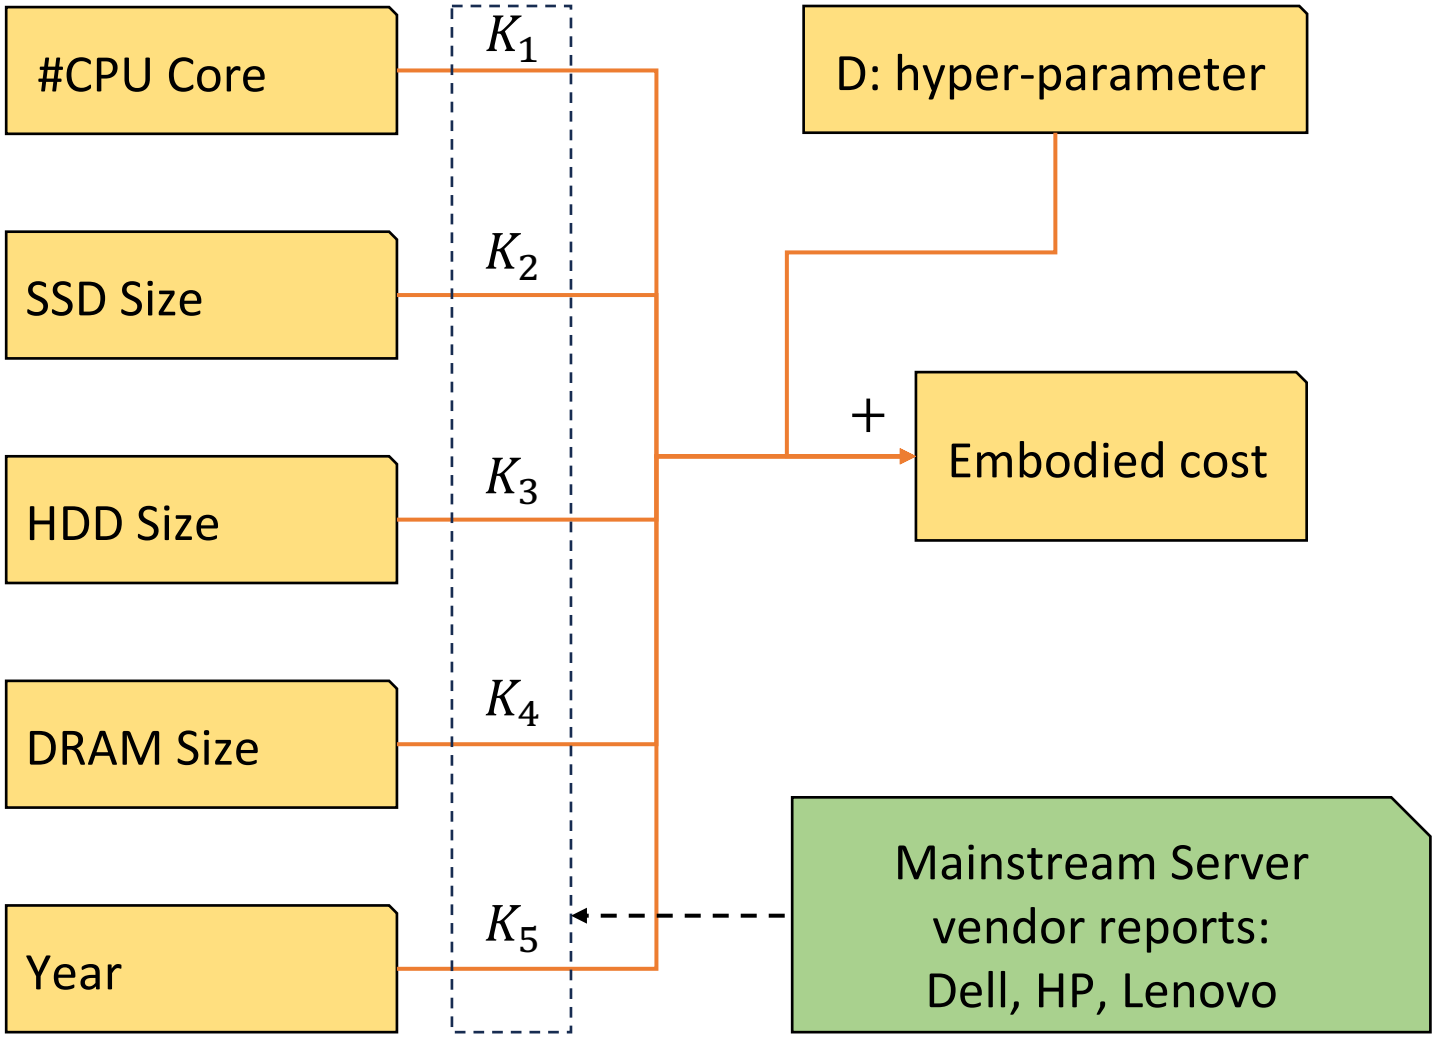

SCARIF: Validation Across Vendors

- Extract features on HP's reports:
- Transfer to Dell's Reports: average error < 0.15 standard deviation (std) from the vendor
- Transfer to Lenovo's Reports: average error < 0.5 std

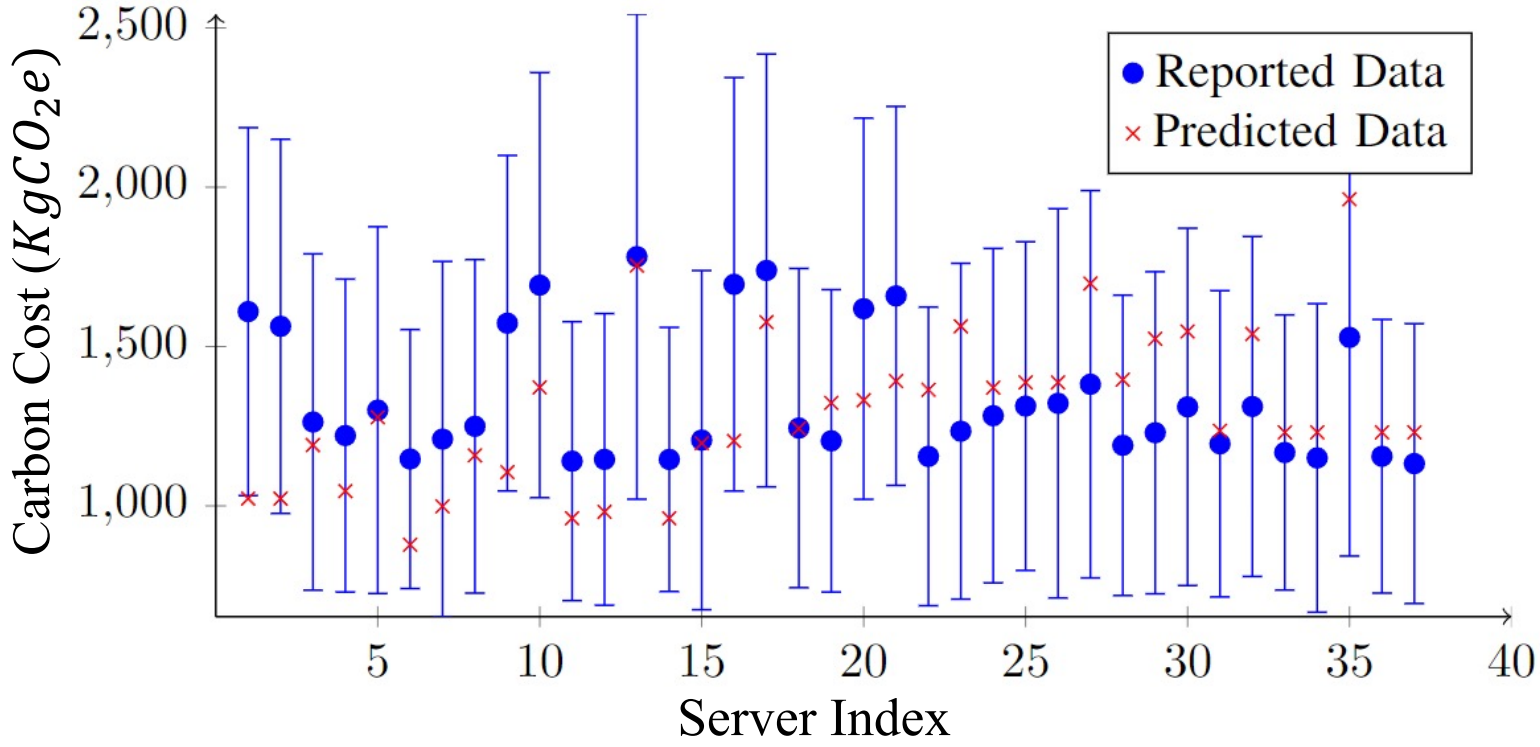

SCARIF: Modeling on Servers with Accelerators

- Basic assumption: the CPU-related part(chip + peripheral) shares the same pattern as the accelerator-related part (chip + peripheral).

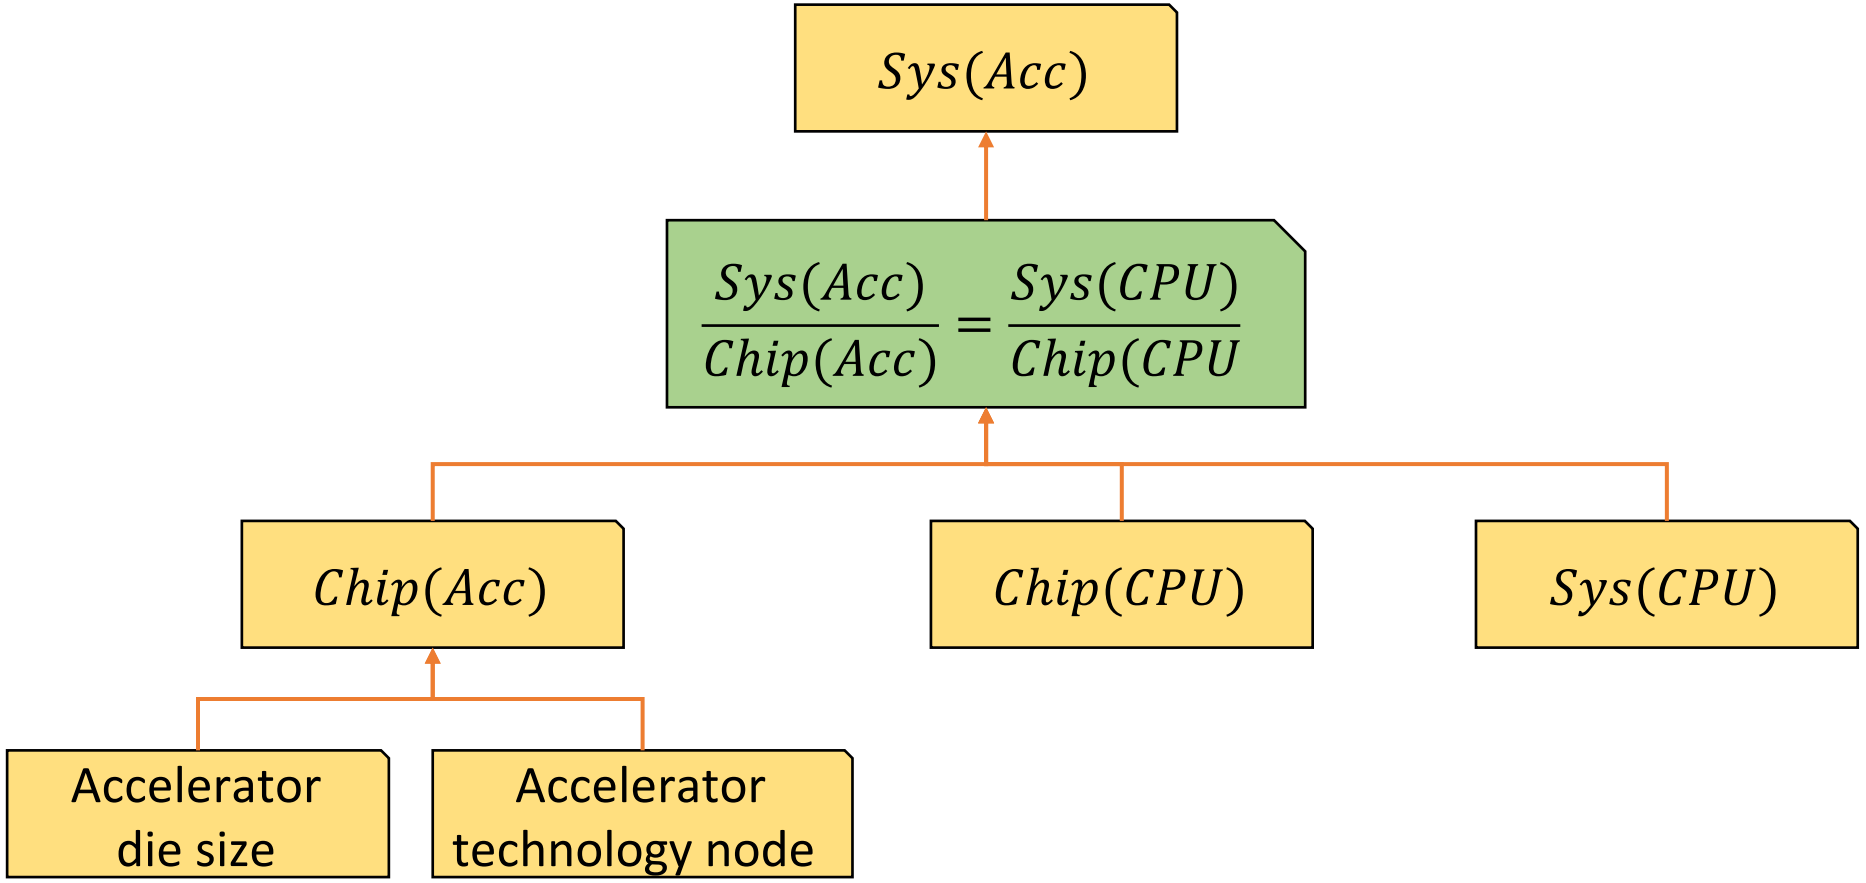

Analysis from SCARIF: Break-Even Analysis

- Research question: Is it worthwhile to pay additional embodied costs to upgrade servers for better energy efficiency? Which one can reduce overall carbon cost

- Experiment setup
  - System 1: 2017 server + 2017 GPU
    - 2x Xeon 8180 CPU
    - 64 GB DDR4 Memory
    - 1 TB HDD
    - Nvidia V100 GPU
  - Application:
    - DeiT-T model inference
    - System 1 runs at 100% utilization
    - System 2 runs at 62.2% utilization
  - System 2: 2020 server + 2020 GPU
    - 2x Xeon 8375 CPU
    - 64 GB DDR4 Memory
    - 1 TB HDD
    - Nvidia A100 GPU
  - Carbon intensity

State	AZ	CA	TX	NY
Carbon intensity $(KgCO_2e/kWh)$	0.395	0.234	0.438	0.188

Analysis from SCARIF: Break-Even Analysis

① The embodied cost of system 2(2542 $KgCO_2e$)

② System 2 uses less energy than system 1(save 1060 kWh yearly)

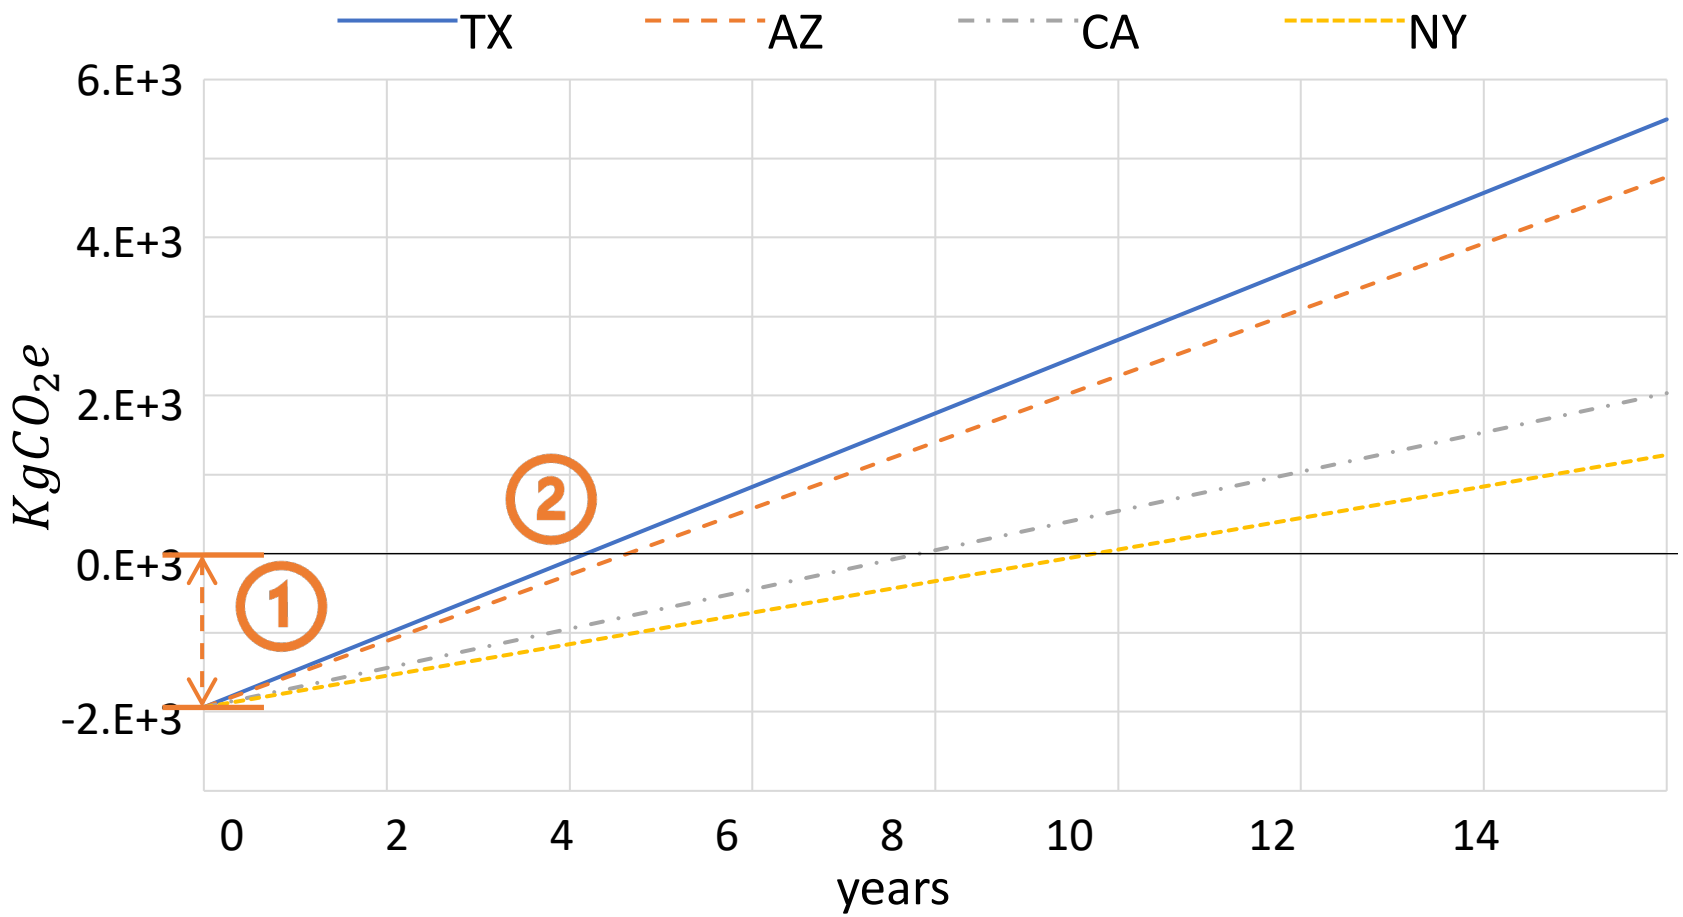

Analysis from SCARIF: Break-Even Analysis

Breakeven Time:
4.2 years (TX)
4.6 years (AZ)
7.8 years (CA)
9.8 years (NY)

Device may not reach the break-even time in its lifetime

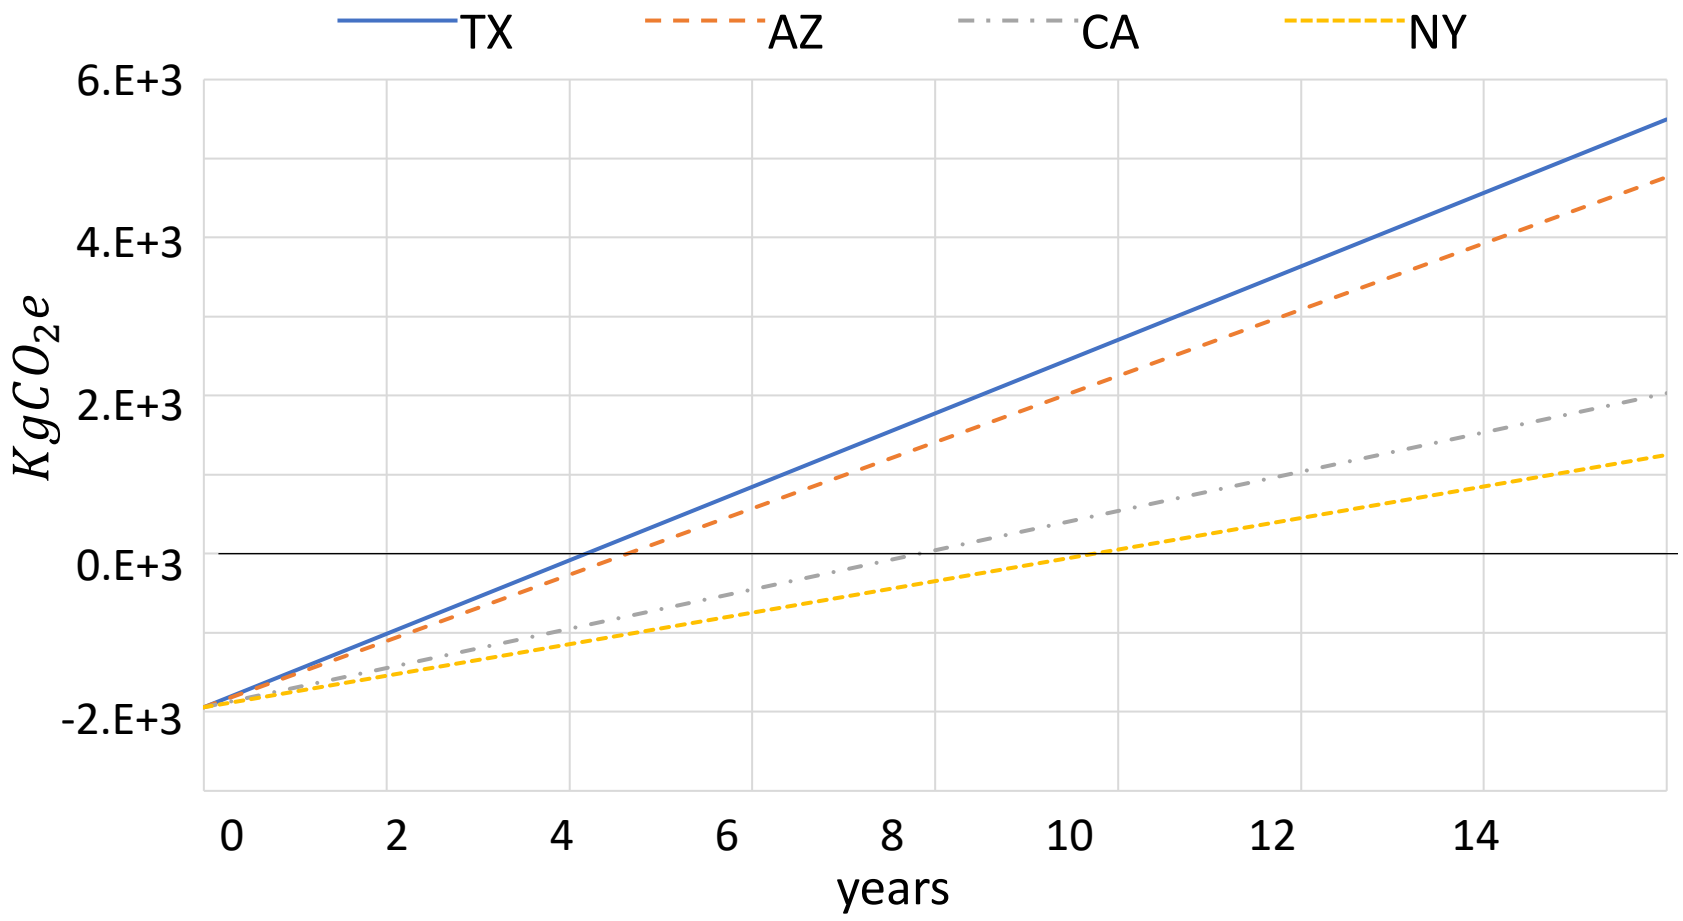

Thank You

SCARIF: Towards Carbon Modeling of Cloud Servers with Accelerators

SCAN ME

Shixin Ji*, Zhuoping Yang*, Xingzhen Chen*, Stephen Cahoon*, Jingtong Hu*, Yiyu Shi§, Alex K. Jones*, and Peipei Zhou*

University of Pittsburgh*; University of Notre Dame§

peipei.zhou@pitt.edu
shixin.ji@pitt.edu

https://peipeizhou-eecs.github.io/

https://github.com/arc-research-lab/SCARIF

National Science Foundation

AMD XILINX