Analyzing the Power Consumption Behavior of a Large Scale Data Center

KASHIF NIZAM KHAN, AALTO UNIVERSITY, FINLAND.

SANJA S., TAPIO N., JUKKA K. N., SEBASTIAN V. A. & OLLI-PEKKA L.

Aalto University
School of Science

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Motivation

- Data center energy spending is ever increasing
  - System power draw is increasing substantially without a breakthrough in energy efficiency
  - Increased economic, social and environmental pressure to decrease the energy cost
  - Performance of future HPC systems will be constrained by power cost

Motivation

- Data center energy spending is ever increasing
- Data center power consumption log analysis is relatively less studied

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Contributions

- Investigate the impact/relation of OS counters and RAPL on total power consumption
- Analyse ‘unsuccessful’ jobs and their influence in energy spending
- Cluster the nodes based on the OS counter and RAPL values
- Model/estimate the total power consumption using OS counters and RAPL value.

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Dataset Description

- 900 nodes **Taito** computing cluster – 460 Sandybridge, 397 Haswell
- Approximately 2 days of production data captured in June 2016
- vmstat , RAPL, plug power and job info.
- Sampled at 0.5Hz

Type	*Haswell*	*Sandy Bridge*
Number of nodes	397	496
Node model	HP XL230a G9	HP SL230s G8
Cores / node	24	16
Memory / node	128GB	64GB

https://research.csc.fi/taito-supercluster

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Power Consumption of Computing Nodes

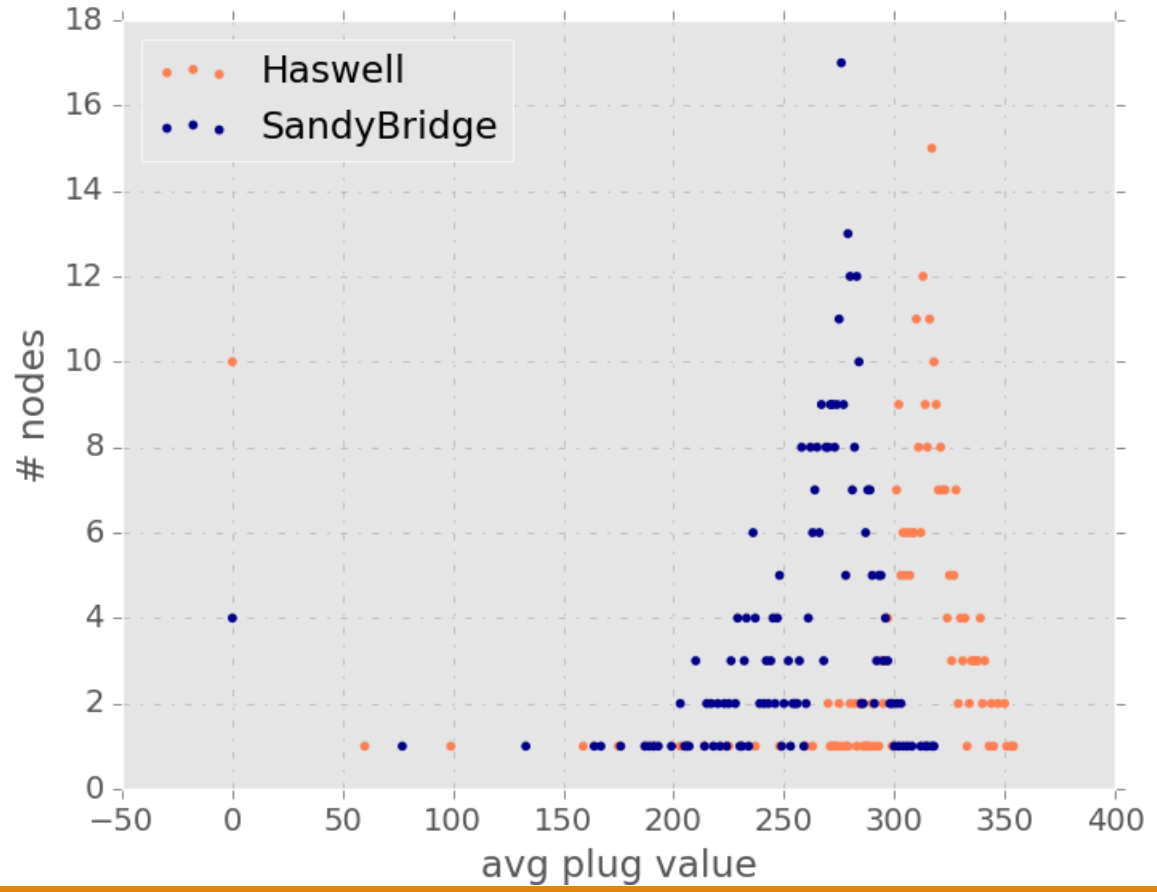

Power Consumption of Computing Nodes

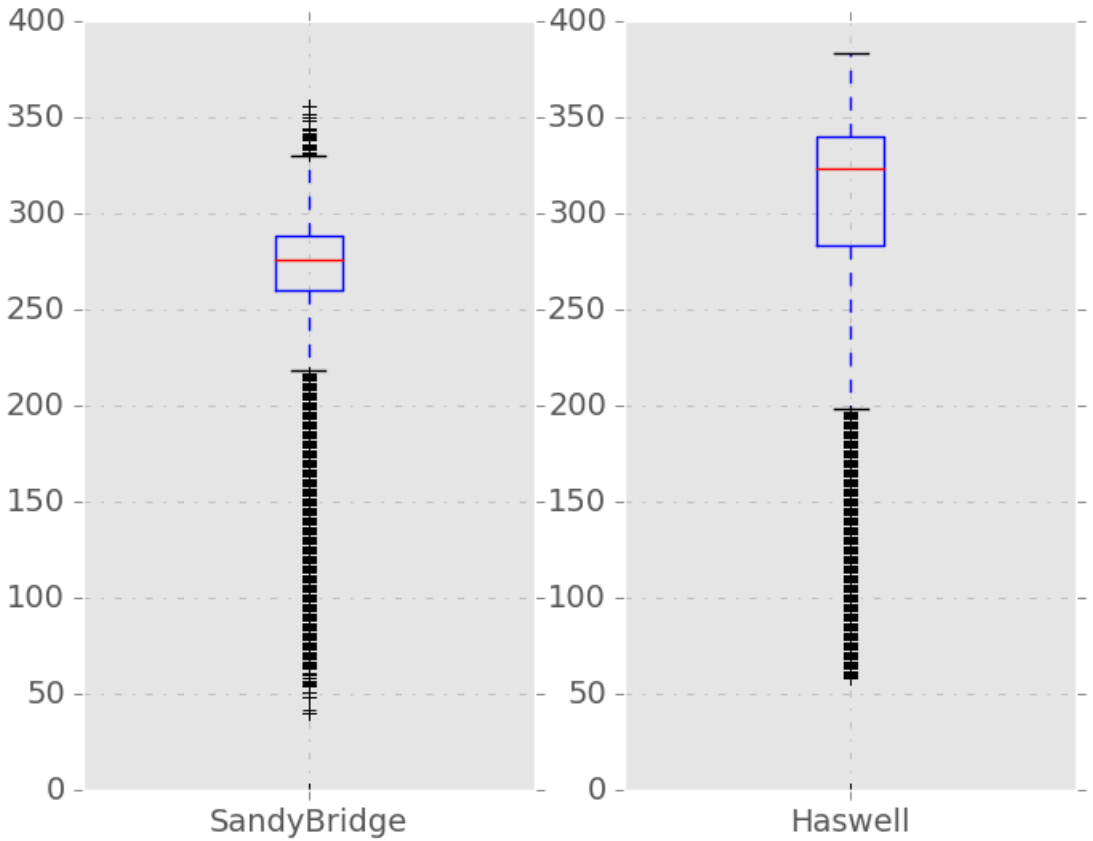

Power Consumption of Computing Nodes

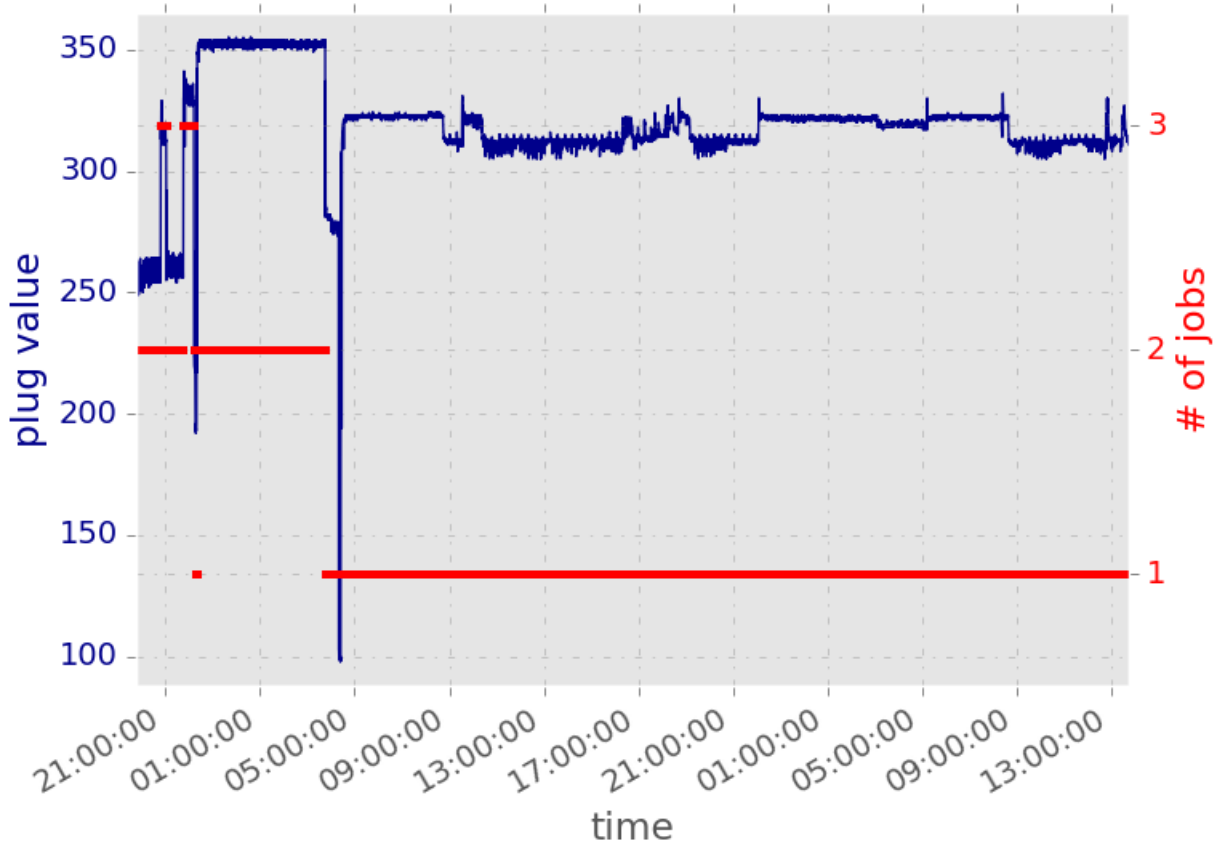

NODE C581

Power Consumption of Computing Nodes

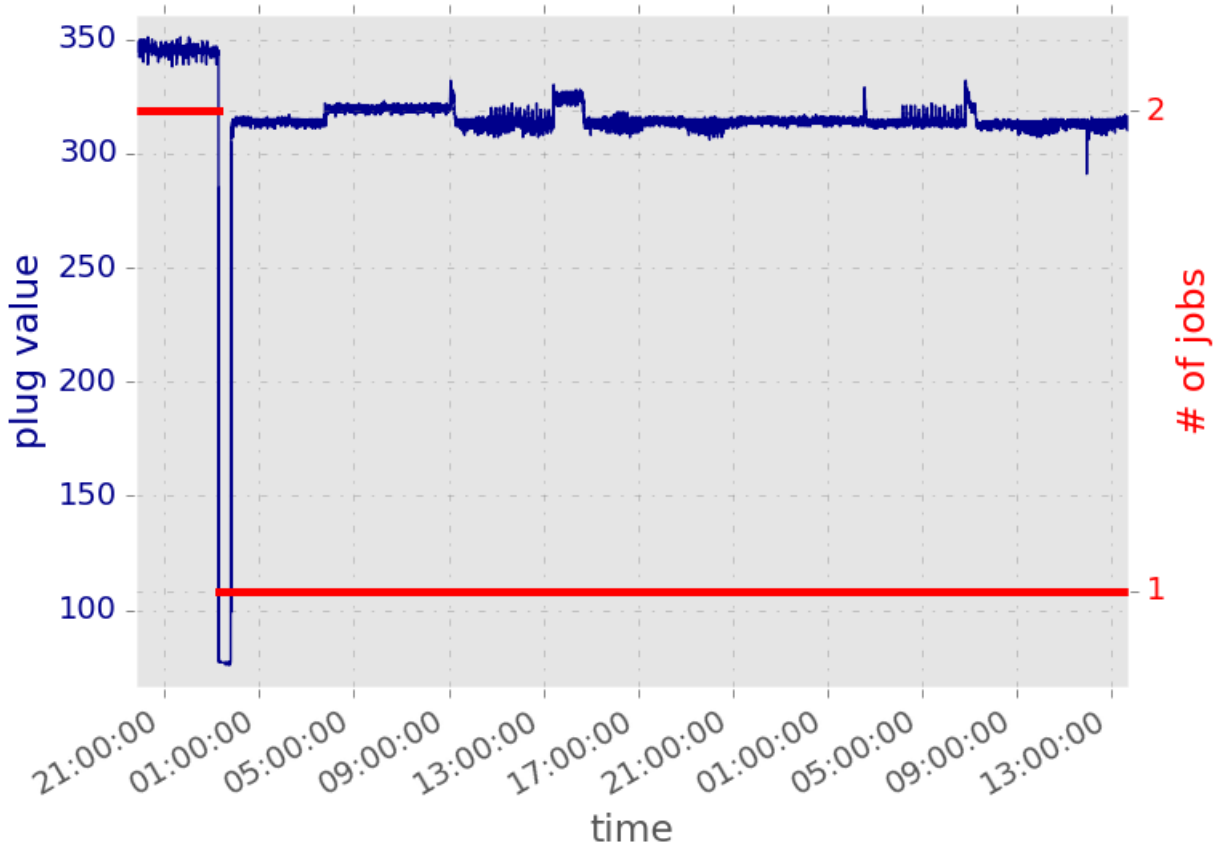

NODE C749

Power Consumption of Computing Nodes

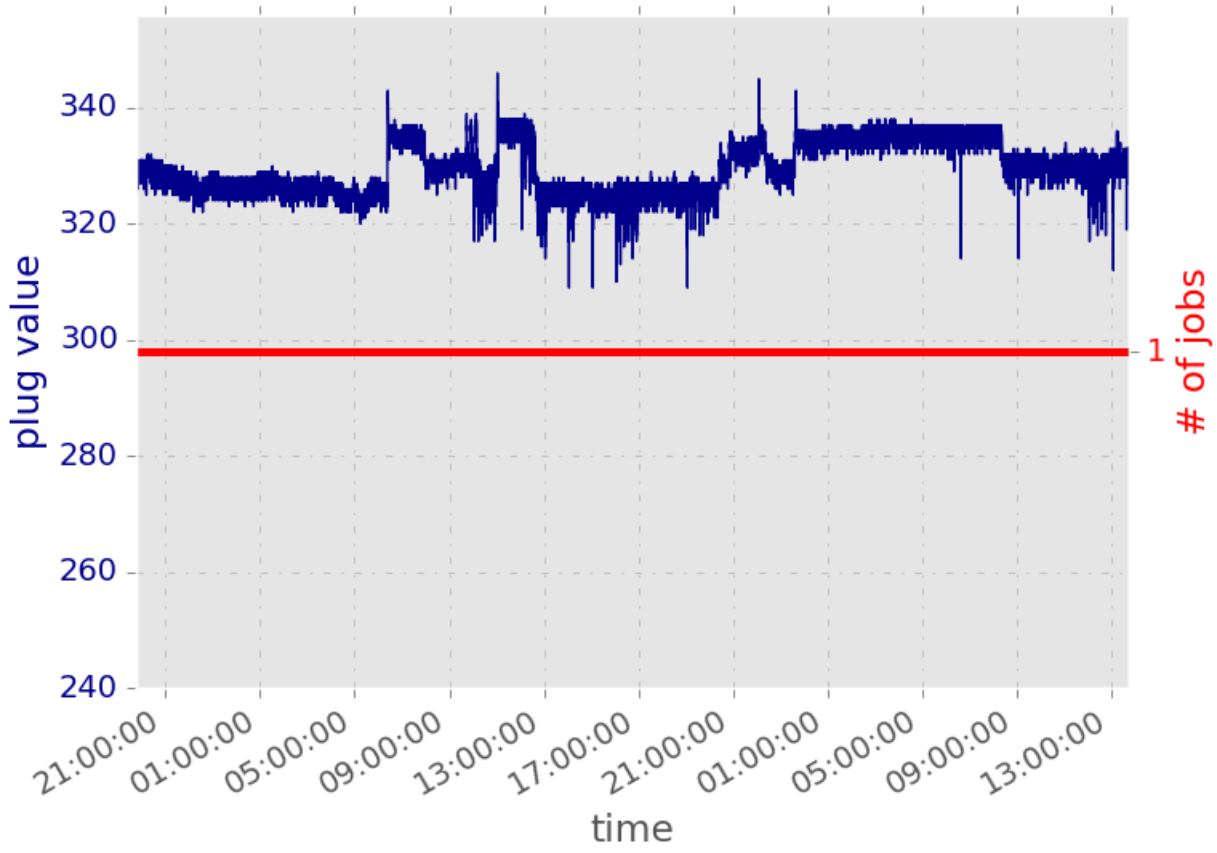

NODE C836

Power Consumption of Computing Nodes

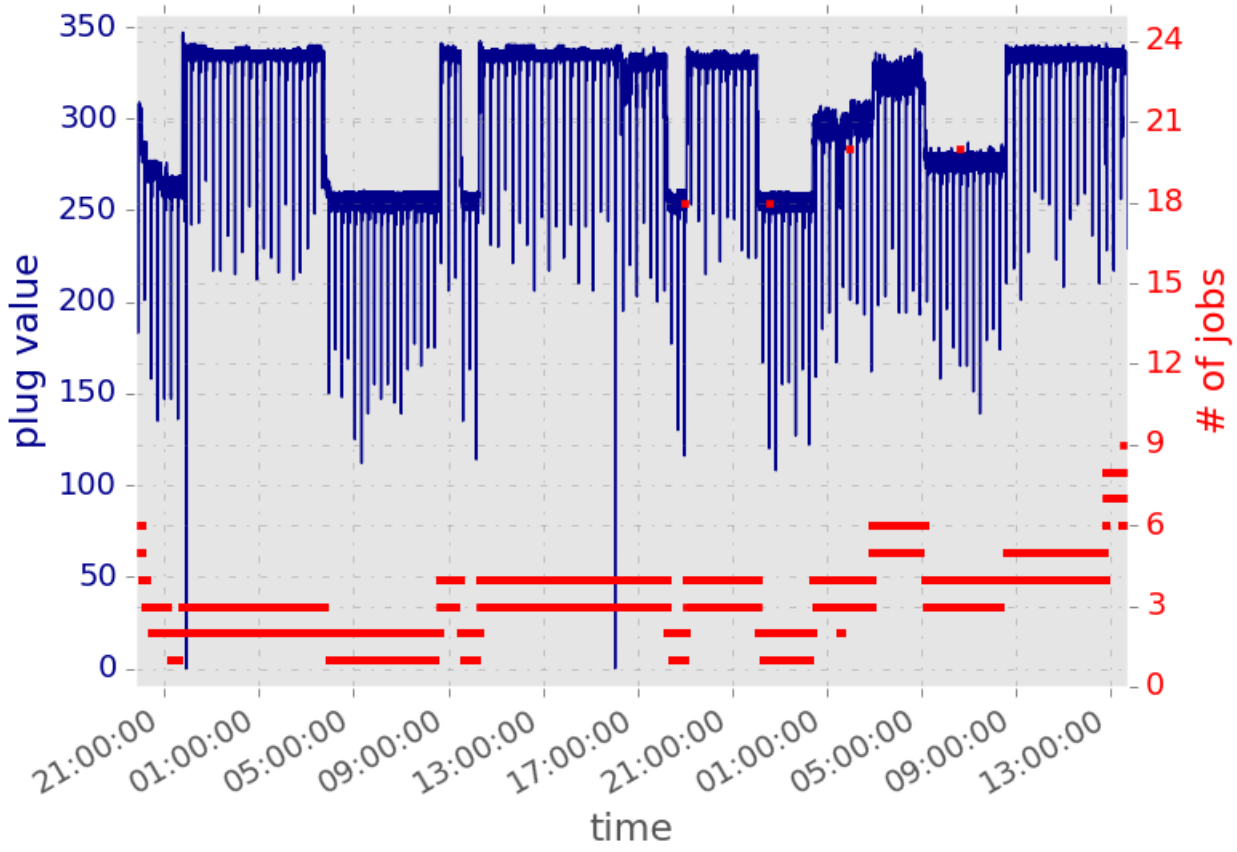

NODE C585

Power Consumption of Computing Nodes

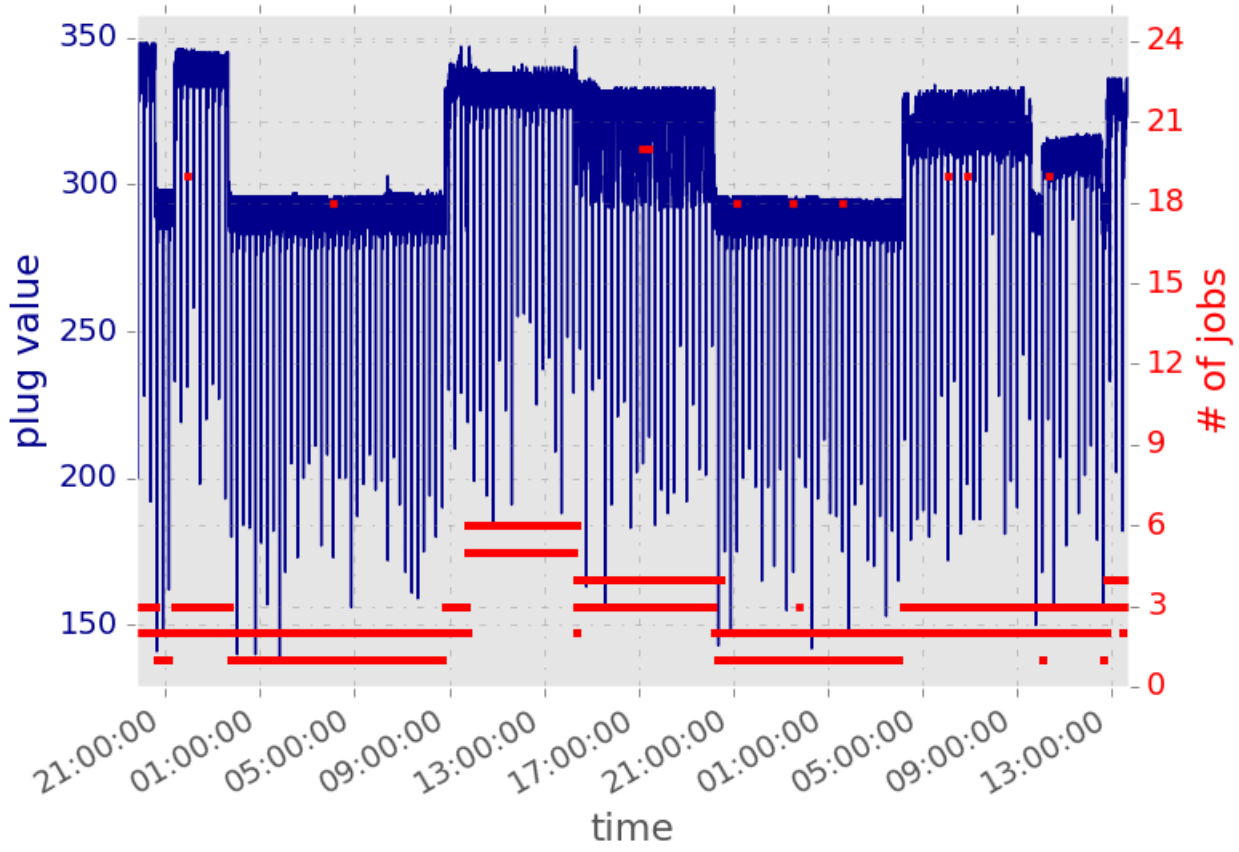

NODE C626

Power Consumption of Computing Nodes

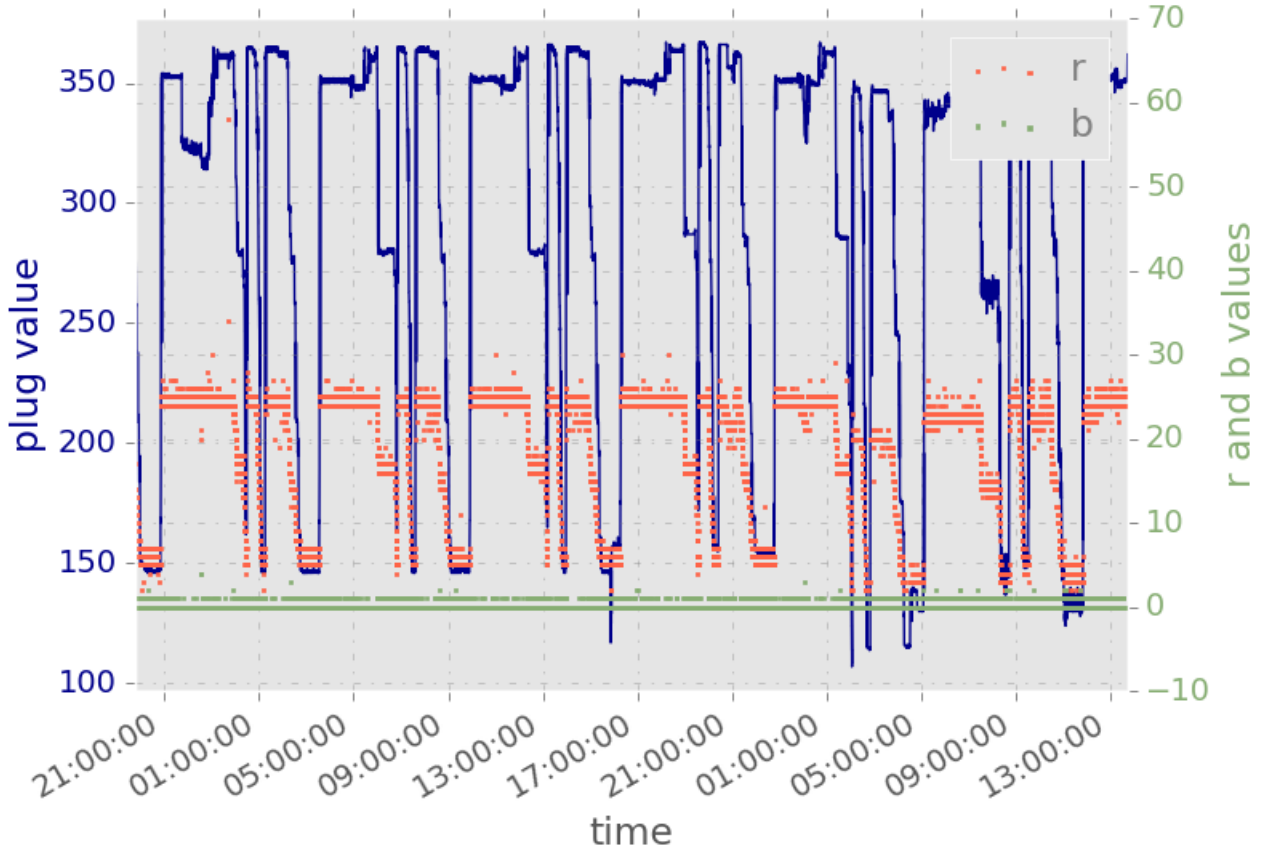

NODE C775

Power Consumption of Computing Nodes

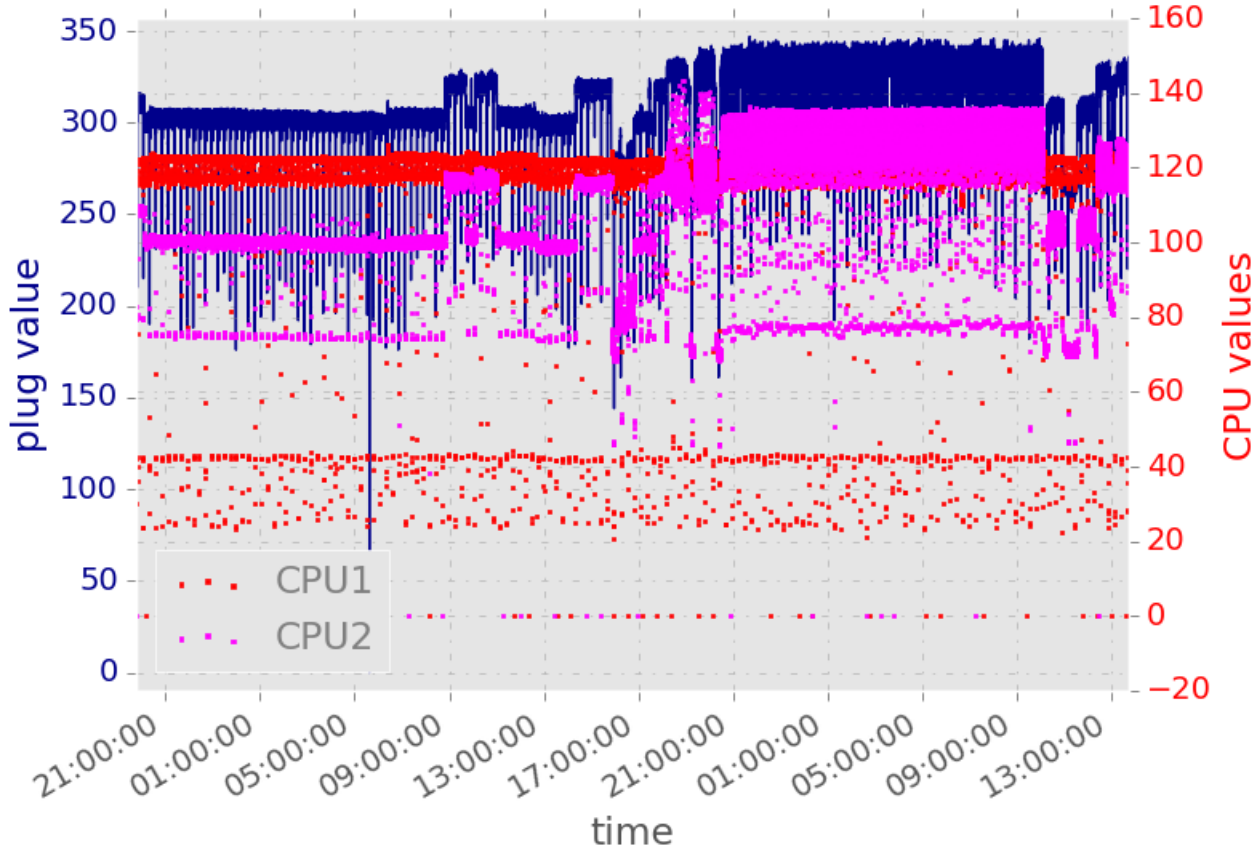

NODE C819

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Analysis of Unsuccessful Jobs

- **Completed** - jobs that ran to completion
- **Failed -** jobs that failed to complete successfully
- **Cancelled-** jobs that are cancelled by their users
- **Timeout-** jobs that did not run to successful completion within a given time limit.

Analysis of Unsuccessful Jobs

Job Status	*Nr. of Jobs (%)*	*Elapsed Time/Job (hrs)*	*CPU Time (%)*
Completed	84.0%	1.0	56.95%
Failed	12.5%	0.7	14.75%
Cancelled	3.0%	8.0	8.96%
Timeout	0.5%	25	19.34%

16% → **43.5%**

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Power Consumption Estimation

- Sample 2% of data from all the nodes (251,244 data samples)
- First 2/3rd of the data is used as historical data and train ML models
- Last 1/3rd of the data is used to validate
- Random Forest gives the best result

Correlation coefficient (corrcoef)	0.97
Mean absolute error (MAE)	3.12
Root mean squared error (RMSE)	9.11

Outline

- Motivation
- Contributions
- Dataset Description
- Power Consumption of Computing Nodes
- Analysis of Unsuccessful Jobs
- Power Consumption Estimation
- Plug Power Modeling

Plug Power Modeling

- Aim - Model the plug power using OS counters and RAPL measurements
- 30.000 measurements from 'Haswell' type computing nodes.

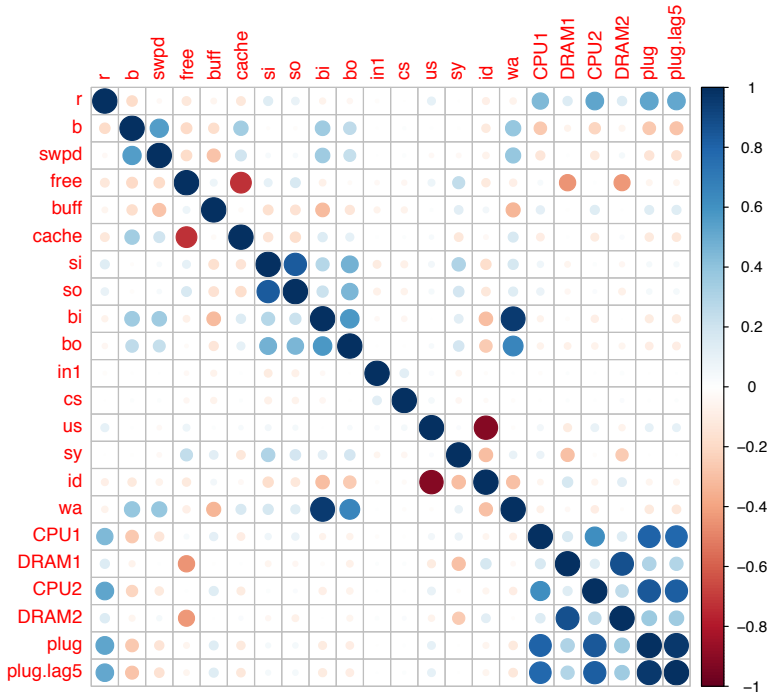

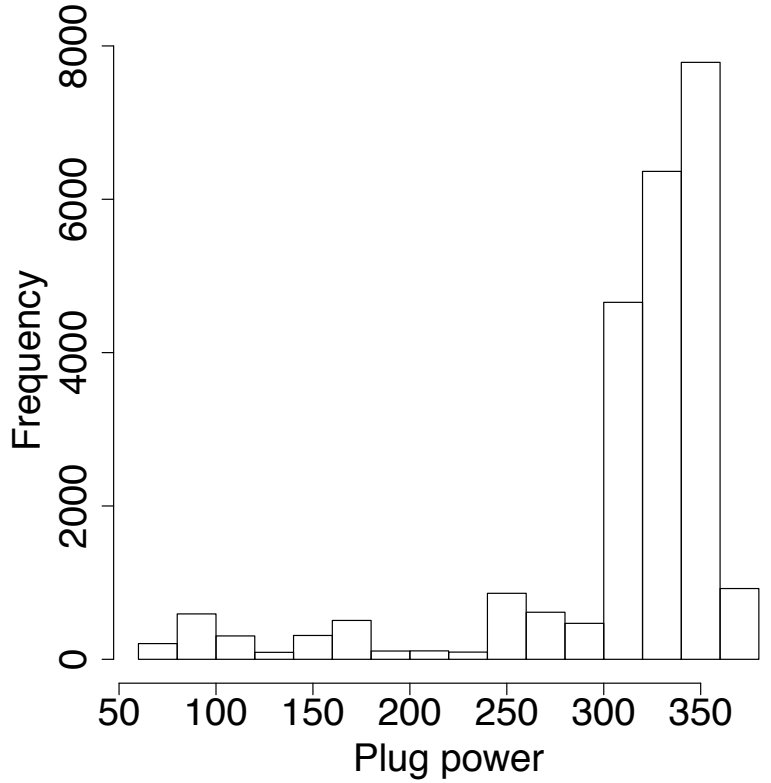

Plug Power Modeling

$$f(x) = a_0 + a_2 CPU1 + a_3 CPU2 + a_4 DRAM1 + a_5 DRAM2 + e$$

MAPE: 2.10%

Plug Power Modeling

$$g(u) = \beta_0 + f_1(x_1) + f_2(x_2) + ... + f_n(x_n) + e$$

MAPE: 1.97%

Clustering

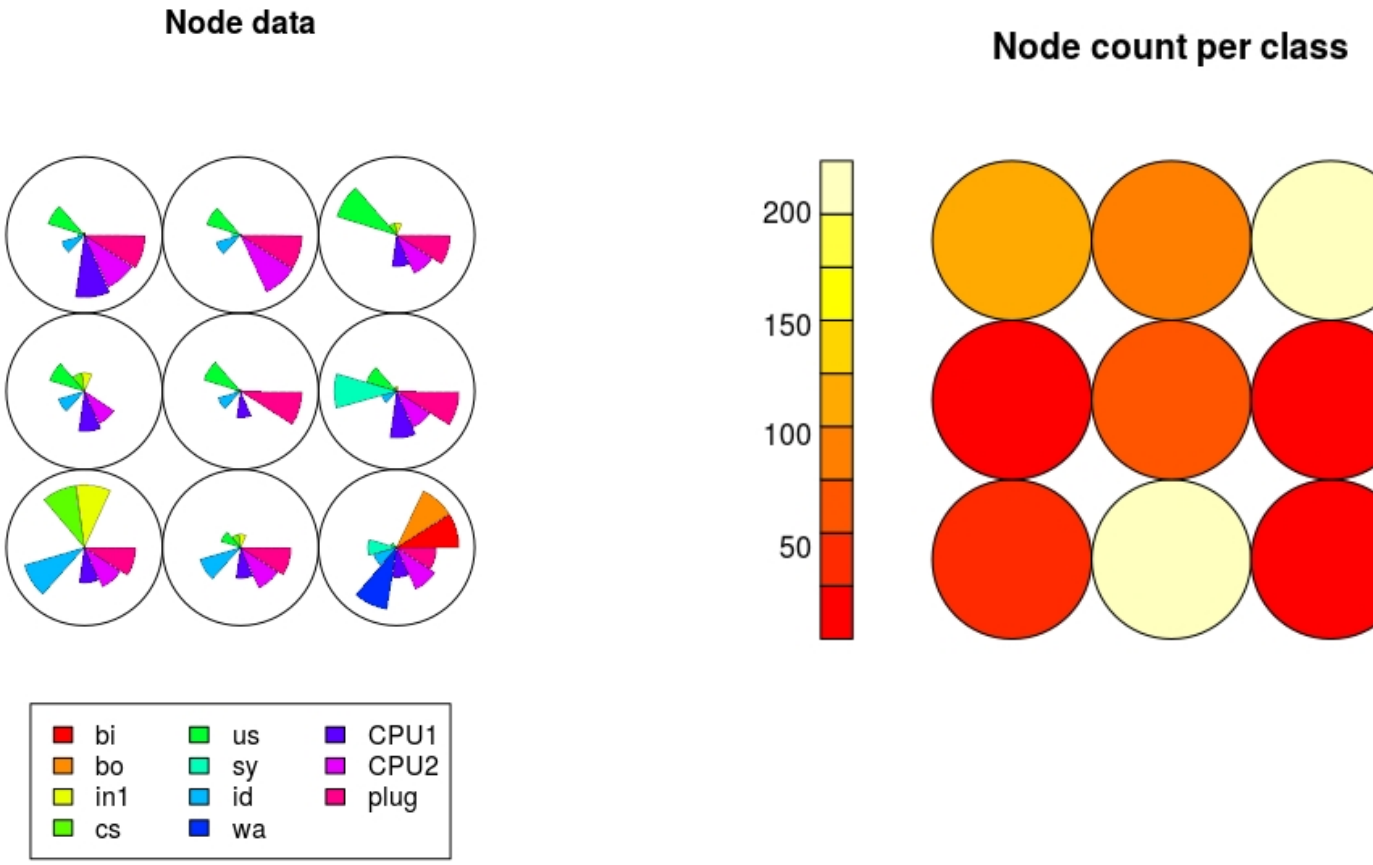
Node data
bi
bo
in1
cs
us
sy
id
wa
CPU1
CPU2
plug
Node count per class
200
150
100
50

Conclusion

- Estimating plug power from utilization metrics is promising
- RAPL add to the accuracy of the models by providing real time power consumption data
- Considering interactions among RAPL variables the error reduces to **1.87**%
- ‘Unsuccessful’ jobs can consume significant resources and power
- In future, we aim to utilize such data center logs to produce job specific power consumption models

Thank You!