

AI4IO: A SUITE OF AI-BASED TOOLS FOR IO-AWARE HPC RESOURCE MANAGEMENT

Michela Taufer

University of Tennessee Knoxville

Abstract:

High performance computing (HPC) is undergoing many changes at the system level. While scientific applications can reach petaflops or more in computing performance, potentially resulting in larger data generation rates and more frequent checkpointing, the data movement to the parallel file system remains costly due to constraints imposed by HPC centers on the IO bandwidth. In other words, the bandwidth to file systems is outpaced by the rate of data generation; the associated IO contention increases job runtime and delays execution. This situation is aggravated by the fact that when users submit their jobs to a HPC system, they rely on resource managers and job schedulers to monitor and manage the computing resources (i.e., nodes). Both resource managers and job schedulers remain blind to the impact of IO contention on the overall simulation performance.

In this talk we discuss how Artificial Intelligence (AI) can augment HPC systems to prevent and mitigate IO contention while dealing with IO bandwidth constraints. Our solution, called Analytics for IO (AI4IO), consists of a suite of AI-based tools that enable IO-awareness on HPC systems. Specifically, we present two AI4IO tools: PRIONN and CanarIO. PRIONN automates predictions about user-submitted job resource usage, including per-job IO bandwidth; CanarIO detects, in real-time, the presence of IO contention on HPC systems and predicts which jobs are affected by that contention (e.g., because of their frequent checkpointing). By working in concert, PRIONN and CanarIO predict the a priori knowledge necessary to prevent and mitigate IO contention with IO-aware scheduling. We integrate AI4IO in the Flux scheduler and show how AI4IO produce improvements in simulation performance: we observe up to 6.2% improvement in makespan of HPC job workloads, which amounts to more than 18,000 node-hours saved per week on a production-size cluster. Our work is the first step to implementing IO-aware scheduling on production HPC systems.

Bio:

Michela Taufer is an ACM Distinguished Scientist and holds the Jack Dongarra Professorship in High Performance Computing in the Department of Electrical Engineering and Computer Science at the University of Tennessee Knoxville (UTK). She earned her undergraduate degrees in Computer Engineering from the University of Padova (Italy) and her doctoral degree in Computer Science from the Swiss Federal Institute of Technology or ETH (Switzerland). From 2003 to 2004 she was a La Jolla Interfaces in Science Training Program (LJIS) Postdoctoral Fellow at the University of California San Diego (UCSD) and The Scripps Research Institute (TSRI), where she worked on interdisciplinary projects in computer systems and computational chemistry. Michela has a long history of interdisciplinary work with scientists. Her research interests include scientific applications on heterogeneous platforms (i.e., multi-core platforms and accelerators); performance analysis, modeling and optimization; Artificial Intelligence (AI) for cyberinfrastructures (CI); AI integration into scientific workflows, computer simulations, and data analytics. She has been serving as the principal investigator of several NSF collaborative projects. She also has significant experience in mentoring a diverse population of students on interdisciplinary research. Michela's training expertise includes efforts to spread high-performance computing participation in undergraduate education and research as well as efforts to increase the interest and participation of diverse populations in interdisciplinary studies.