SAND2020-2367C

Sandia National Laboratories

Machines Learning about Machines – ML for Analysis & Control of HPC Infrastructure

Felix Wang, Sam Green, Kevin Pedretti, Craig M. Vineyard, Andrew Younge - Sandia National Laboratories, Albuquerque, NM

HPC Infrastructure Challenge

- Motivation: the growing complexity of computing systems, ranging from cell phones to supercomputers is becoming difficult for developers to manage, and power utilization is expected to be a major component to regulate going forward
  - By 2025, data centers are forecasted to be using 20% of all available electricity in the world
  - A cloud provider used the equivalent energy consumption of ~366,000 U.S. households in 2014
  - As the world moves towards Exascale, this is not a scalable trajectory
- Goal: employ machine learning (ML) techniques to make more efficient use of HPC infrastructure
  - More intelligent, automated mechanisms are needed to optimize use of available resources, from the subsystem level (e.g. memory management) to the human-interaction level (e.g. job scheduling)
- Experiment: adjust relevant power management states (P-states) and node-level power caps such that compute time is not negatively impacted and power is saved

Figures: manually analyzed power-time profiles for example mini-apps demonstrating the potential for significant power savings with minimal run-time impact

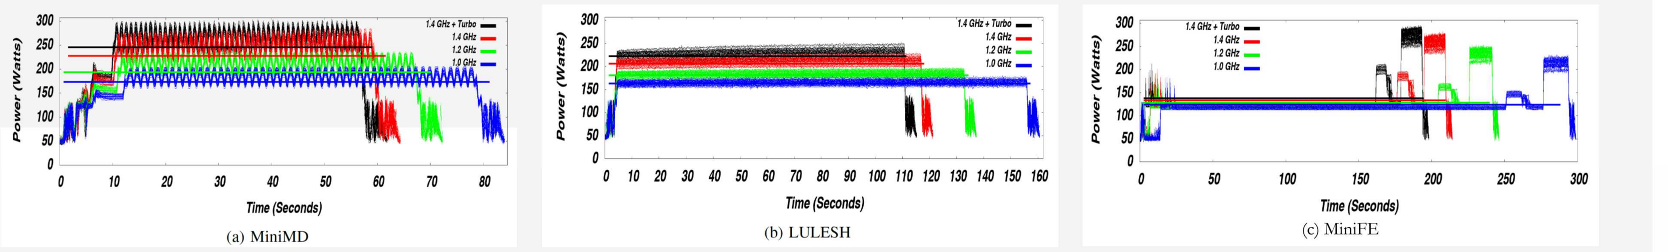

Historical Data Analysis

- Goal: predict the average power of a job given its submission information
  - Categorical/tabular information (e.g. username, job name, layout) makes feature space embedding necessary for standard ML algorithms
  - From power-time profiles, we expect similar jobs to consume similar energy
- Approach: regression analysis of historical job data to enable more accurate forecasts
  - Embedding spaces explored: uncorrelated (binary – 41D), contrast coding (backward difference – 2341D), and neural-based (cat2vec – 142D)
  - Algorithms used: linear, random forest, and support vector regression
  - Contrasted with default settings (e.g. no power cap)

Power Regression r^2 score

	Binary	Cat2Vec	BDE
Lin	0.700	0.806	0.773
SVR	0.657	0.787	N/A
RF	0.813	**0.816**	0.814

Power Regression RSME (W)

	Binary	Cat2Vec	BDE
Lin	42.27	35.52	35.24
SVR	47.22	37.02	N/A
RF	34.87	**34.65**	34.78

Constant Predictor RSME (W)

None (450W)	Mean (~350W)
127.5	80.73

Figures: exemplar power prediction comparisons for test set of different embeddings and regression algorithms (left); measured vs predicted correlation plot of best performing model (below)

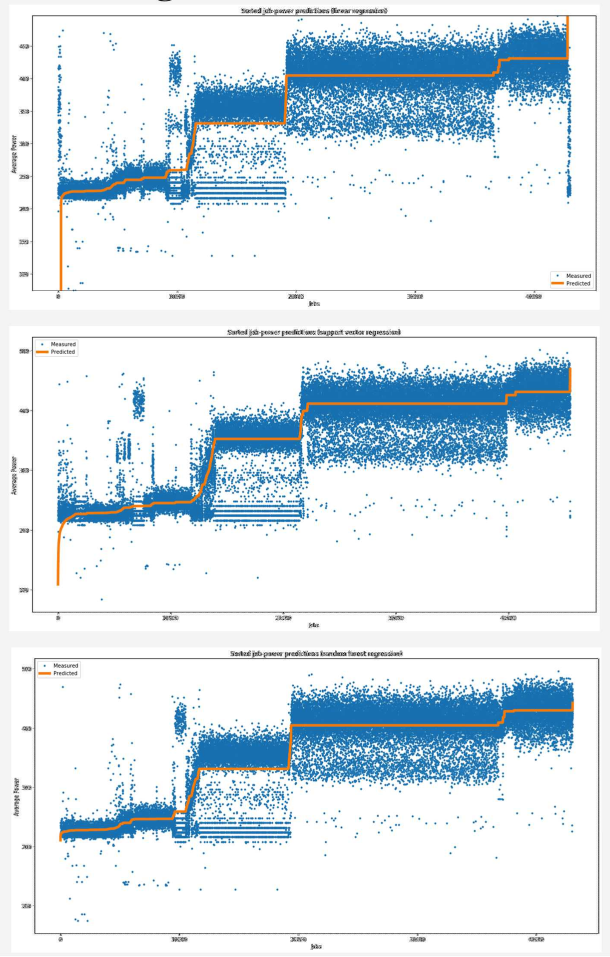

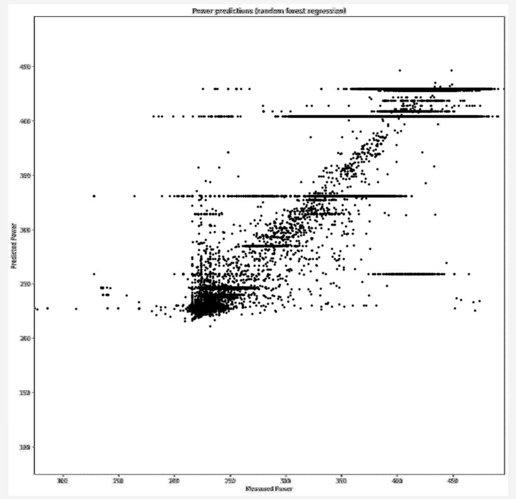

Job-Power Visualization

- Astra: large-scale prototype system under the Sandia Vanguard program
  - 2,592 compute nodes
  - Dual 28-core Cavium ThunderX2 64-bit Arm-v8 processors per node
  - 1.2 MW power consumption

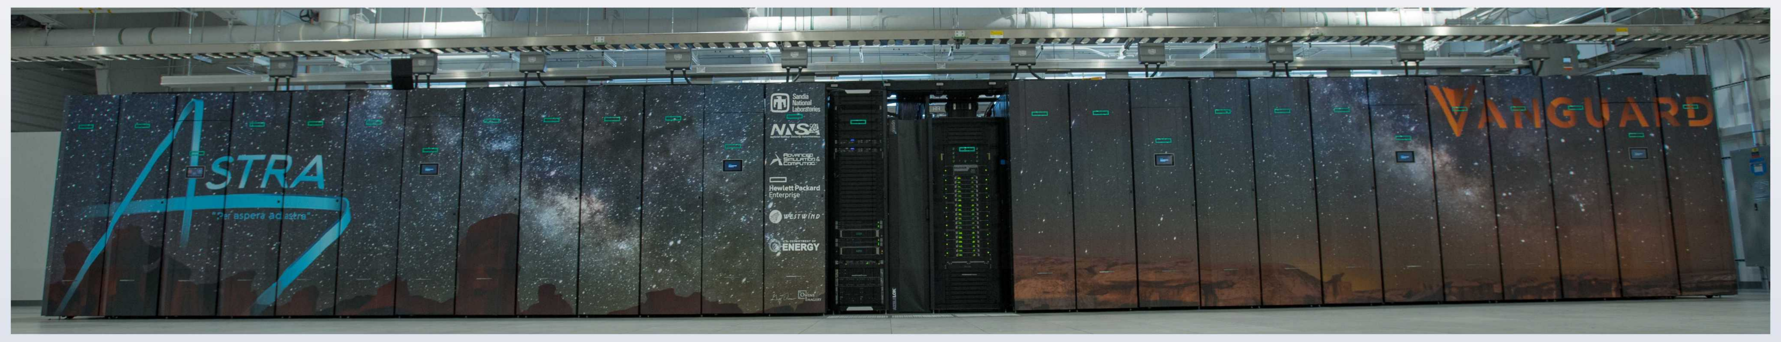

Visualization: scheduled jobs (top row) and power consumption (bottom row) over a five month data collection

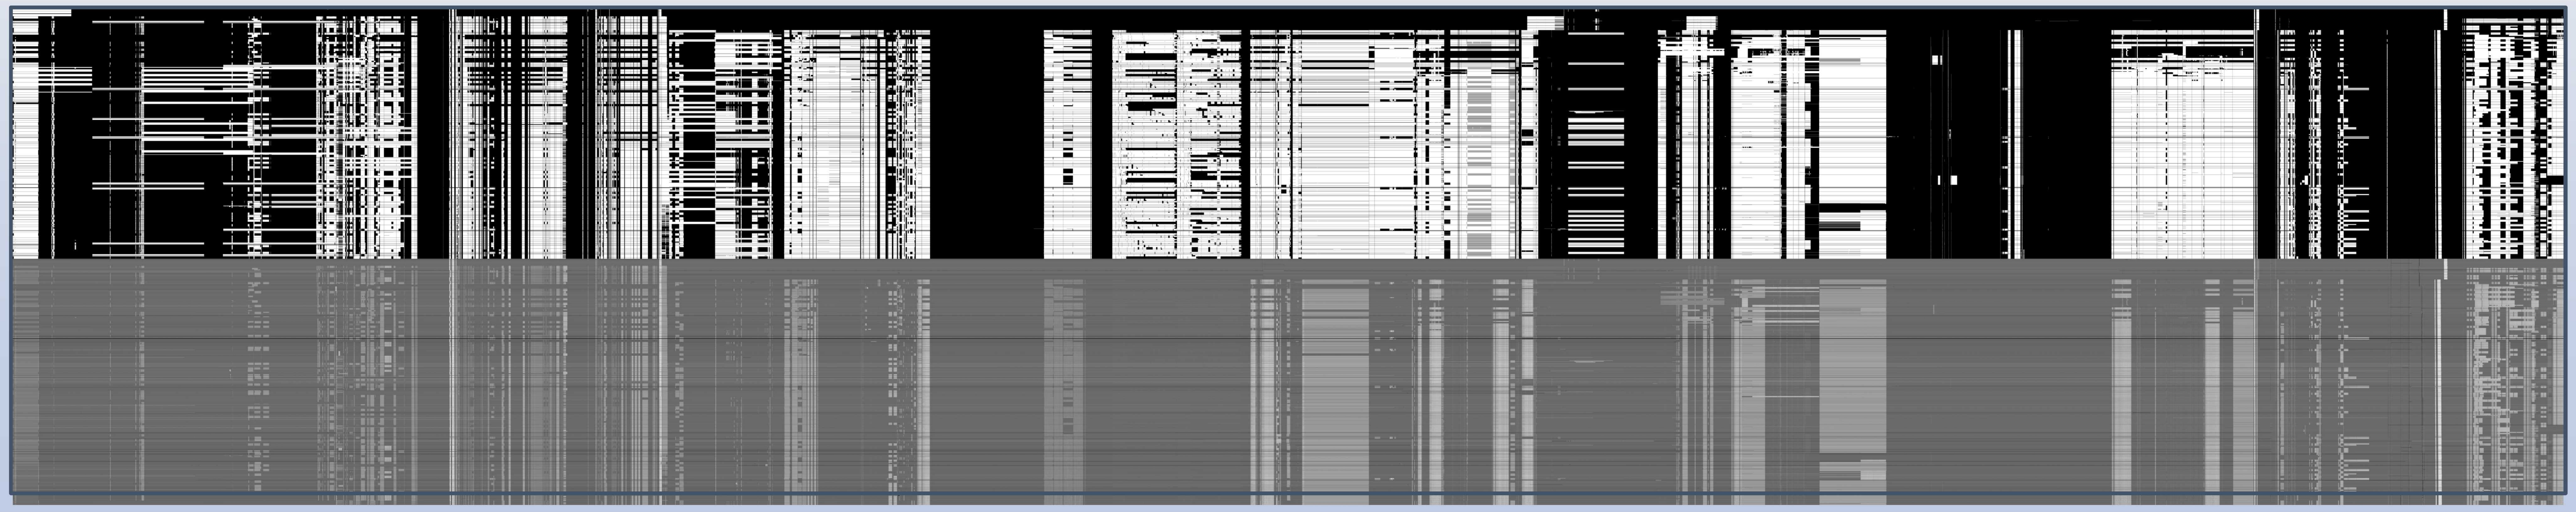

Reinforcement Learning Control

- Reinforcement learning (RL) is a machine learning paradigm where an agent makes actions given state observations from the environment; the environment emits rewards and new state observations, based on the agent's actions

RL Prerequisites	RL for HPC
❑ Sequential: control problem should require sequential decision making	❑ Sequential: power management state (P-state) of each core
❑ Features: observations about the state of the environment	❑ Features: performance counters representing system state
❑ Rewards: system metrics which we care about and can collect during training	❑ Rewards: job run time and power consumption
❑ Actions: parameters of the system which the agent can control	❑ Actions: set P-state

- Goal: learn dynamic P-state policy that balances power utilization and runtime across the execution of applications
  - Although P-states for all cores set identically, running all cores at either lowest or highest P-state settings will be non-optimal
  - Scenario: serial phases of code will not make use of all cores, so they should not be running at maximum P-state, but parallel phases should make use of high P-states

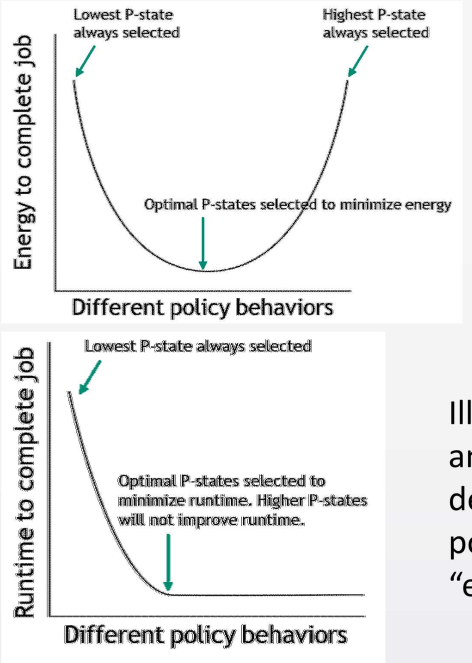

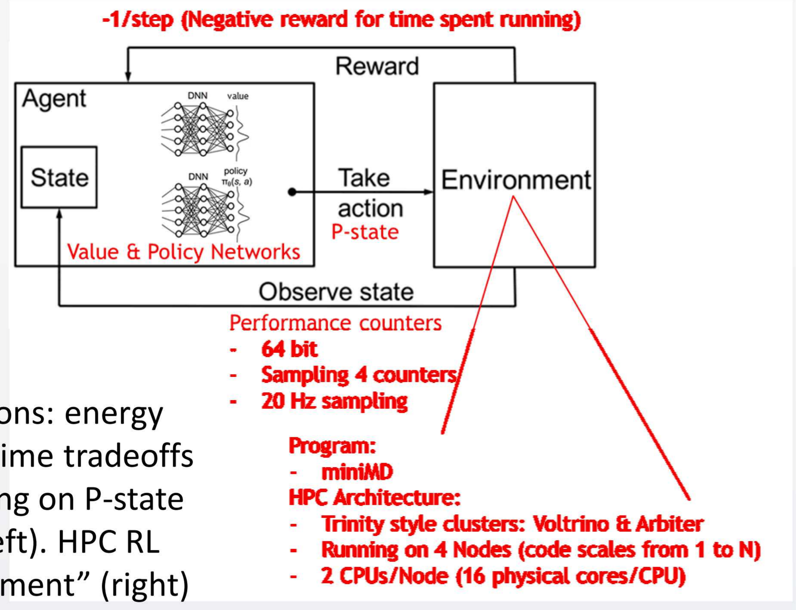

Illustrations: energy and runtime tradeoffs depending on P-state policy (left). HPC RL "environment" (right)

Summary

- Regularities in job-power profiles enables better power utilization management strategies through machine learning approaches
- Power forecasts from job submission information supports increased resource saturation within given power budget
- P-state regulation can further improve power utilization metrics

References

- Andrae, Anders. "Total consumer power consumption forecast." Nordic Digital Business Summit 10 (2017).
- Grant, Ryan E., et al. "Evaluating energy and power profiling techniques for HPC workloads." *2017 (IGSC)*
- https://vanguard.sandia.gov/astra/

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology &Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND No. ___________

U.S. DEPARTMENT OF ENERGY
National Nuclear Security Administration
CCR Center for Computing Research
Sandia National Laboratories