

Towards High-Order Performance Objectives for HPC System Scheduling

Title: [Towards High-Order Performance Objectives for HPC System Scheduling](#)

Authors: C.Lee, A.Snavely, B.Leary, L.Carrington, H.Casanova, R.Bohn, R.Carson, J.Hardy, Y.Schwartzman

Abstract: High Performance Computing (HPC) system schedulers are commonly designed to optimize utilization, average job turnaround and throughput. However, these system-centric metrics do not capture aspects of the system that are important to users. Two examples of these are the scheduler's ability to accurately predict a job's start time and its responsiveness to a range of priority levels. We show that users attach a value to a job's timely completion, and this value typically decreases with time. We present a tool that users can use to improve their customer surplus (i.e. get a better deal). We also show that user estimates can be more accurate and we present a tool to motivate them to do so. Then we show via simulation that an HPC system scheduled to minimize Global Delay Cost explicitly optimizes this metric.

Reference: @techreport{lee05highorder, Author = {C.Lee, A.Snavely, B.Leary, L.Carrington, H.Casanova, R.Bohn, R.Carson, J.Hardy, Y.Schwartzman}, Institution = {SDSC}, Title = {The PMAc Binary Instrumentation Library for PowerPC}, Year = {2004}}