

A Fast Simulator to Enable HPC Scheduling Strategy Comparisons

Alex Wilkinson ^{1,2}, Jess Jones ², Harvey Richardson ², Tim Dykes ², and
Utz-Uwe Haus ²


¹University College London, United Kingdom

²HPE HPC/AI EMEA Research Lab, United Kingdom

25 May 2023

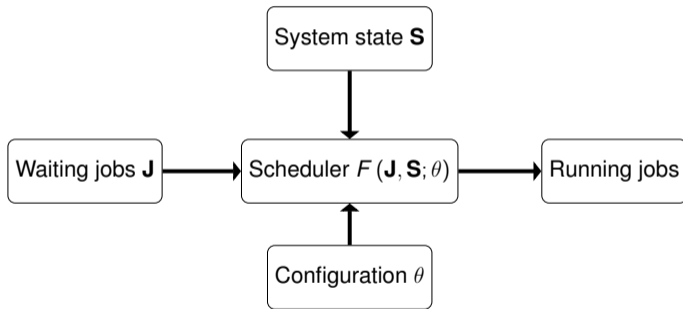
Introduction

- ▶ Job schedulers are a vital part of running an efficient HPC system



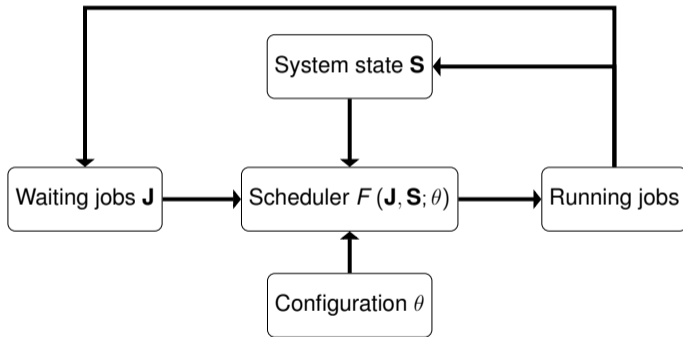
Introduction

- ▶ Job schedulers are a vital part of running an efficient HPC system



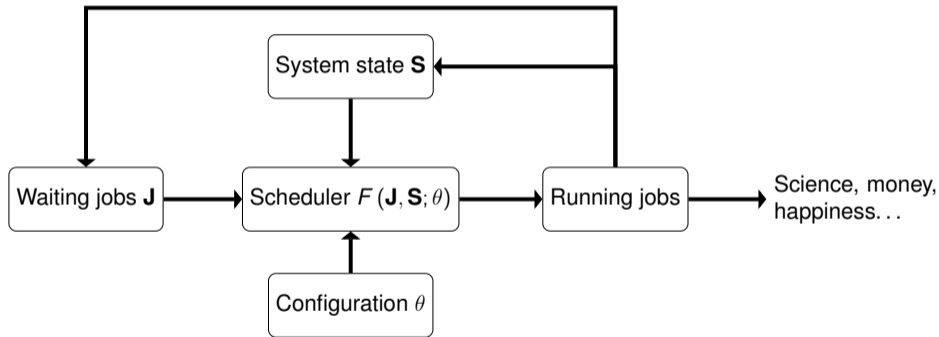
Introduction

- ▶ Job schedulers are a vital part of running an efficient HPC system



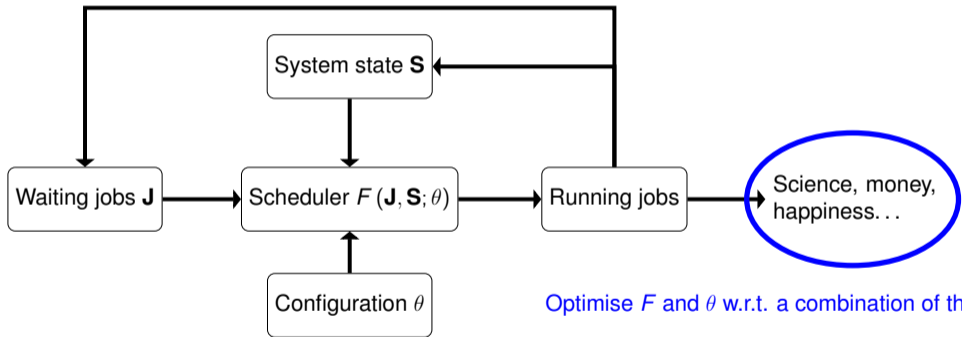
Introduction

- ▶ Job schedulers are a vital part of running an efficient HPC system



Introduction

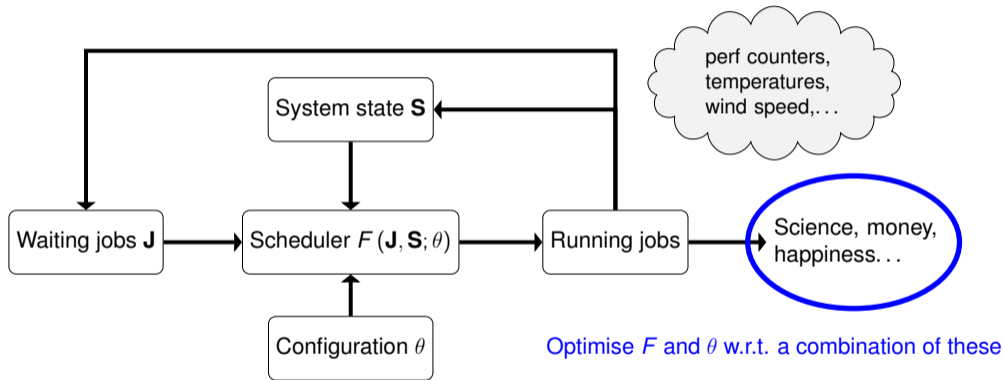
- ▶ Job schedulers are a vital part of running an efficient HPC system



Optimise F and θ w.r.t. a combination of these

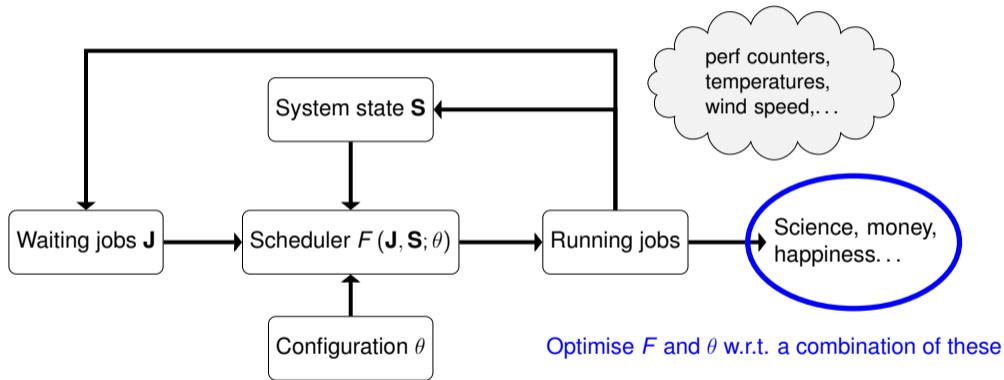
Introduction

- ▶ Job schedulers are a vital part of running an efficient HPC system



Introduction

- ▶ Job schedulers are a vital part of running an efficient HPC system



- ▶ Simulation allows exploration of configurations and scheduling algorithms without risking system efficiency — we will focus on the popular workload manager Slurm

A Lightweight Simulation

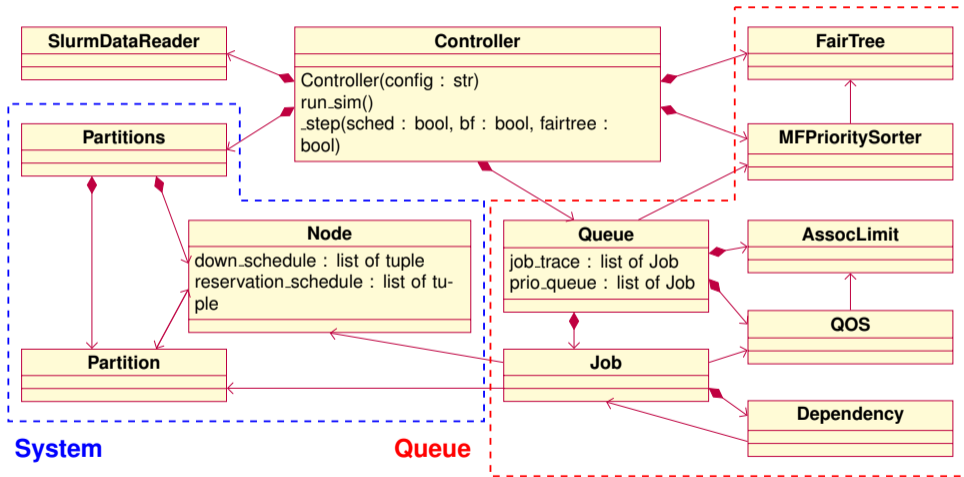
- ▶ Development of a simulation mode for Slurm started with A. Lucero in 2011¹ and has been iterated on in some other excellent works^{2,3}
 - Modify the Slurm source code to emulate communication from nodes and skip through time
 - Limited speed up and extensibility
- ▶ Current research into HPC scheduling often uses custom simulations to evaluate algorithms
 - These can be simplistic and not replicate the configuration of a real system
- ▶ We propose a fast simulation that can accurately reproduce the dynamics of real Slurm without trying to reproduce the specific design
 - Implement features directly relevant to scheduling from scratch

¹ Lucero, A.: Simulation of batch scheduling using real production-ready software tools (2011)

² Jekanovic, A. et al.: Evaluating slurm simulator with real-machine slurm and vice versa (2018)

³ Simakov, N. et al.: A Slurm Simulator: Implementation and Parametric Analysis (2018)

Simulation Structure



Key Simulation Features

Backfilling

Conservative backfilling algorithm that simulates backfilling thread lock release

Resource Limits

Resource limits tracked at quality of service and association level

MultiFactor Priority and Fairshare

Queue sorted using a hierarchy of job features including a fairshare factor. Fairshare is implemented by sorting a rooted ordered tree of users association by usage and system allocation (Slurm's Fair Tree).

- ▶ Recovering full job and system information from Slurm accounting database
 - Information such as dependencies and requested nodes can only be recovered if submitted via command line rather than in batch script
 - Completed reservations not stored

- ▶ Some scheduling features missing from simulation
 - Nodes are the only consumable resource
 - Advanced features: job preemption and heterogeneous jobs

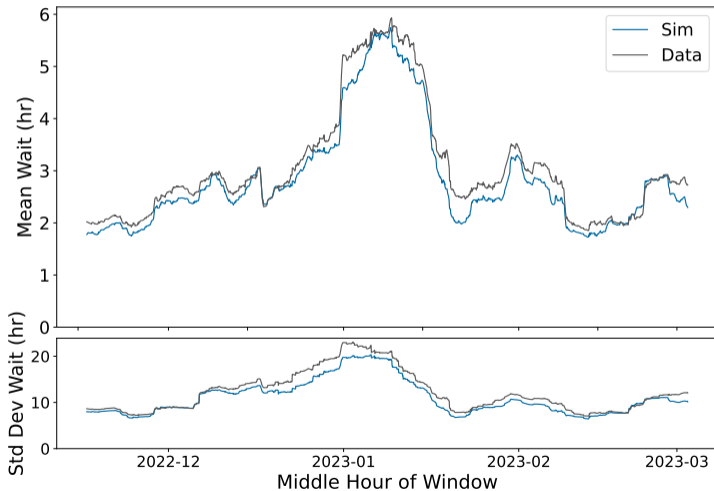
- ▶ Development of simulation was closely tied to ARCHER2
- ▶ ARCHER2 is the UK's national supercomputer consisting of 23 HPE Cray EX cabinets forming a network of 5,860 CPU compute nodes, 28 in TOP500
- ▶ 4 month job trace with $\sim 600,000$ jobs used to validate simulation accuracy
 - Dependencies, overlapping partitions, multiple QoS, advanced reservations, record of down nodes,...



Wait Times

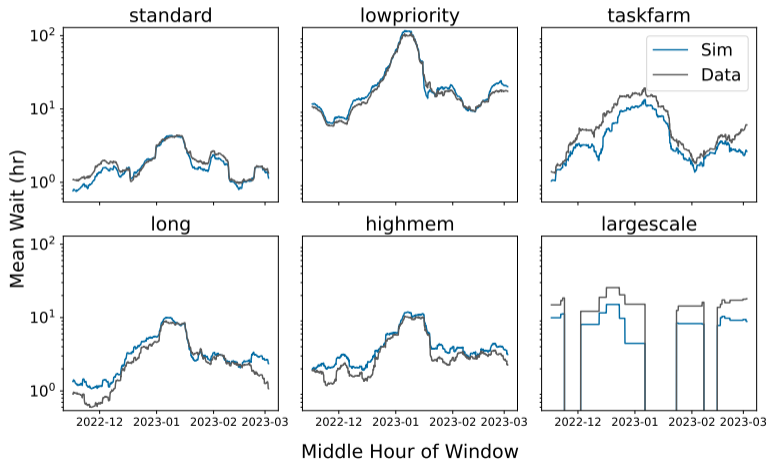


Wait Time 2 Week Moving Window



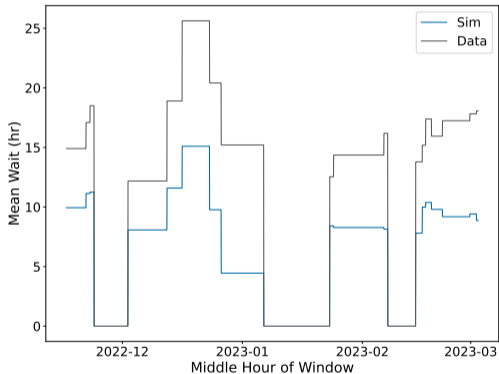
QoS Wait Times

Wait Time 2 Week Moving Window by QoS

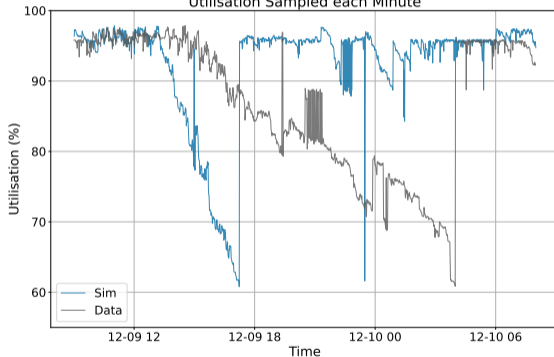


Largescale Jobs Discrepancy

Wait Time 2 Week Moving Window Largescale QoS




Utilisation Sampled each Minute

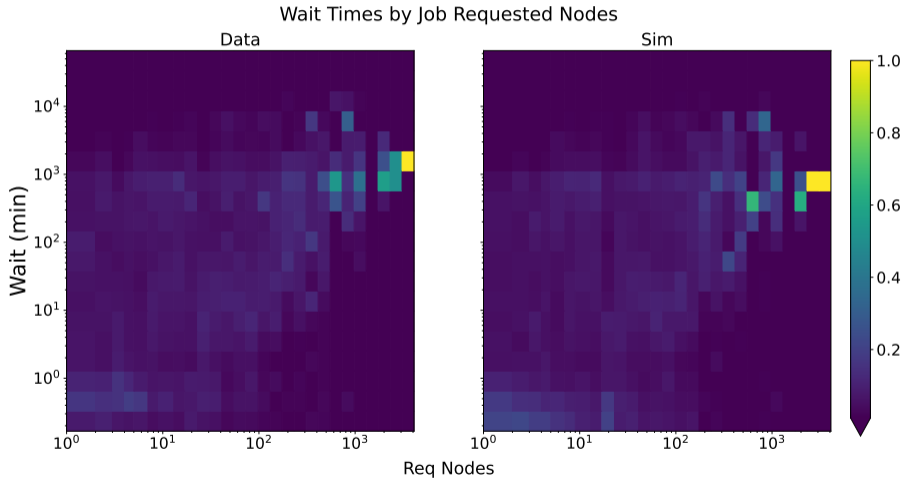


User Wait Times

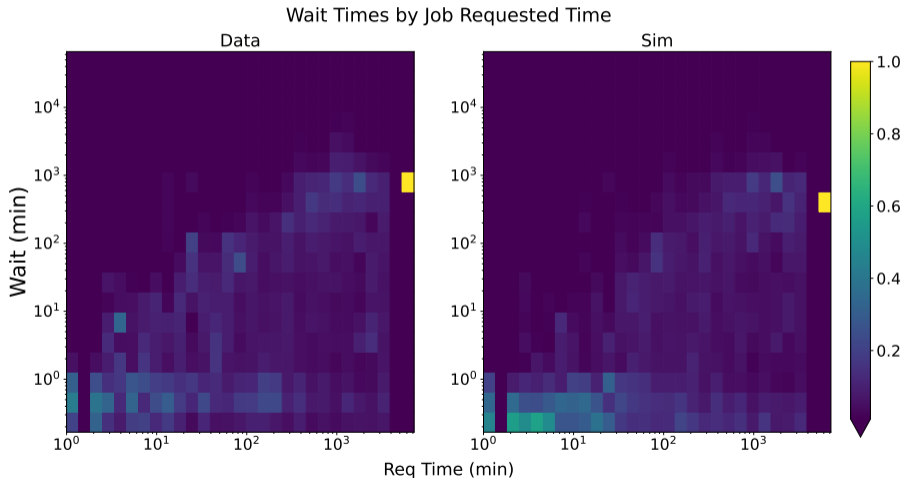
Wait Times for Users with Highest Usage



Job Size Response



Job Length Response



Other Systems: LUMI

- ▶ Important to check that the simulation is not tuned to ARCHER2
- ▶ LUMI is Europe's fastest supercomputer and part of the EuroHPC Joint Undertaking
 - We consider the standard partition consisting of 1,022 CPU nodes
- ▶ 3 month job trace numbering $\sim 25,000$ jobs
- ▶ Difficulty recovering past reservations
 - Approximate using the maximum utilisation from jobs without reservations in a 2 day window

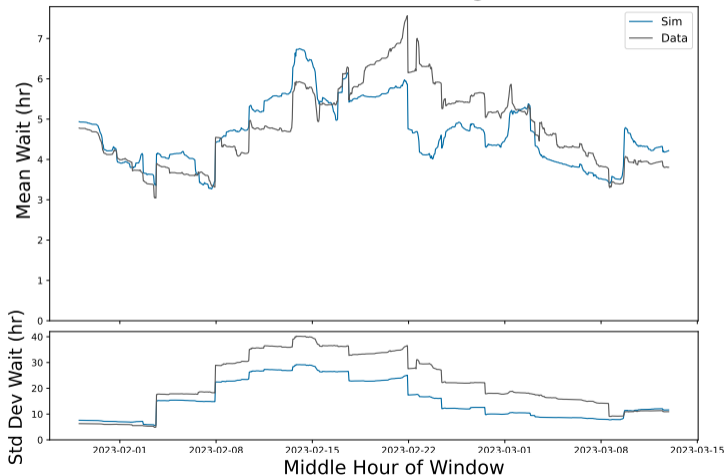


LUMI

The logo for LUMI, consisting of the letters 'L', 'U', 'M', and 'I' in a bold, sans-serif font, enclosed within a rectangular border.


LUMI Wait Times

Wait Time 2 Week Moving Window



LUMI User Wait Times

Wait Times for Users with Highest Usage



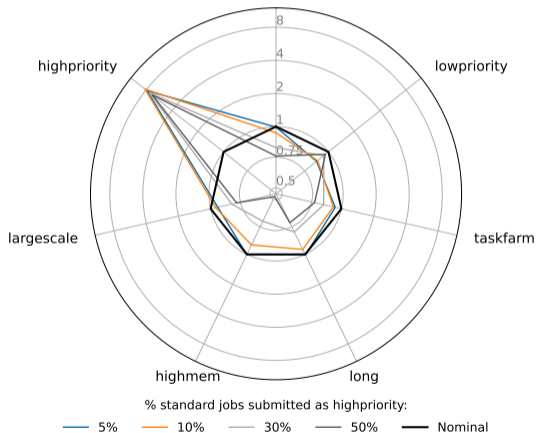
- ▶ ARCHER2 simulation takes approximately 7 hours 20 minutes, LUMI 25 minutes
 - Speed up of ~ 400 for ARCHER2 (400 simulation minutes takes 1 minute)
 - Single threaded, memory usage ~ 2 Gb depending on job trace size
 - Processing time dominated by backfilling
- ▶ Speed ups from simulators in literature are typically between 10 and 25
- ▶ Exception is work from Barcelona Supercomputing Center¹ which achieves a 220 speed up with the CAE Curie log from the Parallel Workloads archive
 - $\sim 200,000$ jobs over an 8 month period running on 5,040 nodes
 - Archive states 62% utilisation
 - Unclear how performance would translate to modern 90+% utilisation workloads
- ▶ Direct comparisons between simulators is important future work

¹Jokanovic, A. et al.: Evaluating slurm simulator with real-machine slurm and vice versa (2018)

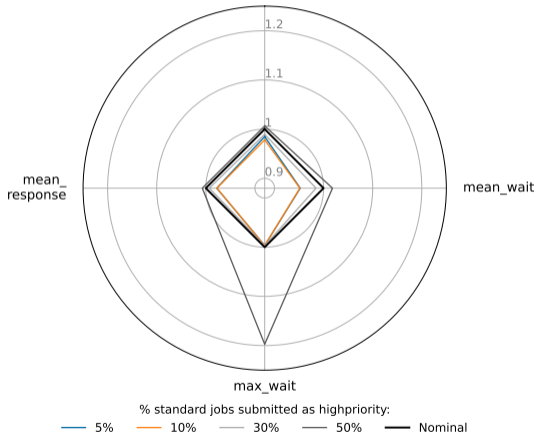
- ▶ The simulator can be used to understand the effect of changes in scheduler behaviour on a production system
- ▶ Start with a simple change to ARCHER2's QoS configuration: adding a high priority QoS
- ▶ Consider scenarios with increasing proportions of *standard* QoS jobs being submitted as *highpriority* in the historical job trace

High Priority

Mean QoS Wait Time Relative to Nominal Simulation
standard



Performance Metrics Relative to Nominal Simulation
mean_slowdown

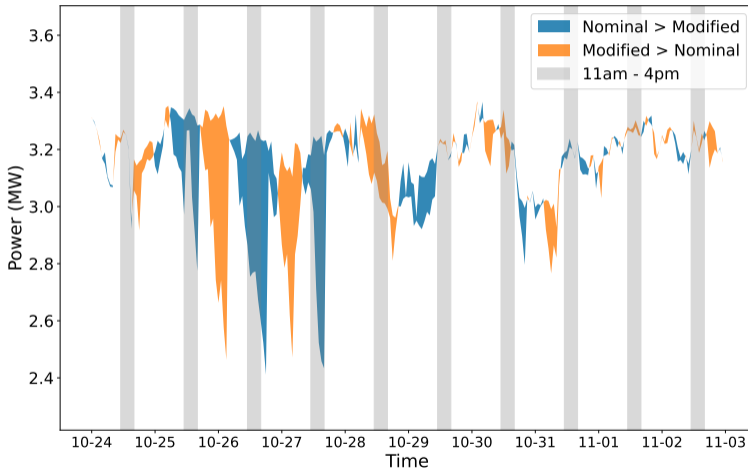


- ▶ Slurm can be configured to associate energy counters from nodes with the jobs running on them
 - System power usage can then be estimated from the jobs running at any given time in the simulation
- ▶ Consider scheduling jobs to minimise power usage during peak times of day
 - Even with backfilling large jobs will require the system to partially drain in order to be scheduled
 - Hold *largescale* jobs until morning, specific time depending on size

Power Usage

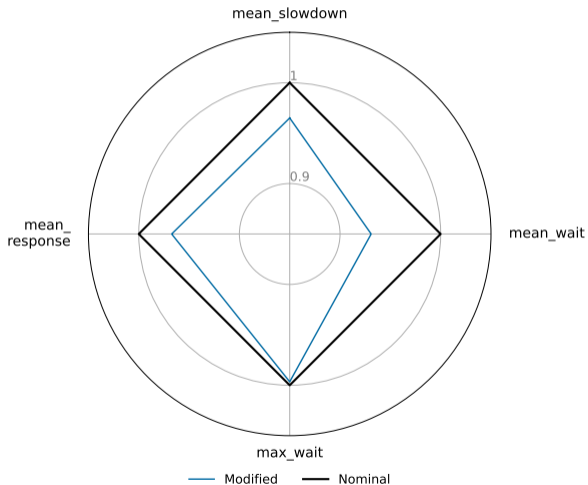


Power Usage Difference for Modified and Nominal Simulation Sampled Hourly



System Efficiency


Performance Metrics Relative to Nominal Simulation



- ▶ A fast and easily extendable scheduling simulation that incorporates many features of Slurm
- ▶ Validated with modern production systems
- ▶ Potential of simulation to provide insight into scheduling strategies demonstrated
- ▶ Future work:
 - Direct comparisons with existing simulations
 - Improving feature coverage of simulation to validate with a wider range of HPC systems

Backup

Job Size Response LUMI



Job Length Response LUMI

