Power-Aware Scheduling for Multi-Center HPC Electricity Cost Optimization

Abrar Hossain, Abubeker Abdurahman,
Mohammad A. Islam, Kishwar Ahmed

This work is supported in part by NSF under grants CNS-2300124, OAC-2411456, CCF-2324915, and ECCS-2152357

Overview of talk

- The HPC Energy Challenge
- Limitations of Existing Approaches
- Introducing TARDIS: Our Solution
- Evaluation & Key Results
- Conclusion & Contributions
- Q&A

The Growing Energy Demands of HPC

- HPC centers can consume up to **30 MW**, costing over **$40M/year**.
- **Electricity prices vary** by 200–300% (peak vs. off-peak).
- **Shifting workloads** can cut costs by ~40%.

Goal: Reduce electricity costs across HPC centers **with minimal performance impact**.

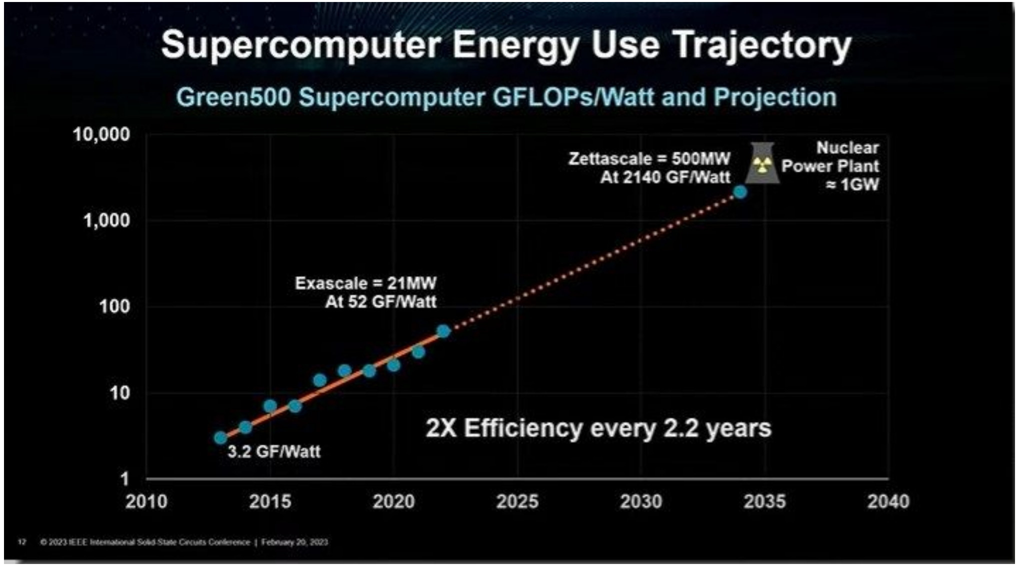

Reference:
https://www.hpcwire.com/2023/02/21/a-zettascale-computer-today-would-need-21-nuclear-power-plants/#foobox-3/0/AMD-Su-ISSCC-slide-supercomputer-energy-use-trajectory.jpg

Challenge: Fluctuating System-Wide Energy Use

- Power usage ranges from ~2 MW (idle) to over 20 MW (peak).
- While baseline PUE can be good (~1.1), sporadic spikes (>1.2) indicate inefficiencies (e.g., cooling, workload distribution).

Implication: These fluctuations highlight periods of high demand and potential inefficiencies, offering opportunities for optimization.

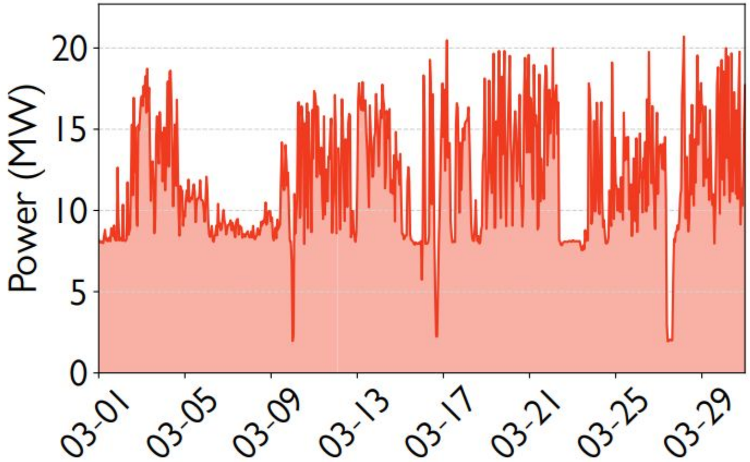

(a) Power consumption.

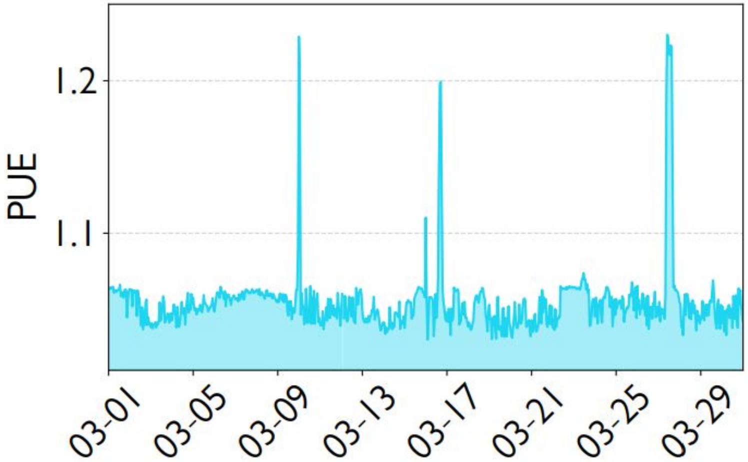

(b) Power usage effectiveness.

[1]J. Sun, Z. Gao, D. Grant, K. Nawaz, P. Wang, C.-M. Yang, P. Boudreaux, S. Kowalski, and S. Huff, "Energy dataset of Frontier supercomputer for waste heat recovery."

Challenge: Unpredictable Job Submission Patterns

- Job submissions vary **daily and hourly**, with **no clear trend or seasonality**.
- Marconi100 data[1] (May–Oct 2020) shows **irregular patterns**, with occasional **midday peaks (~2 PM)**.

Implication: Schedulers must handle **highly variable workload patterns** to optimize energy use.

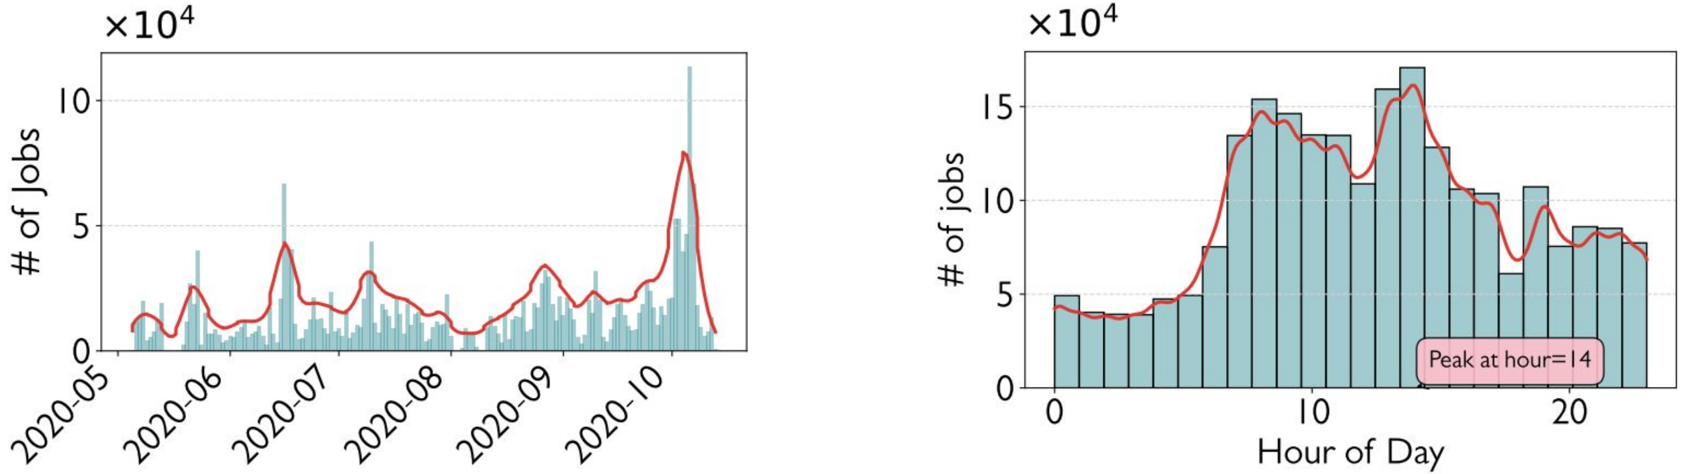

(a) Job submission trends from May - October 2020.

(b) Hourly distribution of job submissions.

[1] F. Antici, M. Seyedkazemi Ardebili, A. Bartolini, and Z. Kiziltan, "PM100: A Job Power Consumption Dataset of a Large-scale Production HPC System,"

Challenge: Diverse Power Needs of Jobs/Components

- **Compute nodes** (GPU-driven) show **high variability** and peak power (>2000W) — key target for optimization.
- **CPUs** and **memory** are more stable, with **minimal variation**.
- **Job-level:** Most jobs use **low to moderate power (10^4–10^6 W)**, but a **long tail of high-power jobs** skews energy use.

Implication: Effective energy management requires identifying and handling **diverse job power profiles**.

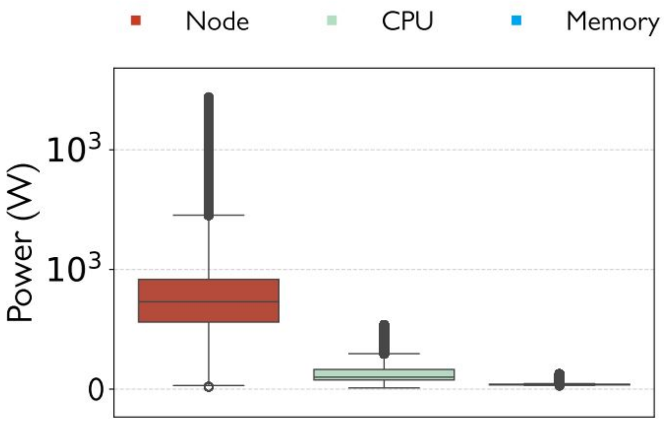

(a) HPC components.

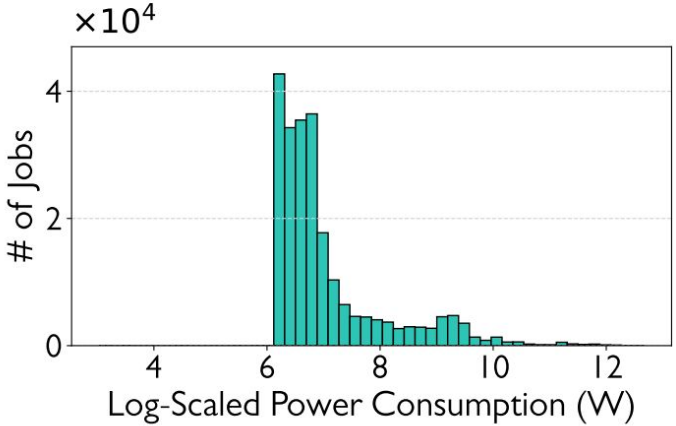

(b) HPC jobs.

Challenge: Job States & Energy Wastage

- **Completed jobs**: Consistent duration, efficient power use.
- **Failed/Timeout jobs**: High variability, often waste energy.

<u>Implication</u>: Dynamic scheduling must **balance load**, **allocate resources smartly**, and **align provisioning with actual workload** to reduce was

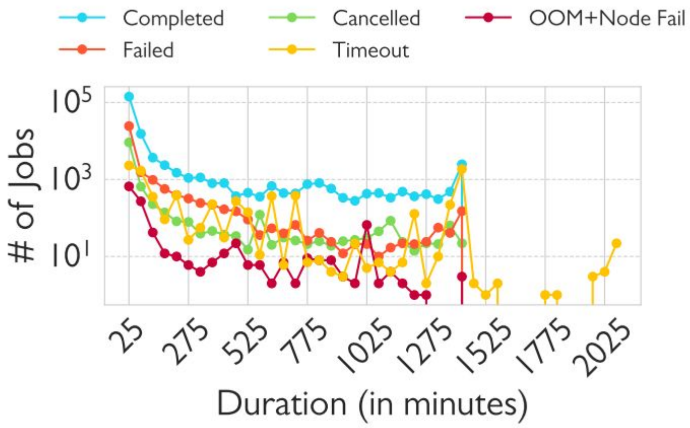

(a) Job trends.

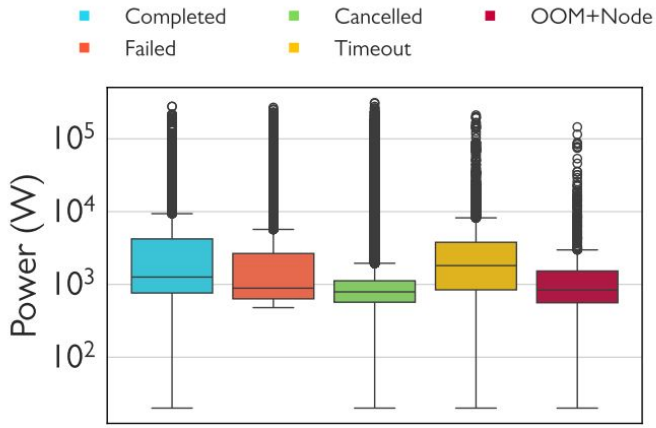

(b) Power consumption.

Why Current Schedulers Fall Short

Limitations of Existing Schedulers

- **Traditional schedulers** (FCFS, backfilling): Focus on throughput, ignore energy costs.
- **Power-aware schedulers**: Use **simplified models** or **static caps**, often neglect **dynamic pricing**.
- Lack **accurate, fine-grained job power prediction**
- Focus on **temporal optimization** (single center); **spatial optimization** is rare
- **Rigid power capping** can hurt performance

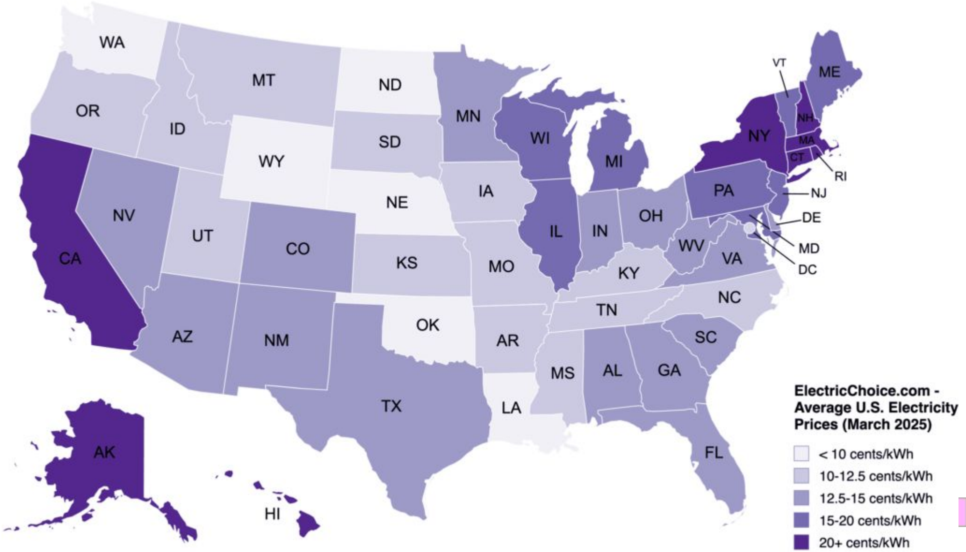

TARDIS: Power-Aware Scheduling for Multi-Center HPC

A novel scheduler that **minimizes electricity costs** via **temporal and spatial optimization**.

- **Power Prediction**: Uses a **Graph Neural Network (GNN)** to estimate job-level power.

- **Holistic Scheduling**: Allocates jobs across centers based on:
  - Predicted power
  - **Dynamic electricity prices**
  - System load & job characteristics

TARDIS: System Workflow

1. **Power Prediction**: GNN estimates job-level power.

2. **Job Scoring**: Jobs ranked by **cost**, **efficiency**, and **wait time** across sites/times.

$$Score_{j,k,t} = w_c C_{j,k,t} + w_p P_{j,k} + w_u U_{k,t} + w_w W_{j,t}$$

3. **Spatio-Temporal Dispatch**: Assigns jobs to HPC centers to **minimize global cost** while meeting system constraints.

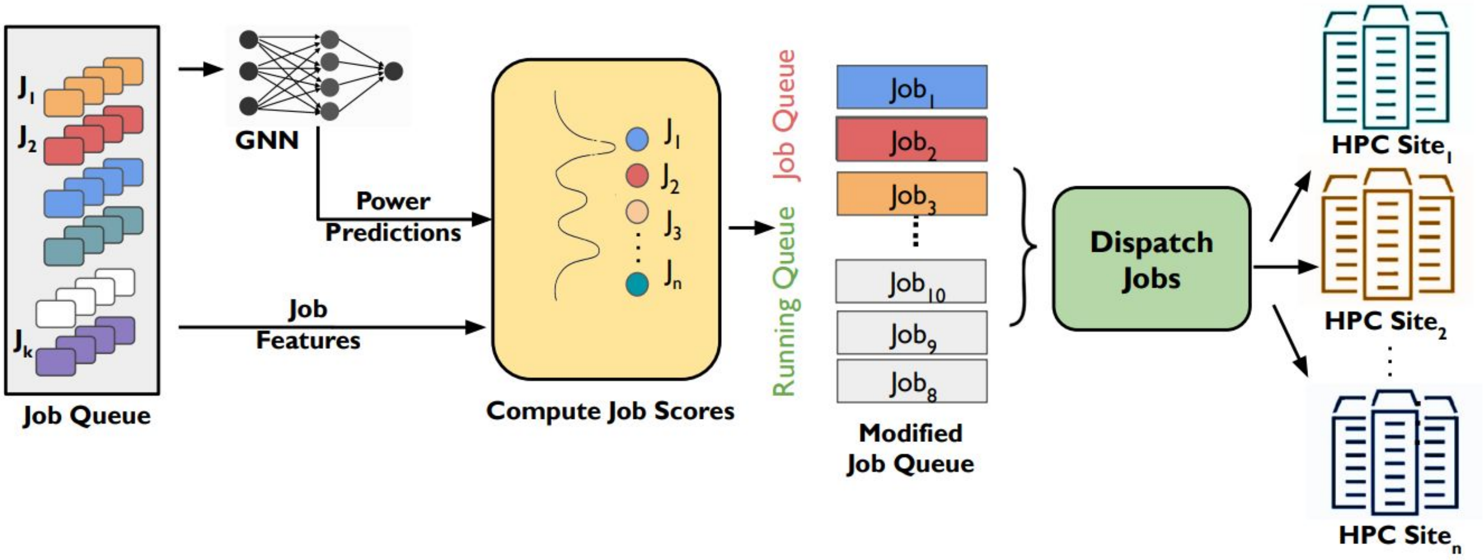

Accurate Power Prediction with GNNs

Key Details:

- **Inputs:** Node count, cores/task, memory, runtime, job type, etc.
- **Graph:** Built via k-Nearest Neighbors on job features.
- **Architecture:** Embedding + GCNConv layers + BatchNorm + ReLU + Residual connections.

Performance:

- Prediction errors **normally distributed around zero**.
- Median error ranges from ~30W (low power) to 120W (high power) — accurate across job sizes.

Component	**Power Prediction GNN**
Input	$[batch\_size, 8]$
Embedding Layer	Linear($8 \to 128$) + BatchNorm + ReLU + Dropout
GCN Layer 1	GCNConv($128 \to 128$) + BatchNorm + ReLU
GCN Layer 2	GCNConv($128 \to 128$) + BatchNorm + ReLU + Residual
FC Layer 1	Linear($128 \to 64$) + ReLU + Dropout
Output Layer	Linear($64 \to 1$)
Trainable Parameters	**43,265**

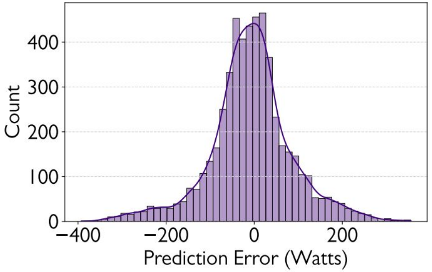

(a) Overall Error distribution.

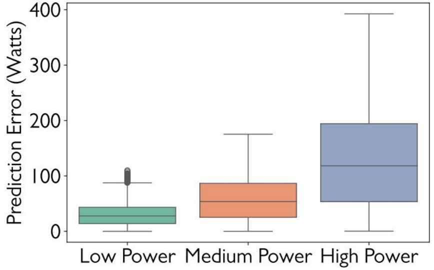

(b) Across three job categories.

Evaluation Methodology

- **Job Trace:** PM100 dataset (~230K jobs)
- **GNN Training:** 30% of data (temporal split)
- **Scheduling Evaluation:** Remaining 70% under 3 workload scenarios (Low, Avg, High)

Dynamic Pricing Simulation:

- 3 HPC sites (A, B, C) with different time zones and peak/off-peak rates (3x difference, e.g., $0.12 / $0.36 per kWh)
- Power budgets varied from 25% to 100% of historical peak power

Baselines Compared: FCFS, SJF, Backfilling, Random (multi-site)

Temporal Optimization: Cost Savings

- **Electricity Cost Reduction:** Up to 18% in temporal-only scenarios
- **Consistent Outperformance:** Across power budgets and workload months
- **Wait Times:** Slight, controlled increase under tight power limits
- **Trade-off:** Effective balance between cost savings and job turnaround

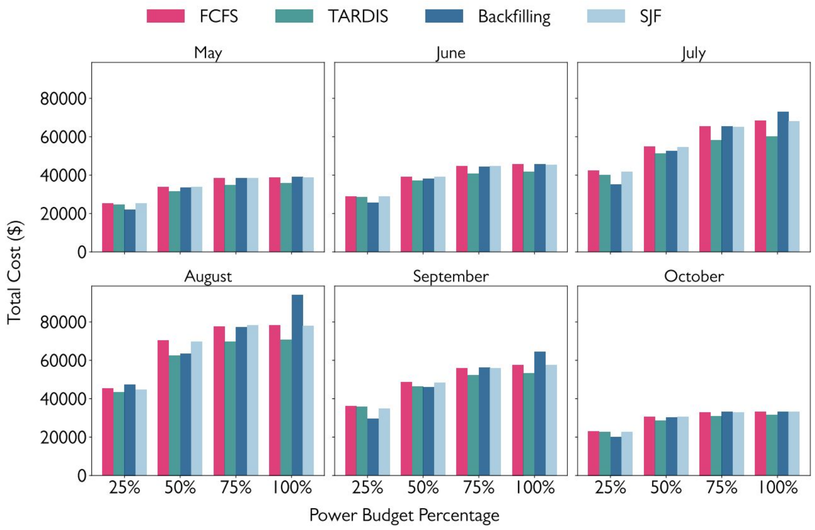

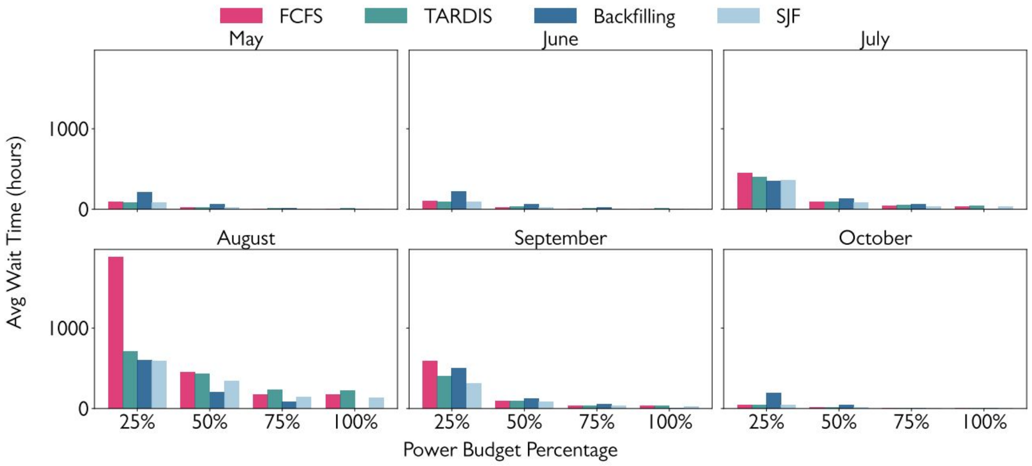

Multi-Site Optimization: Enhanced Savings

- **Peak hours:** Runs ~18% high-power jobs vs. 22–25% in baselines
- **Off-peak hours:** Runs ~70% high-power jobs vs. 45–50% in baselines
- Maintains balanced execution of low/medium power jobs throughout

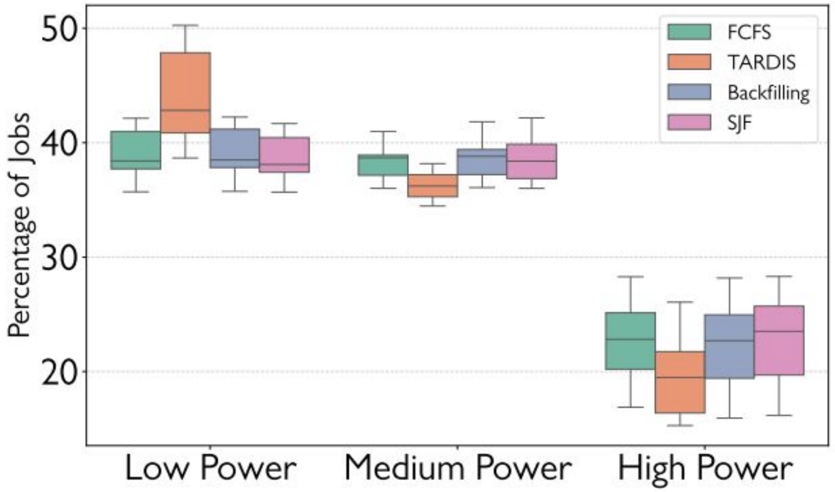

(a) Peak hours.

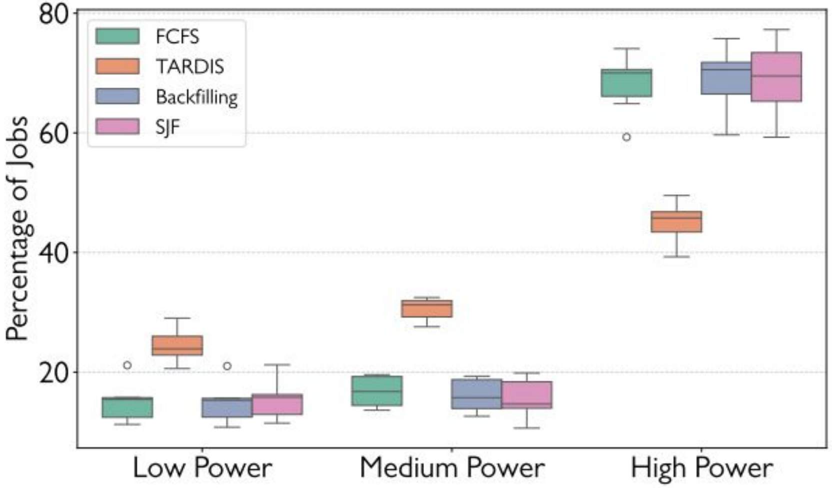

(b) Off-peak hours.

Mechanism: Strategic Peak Shifting

- **TARDIS:** $100–$650 daily
- **Single-site schedulers:** $200–$2,200 daily
- **Random assignment:** Often > $1,800 daily

Aggregate Savings vs. Single-Site:

- 10–15% lower total electricity cost
- 55% workload during peak hours (vs. 65–70%)
- Lower cost per job (~$5.5 vs. $6.5–$7.0)

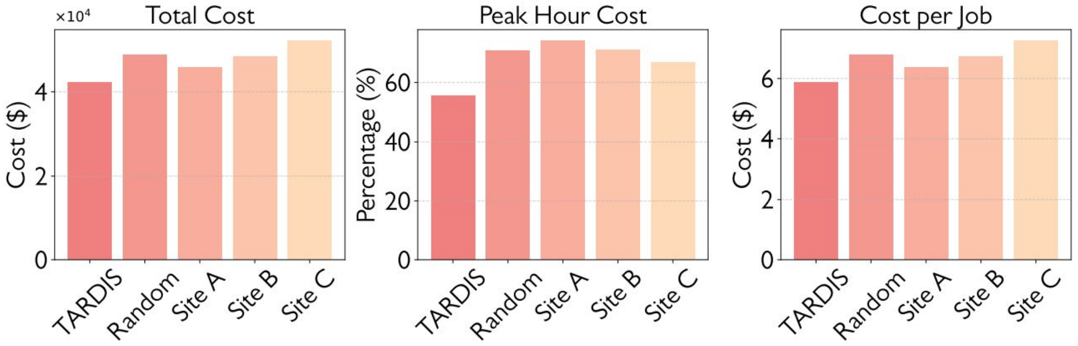

Multi-Site Dynamics & Utilization

Dynamically **shifts jobs across sites** to exploit non-overlapping peak hours

- E.g., **Site C picks up load** when A & B are in peak, and vice versa

System Utilization:

- Lower average (~30%) vs. single-site (80–95%)
- Similar peak utilization (~80%)

<u>Takeaway</u>: This **"deliberate underutilization"** allows flexible, cost-aware scheduling **without sacrificing throughput**

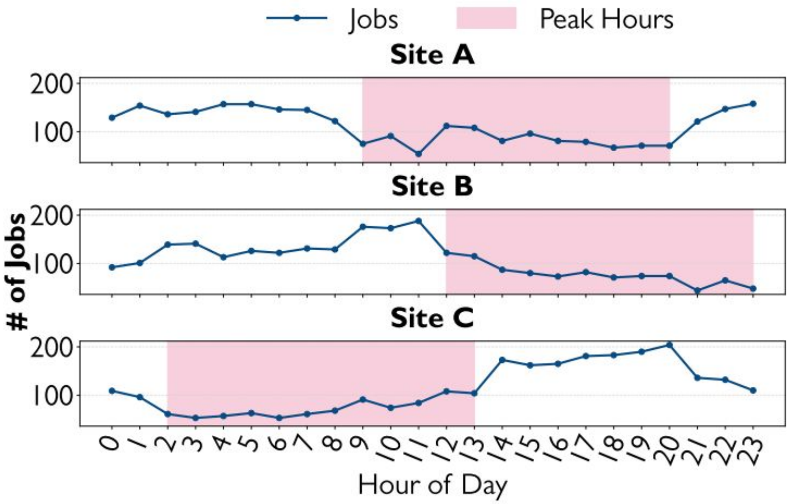

(a) Job distribution.

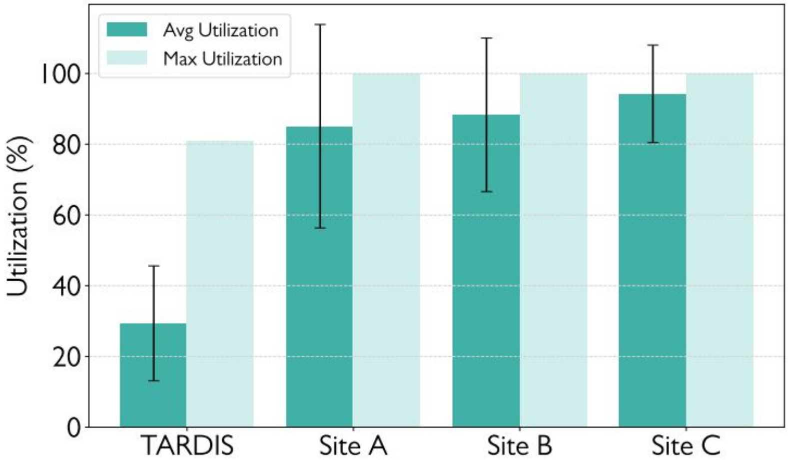

(b) System utilization.

Contributions & Conclusion

- **TARDIS**: GNN-based, spatio-temporal, power-aware scheduler

Key Contributions:

- Accurate **job-level power prediction** via GNN
- **Dynamic, price-aware scheduling** across time and sites
- **Multi-center extension** for cost-efficient workload distribution

Results:

- **Up to 18% cost savings (temporal)**, **10–20% (multi-site)** vs. baselines
- Achieved via **smart job shifting** using predicted power + dynamic pricing
- Maintains high throughput with minimal wait-time trade-offs

Thank You!
Questions?

Contact: abrar.hossain@utoledo.edu