



Contribution ID: 18

Type: Talk

## Evaluating and Characterizing Parallel I/O in HPC Systems: Best Practices and Future Directions

*Thursday, 2 December 2021 14:30 (30 minutes)*

As a recent I/O behaviour analysis [1] has revealed, High Performance Computing (HPC) storage systems may no longer be dominated by write I/O – challenging the long- and widely-held belief that HPC workloads are write-intensive. HPC applications are evolving to include not only traditional scale-up modelling and simulation bulk-synchronous workloads but also scale-out workloads [2] like artificial intelligence (AI), advanced and big data analytics [3], machine learning, deep learning [4], and complex multi-step workflows [5]–[7]. Exascale workflows are projected to include multiple different components from both scale-up and scale-out communities operating together to drive scientific discovery and innovation. With the often conflicting design choices between optimizing for write-intensive vs. read-intensive workloads, having flexible I/O systems will be crucial to support these emerging hybrid workloads. Another performance aspect is the intensifying complexity of parallel file and storage systems in large-scale cluster environments. Storage system designs are advancing beyond the traditional two-tiered file system and archive model by introducing new tiers of temporary, fast storage close to the computing resources with distinctly different performance characteristics. The changing landscape of emerging hybrid HPC workloads along with the ever increasing gap between the compute and storage performance capabilities reinforce the need for an in-depth understanding of extreme-scale parallel I/O and for rethinking existing data storage and management evaluation techniques and strategies. In this talk, an overview and taxonomy [8] of the current state-of-the-art research on large-scale parallel I/O evaluation and characterization techniques in the context of HPC systems is presented. Traditionally, the process of understanding large-scale I/O behaviour and performance for specific applications or storage systems is performed iteratively and empirically in a closed loop fashion, as outlined in Figure 1, and consists of three main phases: (1) Measurements and Statistics Collection, (2) Modelling and Prediction, and (3) Simulation. The overview and broad knowledge base provided by this talk is invaluable to the whole scientific community, as applications often observe poor performance due to bottlenecks in the parallel I/O and storage system. In addition, this talk aims to identify future re-search challenges with regard to emerging exascale computing systems and more complex hybrid HPC workloads.

### **Student?**

No

### **Supervisor name**

### **Supervisor email**

**Primary author:** NEUWIRTH, Sarah (Goethe-University Frankfurt)

**Presenter:** NEUWIRTH, Sarah (Goethe-University Frankfurt)

**Session Classification:** HPC Technology

**Track Classification:** Storage and IO