SizeCap: Efficiently Handling Power Surges for Fuel Cell Powered Data Centers

Yang Li, Di Wang, Saugata Ghose, Jie Liu,
Sriram Govindan, Sean James, Eric Peterson,
John Siegler, Rachata Ausavarungnirun, Onur Mutlu

SAFARI Carnegie Mellon Microsoft

Executive Summary

- Fuel cells: efficient power source for data centers
- Problem: **limited load following capability**
  - Fuel cells only *gradually* increase output power when load increases
  - Power surges may lead to a **power shortfall** → server shutdown or damage
- Existing Approaches
  - Power capping: hurts performance
  - Energy storage device (ESD): increases cost
- Our Approach: SizeCap
  - **Our goal:** low cost, still guarantee workload performance
  - Key Idea 1: Size the ESD to cover only **typical-case** power surges
  - Key Idea 2: Use **smart power capping**, which is aware of fuel cell and workload behavior, to handle remaining power surges
- SizeCap safely **reduces ESD size by 50 – 85%**

Outline

- Background
- Problem
- Existing Approaches
- Key Ideas
- Detailed Design
- Evaluation
- Conclusion

Fuel Cell Powered Data Centers

- Data center power consumption continues to grow
  - In USA alone: 91 billion kWh @2013 → 140 billion kWh @2020
  - We need more energy-efficient power sources
- Fuel cells
  - Convert fuel (e.g., hydrogen, natural gas) into electricity
  - Advantages: high energy efficiency, low CO_2 emission, highly reliable delivery infrastructure

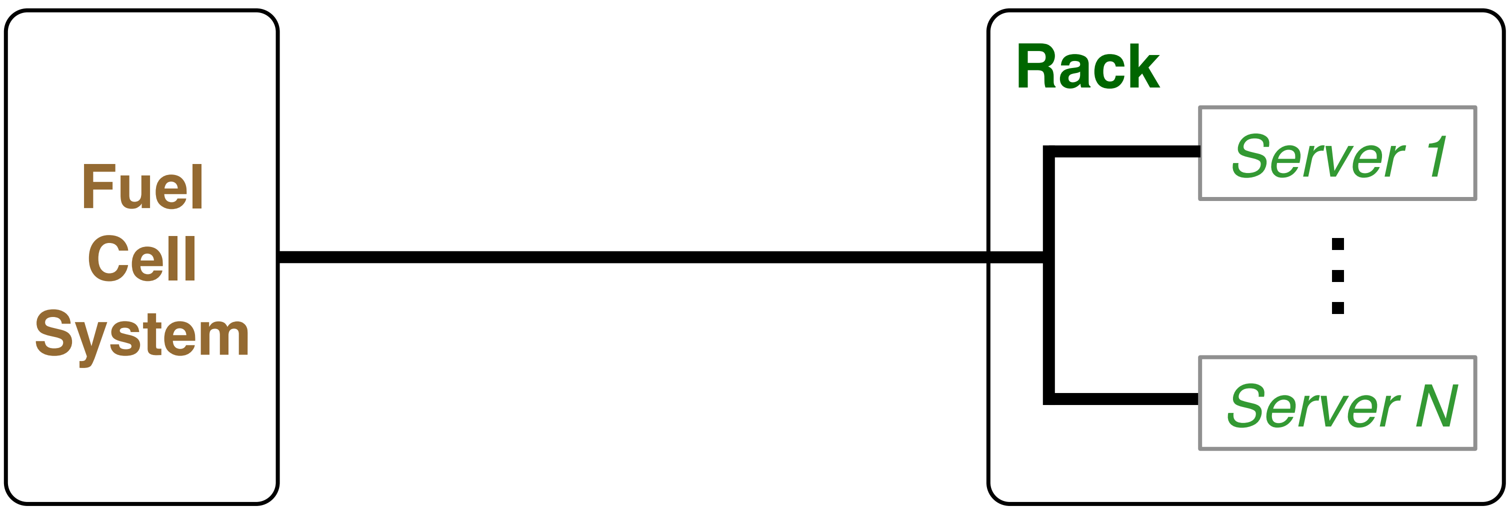

Outline

- Background
- Problem
- Existing Approaches
- Key Ideas
- Detailed Design
- Evaluation
- Conclusion

Problem: Limited Load Following Capability

- Fuel cell power output only gradually increases when power demand increases

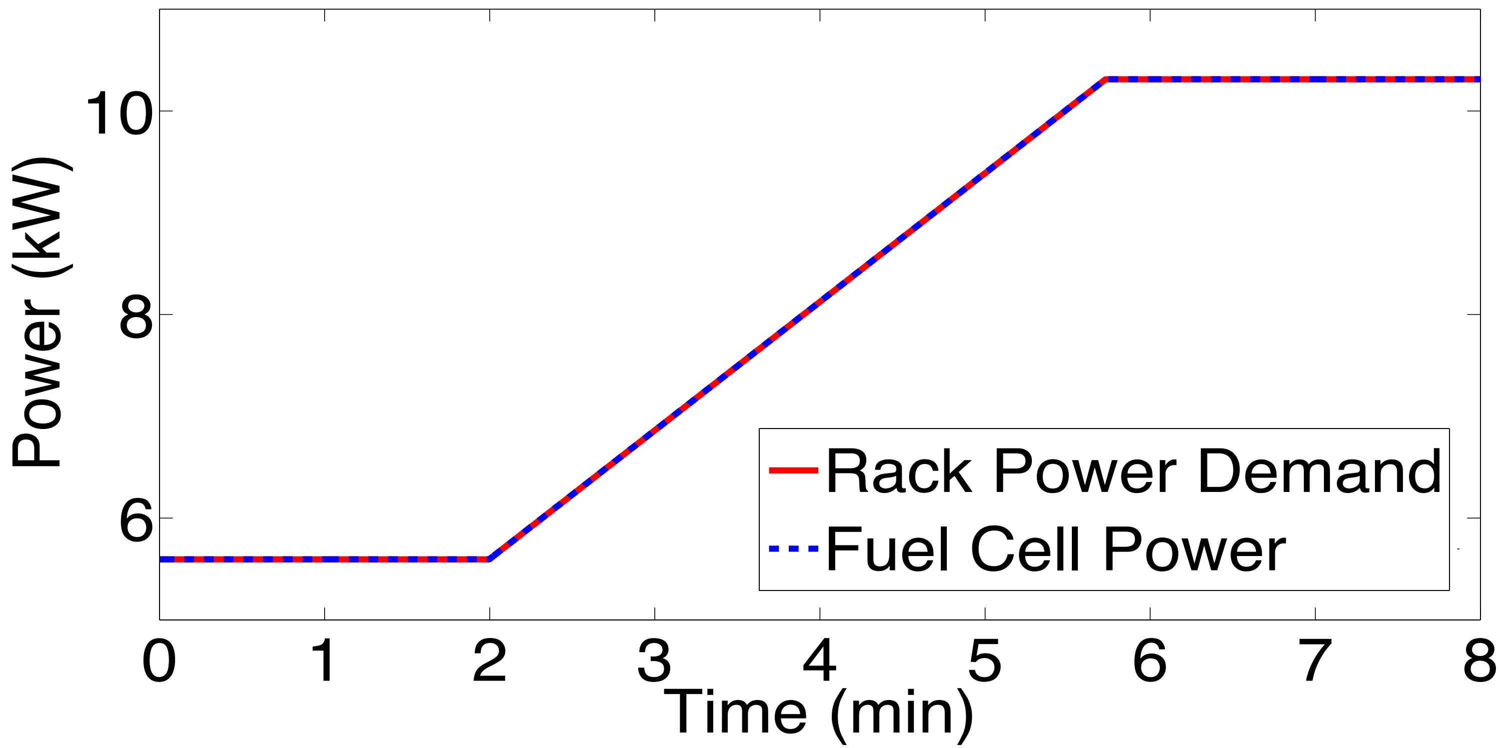

Problem: Limited Load Following Capability

- Fuel cell power output only gradually increases when power demand increases

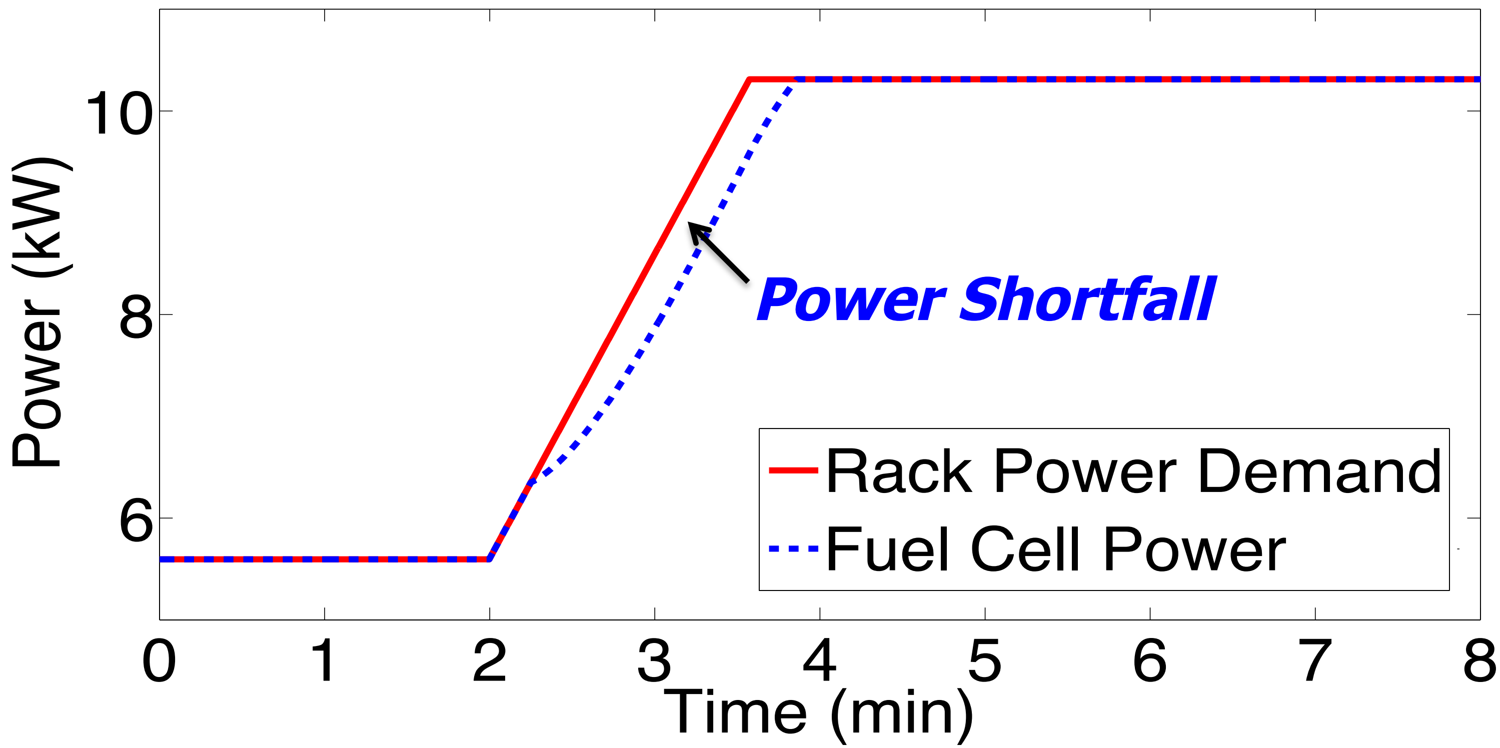

Problem: Limited Load Following Capability

- Fuel cell power output only gradually increases when power demand increases

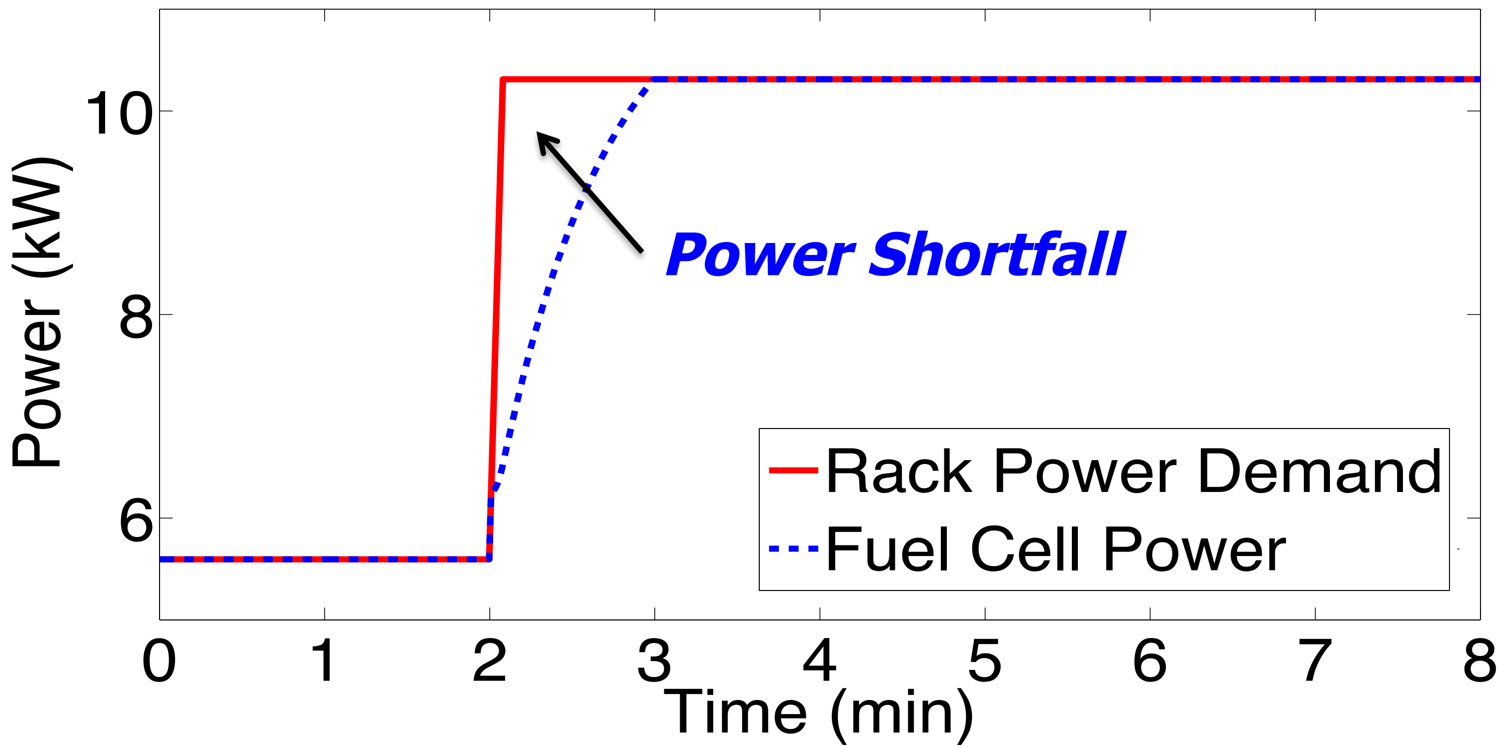

- Can lead to server damage or shut down

Problem: Limited Load Following Capability

- Fuel cell power output only gradually increases when power demand increases

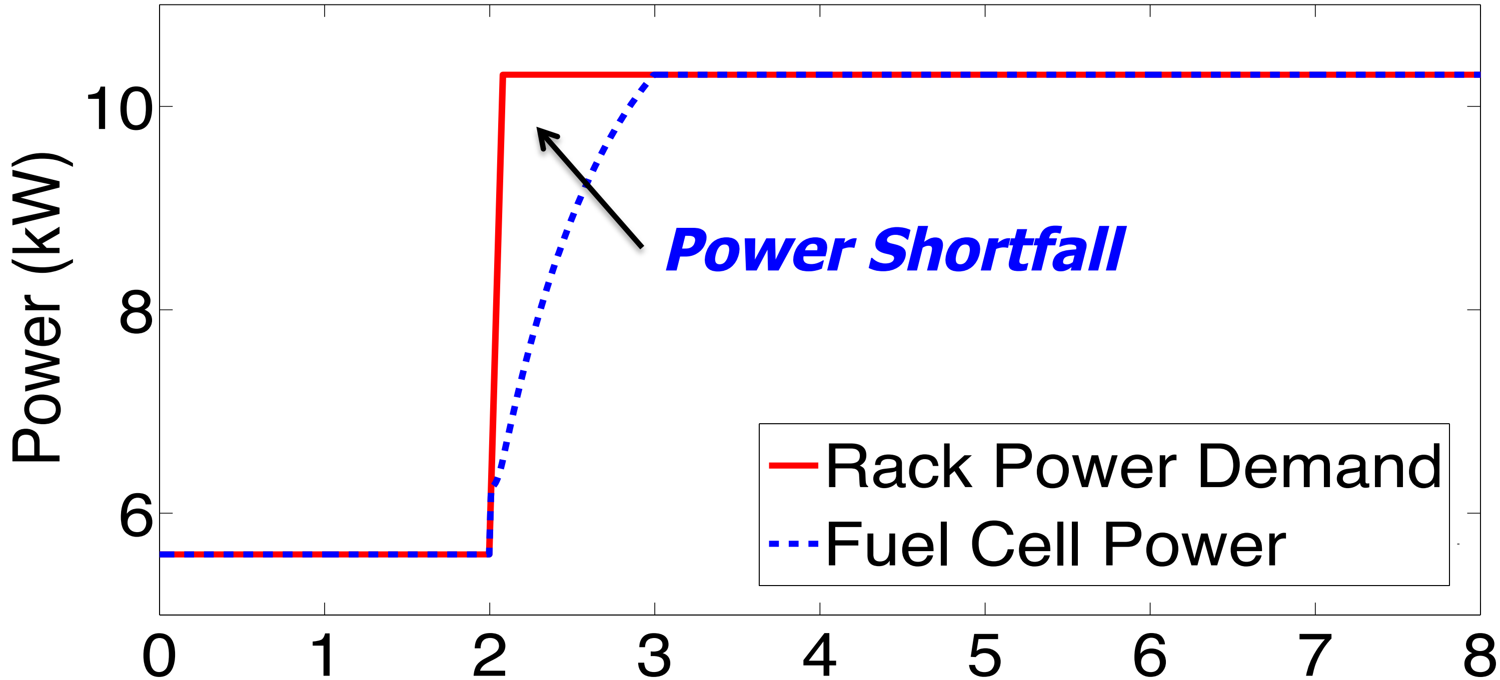

How can we efficiently handle power shortfalls?

Outline

- Background
- Problem
- Existing Approaches
- Key Ideas
- Detailed Design
- Evaluation
- Conclusion

Existing Approaches to Handling Power Shortfalls

- Power capping
  - Cuts down the power demand
  - Performs DVFS or shuts down nodes
  - Low cost
  - Hurts performance
- Energy storage device (ESD)
  - Buffers energy
  - Supplies extra energy when needed
  - High performance
  - High cost: ESD is sized to handle ***worst-case power surges***, even though they *rarely* occur
- **Our goal: high performance, low cost**

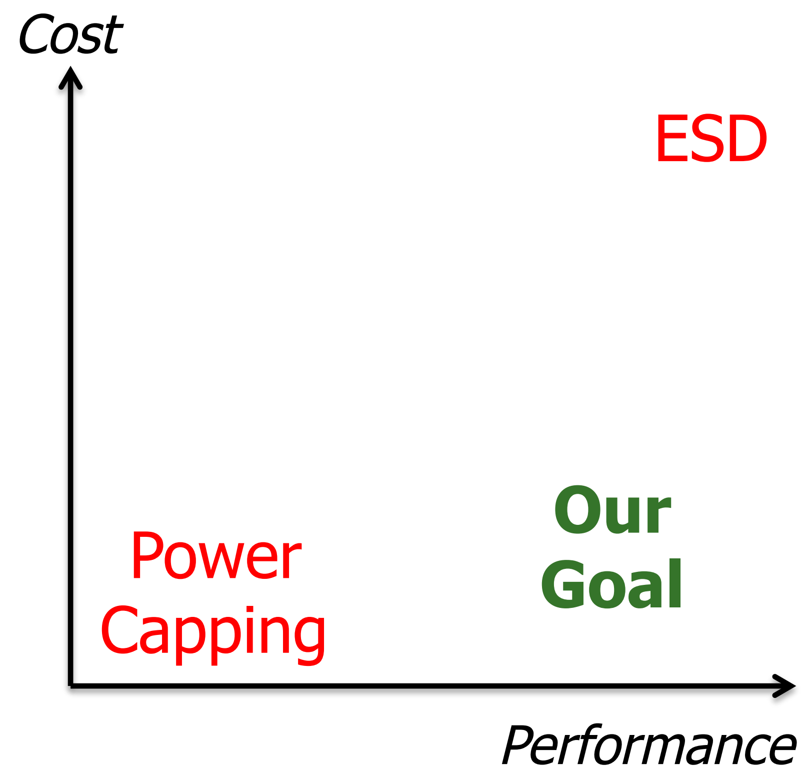

Outline

- Background
- Problem
- Existing Approaches
- Key Ideas
- Detailed Design
- Evaluation
- Conclusion

SizeCap: Key Ideas

Key Idea 1: Size ESD based on **typical-case power surges, not worst-case surges**

Key Idea 2: Use **smart power capping** to handle remaining power surges

Key Idea 1: Size ESD Based on Typical Case

- We study **production data center traces** from Microsoft

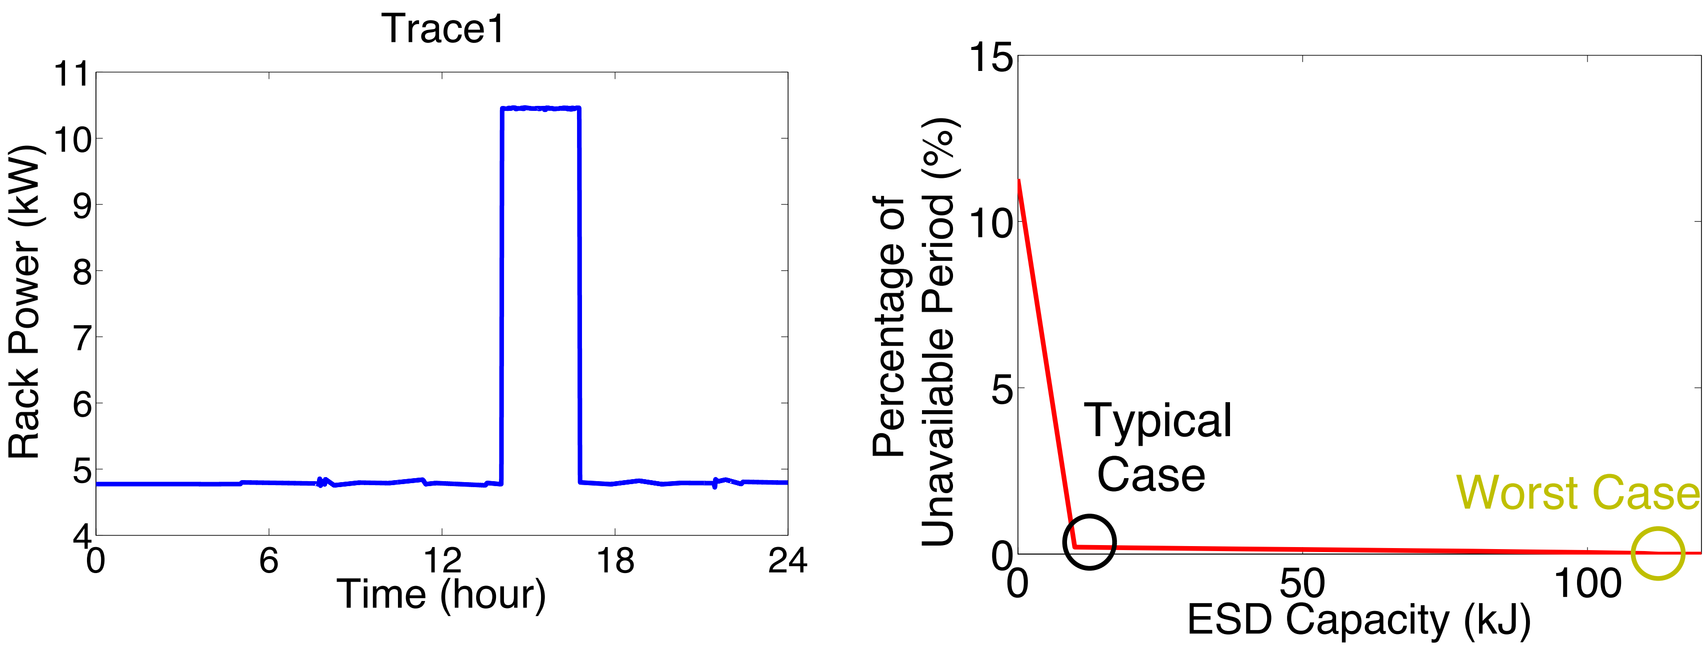

- Unavailable period: time that underprovisioned ESD cannot handle power surges
- Trace 1: reduce ESD size by **85%** → only 0.4% unavailable period

Key Idea 1: Size ESD Based on Typical Case

- We study **production data center traces** from Microsoft

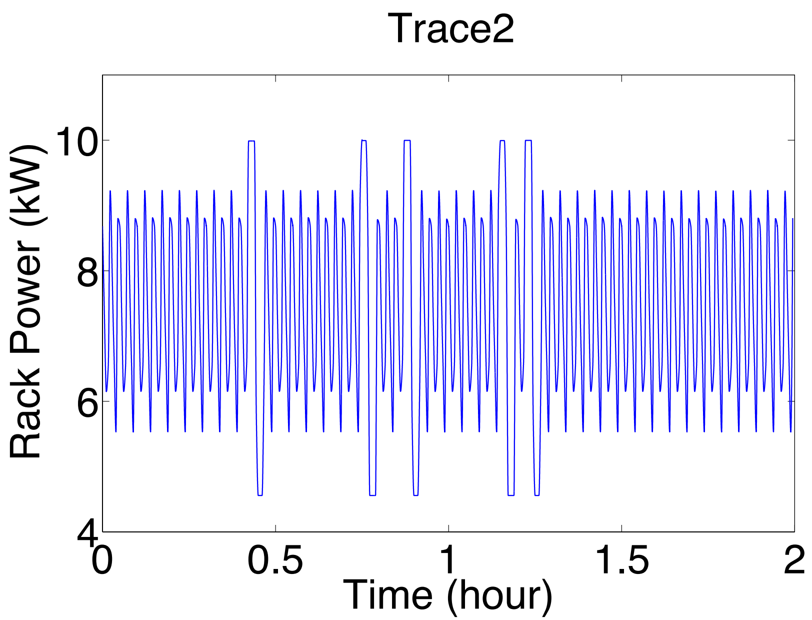

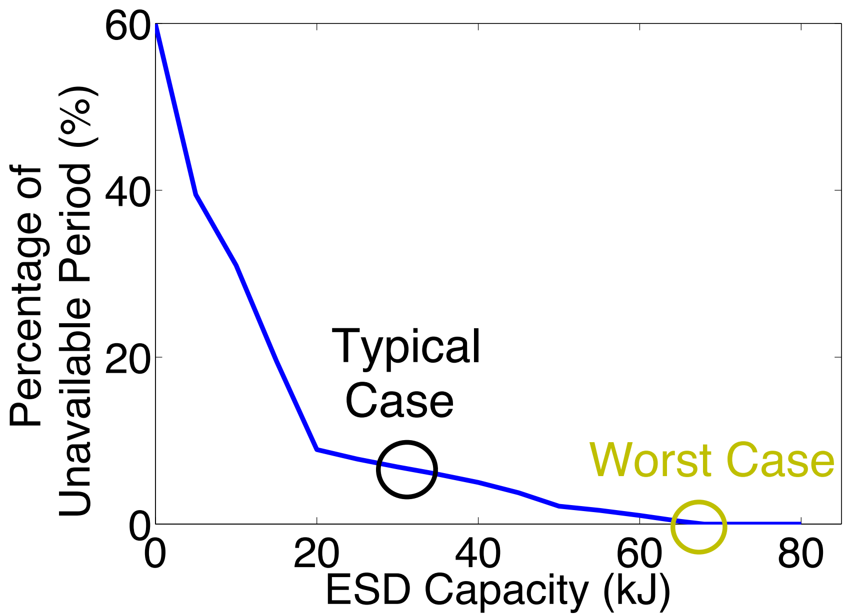

- Trace 2: reduce ESD size by **50%** → 6.2% unavailable period

Sizing ESD based on typical-case power surges does not hurt performance significantly

SizeCap: Key Ideas

Key Idea 1: Size ESD based on **typical-case power surges, not worst-case surges**

Key Idea 2: Use **smart power capping** to handle remaining power surges

Key Idea 2: Smart Power Capping

- Make power capping aware of **fuel cell load following behavior**
  - Fuel cells respond ***differently*** to different power surges
  - With fuel cell load following model, we can know how fuel cell power responds to rack power demand
  - Control the rack power such that it never exceeds sum of fuel cell power and ESD output

- Make power capping aware of **workload behavior**
  - Workload performance is dependent on ***how*** power is allocated over time
  - Allocate power over time to maximize workload performance

Smart power capping uses fuel cell, workload behavior to deliver higher benefits

SizeCap

- A framework to reduce ESD capacity by employing smart power capping policies
- At design time
  - Select **best** power capping policy implementable in system
  - Find **minimum** ESD size that still meets service level agreement (SLA) under the selected policy
- At runtime
  - Period-based power control
  - Every period: use power capping policy to determine power used by each server in next period

Outline

- Background
- Problem
- Existing Approaches
- Key Ideas
- Detailed Design
- Evaluation
- Conclusion

Design Time: Policy Selection & ESD Sizing

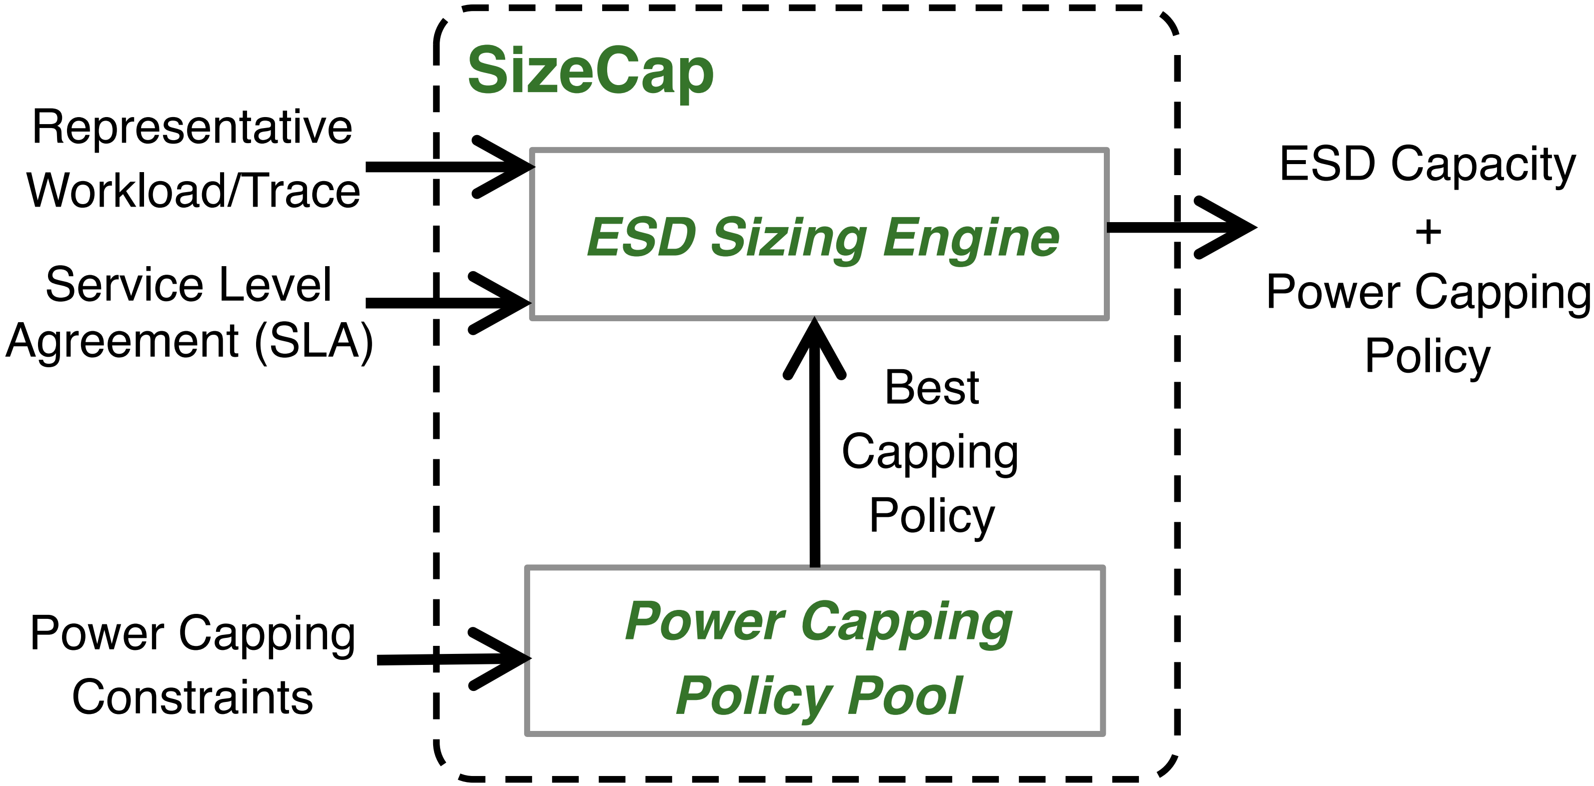

Runtime: Execute Power Capping Policy

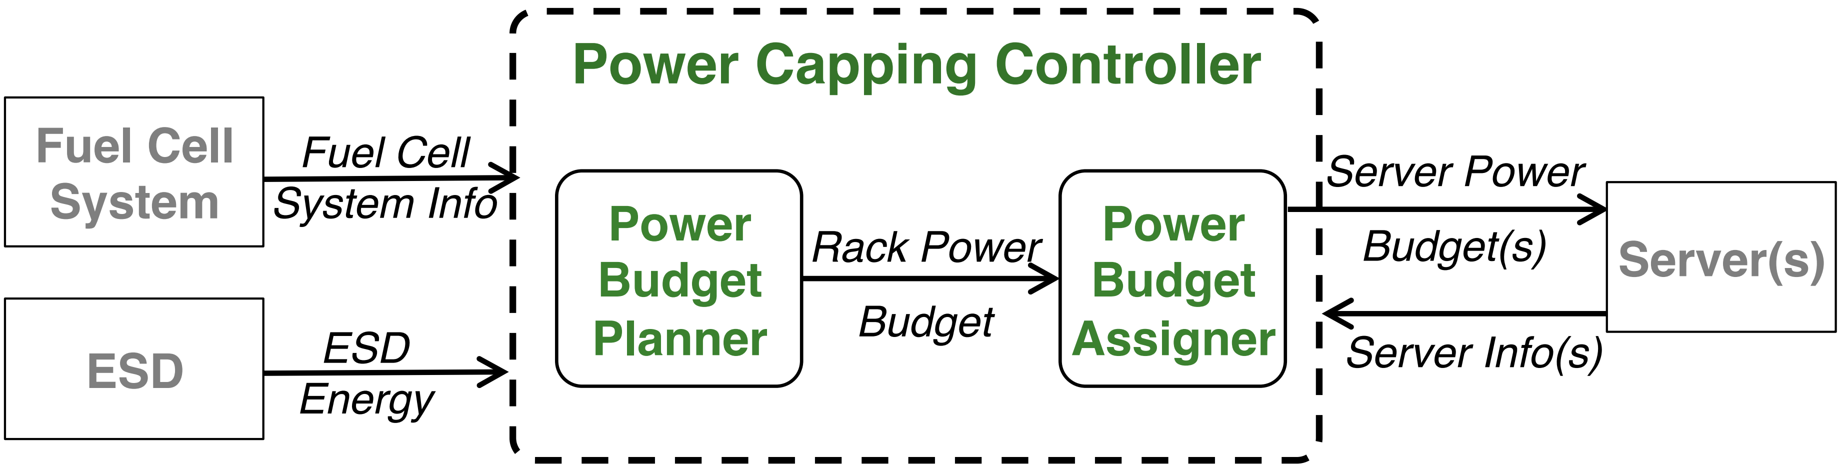

- Power Budget Planner: Plan total rack power budget for next period
- Power Budget Assigner: Distribute rack power among the servers for next period
- Controller can be centralized or decentralized

Power Capping Policy Taxonomy

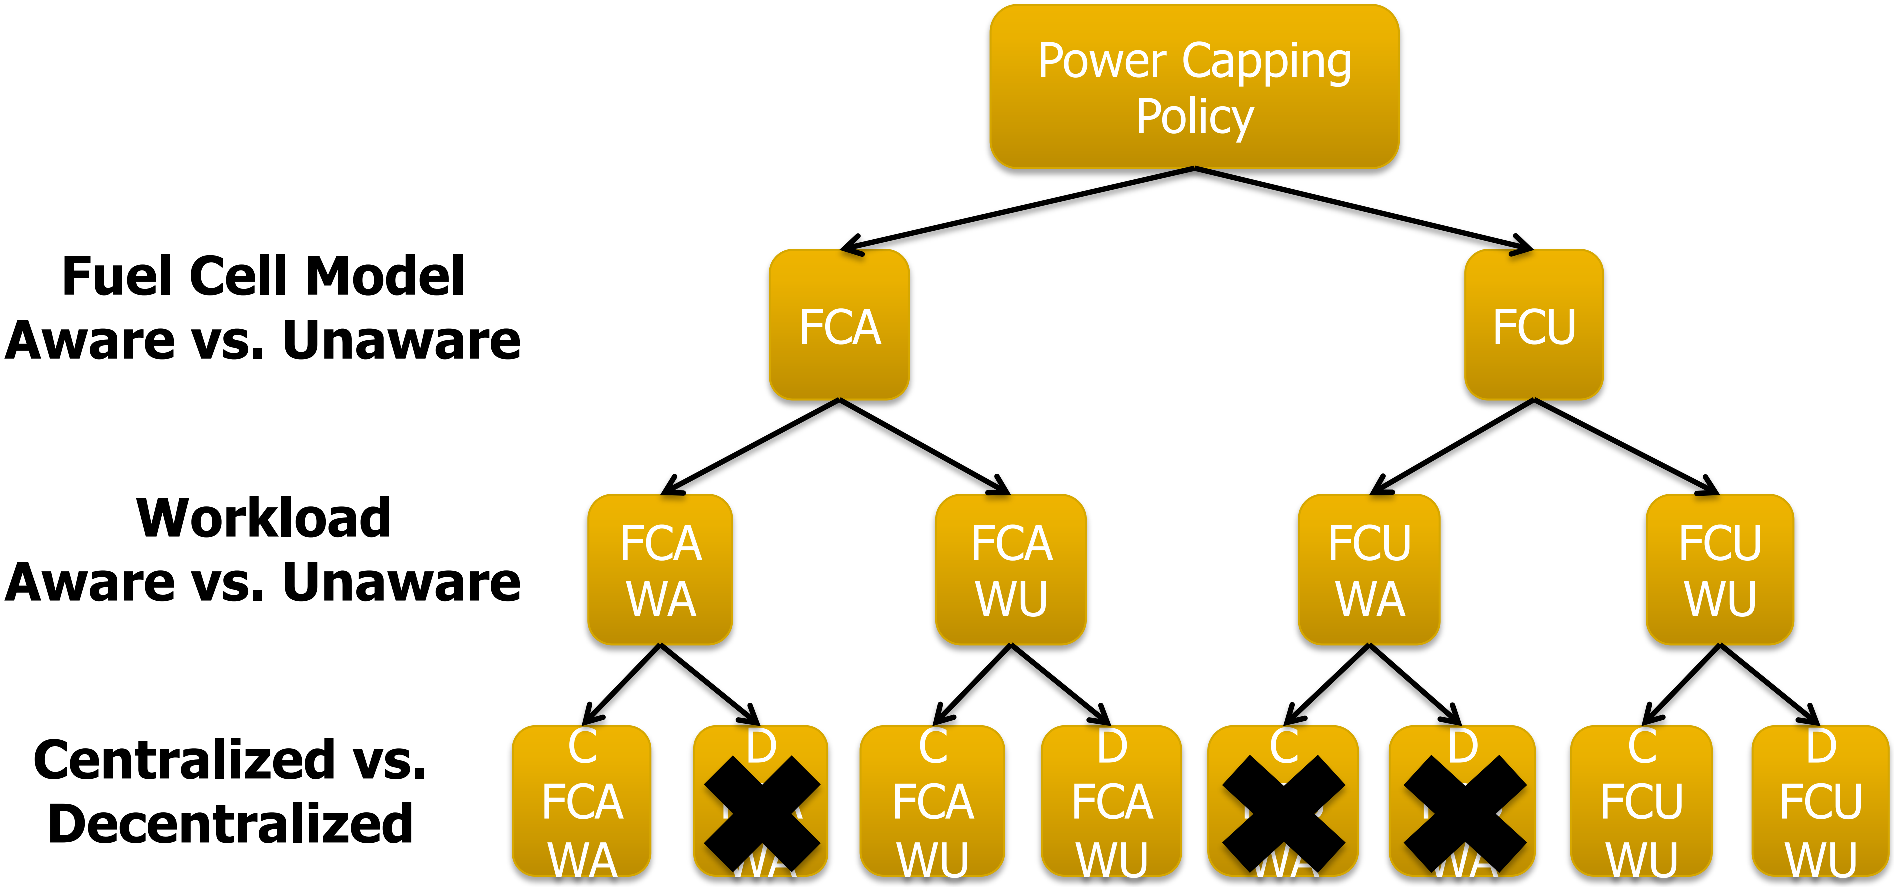

Power Capping Policy Taxonomy

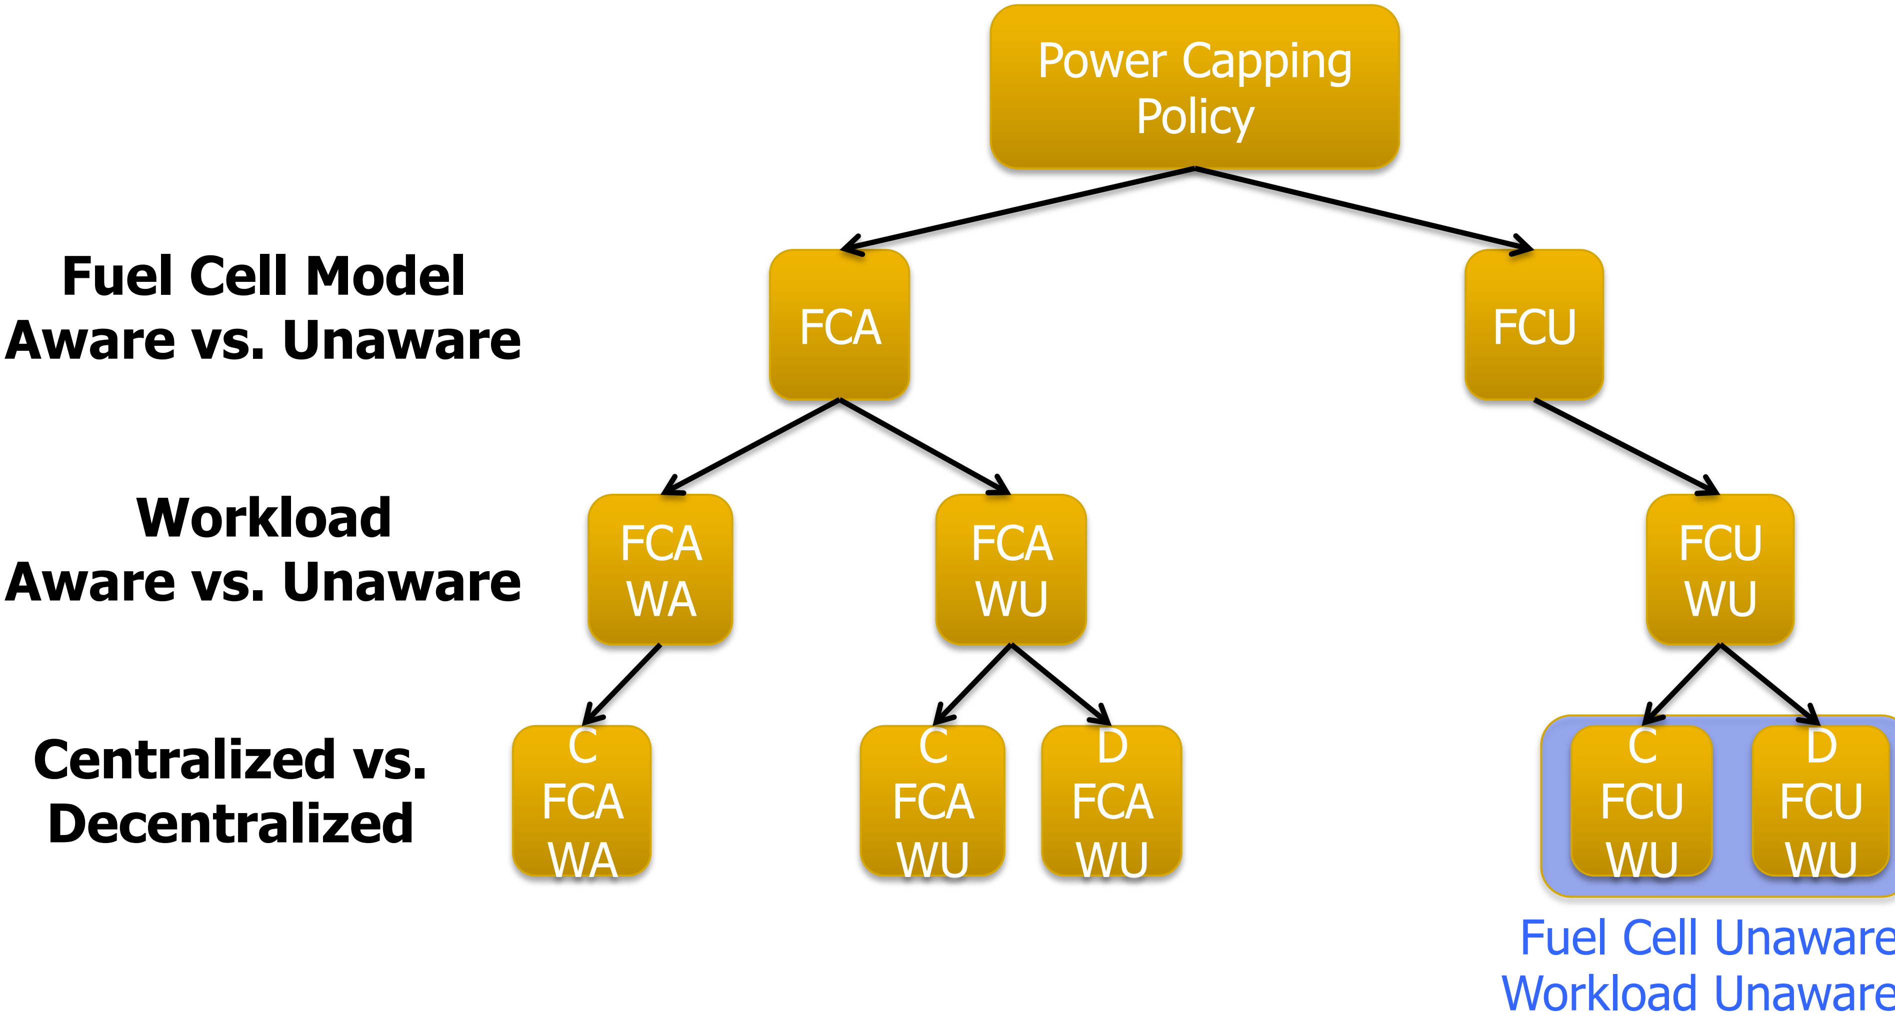

Fuel Cell and Workload Unaware Policies

- Goal: ***No power shortfalls***, optimize performance in ***next period***

- Power Budget Planner
  - Use ESD first
  - When ESD is used up
    - Ramp up rack power with ***conservative but safe rate***
    - **Static rate** that guarantees no shortfalls in entire fuel cell operating range

- Power Budget Assigner
  - Assign power to each server proportional to each server’s workload intensity or current power consumption

Power Capping Policy Taxonomy

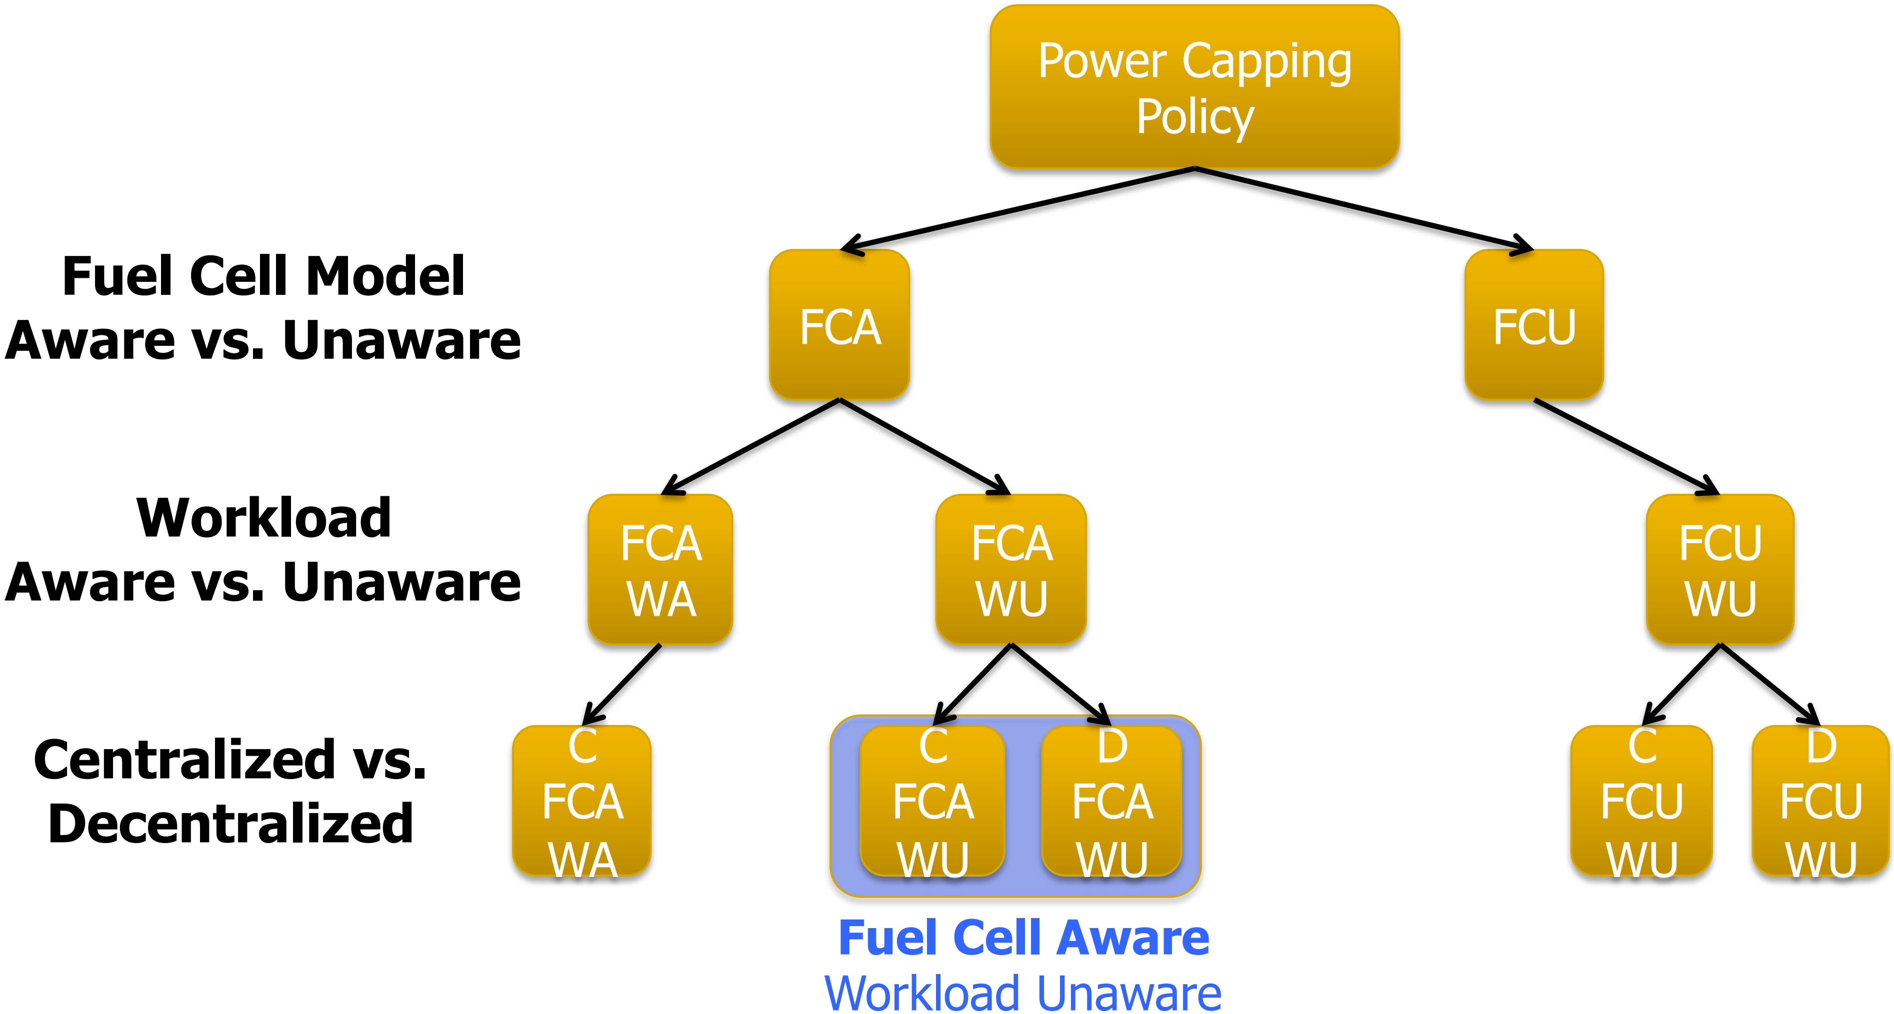

Fuel Cell Aware, Workload Unaware Policies

- Goal: ***No power shortfalls***, optimize performance in ***next period***
- Power Budget Planner
  - Use ESD first
  - When ESD is used up
    - Ramp up rack power with **maximum safe ramp rate**
    - **Dynamically adapted rate** to guarantee no shortfalls only under current conditions, derived from fuel cell model
- Power Budget Assigner
  - Same as fuel cell and workload unaware policies

Power Capping Policy Taxonomy

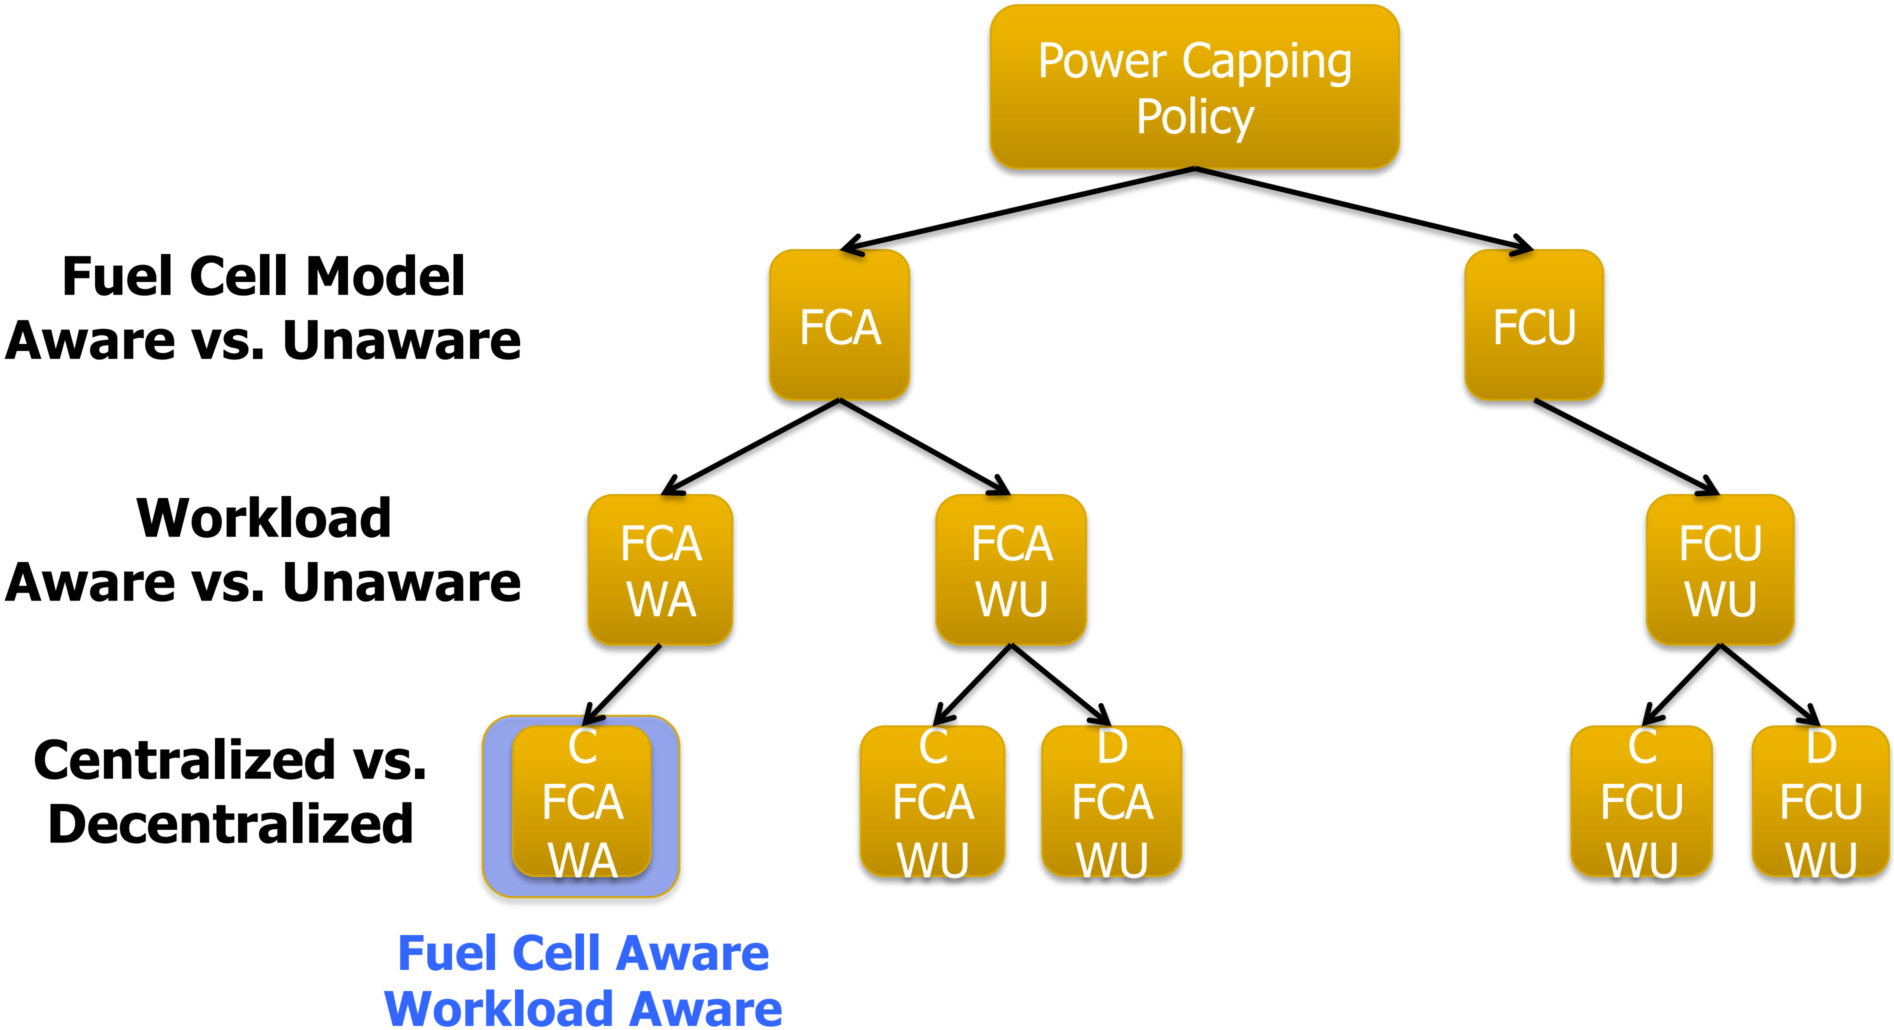

Fuel Cell and Workload Aware Policy

- Goal: ***No power shortfalls***, optimize performance ***over multiple periods***
  - Spend max ESD power now, cap aggressively in later periods
  - Save some power now and cap more, use power in later periods
- Power Budget Planner
  - Use fuel cell model to find ***all safe power capping settings***
  - Use workload behavior to assign power
    - Look at how workload performs over next several periods under different power allocations
    - Pick power capping setting that ***maximizes long-term performance***
- Power Budget Assigner
  - Similar to previous policies

Outline

- Background
- Problem
- Existing Approaches
- Key Ideas
- Detailed Design
- Evaluation
- Conclusion

Evaluation Methodology

- Simulation Configuration
  - Rack with **45 production servers**
  - Each server runs power capping driver developed in-house
- Traces
  - **Production traces** collected from Microsoft data centers
  - WebSearch workload
- Metrics
  - **Success rate**: Percentage of requests completed within the maximum allowable service time
  - **Average latency**: Average service latency of all requests
  - **P95 latency**: 95th percentile (tail) latency

Key Evaluation Results

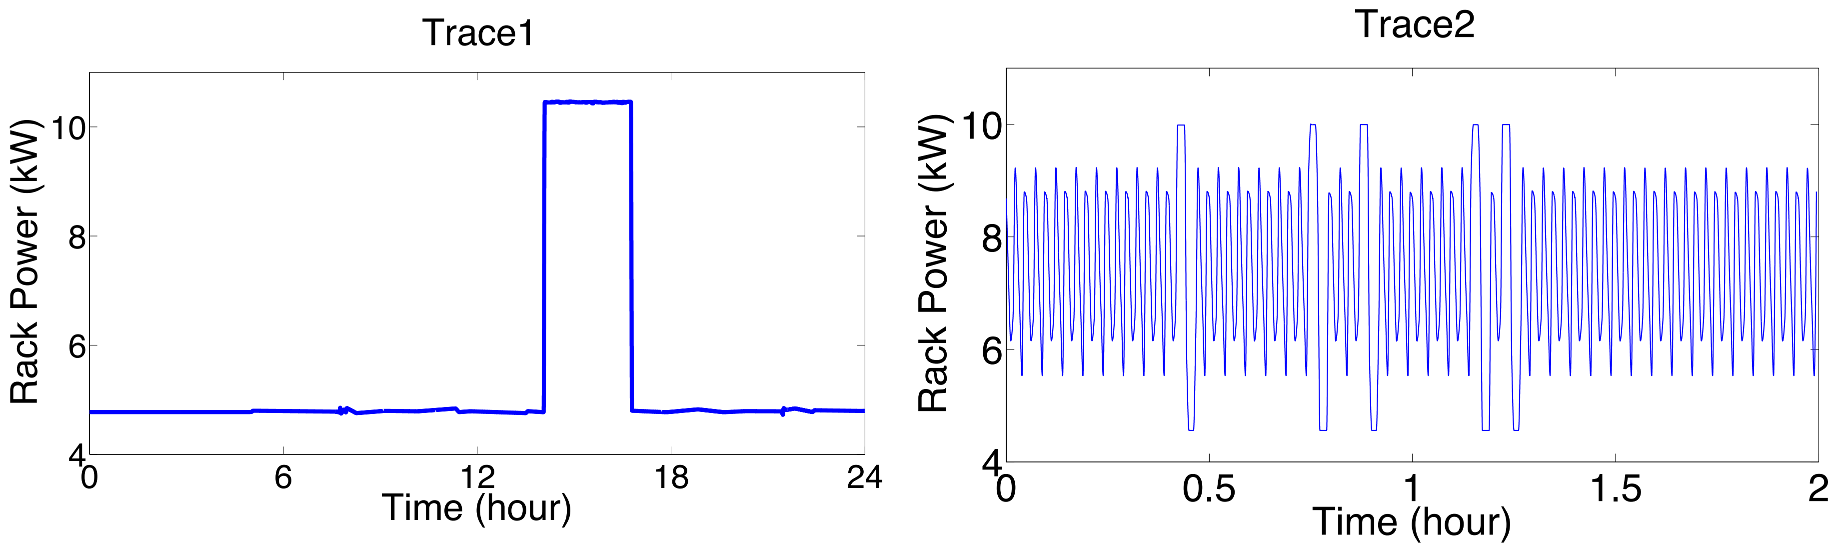

- SLA: Assume margins of 0.1% success rate, 3% average latency, and 10% P95 latency under fully-provisioned ESD
- **D-FCA-WU:** ***Safely reduces ESD size by 85% for Trace 1, by 50% for Trace 2, and meets SLA***
- **Policies with awareness** of fuel cell and/or workload behavior **reduce ESD size 10–20% more** than unaware policies

Conclusion

- Fuel cells: efficient power source for data centers
- Problem: **limited load following capability**
  - Fuel cells only *gradually* increase output power when load increases
  - Power surges may lead to a **power shortfall** → server shutdown or damage
- Existing Approaches
  - Power capping: hurts performance
  - Energy storage device (ESD): increases cost
- Our Approach: SizeCap
  - **Our goal:** low cost, still guarantee workload performance
  - Key Idea 1: Size the ESD to cover only **typical-case** power surges
  - Key Idea 2: Use **smart power capping**, which is aware of fuel cell and workload behavior, to handle remaining power surges
- SizeCap safely **reduces ESD size by 50 – 85%**

SizeCap: Efficiently Handling Power Surges for Fuel Cell Powered Data Centers

Yang Li, Di Wang, Saugata Ghose, Jie Liu,
Sriram Govindan, Sean James, Eric Peterson,
John Siegler, Rachata Ausavarungnirun, Onur Mutlu

Impact of ESD Cost on TCO

- ESD cost
  - Supercapacitor: $5.6 per kJ [McCawley, Fung Institute 2014]

- ESD sizes for our traces
  - Trace 1
    - Fully-provisioned: 112.5 kJ per rack → $630.00
    - After SizeCap: 16.9 kJ per rack → $94.50
  - Trace 2
    - Fully-provisioned: 68.0 kJ per rack → $380.80
    - After SizeCap: 34.0 kJ per rack → $190.40