
Experiences in the design of a malleable and I/O-aware scheduler

Nicolas Vidal , Guillaume Aupy , Alberto Garcia , David Singh*¹, Jesus Carretero , and Emmanuel Jeannot²

¹Carlos III University of Madrid – Spain

²Inria Bordeaux - Sud-Ouest – Institut National de Recherche en Informatique et en Automatique – France

Abstract

In this talk I will describe the execution framework that integrates CLARISSE and FlexMPI into a parallel application scheduler prototype. CLARISSE is a middleware for enhancing I/O flow coordination and control in the HPC I/O stack software. It monitors the I/O activity of each running application and provides inter-application coordination mechanisms for accessing the storage. FlexMPI provides dynamic load balancing and malleability capabilities for MPI applications, as well as CPU and communication monitoring. By the joint use of both frameworks it is possible dynamically characterize each running application and offer a holistic view of the platform. Using this information in combination with the coordination mechanisms provided by CLARISSE and FlexMPI it is possible to implement CPU and I/O scheduling policies aimed at improving overall platform performance. The main points of this talk will cover runtime monitoring, CPU and I/O scheduling, and malleability.

Keywords: Scheduling, Malleability, I/O scheduling, Cross, layer optimizations, Parallel I/O, monitoring

*Speaker