

---

# Power-aware scheduling for Large-scale Supercomputers

Tapasya Patki\*<sup>1</sup>

<sup>1</sup>Lawrence Livermore National Laboratory (LLNL) – Lawrence Livermore National Laboratory 7000 East Avenue • Livermore, CA 94550, United States

## Abstract

Power is a critical problem as the supercomputing community ventures toward exascale. Several system, hardware and application-level challenges to power management make it challenging to schedule jobs effectively with high throughput and utilization on large-scale supercomputers. In this talk, I will present two power-aware scheduling strategies that are being developed at LLNL. The first is an algorithm implemented within the popular resource manager, SLURM, and the second is a variation-aware scheduling algorithm implemented within LLNL's new hierarchical scheduling framework, Flux. Using real-world datasets from large supercomputers at LLNL and University of Tokyo, I will demonstrate the effectiveness of these power-management scheduling algorithms, and discuss how other flow resources (such as network, IO, etc) in HPC clusters can leverage similar techniques.

---

\*Speaker