Characterizing Datacenter Server Generations for Lifetime Extension and Carbon Reduction

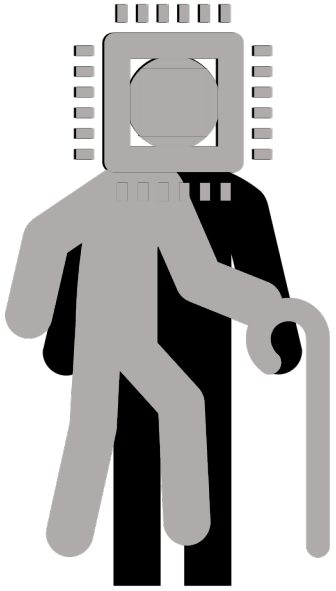

Jaylen Wang, Udit Gupta, Akshitha Sriraman

Carnegie Mellon University

Cornell University Founded A.D. 1865

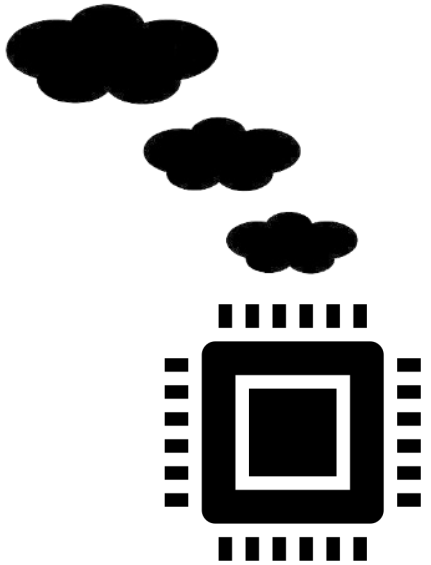

Computing has a carbon catch

If we overshoot 1.5°C (very likely):

Datacenters' electricity usage: Today: ~1%
Tomorrow: ~7% (2030)
[Jones'18]

Urgent need for new ways to reduce datacenter carbon emissions

What contributes to datacenter emissions?

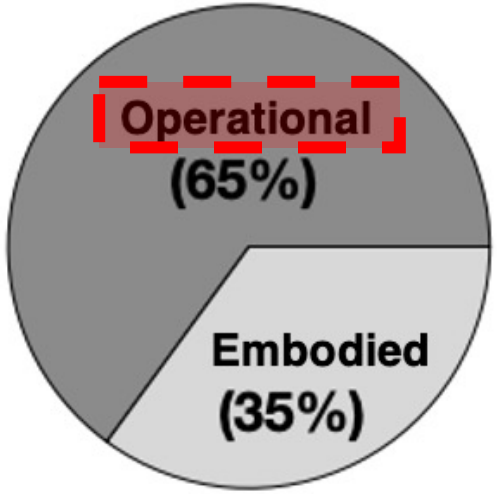

Without renewables

Facebook (2018) carbon footprint

[Gupta'21]

As renewables increase, embodied emissions is the **carbon bottleneck**

How to **lower embodied emissions**?

How to **lower embodied emissions**?

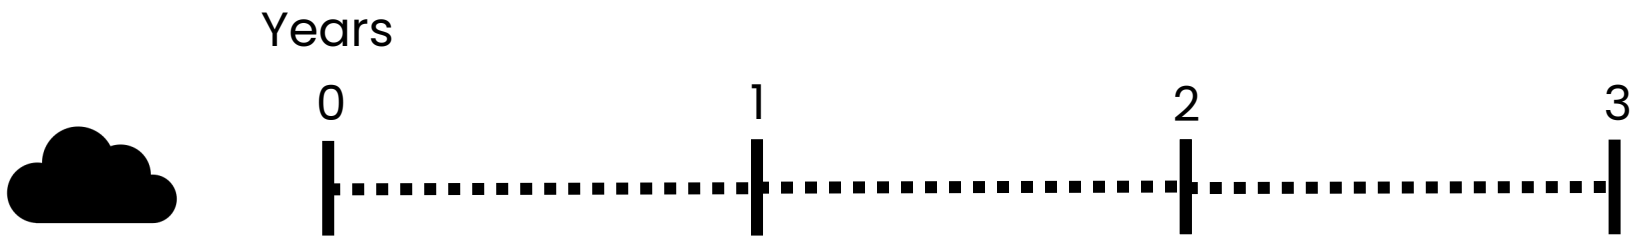

How to **lower embodied emissions**?

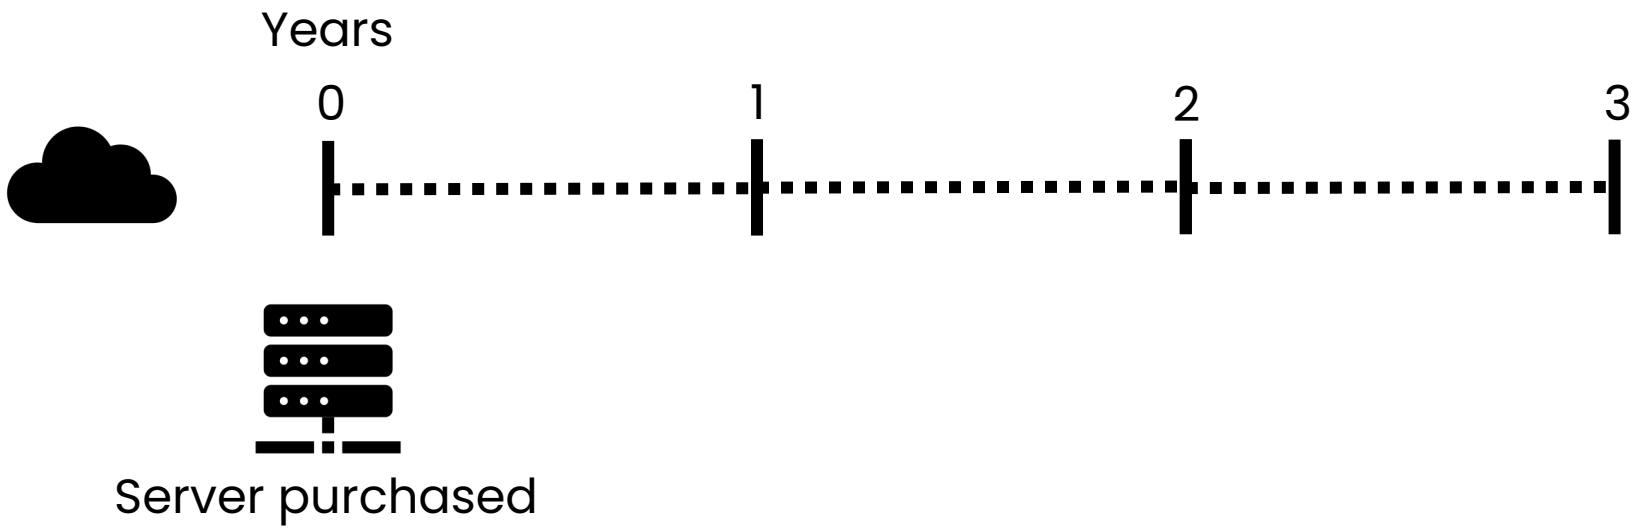

How to **lower embodied emissions**?

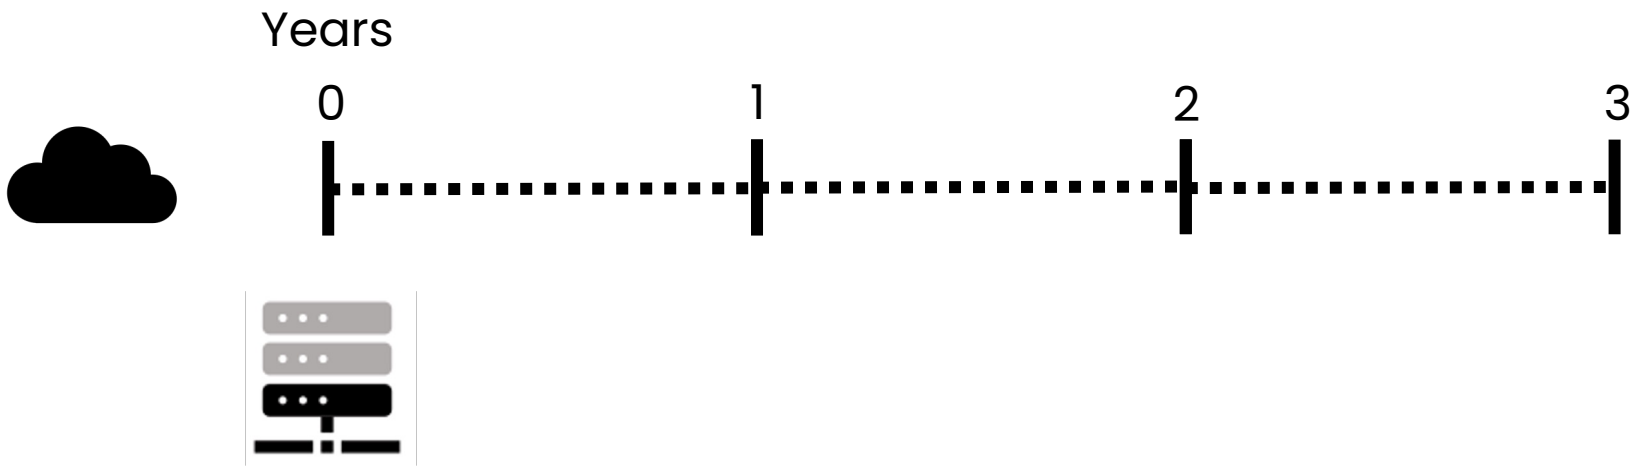

How to **lower embodied emissions**?

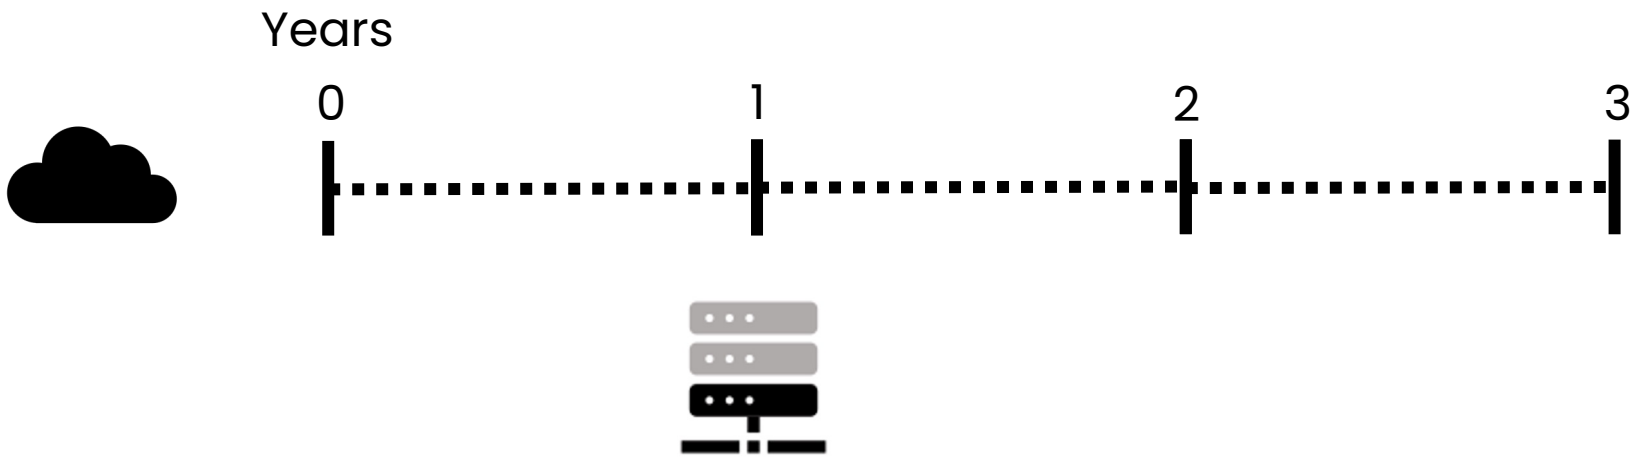

How to **lower embodied emissions**?

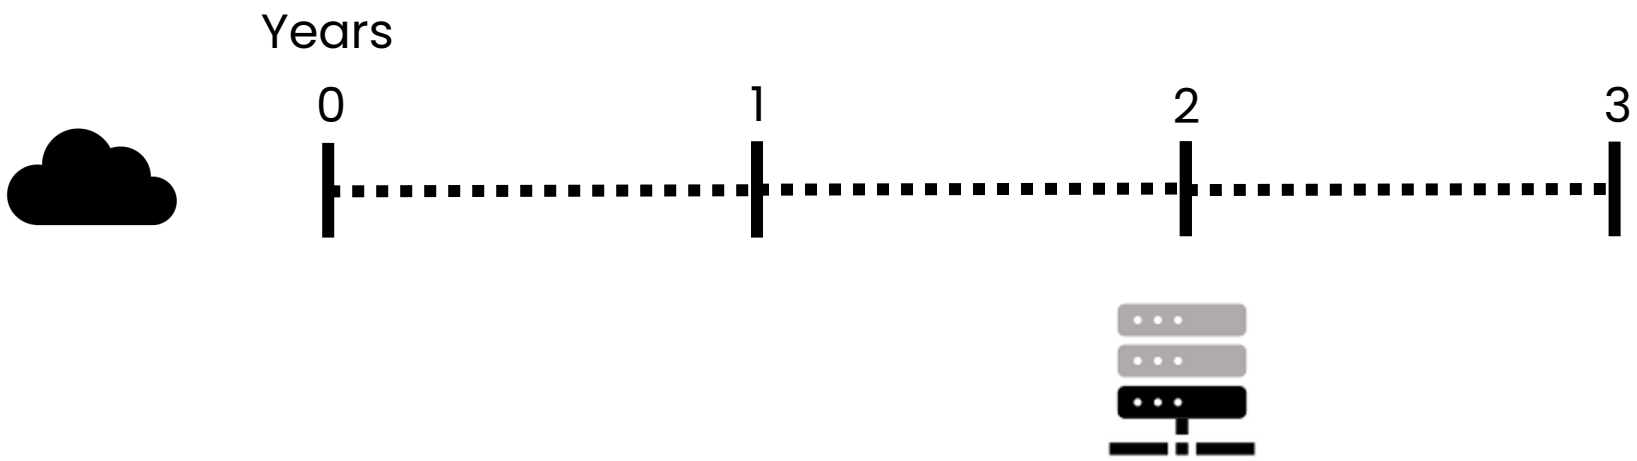

How to **lower embodied emissions**?

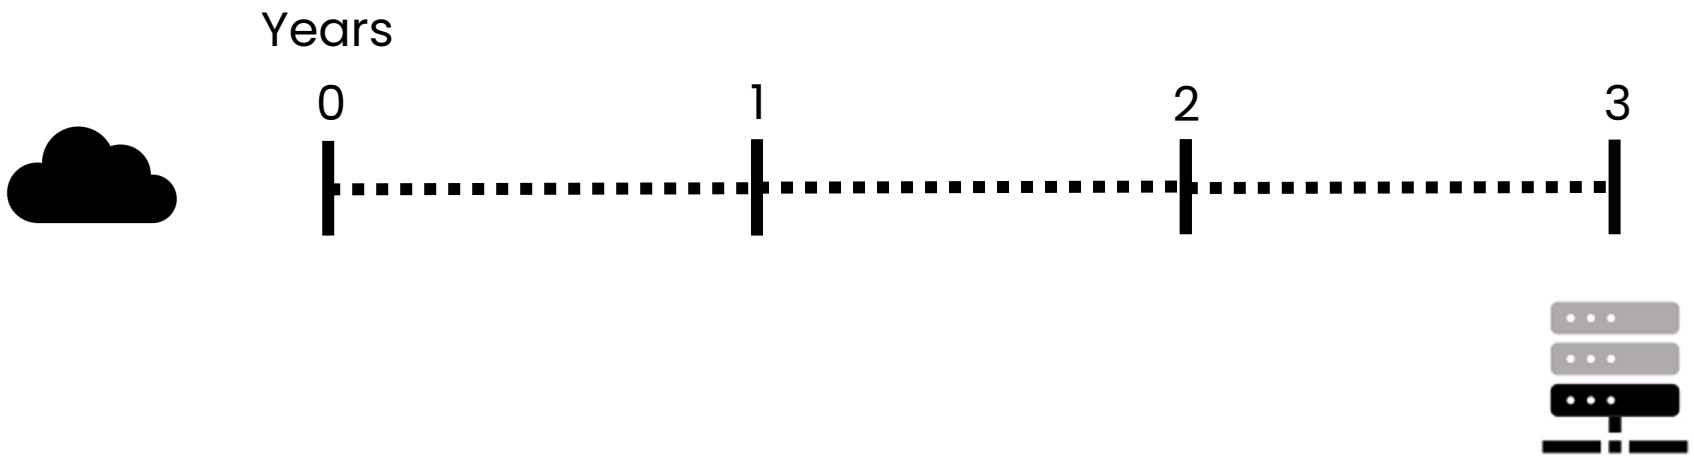

How to **lower embodied emissions**?

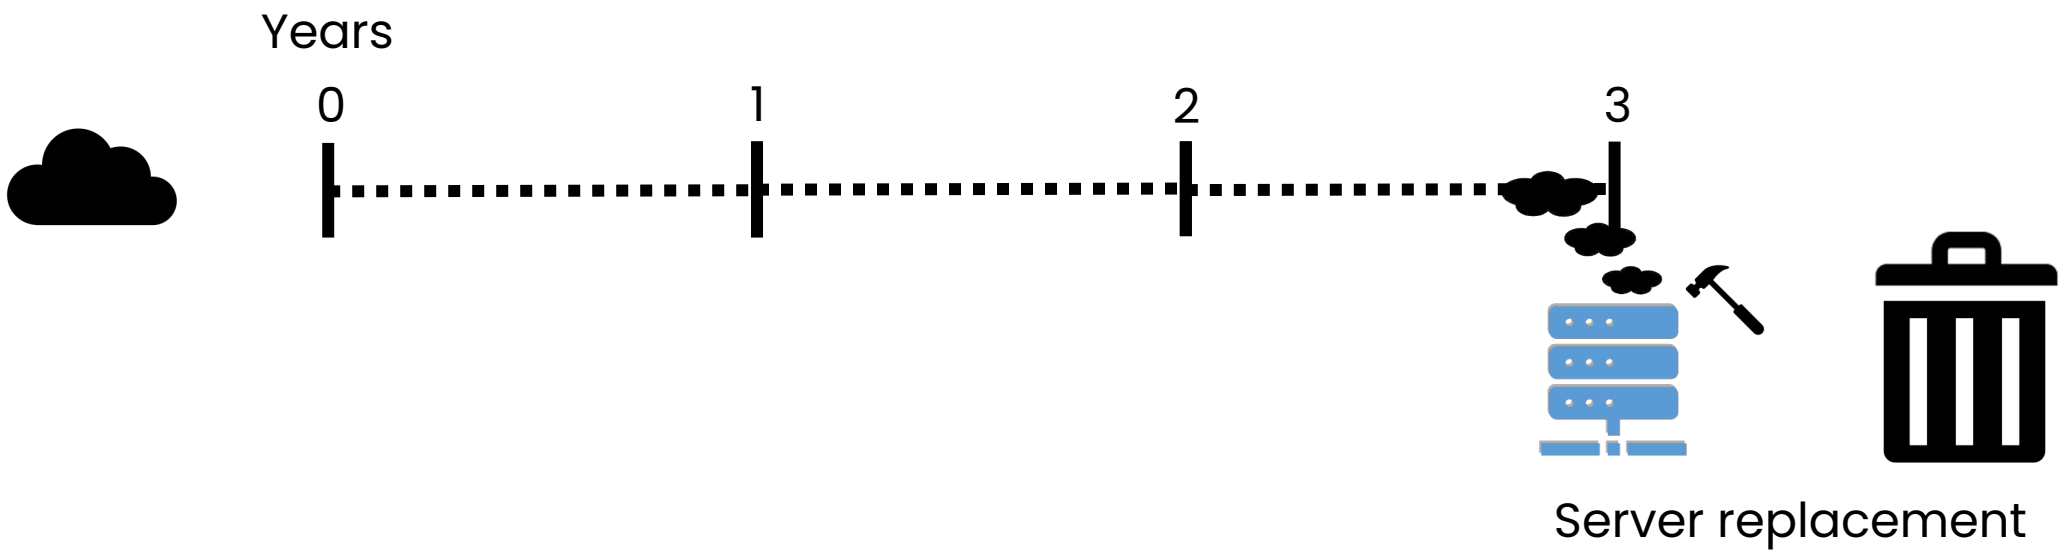

We must **extend server lifetimes** to lower embodied emissions

So, what drives shortened server lifetimes?

Intel Processor **Performance**

Our goal: Achieve acceptable performance on older HW to increase server lifetimes

Our contributions

1. A **characterization** of server generations running web services
2. A **scheduler** to performance-efficiently schedule on old HW
3. A step towards **measuring** carbon savings

This talk

Characterization

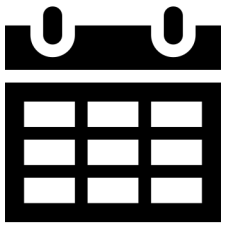

Scheduler

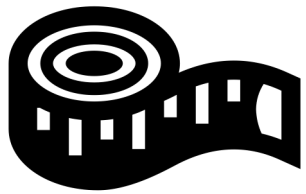

Carbon Measurement

This talk

Characterization

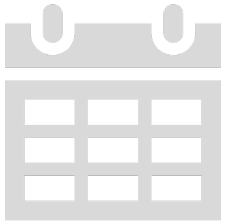

Scheduler

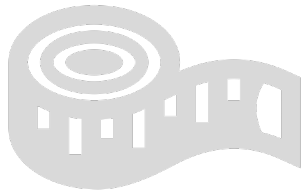

Carbon Measurement

Q1: Do **end-to-end** services perform well on old HW?

DeathStarBench µServices

[Gan'19]

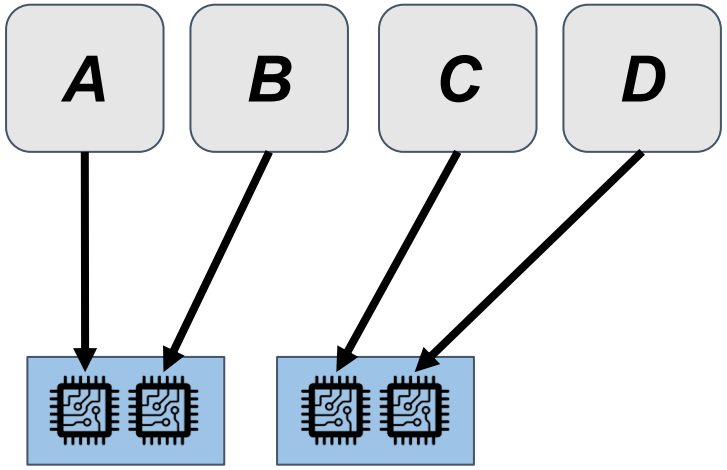

Servers

All running on new

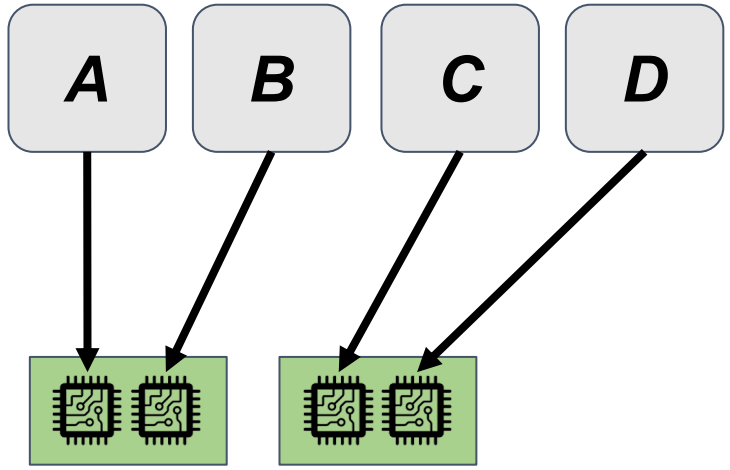

All running on old

Q1: Do **end-to-end** services perform well on old HW?

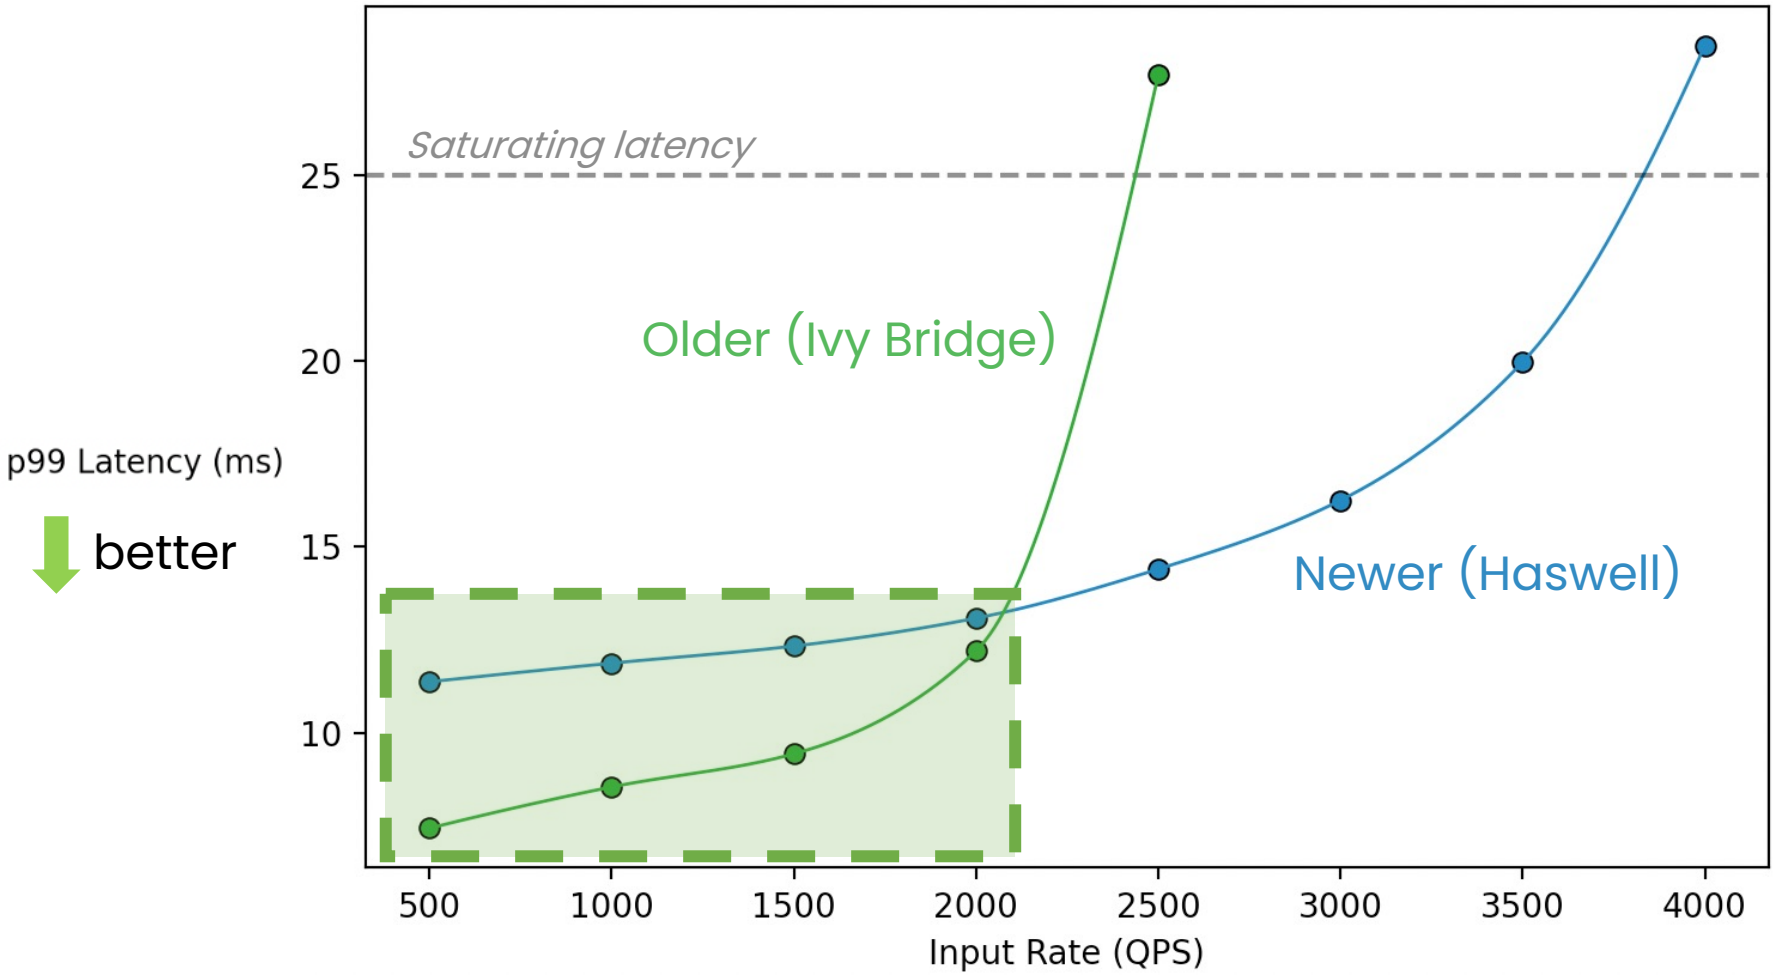

Older servers may be beneficial at **low load conditions**

Q2: Do **specific** microservices perform well on older HW?

DeathStarBench µServices
[Gan'19]

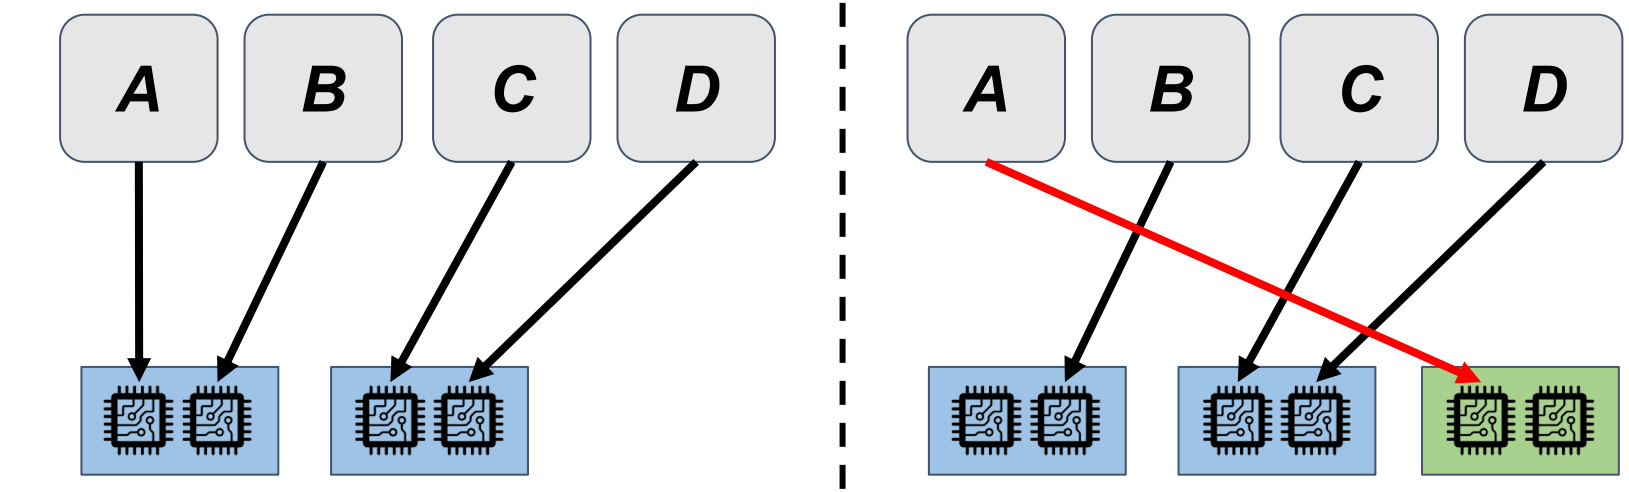

Servers

All running on new

A running on old

Q2: Do **specific** microservices perform well on older HW?

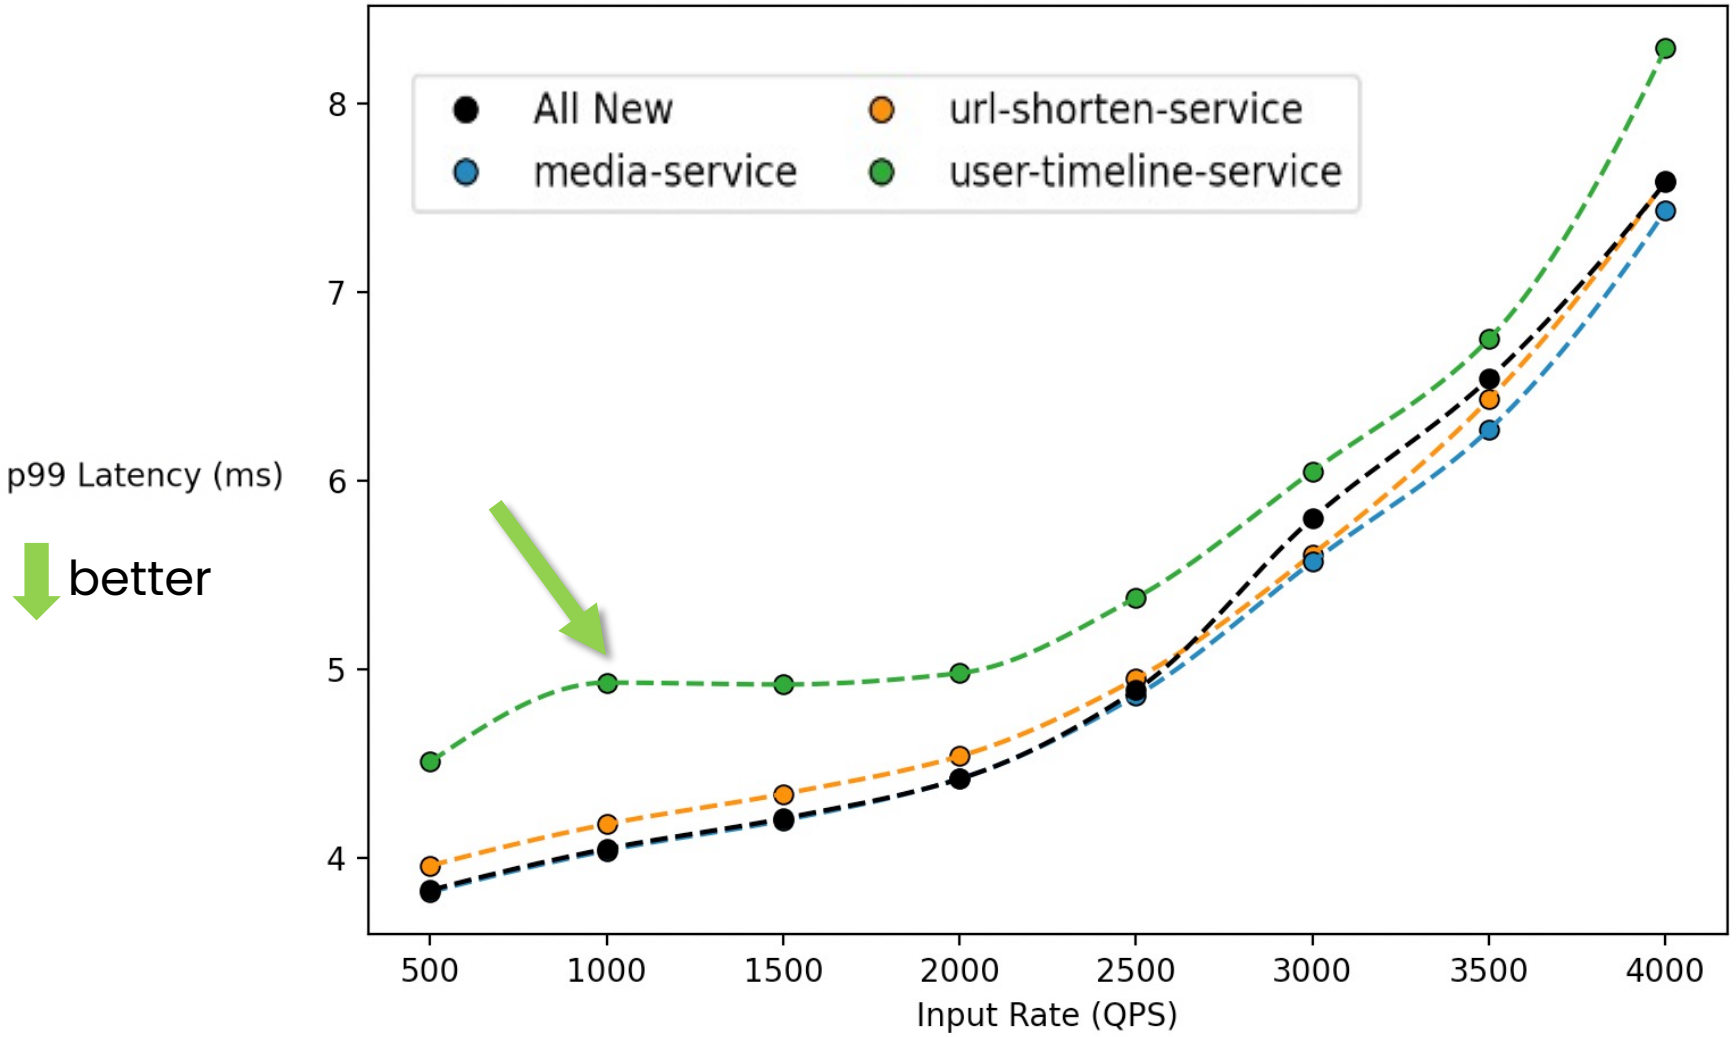

Placing all microservices on new HW is carbon-inefficient

This talk

Characterization

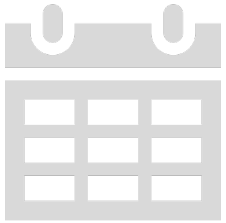

Scheduler

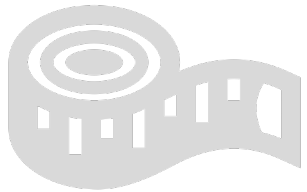

Carbon Measurement

This talk

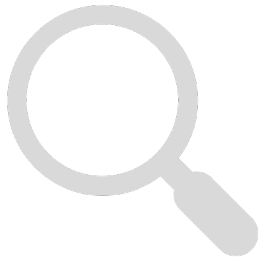

Characterization

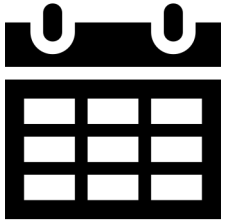

Scheduler

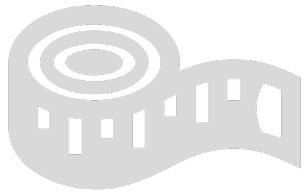

Carbon Measurement

How do we solve these **carbon inefficiencies**?

- Application profiling
- Load condition
- Heterogeneous hardware

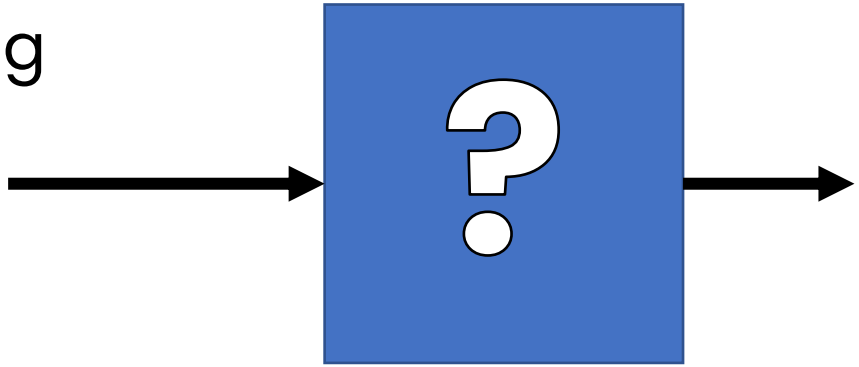

- Hardware/instance capacities
- (μ)Service placement
- Resulting carbon emissions

How do we solve these **carbon inefficiencies**?

- Application profiling
- Load condition
- Heterogeneous hardware

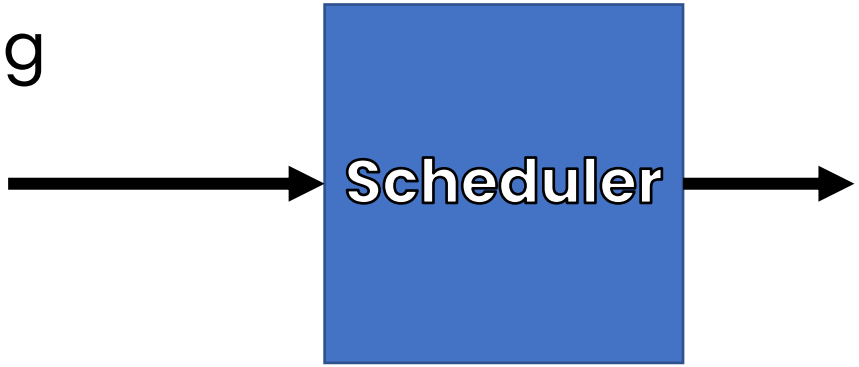

- Hardware/instance capacities
- (μ)Service placement
- Resulting carbon emissions

How do we solve these **carbon inefficiencies**?

- Application profiling
- Load condition
- Heterogeneous hardware

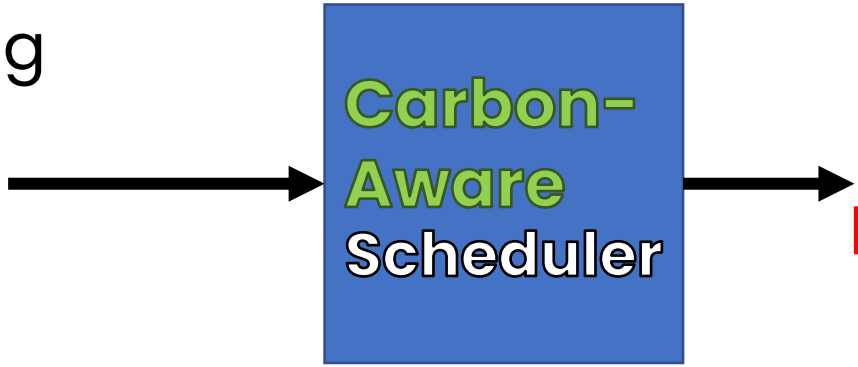

- Hardware/instance capacities
- (μ)Service placement
- Resulting carbon emissions

?

We need **intelligent carbon-aware scheduling**, driven by better metrics

This talk

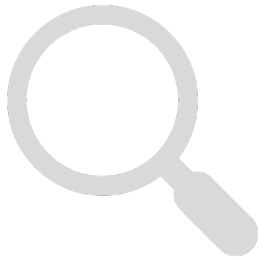

Characterization

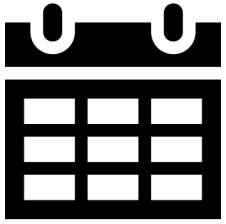

Scheduler

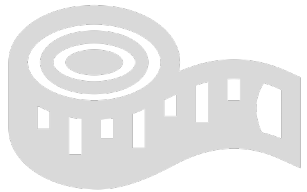

Carbon Measurement

This talk

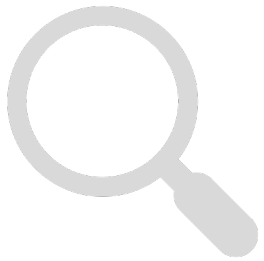

Characterization

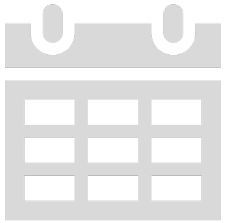

Scheduler

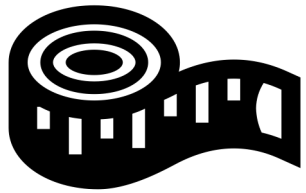

Carbon Measurement

How do we measure the scheduler's **carbon impact**?

and incentivizes reducing :)

^

A good metric measures the carbon impact of:

Embodied cost

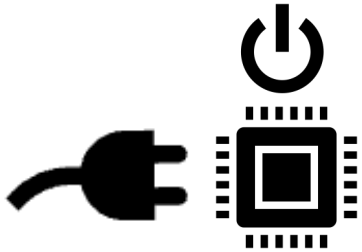

Operational cost

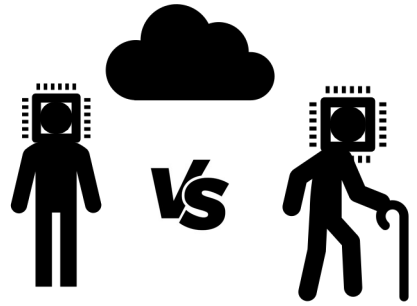

Using **older** vs newer

And...?

We need **better metrics** to measure emissions

Characterizing Datacenter Server Generations for Lifetime Extension and Carbon Reduction

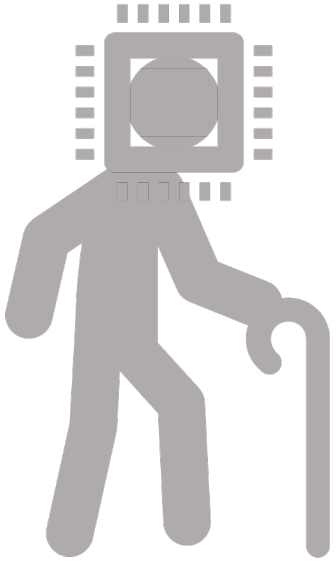

Jaylen Wang, Udit Gupta, Akshitha Sriraman

Carnegie Mellon University

Cornell University Founded A.D. 1865

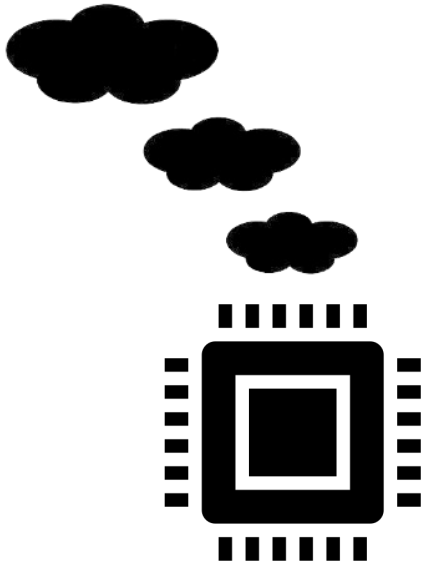

Backup Slides

That’s a great question – good thing I have backup slides :)

Other factors than performance…

Device	Proportion
HDD	81.84 %
Miscellaneous	10.20 %
Memory	3.06 %
Power	1.74 %
RAID card	1.23 %
Flash card	0.67 %
Motherboard	0.57 %
SSD	0.31 %
Fan	0.19 %
HDD backboard	0.14 %
CPU	0.04 %

(Of failures in datacenter over 4 years)

[Wang'17]

We have AMD too!

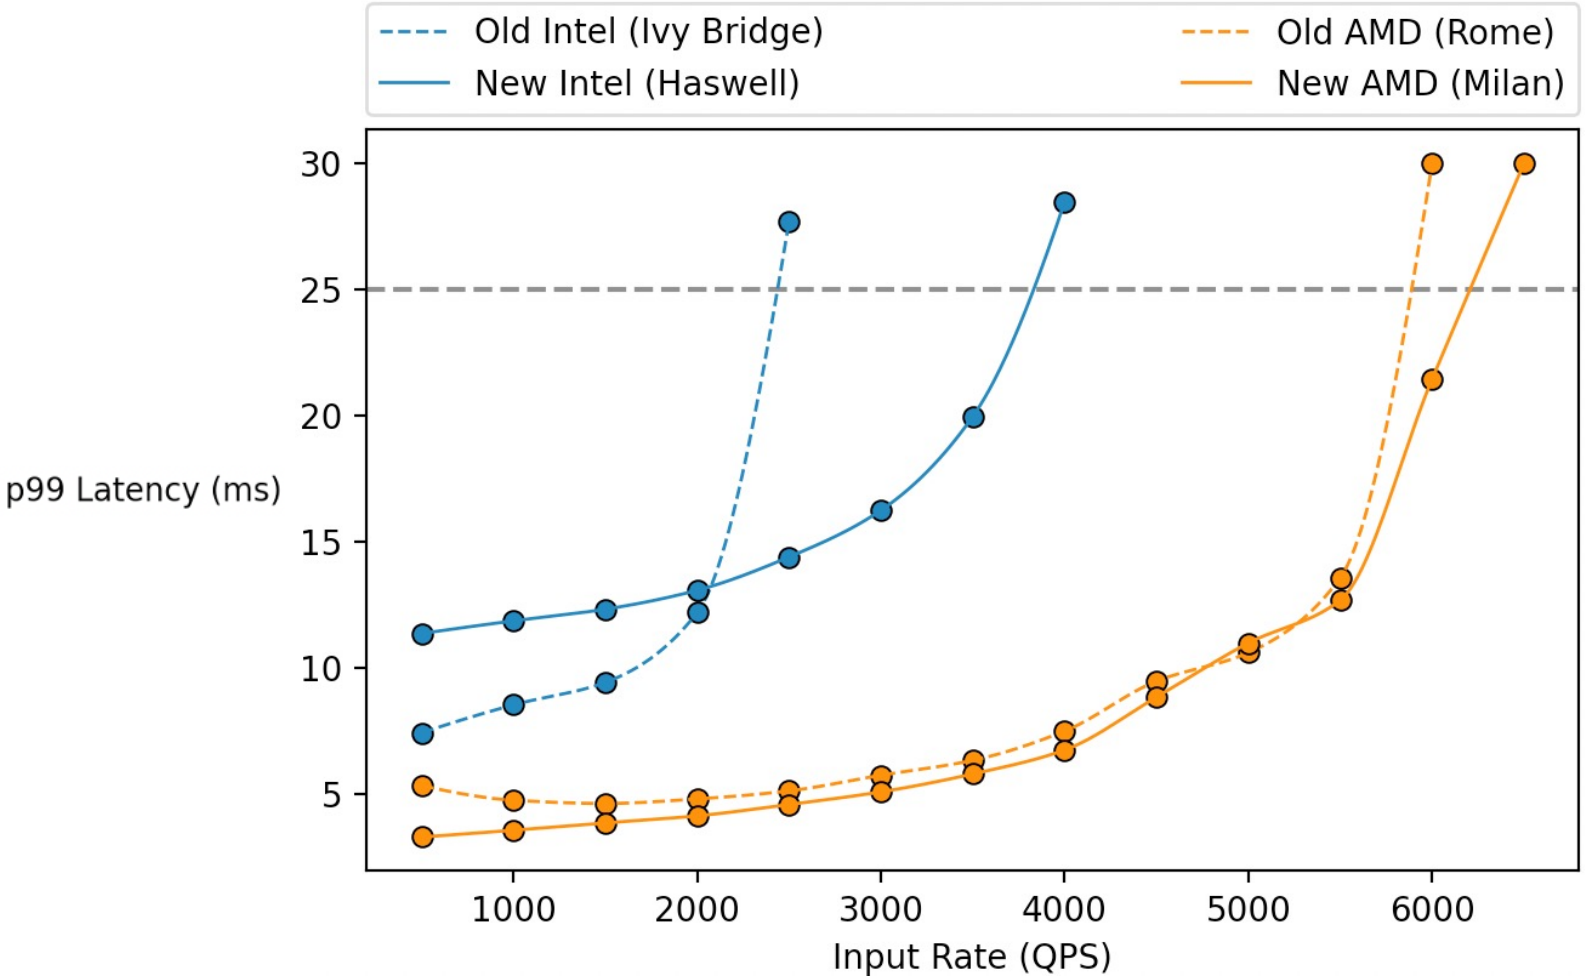

We have HotelReservation too!

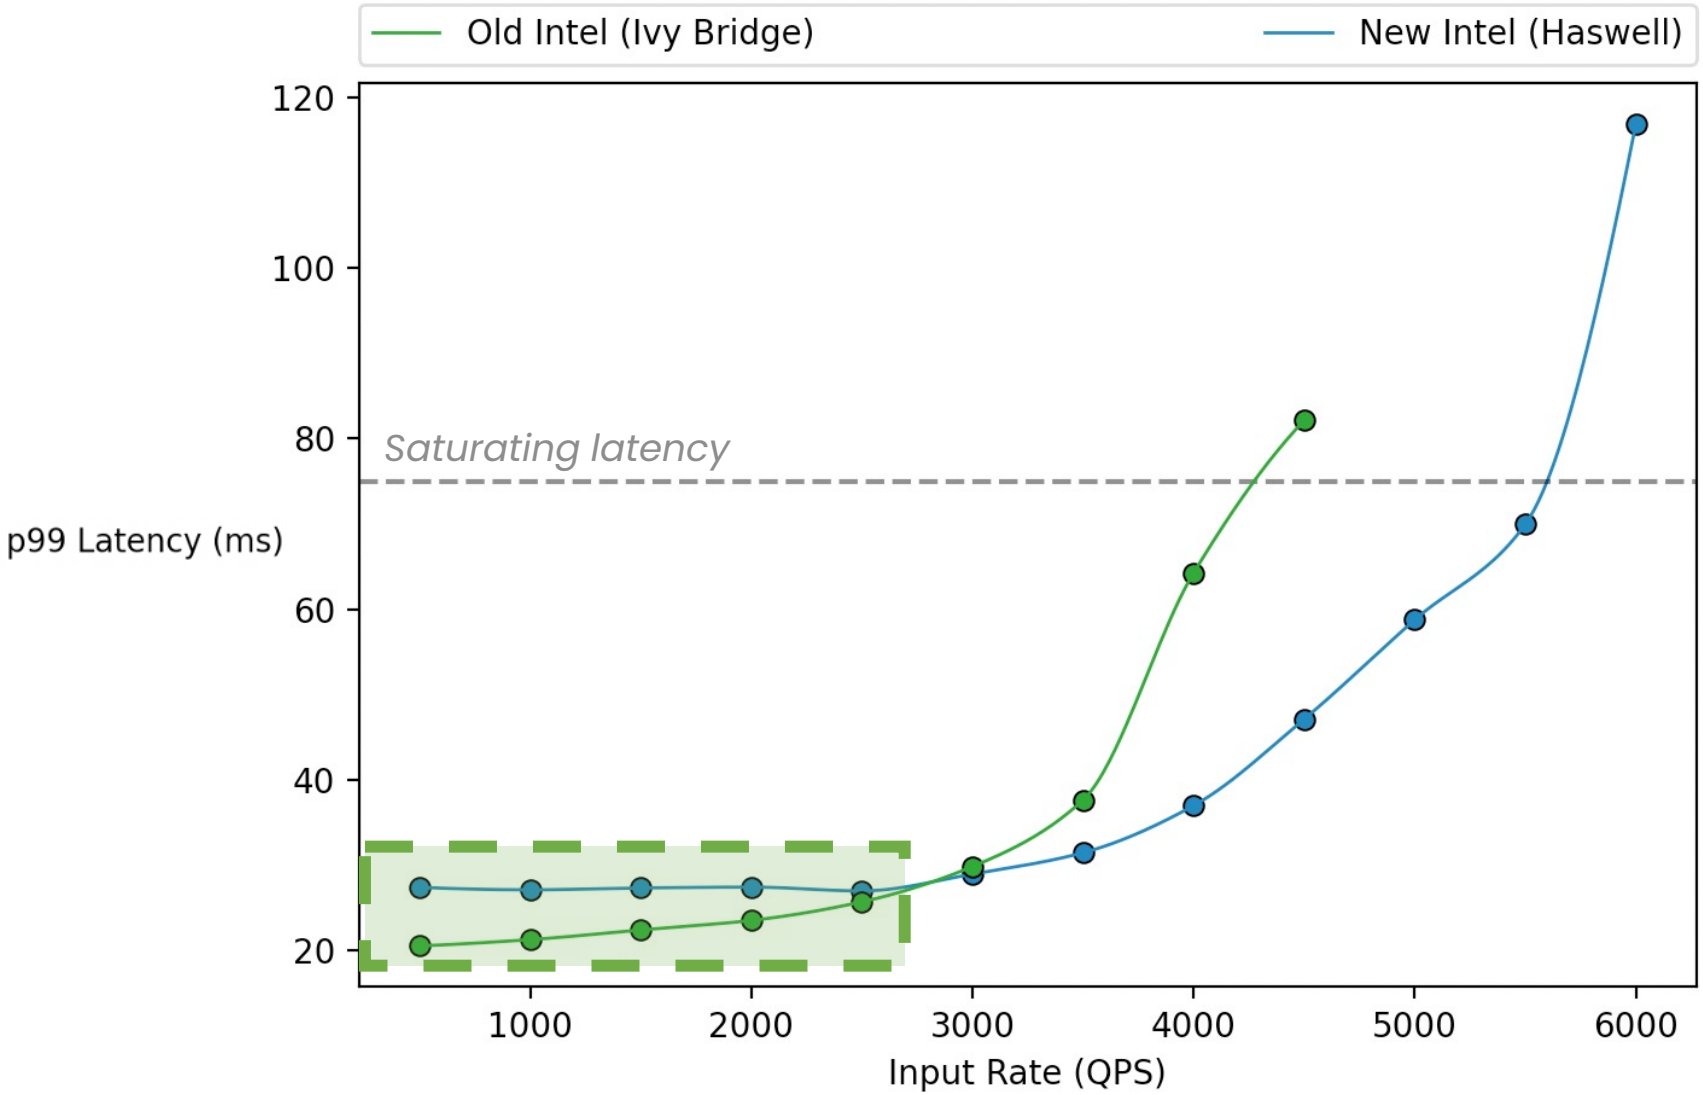

Here's ALL microservices (it's a lot)

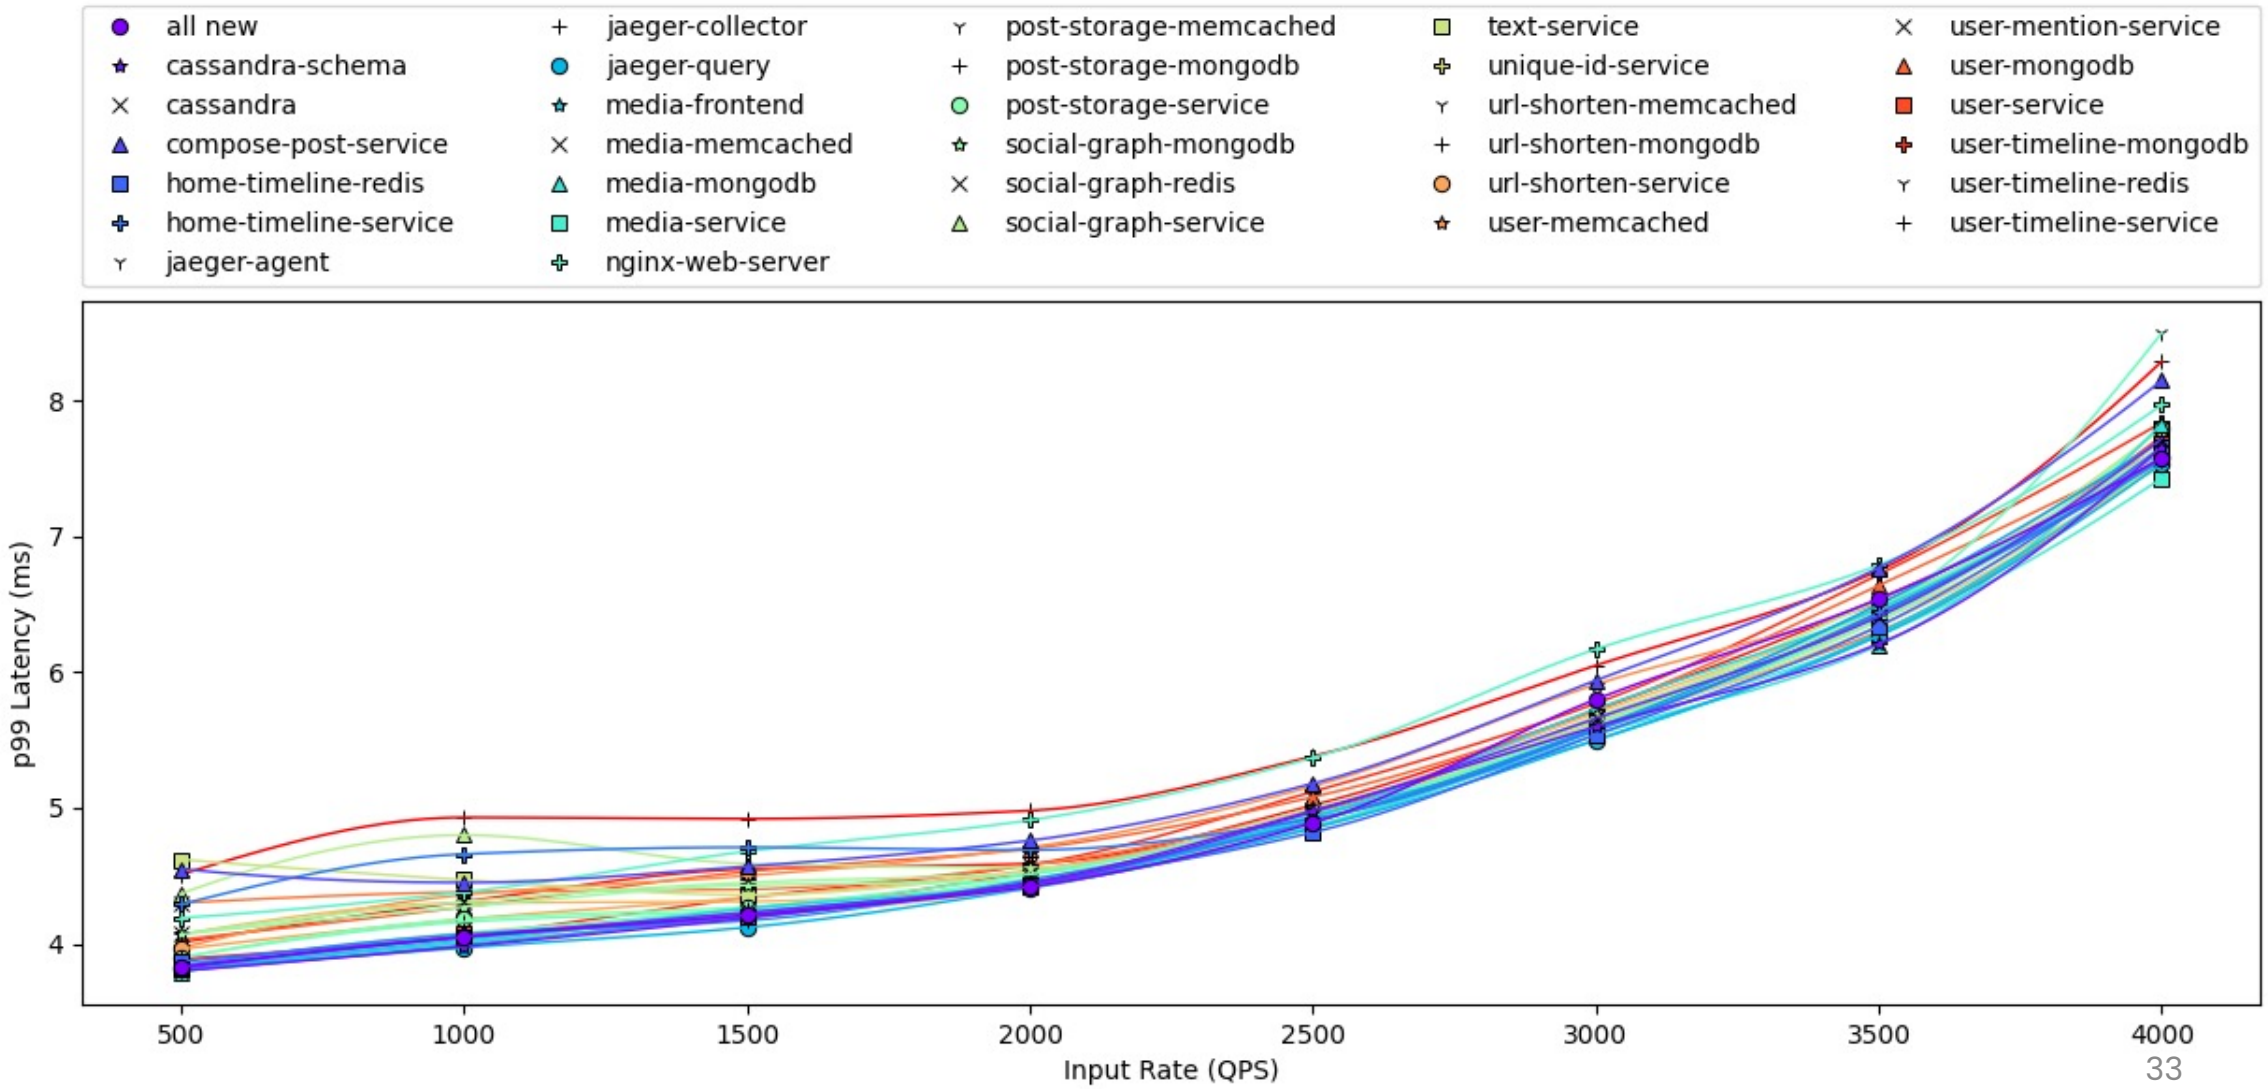

Server Setups

	Intel		AMD	
	Xeon E5-2660 v2	**Xeon E5-2660 v3**	**EPYC 7542**	**EPYC 7543**
Microarchitecture	Ivy Bridge (2012)	Haswell (2013)	Rome (2019)	Milan (2021)
Cores/Threads	10/20	10/20	32/64	32/64
Node	22 nm	22 nm	7 nm	7 nm
Base/Turbo (GHz)	2.2 / 3	2.6 / 3.3	2.9 / 3.4	2.8 / 3.7
LLC Cache Size	25 MB	25 MB	128 MB	256 MB
TDP (W)	95	105	225	225
RAM (DDR4)	256GB (1.6 GHz)	160GB (2.133 GHz)	256GB (3.2 GHz)	512GB (3.2 GHz)
Disk (SATA)	2 TB HDD	480 GB SSD	1.6 TB SSD	2 TB SSD
NIC	10Gb (PCIe v3)	10 Gb (PCIe v3)	25 Gb (PCIe v4.0)	25 Gb (PCIe v4.0)

TABLE I

CHARACTERISTICS OF TWO GENERATIONS (OLD ON THE LEFT, NEW ON THE RIGHT) OF INTEL AND AMD SERVERS USED IN EXPERIMENTS.

What should the scheduler do?

1. Profile the application across a distributed service
2. Identify carbon inefficiencies where older ≈ newer
3. Place services across nodes to favor older use
4. Measure carbon emissions with the metric
5. Iterate on decision until carbon is sufficiently low