SfS^2 (SECURE for Student Success)
First Annual Undergraduate Research Symposium

SfS^2
SECURE for Student Success

Orchard Conference Center
September 11, 2024

SfS2 SECURE for Student Success

Machine Learning-based Energy Prediction for Workload Management in Datacenters

Matthew Smith, Brandon Ismalej

Advisor: Dr. Xunfei Jiang

Department of Computer Science

California State University, Northridge

Outline

- Introduction
- Techniques/Methodologies
- Data/Observations
- Results
- Conclusions

Introduction

- Data centers consume 1% of world-wide energy, growing
  - Is there a way we can limit the energy consumed by data centers?

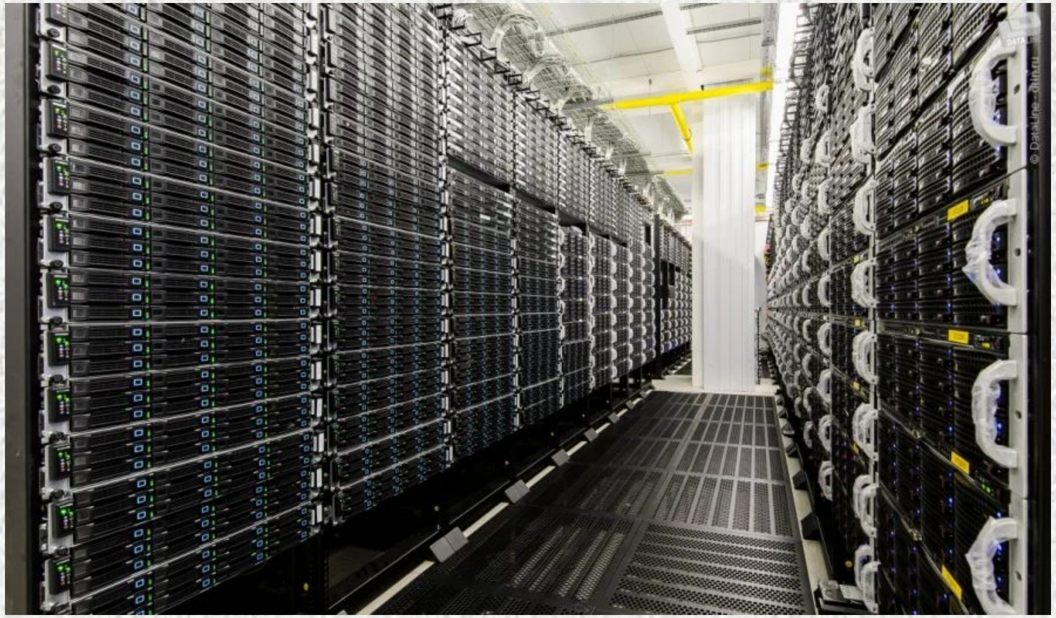

Figure 1: Image of a data center [1].

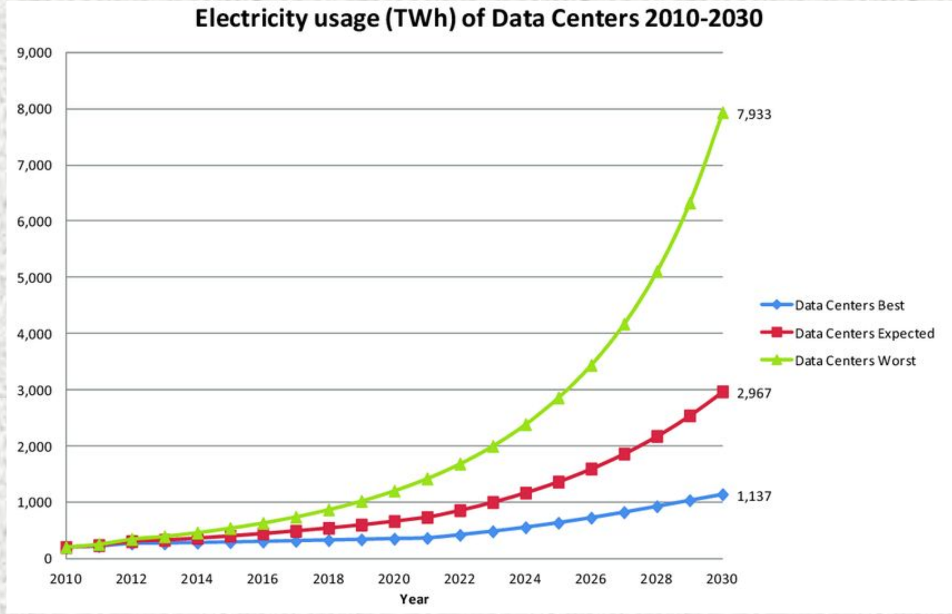

Figure 2: Global electricity demand of data centers for 2010-2030 [2].

Introduction: Load Balancing

- **Load balancing:** how a network distributes tasks to servers

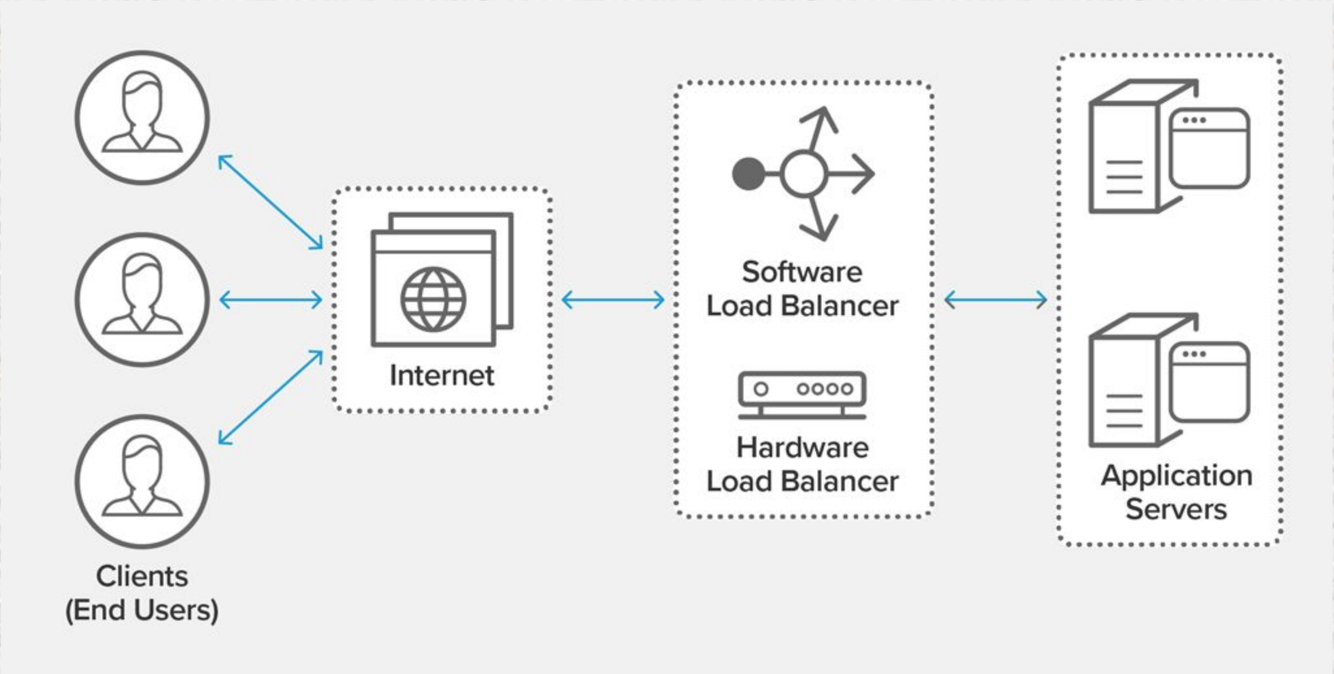

Figure 3: The load balancing process from client to server [3].

SfS2 SECURE for Student Success

Introduction: Virtual Machines and Tasks

- **Virtual machines (VMs):** portion of hardware dedicated to a task
  - Maximize number of programs running in each server

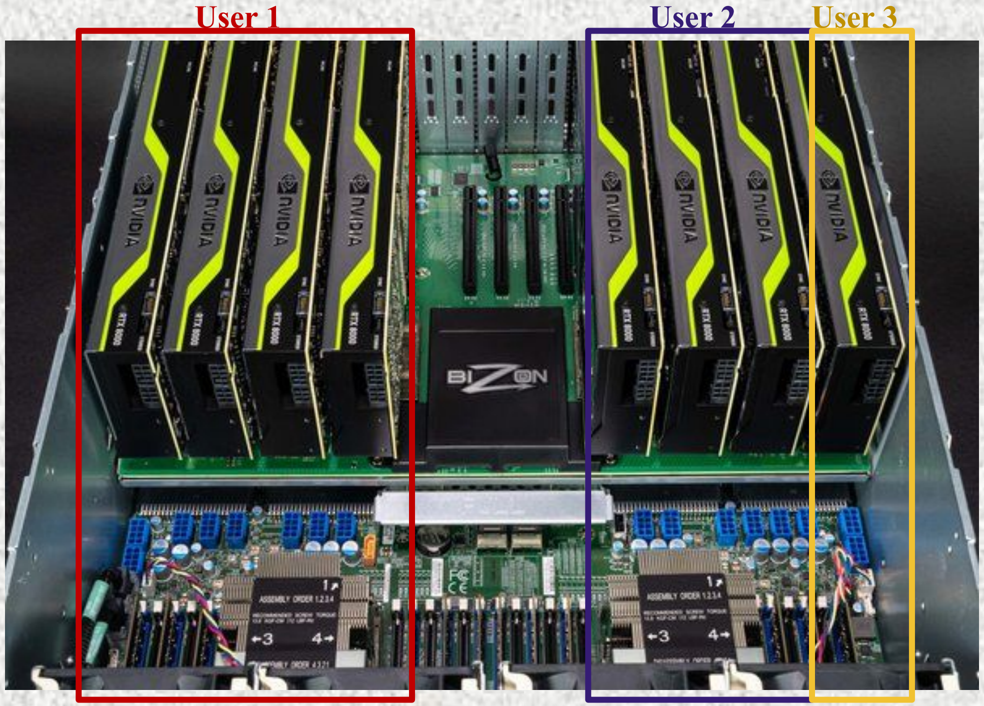

Figure 4: A single server with 8 GPUs and high number of CPUs and RAM [4].

Introduction: Project Objectives

- Part 1: GPU Energy Modeling*
  - Synthesize GPU data that mimics characteristics of real-world workload traces
  - To develop accurate GPU power prediction models to advance the study of thermal-aware workload management
- Part 2: Thermal-Aware Load Balancing
  - Evaluate thermal aware load balancing algorithms in simulations
  - Integrate GPU power prediction models into the simulator
  - ** No way to model GPU power in simulation by default*

Part 1: GPU Energy Modeling

- Objectives:
  - Synthesize GPU data that mimics characteristics of real-world workload traces
  - To develop accurate GPU power prediction models to advance the study of thermal-aware workload management

Emulating Real-World GPU Data

- Statistical analysis of 2 real-world workload traces:
  - Alibaba v2020 GPU Cluster Trace
  - SenseTime Helios GPU Cluster Trace
  - Focus:
    - Inter-task delay times ≥ 0 secs.
    - Tasks running sequentially on individual GPUs
- Goal:
  - To extract inter-task delay times of real-world data
  - To shape the design of our experiments and data synthesis

Emulating Real-World GPU Data

- Findings from analysis:
  - Distribution of delay times are heavily skewed
  - Most delay times were zero seconds
  - We focused on the 95th percentile to shape our experiments

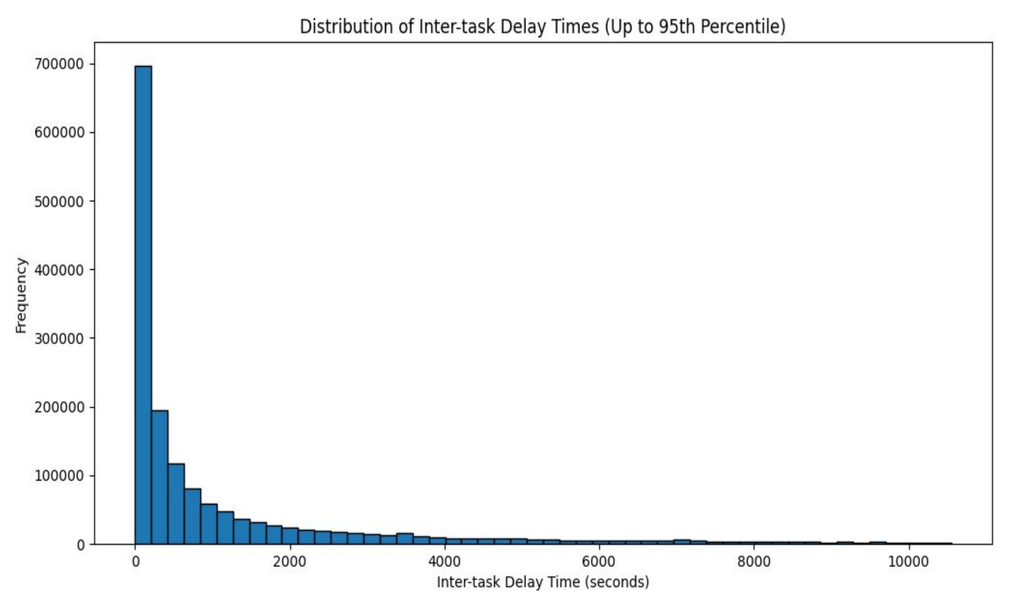

Figure 5: Distribution of Alibaba traces (up to 95th percentile)

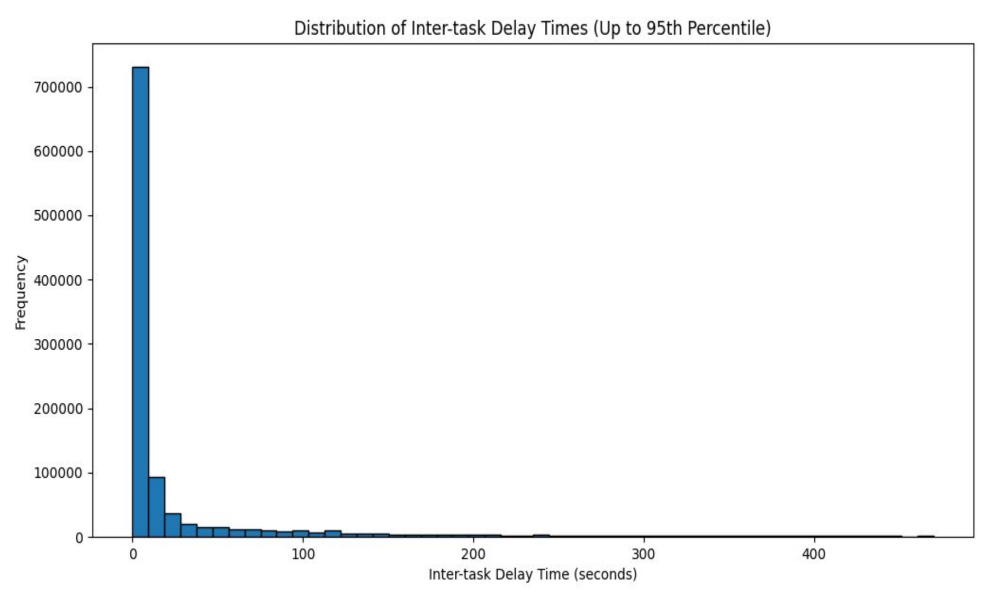

Figure 6: Distribution of Helios traces (up to 95th percentile)

Designing Experiments based on Analysis

- A set of 14 GPU-intensive benchmarks:

Benchmark Application	Domain	Avg. Power [W]	Avg. GPU Util [%]	GRAM Avg. [GiB]
2D_CONVOLUTION	Image Processing	59.05	98.10	0.60
BERT_QA_PREDICTION	Natural Language Processing	60.73	95.05	1.29
BLACKSCHOLES	Computational Finance	22.86	8.83	0.16
DISTILBERT_TRAINING	Natural Language Processing	57.87	98.62	6.47
EUCLIDEAN_DISTANCE	Mathematical Computation	51.08	99.33	1.64
FFT	Signal Processing	55.19	99.45	4.61
GAN_MNIST_DIGITS	Generative Adversarial Network	60.77	97.96	0.79
IMAGE_CLASSIFIER	Image Classification	23.09	6.74	0.29
K_MEANS	Unsupervised Learning	22.86	17.46	0.41
MATMUL	Matrix Multiplication	27.85	13.06	1.36
MONTE_CARLO_PRICING	Financial Simulation	43.01	98.31	3.48
SEPIA_TRANSFORMATION	Image Processing	23.10	5.94	0.26
SPECTROGRAM_TRANSFORMATION	Audio Processing	23.65	5.18	0.18
SIMPLE_CNN	Image Recognition	42.30	97.21	2.02

Table 1: Domain of GPU Benchmark Applications

Designing Experiments based on Analysis

- Designed five distinct experiments to cover a wide range of inter-task delay times, using our benchmarks
  - “No delays” simulates peak usage
  - Short delays simulate high-frequency task switching
  - Longer delays reflect idle periods, like downtime

  - Experiments yielded ~40 hours of data
  - Data was collected second-by-second

CPU	2 * Intel(R) Xeon(R) CPU E5-2690 v2 @ 3.00GHz (10 cores)
Memory	98,304 MB
Disk	2 * 480 GB
GPU	NVIDIA Tesla P4 8GB GDDR5
OS	Ubuntu 22.04.2 LTS

Table 2: Cluster System Configuration

	Delay Times (seconds)	Task Order
Exp 1	No delays	Sequential, reverse, shuffled
Exp 2	0 - 10	shuffled
Exp 3	1 - 20	Sequential, reverse, shuffled
Exp 4	1 - 30	shuffled
Exp 5	300 - 1000	shuffled

Table 3: Experiment Configurations

SfS² SECURE for Student Success

Preprocessing Synthesized Data

- Features extracted for duration of every task-active and idle period
  - To align with available features of real-world workloads that would be implemented in cloud simulator

timestamp	gpu_temp	gpu_power	gpu_GRAM_util	gpu_util
2024-06-19 21:12:14	36	6.74	0	0
2024-06-19 21:12:15	36	6.64	0	0
2024-06-19 21:12:16	36	6.64	0	0
2024-06-19 21:12:17	36	6.55	0	0
2024-06-19 21:12:18	36	6.64	0	0
2024-06-19 21:12:19	35	6.64	0.1220703125	0
2024-06-19 21:12:20	36	22.71	5.60302734375	1
2024-06-19 21:12:21	36	54.98	75.47607421875	36
2024-06-19 21:12:22	38	59.76	82.80029296875	96
2024-06-19 21:12:23	39	43.94	85.14404296875	100
2024-06-19 21:12:24	39	65.01	85.14404296875	100
2024-06-19 21:12:25	40	49.21	85.14404296875	100
2024-06-19 21:12:26	41	50.96	85.14404296875	100
2024-06-19 21:12:28	41	65.26	85.14404296875	100
2024-06-19 21:12:29	42	65.46	85.14404296875	100
2024-06-19 21:12:30	43	61.72	85.14404296875	100
2024-06-19 21:12:31	43	60.59	85.14404296875	100
2024-06-19 21:12:32	43	61.27	85.14404296875	100
2024-06-19 21:12:33	44	X	85.14404296875	100

Start of GPU Task
GPU GRAM Avg.
GPU GRAM Max.
GPU Power. Avg.
GPU Util. Avg.
Prediction Value

Figure 7: Illustration of Data Preprocessing

GRAM Avg. & Max. [GiB] ; Util. Avg. [%] ; Power Avg. [W]

Modelling GPU Power

- Four machine learning algorithms were utilized
  - XGBoost (eXtreme Gradient Boosting)
  - CatBoost (Categorical Boosting)
  - LightGBM
  - LSTM (Long Short-Term Memory)

- Utilized a 60/20/20 split of data for:
  - Training/Testing/Validation

Modelling GPU Power

- To predict GPU power consumption:
  - Training input:
    - GPU GRAM Avg. [GiB]
    - GPU GRAM Max. [GiB]
    - GPU Utilization Avg. [%]
  - Training output:
    - GPU Power Avg. [W]

- Initial ML model training

	CatBoost	LightGBM	XGBoost	LSTM
RMSE Value	1.218	1.224	1.284	1.520

Table 4: RMSE Values of Initial Model Training

- Grid-search hyperparameter tuning performed to improve RMSE

Modelling GPU Power

- Best model - **XGBoost**
  - RMSE = 1.217
  - On average, predictions for GPU power consumption are off by about 1.217 watts from actual measured values

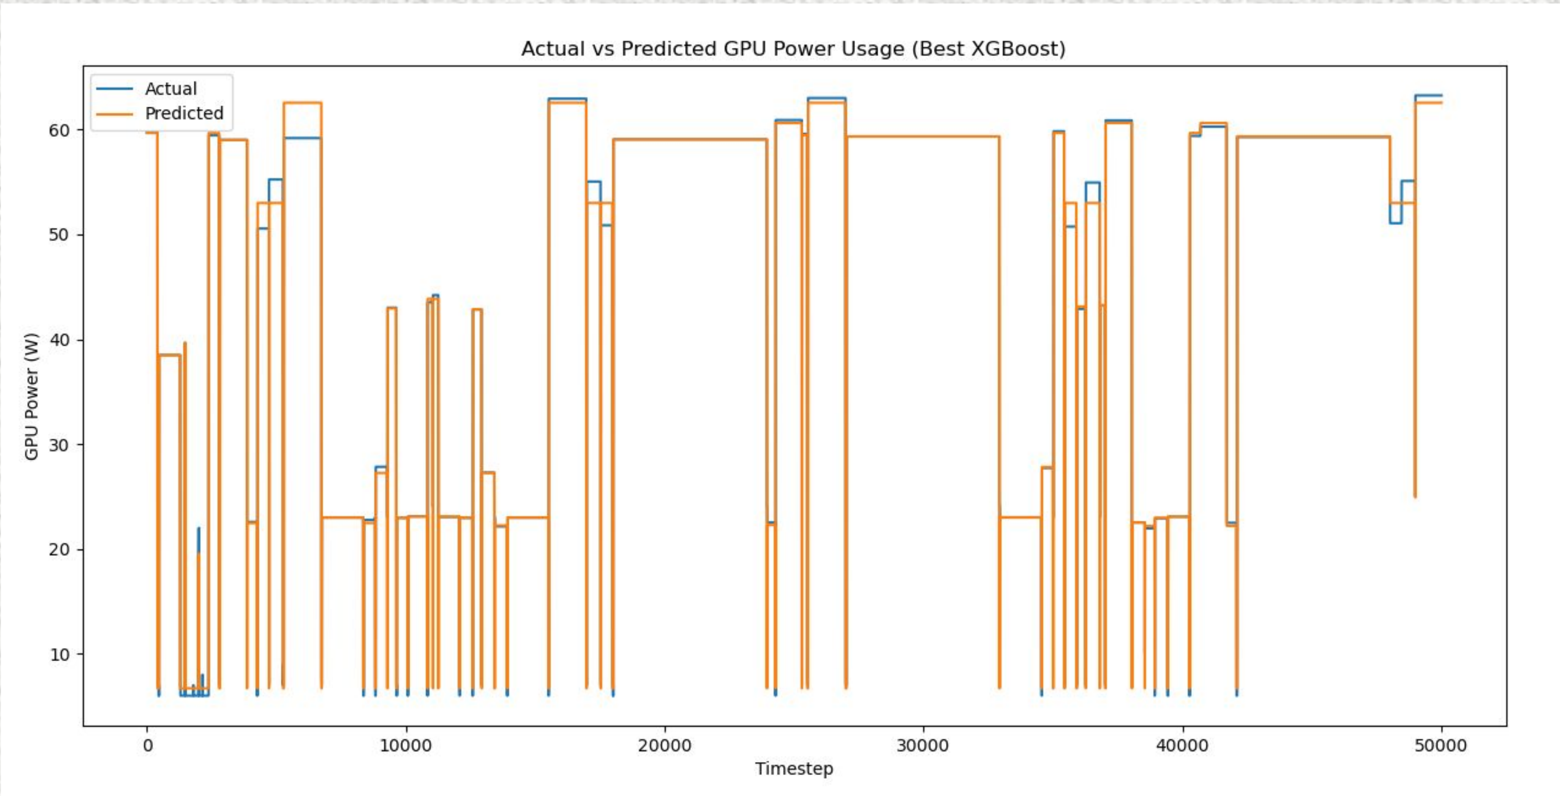

Figure 8: Best XGBoost Model Predicted vs. Actual GPU Power Consumption

Part 2: Thermal-Aware Load Balancing

- Objectives:
  - Evaluate thermal aware load balancing algorithms in simulations
  - Integrate GPU power prediction models into the simulator

Researching Workload Traces

- Investigate workload traces of data centers using GPUs
  - HeliosData, **Alibaba Cluster v2020**, Alibaba Cluster v2023

Figure 9: An Alibaba data center [5].

Processing the Alibaba 2020 Workload Trace

- Alibaba v2020 (5 GBs)
  - We used 4 files: machine spec, task, instance, and sensor
- Filtered out tasks missing:
  - Planned VM resources
  - Sensor, server information
  - Task duration
- Over 66 days
  - 497 hosts, 325k VMs/tasks

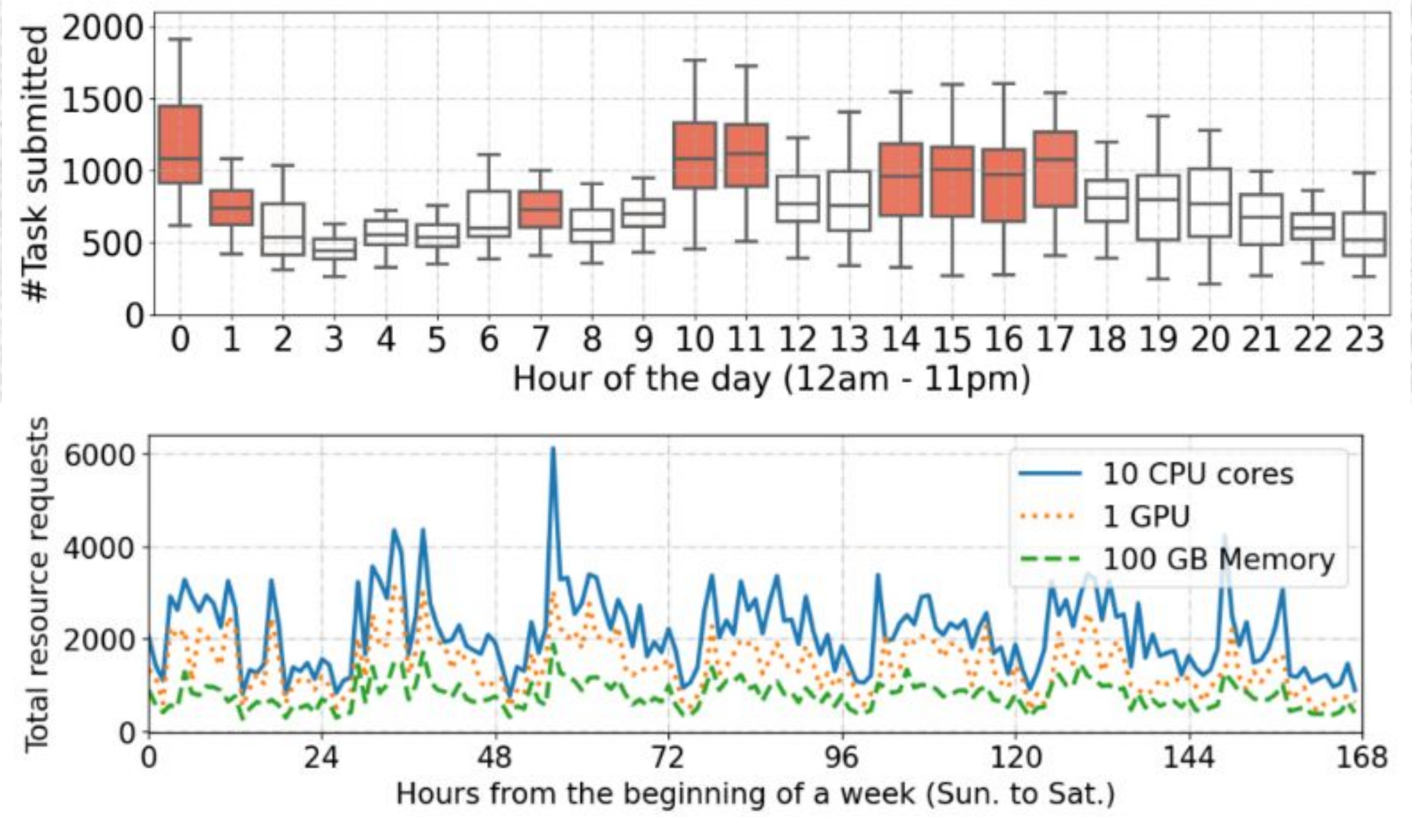

Figure 10: Number of tasks submitted over a day and total resource requests over a week [6].

Modify our Simulations

1. Initialize simulations using Alibaba 2020
2. Update config and output files
3. Test simulation is working correctly
4. Integrate new energy models
5. Rewrite load balancing algorithms
6. Run full duration simulations

SfS2 SECURE for Student Success

Results: Modifying our Simulations

- Initialize simulations
- Update config and output files
- Test simulation:
  - Problem: 56 days to run
  - Optimizations (20x faster!)
    - ML batch prediction
    - Caching calculations
  - Add checkpointing

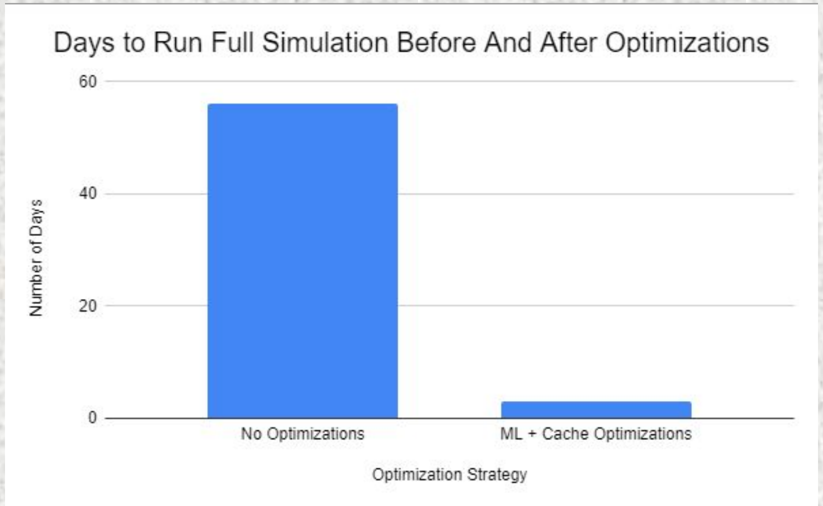

Figure 11: Optimization strategies used to reduce the simulation runtime.

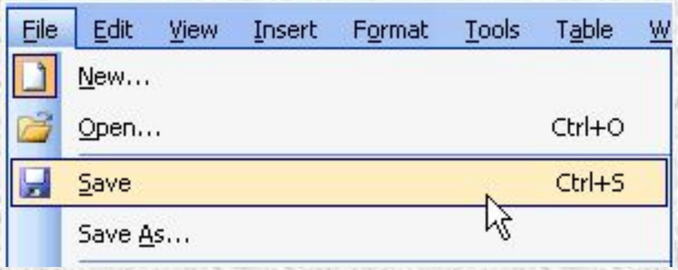

Figure 12: Checkpoint strategy for saving simulation progress midway through [7].

SfS2 SECURE for Student Success

Part 2: Future Work

- Fixing simulation:
  - Trigger events in simulator

- Evaluating energy:
  - Integrate new GPU power model
  - Model energy for servers using multi-CPUs and multi-GPUs
  - Rewrite load balancing algorithms
  - Run full duration simulations

Conclusion

- We learned:
  - Performing statistical analyses from noisy datasets
  - Developing machine learning models to solve a problem
  - Load balancing and applications to job scheduling
  - Efficiency of using ML batch prediction

SECURE for Student Success

Acknowledgement

- Supported by the SfS^2 Program and funded by the United States Department of Education FY 2023 Title V, Part A, Developing Hispanic-Serving Institutions Program five-year grant, Award Number P31S0230232, CFDA Number 84.031S.

- However, the contents of this presentation do not necessarily represent the policy of the US Department of Education, and you should not assume endorsement by the Federal Government.

Image Citations

[1] Online. From: https://songrgg.github.io/images/datacenter.webp
[2] On Global Electricity Usage of Communication Technology: Trends to 2030 - Scientific Figure on ResearchGate. Available from: https://www.researchgate.net/figure/Global-electricity-demand-of-data-centers-2010-2030_fig2_275653947
[3] Online. From: https://www.nginx.com/resources/glossary/load-balancing/
[4] Online. From: https://qph.cf2.quoracdn.net/main-qimg-6a23e1295bed5350309870109a808cdf-lq
[5] Online. From: https://github.com/alibaba/clusterdata/blob/master/cluster-trace-gpu-v2020/README.md
[6] Online. From: https://finance.yahoo.com/news/alibaba-close-data-centres-australia-093000629.html
[7] Online. From: https://edu.gcfglobal.org/en/word2003/save-and-save-as-/1/

Questions?