Microsoft Research Faculty Summit 2014

15TH ANNUAL

Microsoft Research
Faculty Summit
2014 15TH ANNUAL

Capacity Provisioning Problems in Geo-distributed Data Centers

Bhuvan Urgaonkar
Dept. of CSE
The Pennsylvania State University

Microsoft

Geo-distributed Data Centers

- Reasons for geo-distribution:
  - Latency
  - Availability
- What are the **cost** implications?

What's New?

- What is well-understood:
  - How to build **single** data centers cost-effectively?
  - How to create distributed applications using an **existing** pool of data centers (that were built separately)?

How do costs change when we build a geo-distributed version of a centralized DC?

- Approach: specific case studies -> general insights & challenges

A Simple Thought Experiment

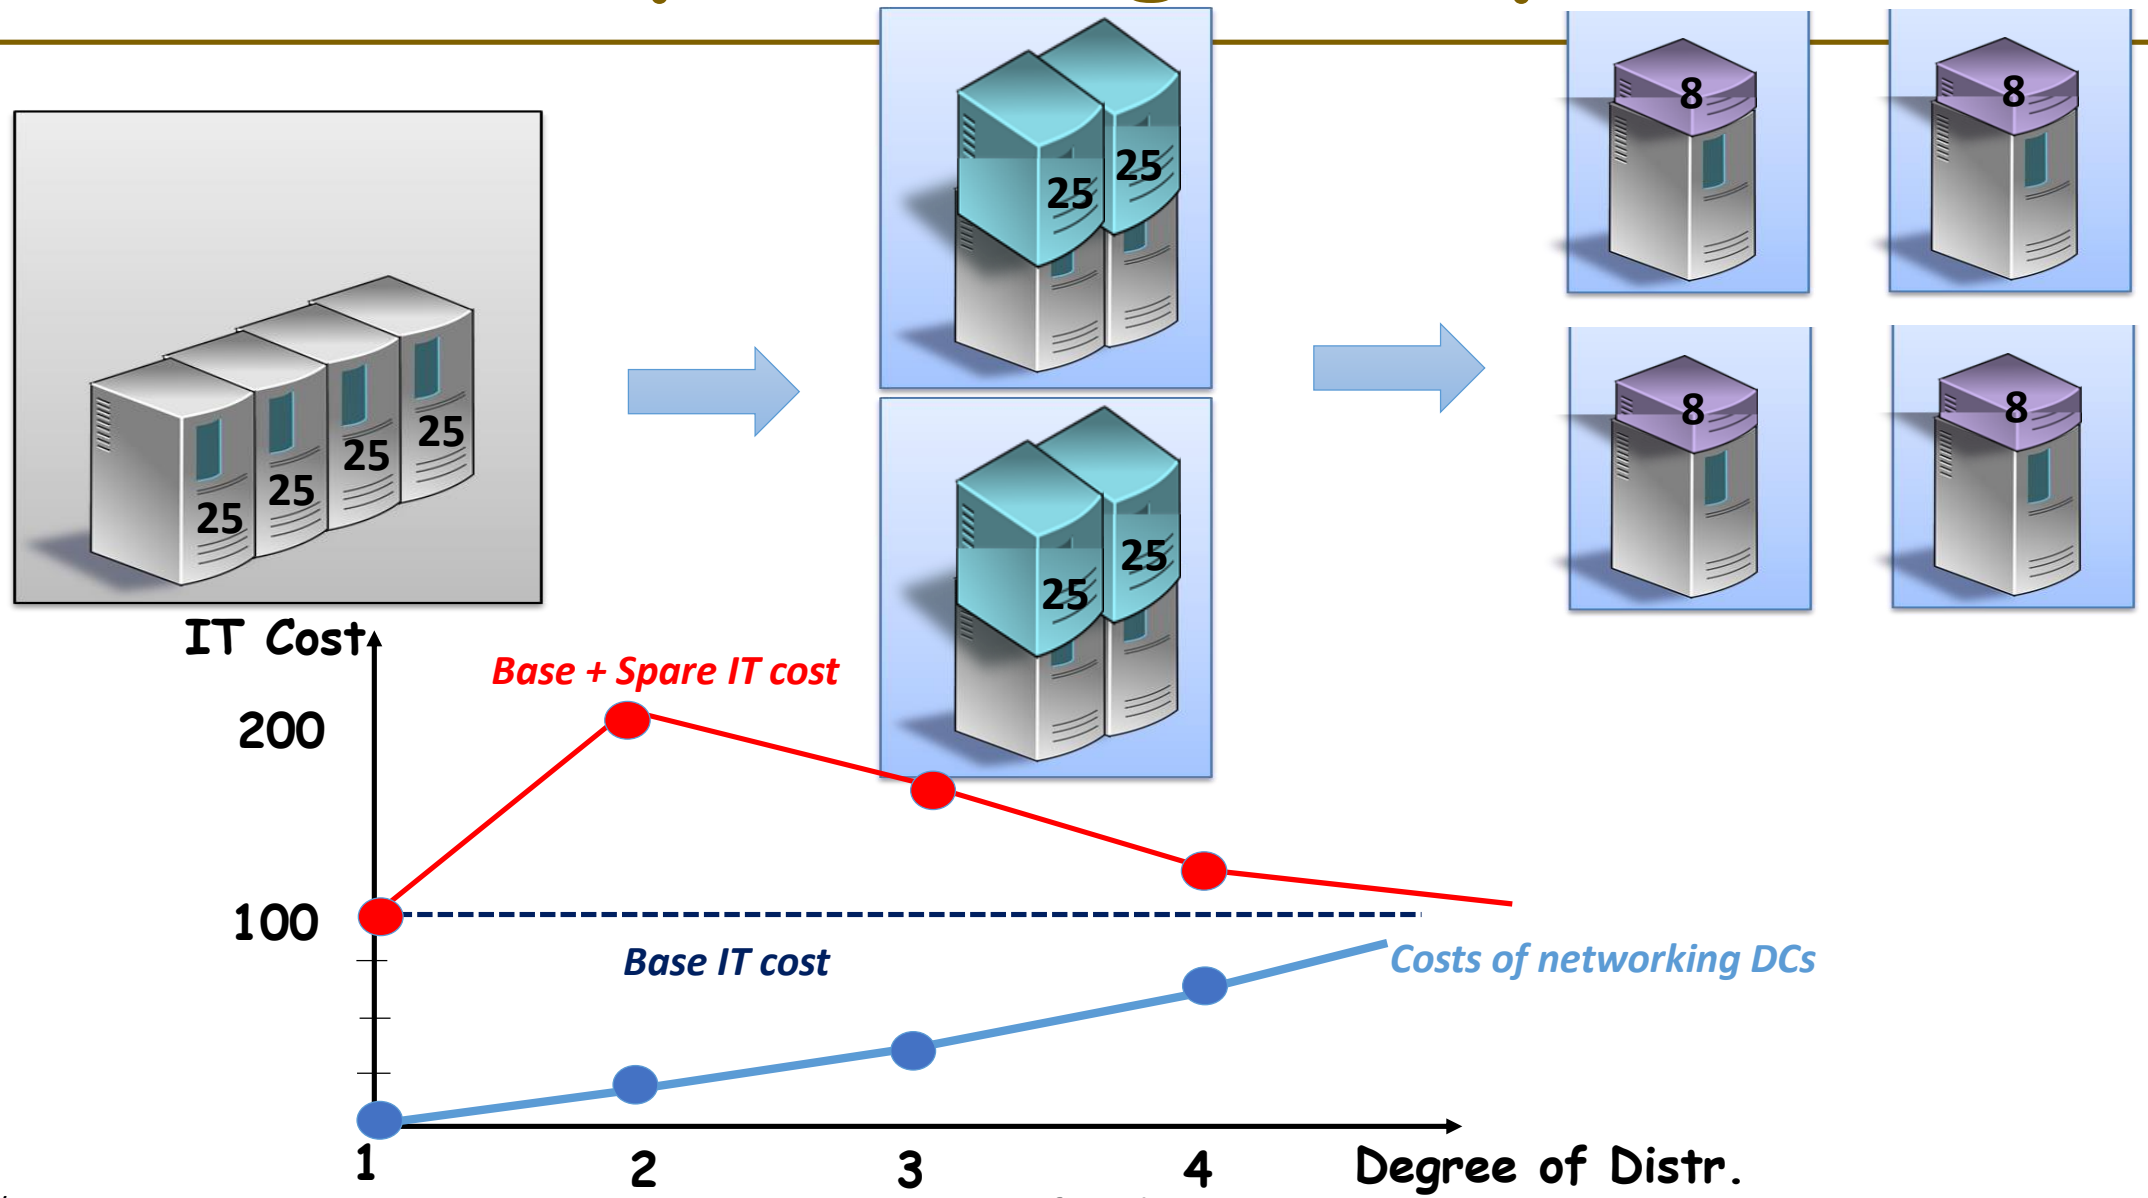

Costs: What have we made worse?

- Networking infrastructure to connect DCs
- Larger overall IT capacity
  - Redundancy for availability
    - Higher for heterogeneous collection of DCs
  - **Poorer statistical multiplexing**

Costs: What have we made worse?

- Networking infrastructure to connect DCs
- Larger overall IT capacity
  - Redundancy for availability
    - Higher for heterogeneous collection of DCs
  - **Poorer statistical multiplexing**

How do we keep this "small"?

Costs: What have we made worse?

- Networking infrastructure to connect DCs
- Larger overall IT capacity
  - Redundancy for availability
    - Higher for heterogeneous collection of DCs
  - Poorer statistical multiplexing
- Non-IT infrastructure (power+cooling) costs
  - To support higher IT capacity

Costs: What have we made worse?

- Networking infrastructure to connect DCs
- Larger overall IT capacity
  - Redundancy for availability
    - Higher for heterogeneous collection of DCs
  - Poorer statistical multiplexing
- Non-IT infrastructure (power+cooling) costs
  - To support higher IT capacity

Can we keep non-IT Infra. "size" small?

Costs: What has improved?

- Revenue due to better latency improvements

Costs: What has improved?

- Revenue due to better latency improvements
- Aspects of availability

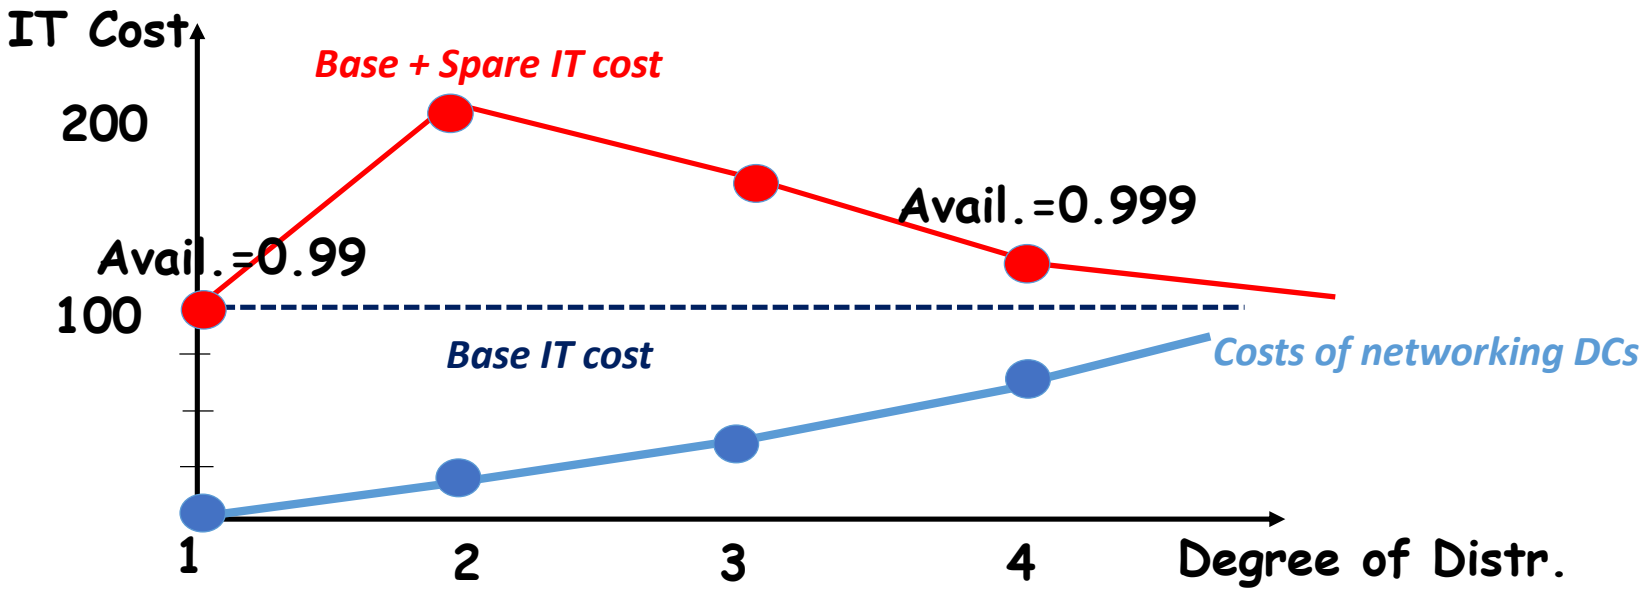

Costs: What has improved?

- Revenue due to better latency improvements
- Aspects of availability

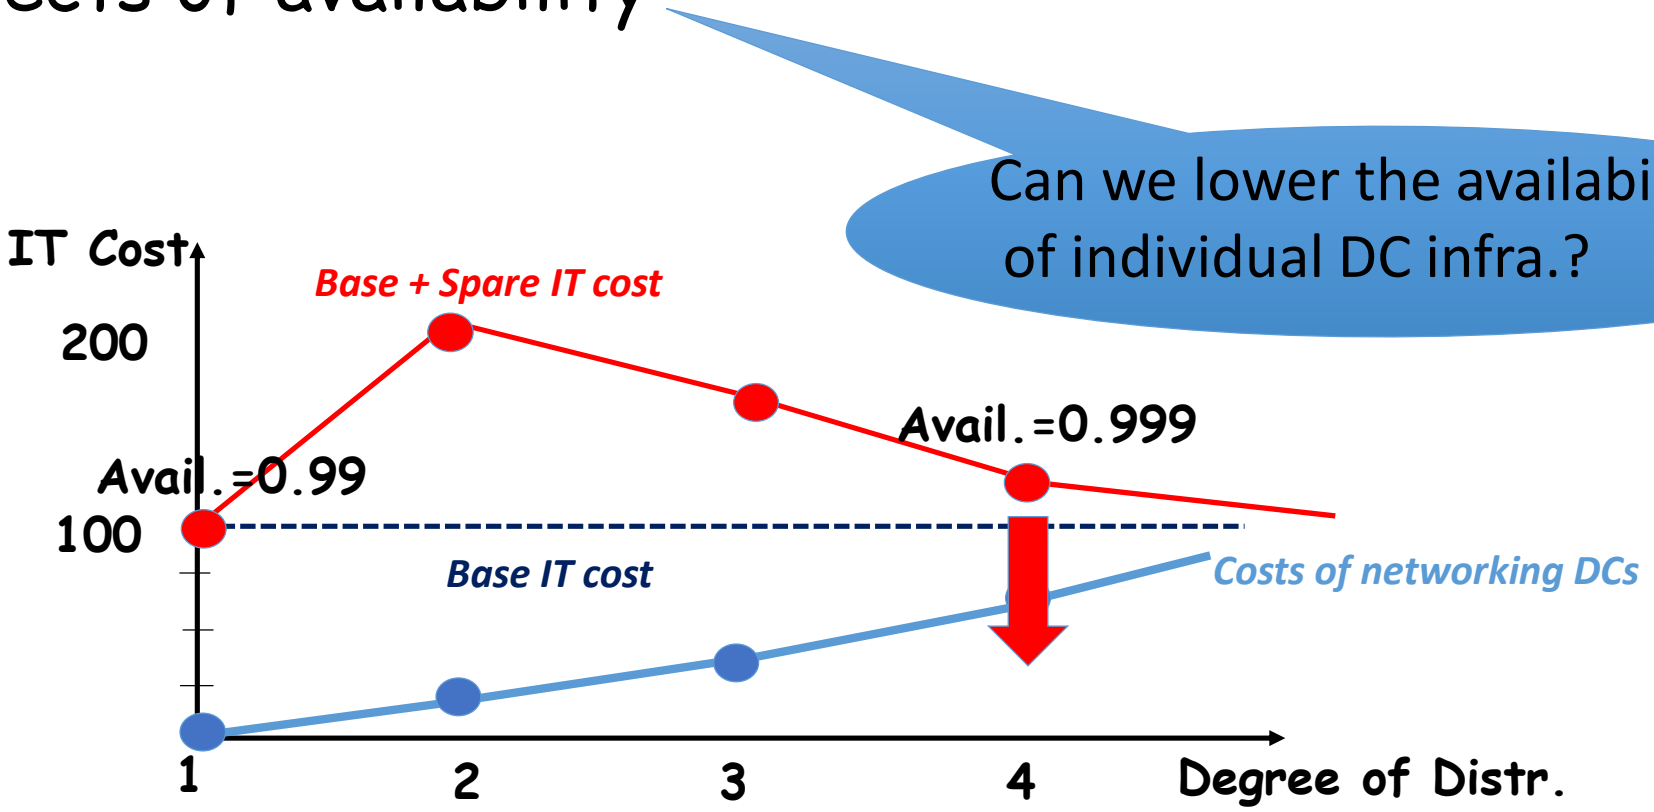

Outline

- An example of cost-effective IT provisioning
- Keeping non-IT infrastructure costs low
  - Lowering peak power related costs using batteries
- Conclusions

Problem Setting

- DC locations given
- Client demands known, time-varying
- Goal: determine total capacity at each DC
  - To meet latency constraints, and
  - To allow for one DC to fail
- Our optimizer: An LP
  - Generally, NP-hard facility location problems

Results

- DC locations
  - 6 MS data centers in the US
- Client demand model
  - Exhibits time zone specific variation
  - Proportional to population

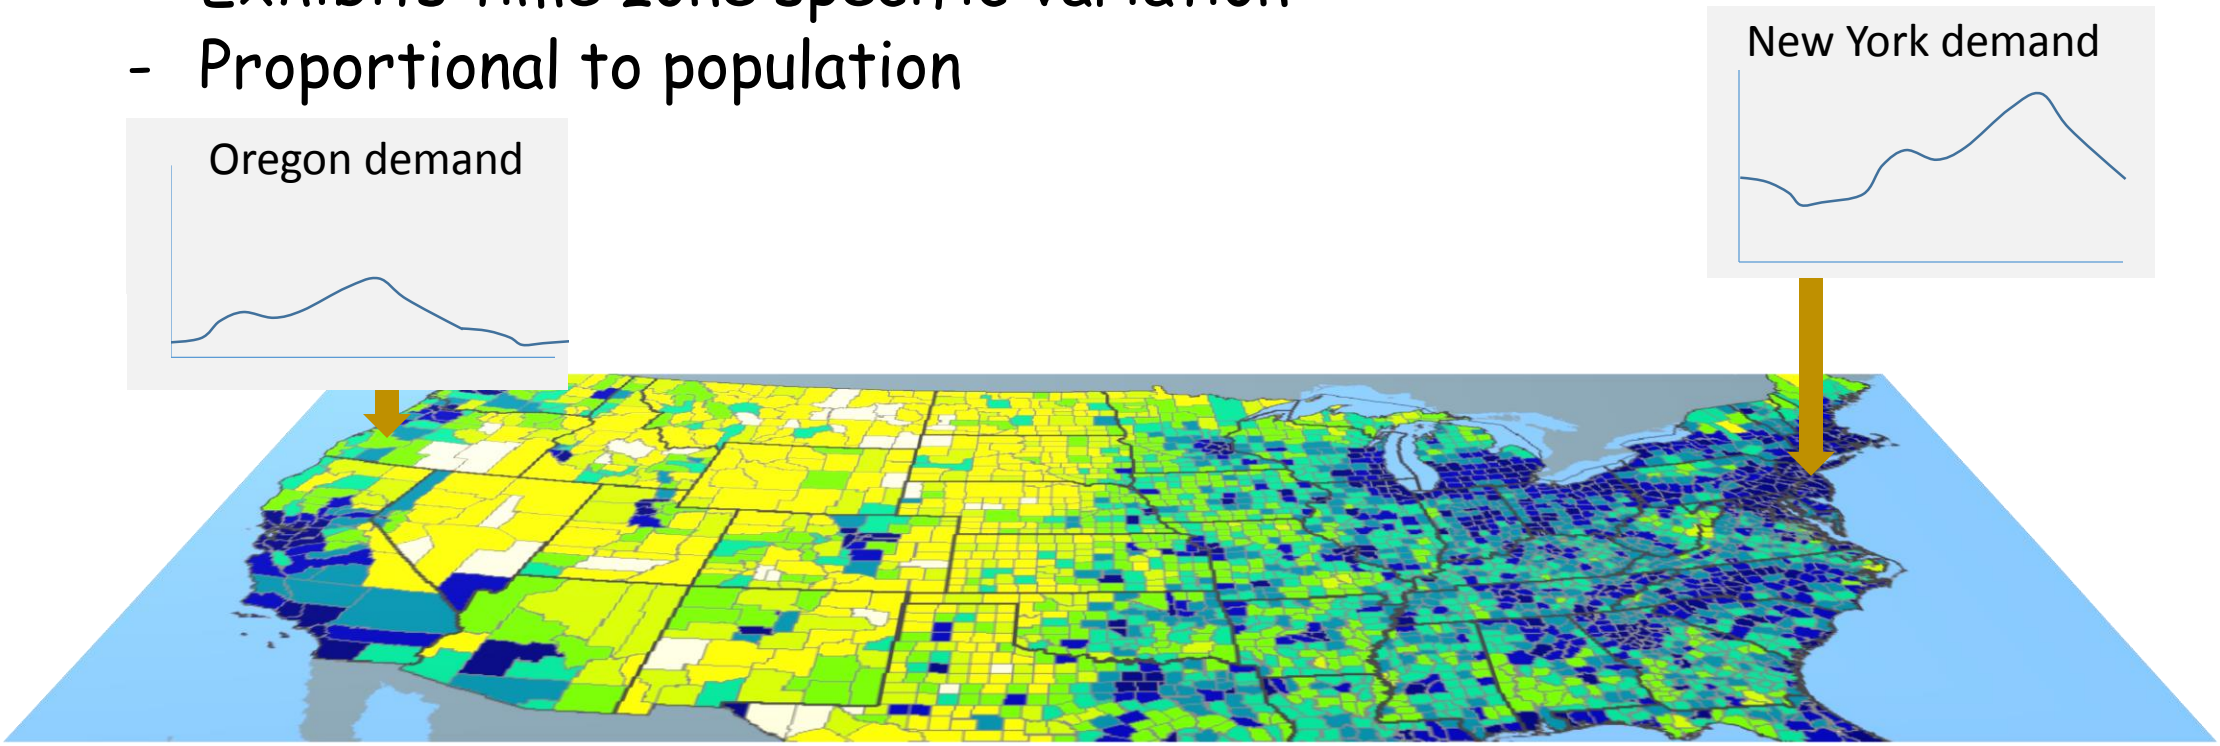

Results

Experiments using demand measured for one Microsoft cluster, and 6 MS DC locations within US. L'= L

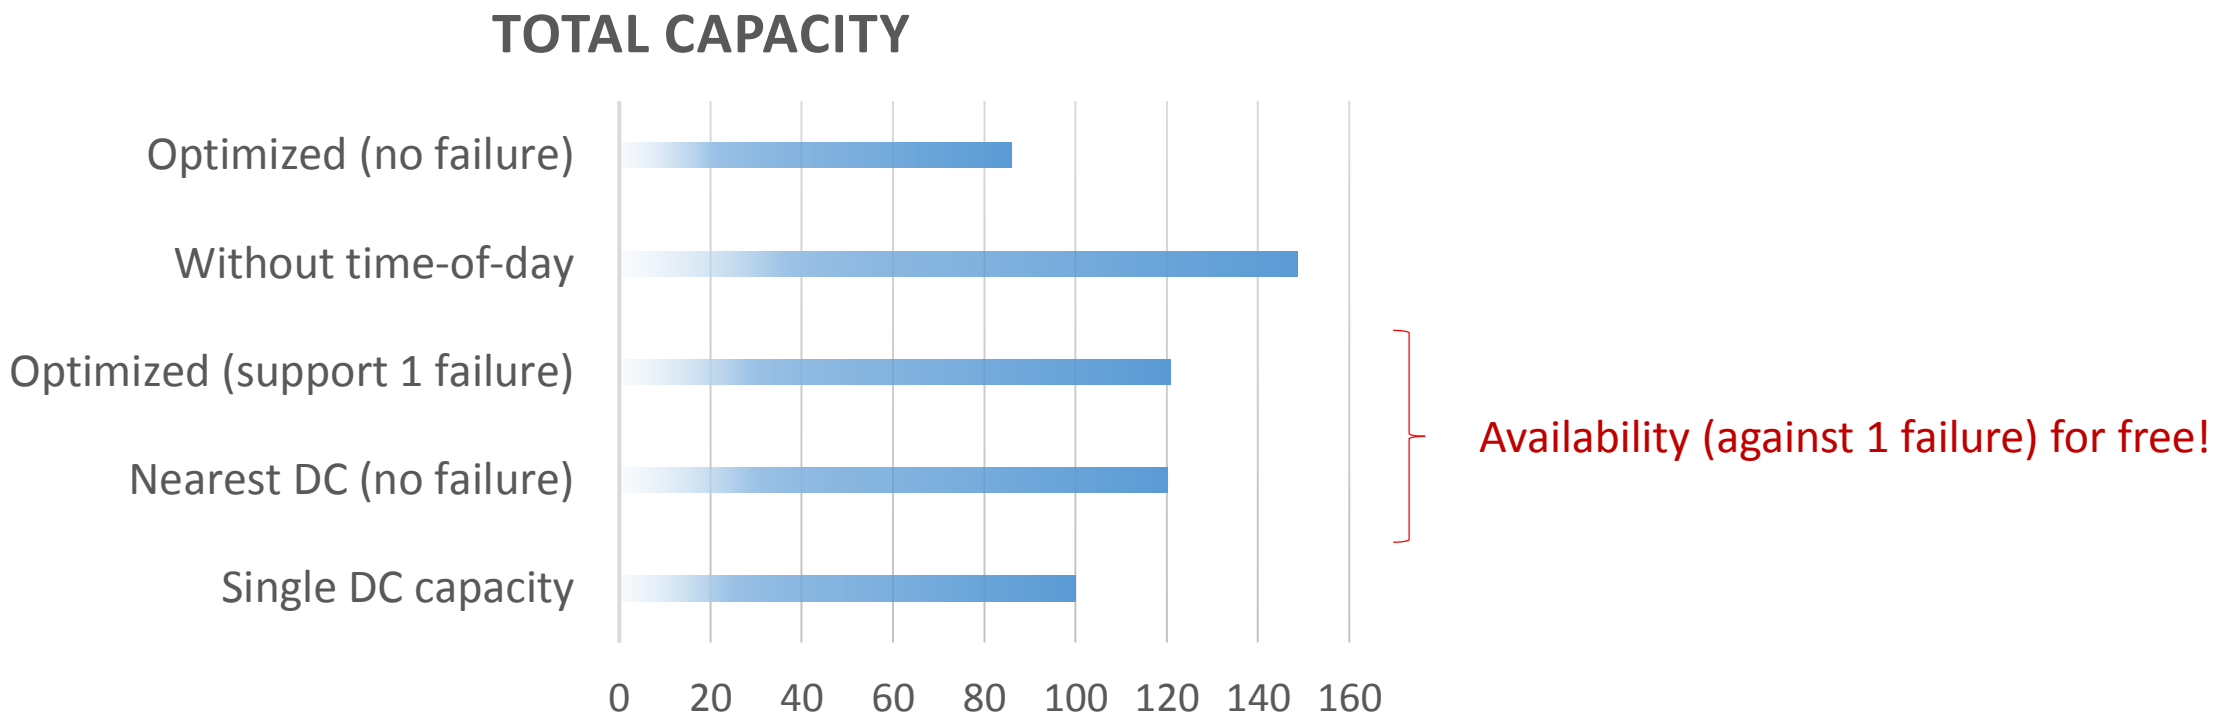

Results

Experiments using demand measured for one Microsoft cluster, and 6 MS DC locations within US. L=L'

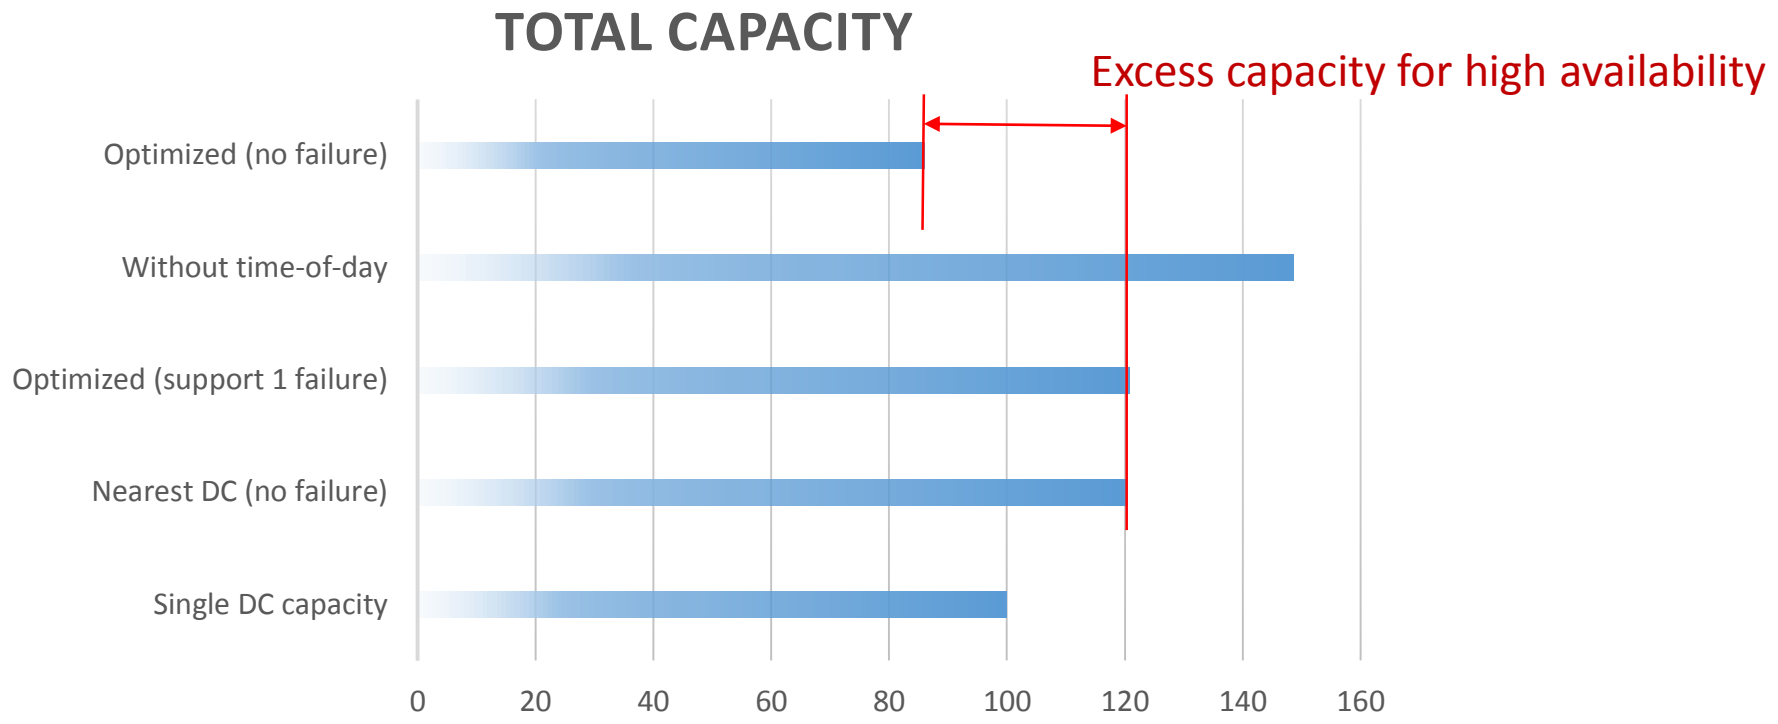

Details: Narayanan et al., "Towards leaner geo-distributed cloud infrastructure," Proc. HotCloud 2014

Outline

- ✓ An example of cost-effective IT provisioning
- ➢ Keeping non-IT infrastructure costs low
  - Lowering peak power related costs using batteries
- Conclusions

A Closer Look at Power Infrastructure

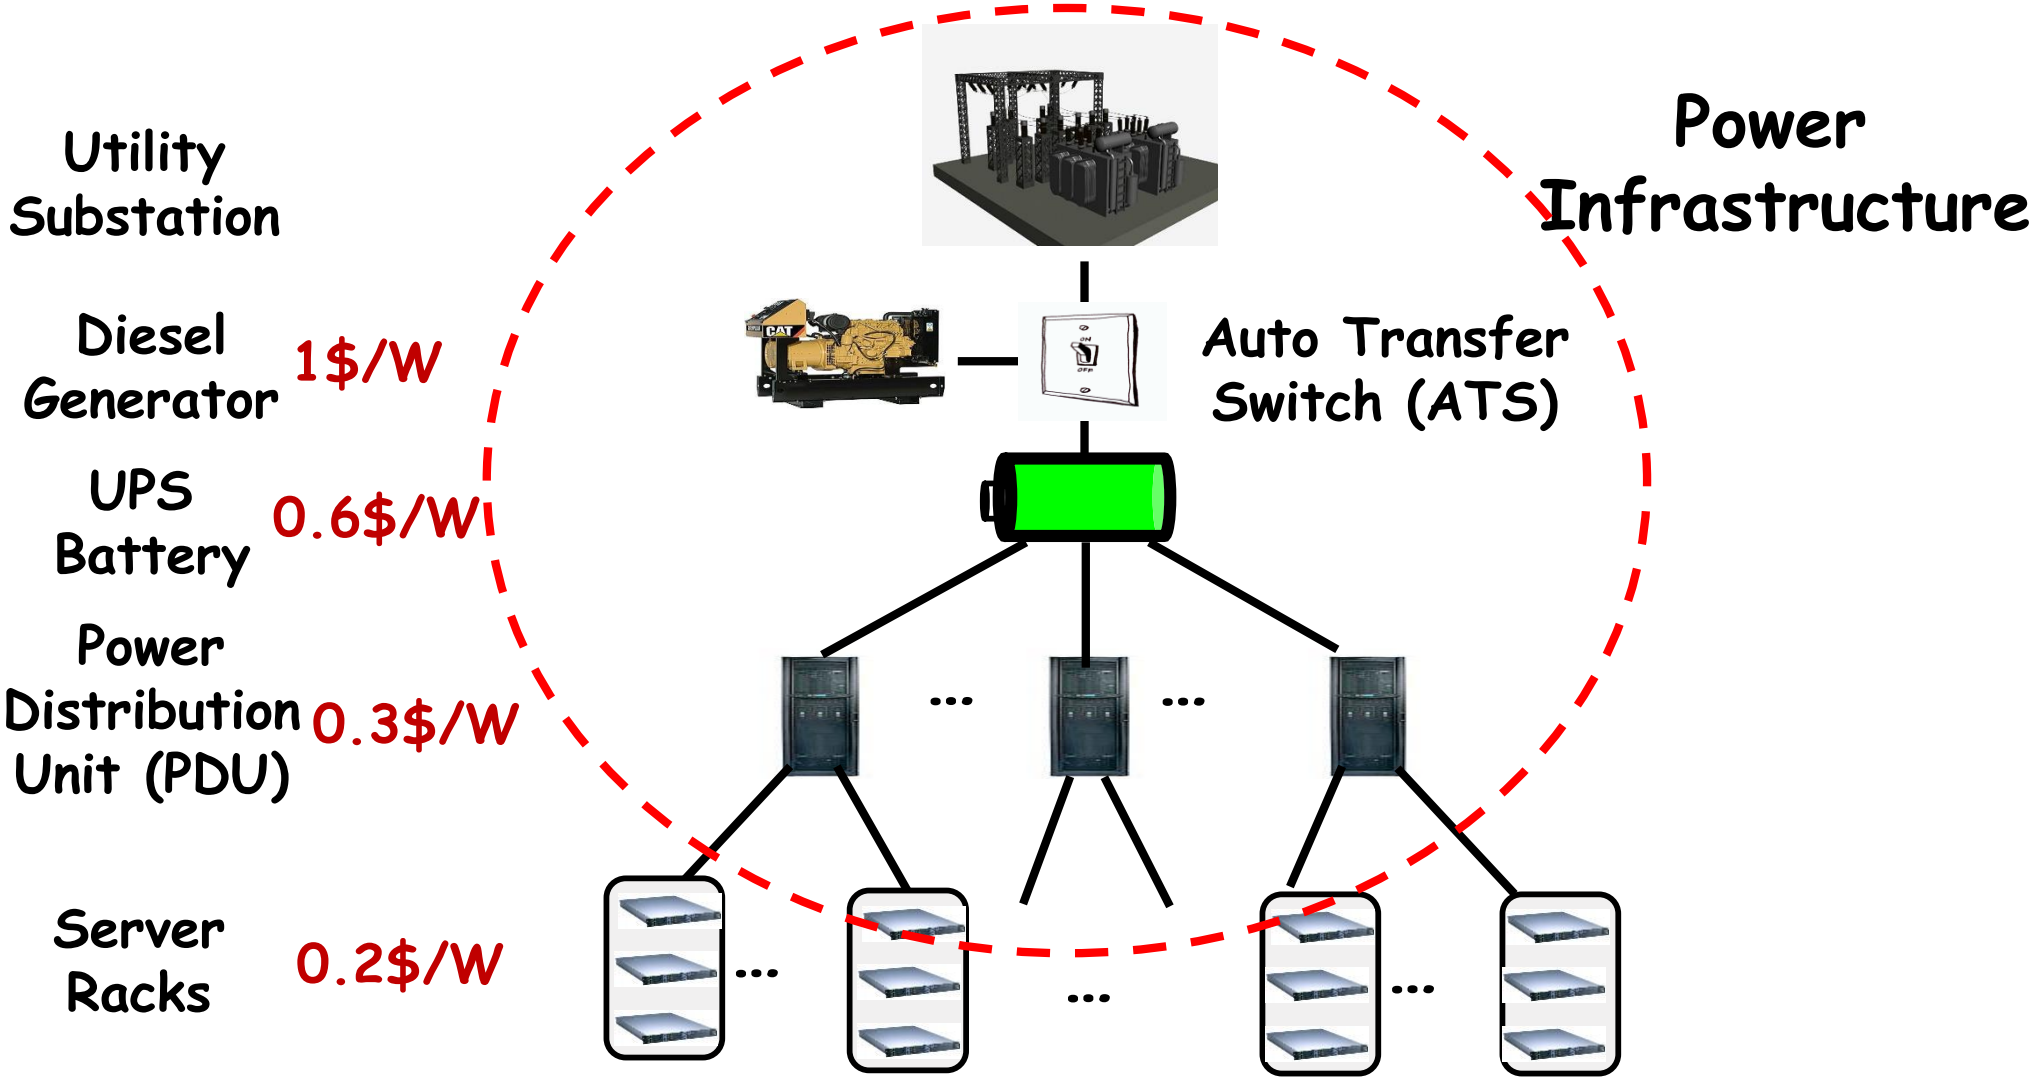

Lowering Peak Draw

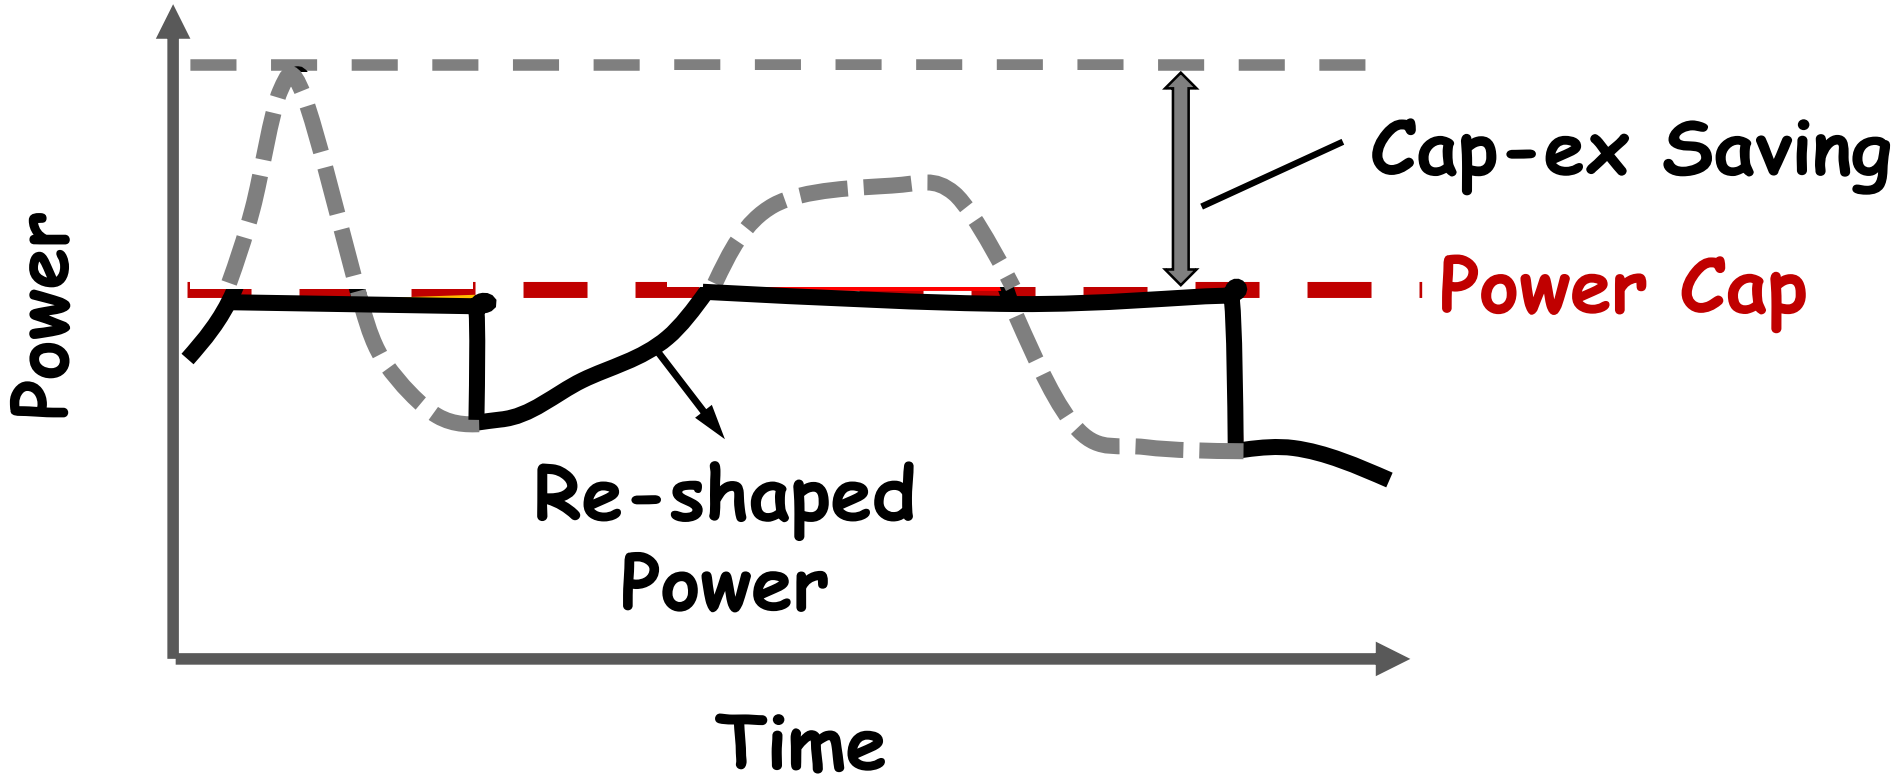

Using Energy Storage

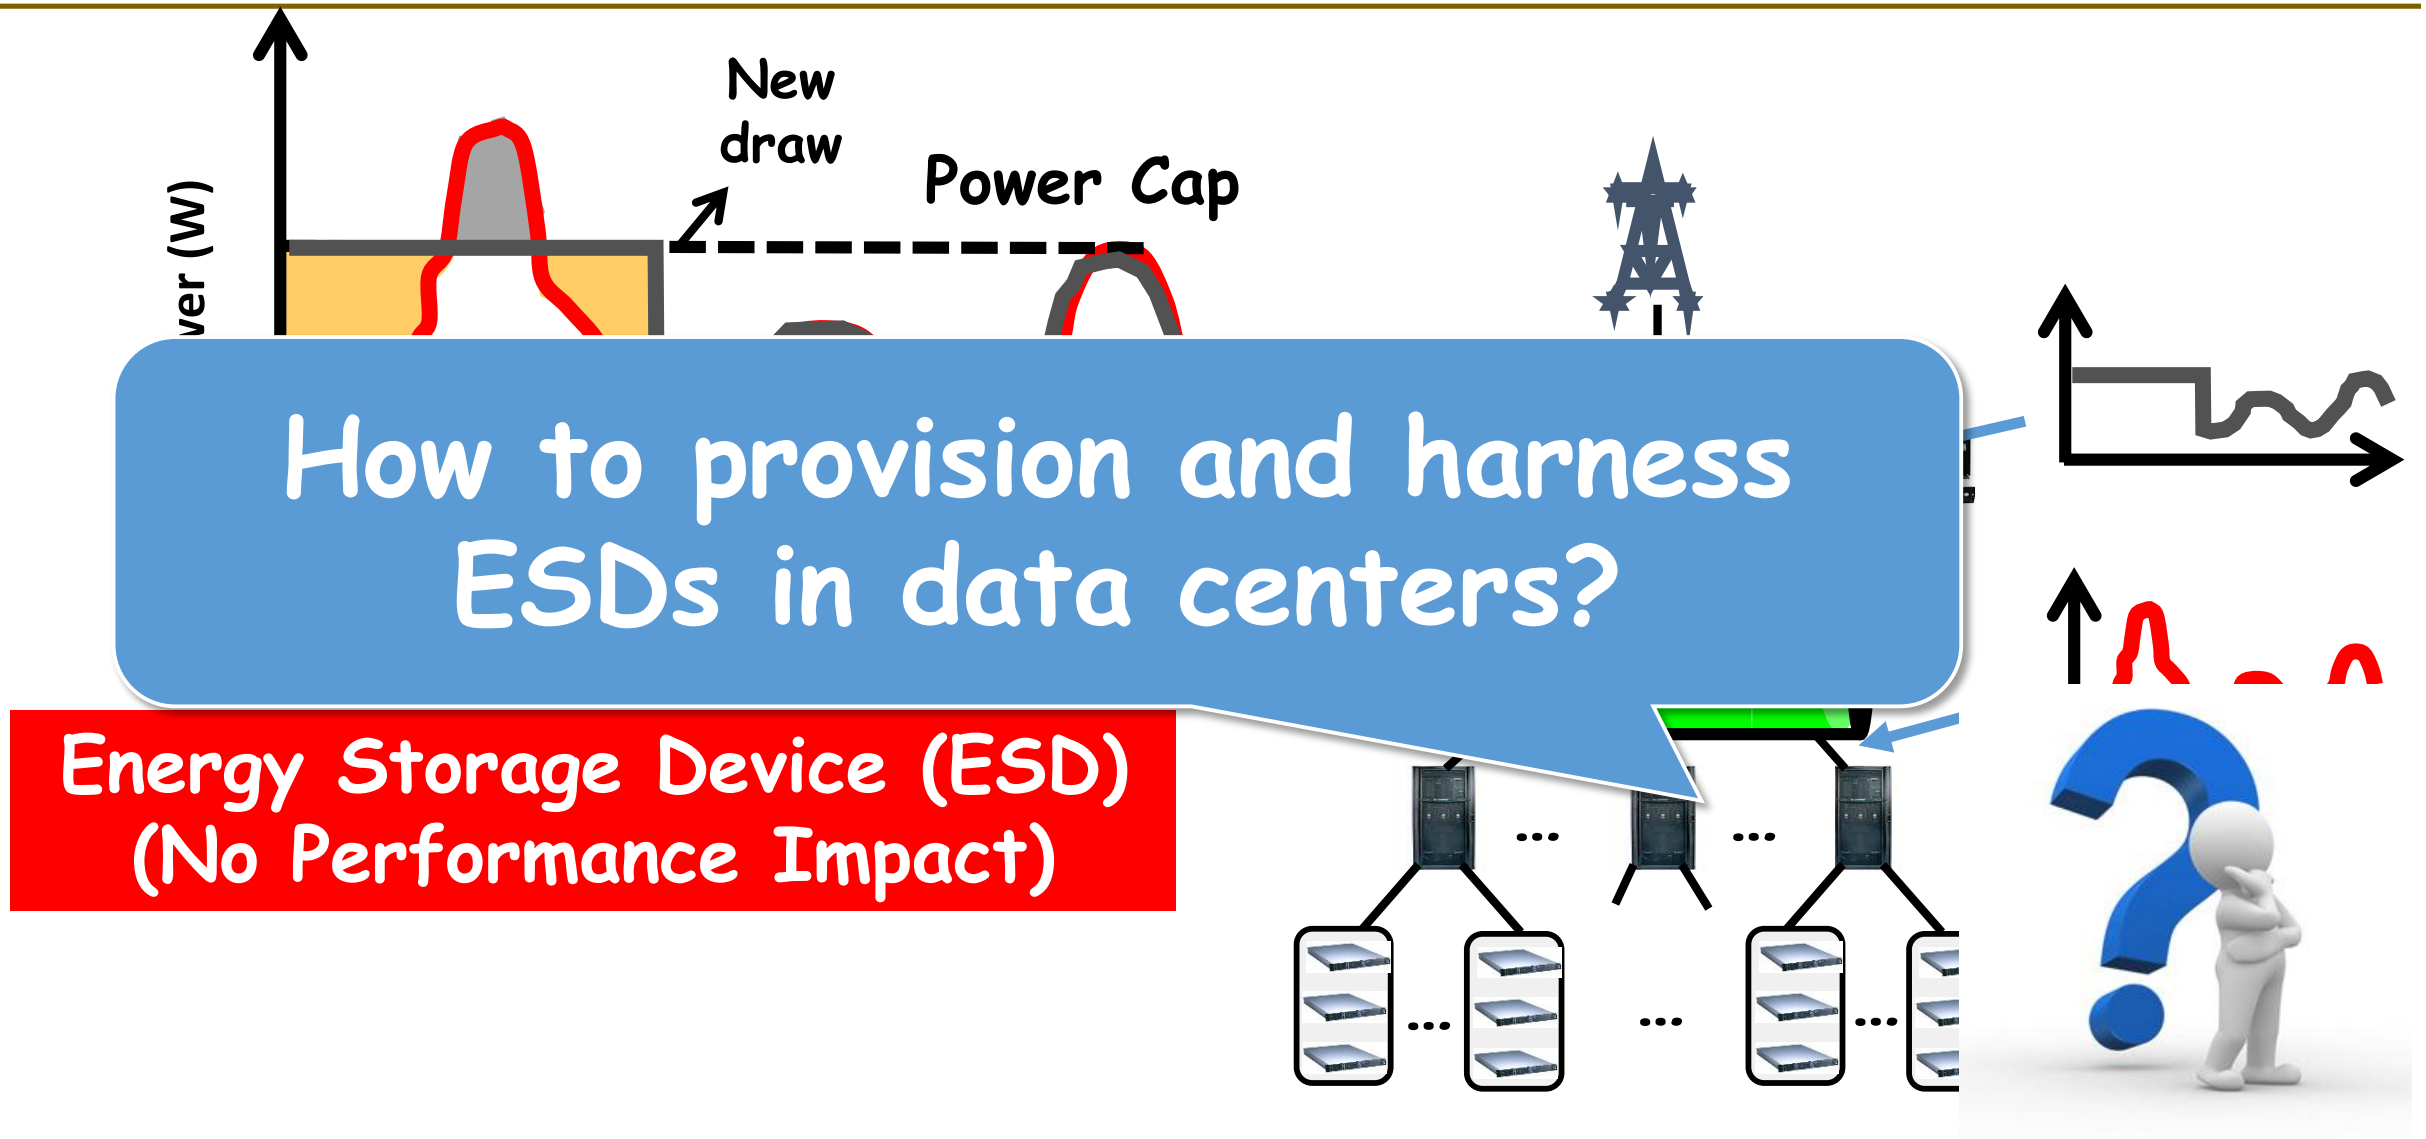

ESDs in Current Data Centers

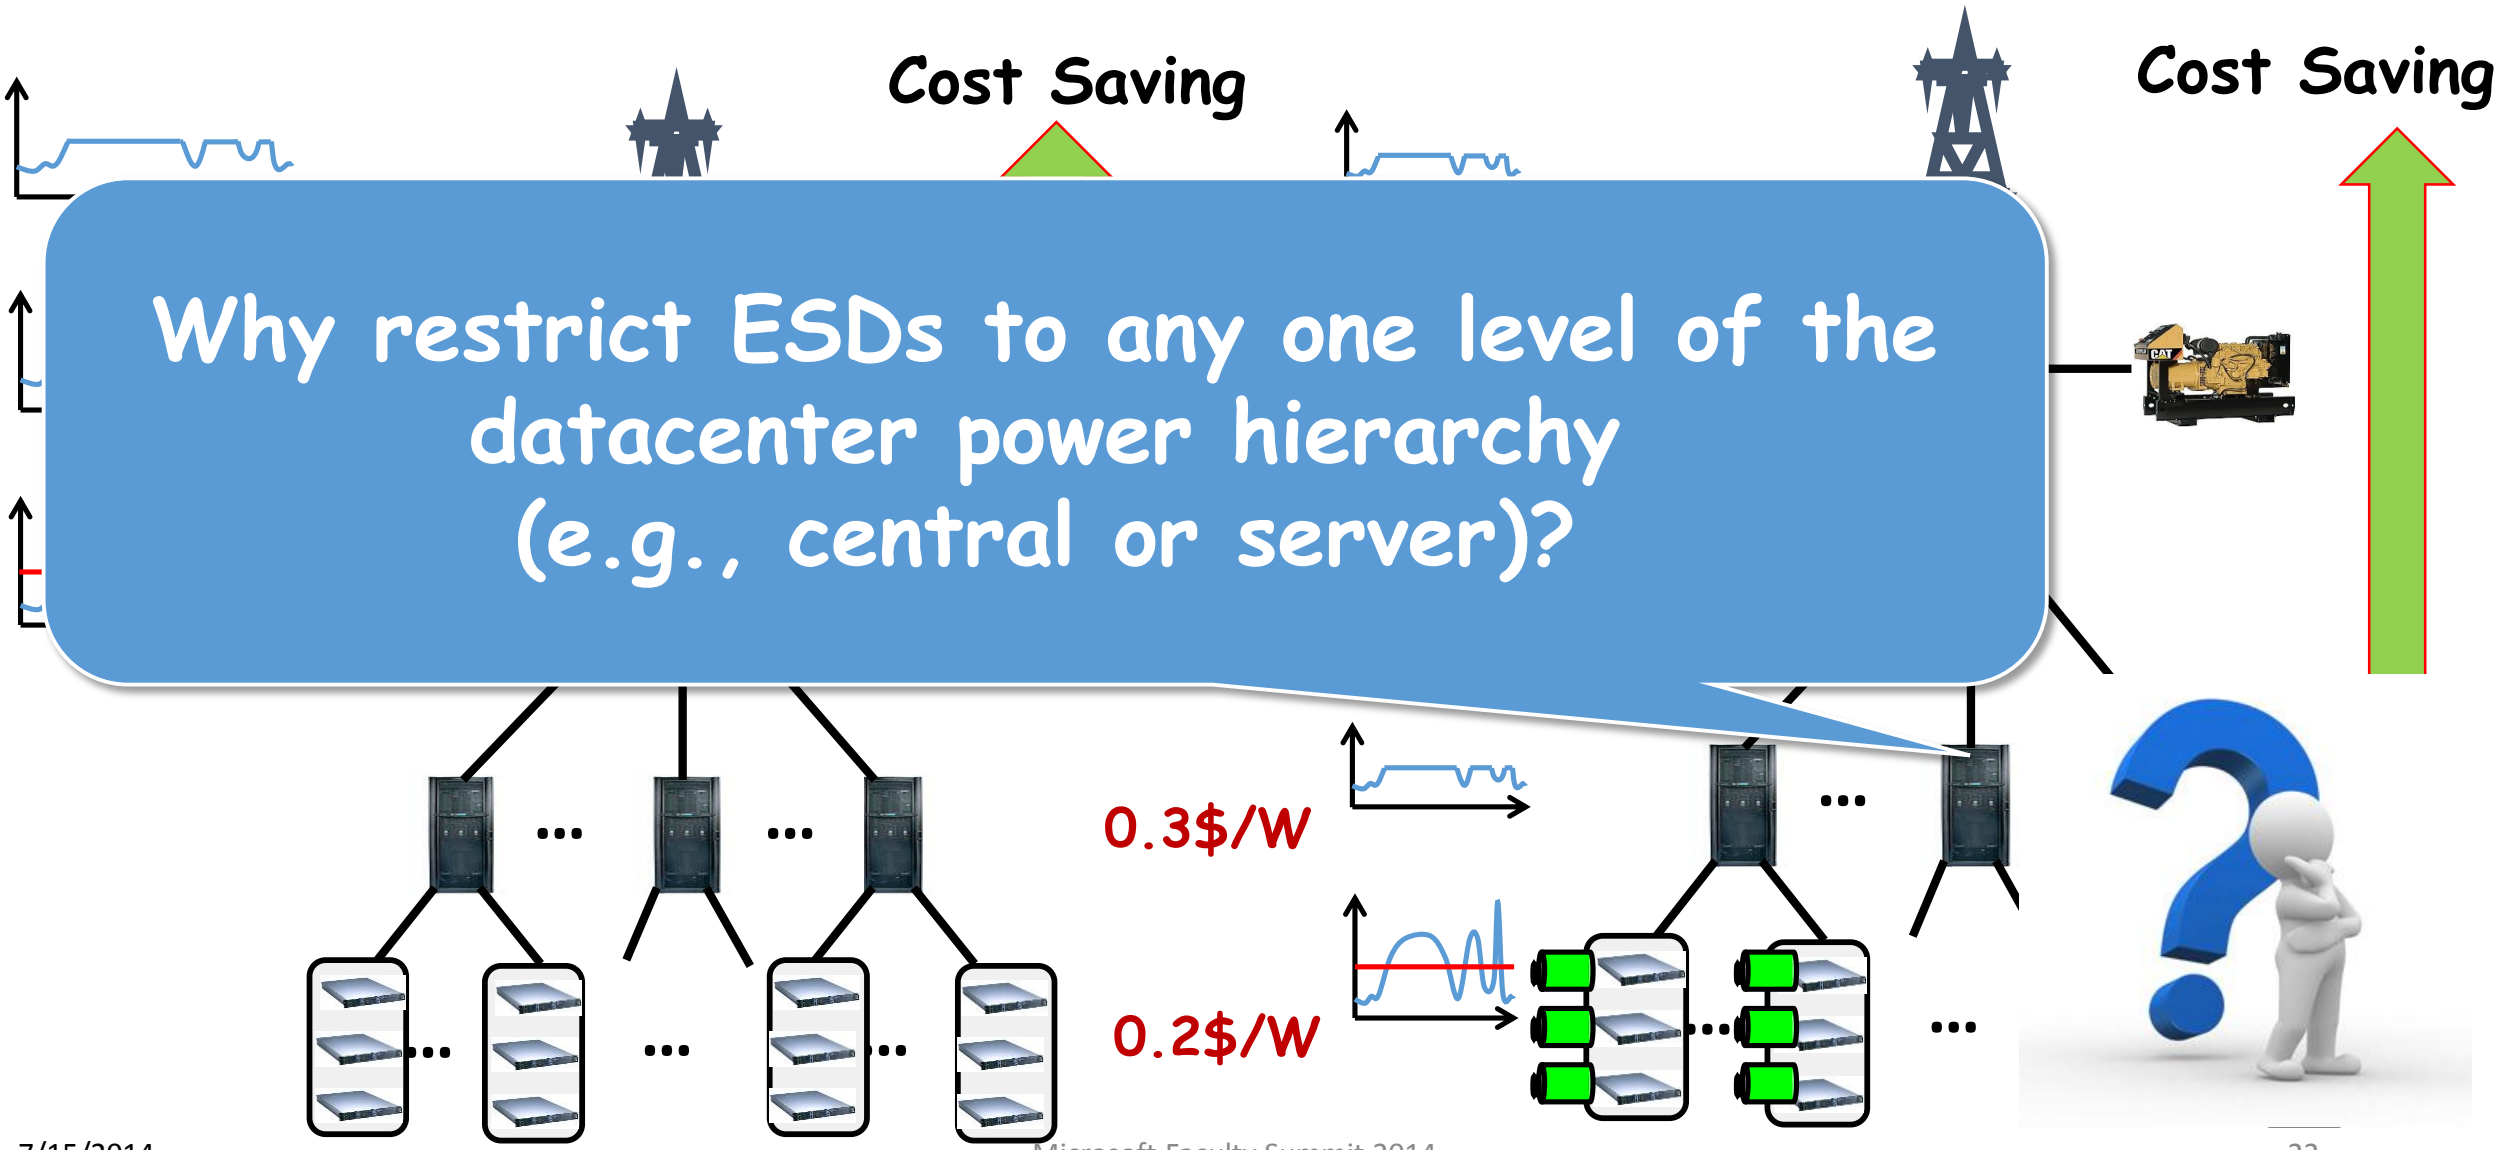

Ragone Plot

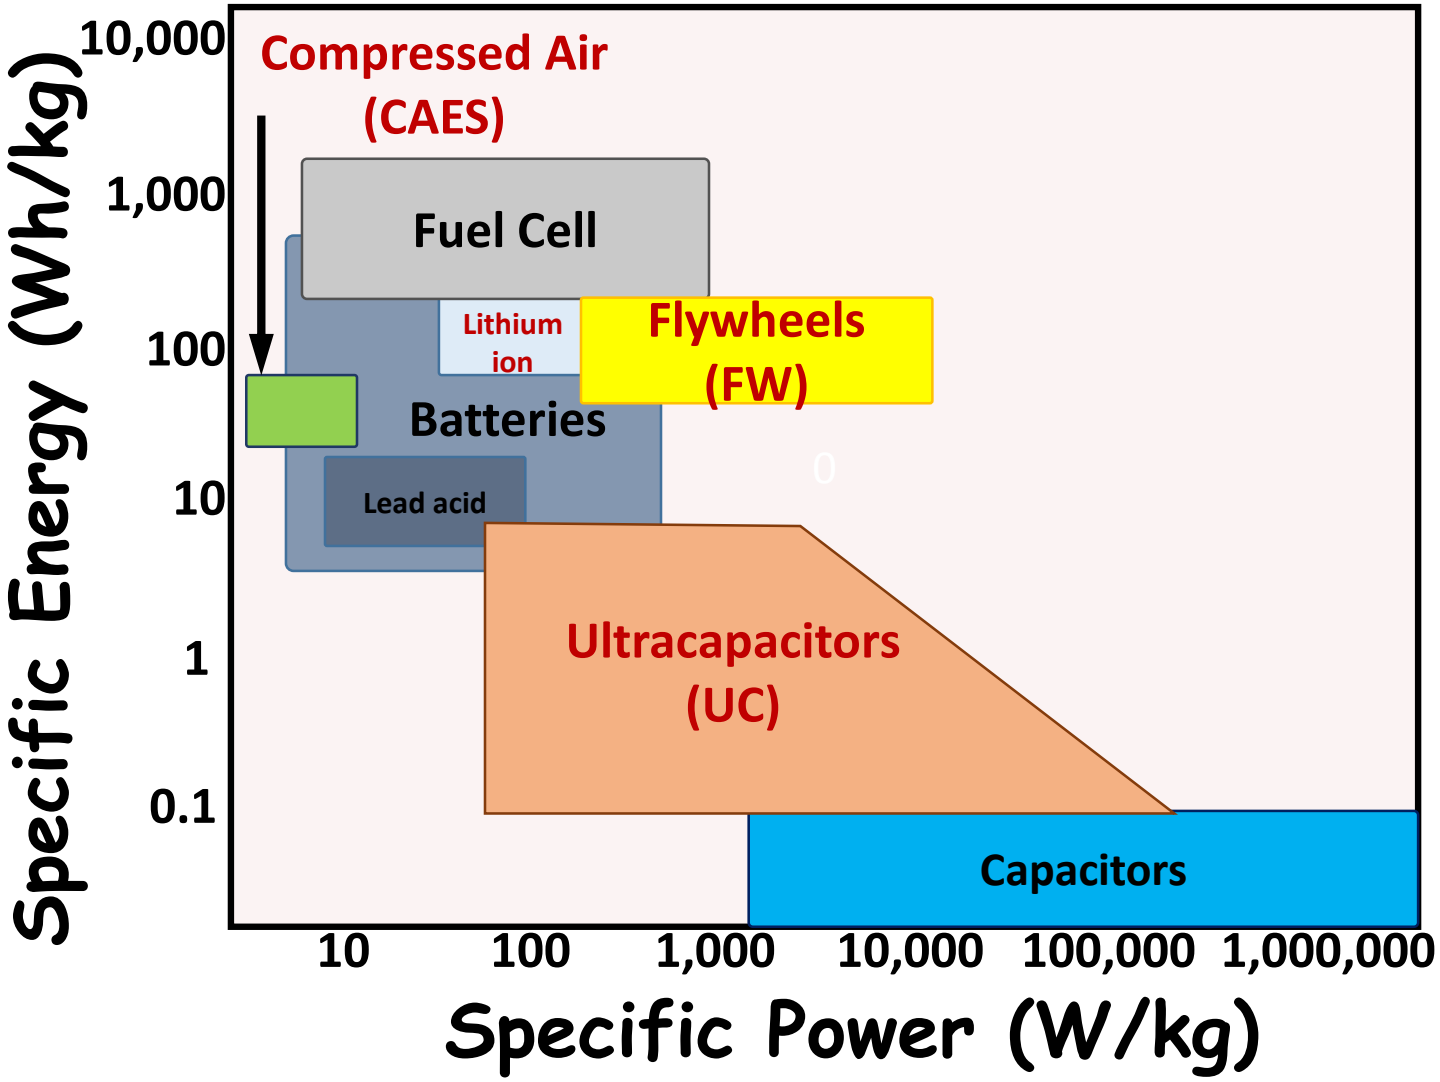
Specific Energy (Wh/kg)
10,000
1,000
100
10
1
0.1
Compressed Air
(CAES)
Fuel Cell
Lithium
ion
Flywheels
(FW)
Batteries
Lead acid
Ultracapacitors
(UC)
Capacitors
10
100
1,000
10,000
100,000
1,000,000
Specific Power (W/kg)

Capital Cost (Energy and Power)

	Ultracapacitor	Flywheel	Lithium ion battery	Lead-acid battery	Compressed air
Cost ($/kW)					

Why restrict to single ESD technology (e.g., Lead acid battery)?

Multi-level Multi-technology ESDs

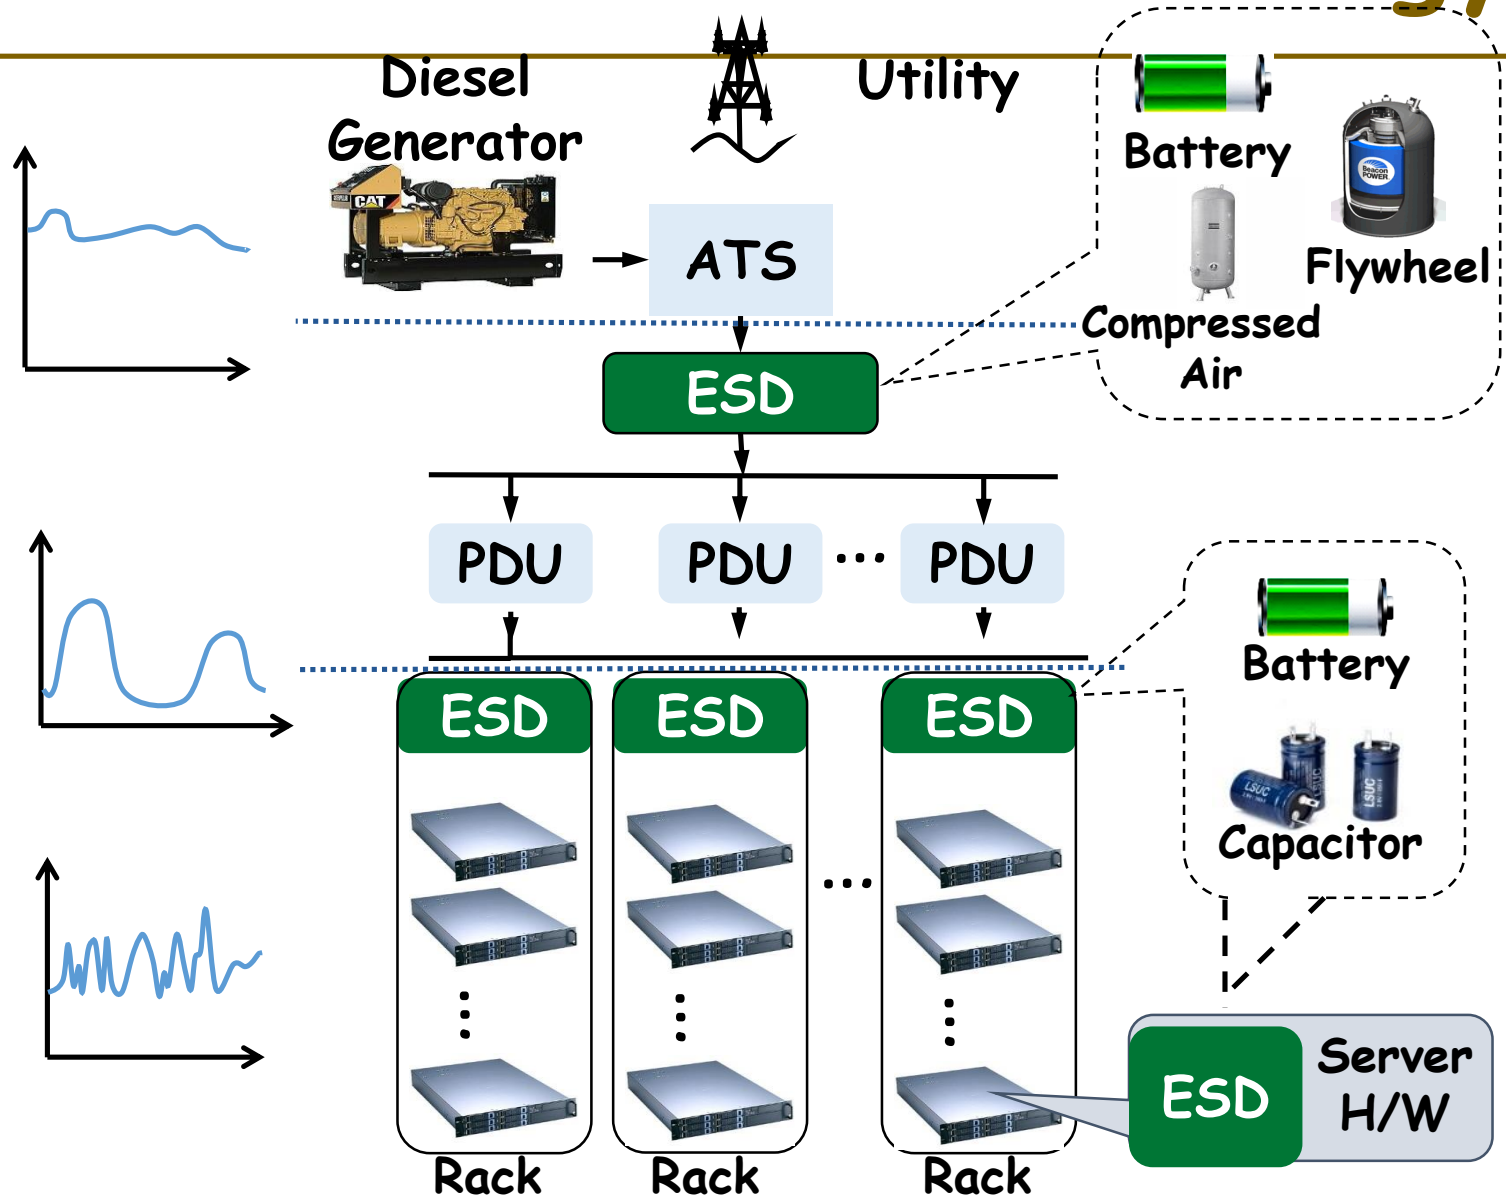

Cost Savings for Google Workloads

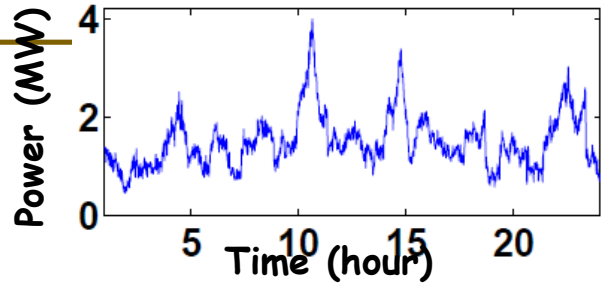

(Savings, ESD cost) Total cost without ESD is $12k/day

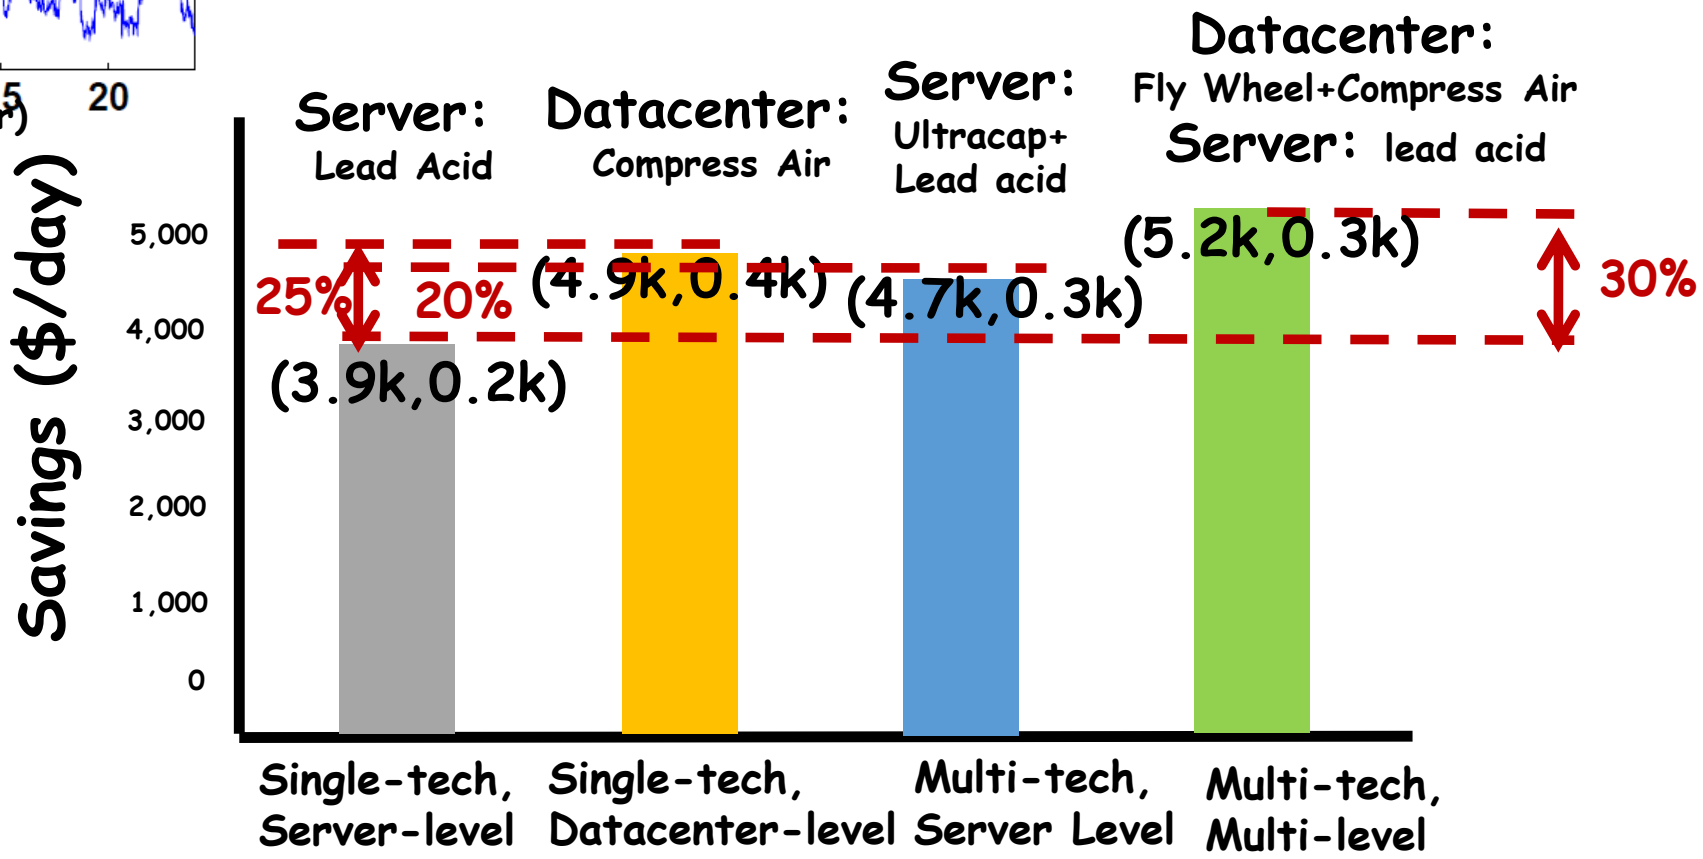

Details: Wang et al., "Energy Storage in the Datacenter: What, Where, and How Much?," Proc. ACM Sigmetrics 2012

Outline

- ✓ An example of cost-effective IT provisioning
- ✓ Keeping non-IT infrastructure costs low
  - ✓ Lowering peak power related costs using batteries
- • Conclusions

Related Work

- IT capacity provisioning
  - Capacity planning [Goiri et al. ICDCS'11]
    - Showed that more DCs, where each is lower availability (lower cost) but extra geo-spares, better
    - Computed optimal capacity placements
- Lowering infrastructure availability/cost
  - Reducing the "size" of power infrastructure
    - Under-provisioning backup generators [Wang14]
    - Reducing component redundancy [Govindan11,Kansal13]
  - Less aggressive cooling design
    - Has similarity in offering an availability vs cost trade-off [Schroeder@Sigmetrics12]
    - Related work in geo-distributed setting: [Wierman]
  - Lower availability IT

Conclusions

- Cost-effective capacity provisioning of geo-distributed data centers presents opportunities for novel problems in optimization and system design
  - Putting together lower availability data centers with appropriate fault tolerance mechanisms during subsequent operation
  - Key source of difficulty is uncertainty of subsequent workload evolution
    - Typical facility location based formulations might be inadequate
    - Stochastic optimization? Robust optimization?
- More information: http://www.cse.psu.edu/~bhuvan
- Joint work with: Anand Sivasubramaniam, Aman Kansal, Di Wang, Sriram Govindan, Hosam Fathy, Iyswarya Narayanan

Save the planet and return
your name badge before you
leave (on Tuesday)