An Exascale Vision of Scientific Data Management

Achievement: Computer scientists articulate how to use new technologies to reduce latency to scientific insight in big data.

Significance and Impact: I/O is worsening as a bottleneck for exascale; advances in compression, progressive refinement, and intelligent storage system usage are all needed to address this.

Research Details:

- Implemented two methods to access data at different qualities progressively. Refactored data can be split over multiple storage tiers (SSD, parallel file system, tape, etc.) for better performance
- Developed mathematically motivated approach for "lossless decimation" with comparable compression to other state of the art techniques
- Exploring feasibility of hardware acceleration for compression: GPUs, FPGAs, etc.

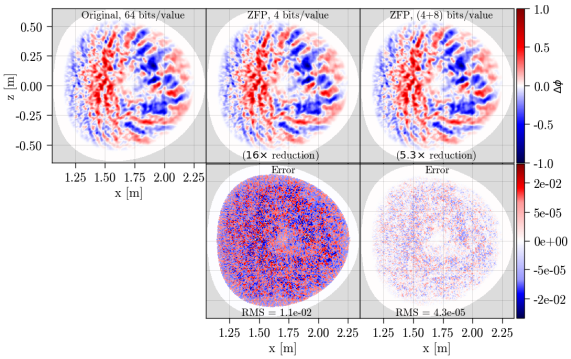

Progressive refinement with plasma physics data. The top panels show how the electric potential in a simulated tokamak deviates from background. The data have been saved at different precision-levels in each column, compressing with ZFP. The bottom row shows the errors incurred for that precision-level. Low precision can be saved to faster storage (SSD) with additional accuracy added from larger, slower storage (parallel file system) if enough bandwidth is available.

Sponsor/Facility: Advanced Scientific Computing Research (ASCR) program.

PI and affiliation: Scott Klasky, Oak Ridge National Laboratory

Team: S. Klasky, E. Suchyta, M. Ainsworth, Q. Liu, et al.

Publication: S. Klasky, E. Suchyta, M. Ainsworth, Q. Liu, et al., "Exacution: Enhancing Scientific Data management for Exascale," in ICDCS'17, Atlanta, GA, 2017.

Overview:

The computational power of leadership class computing facilities is growing much faster than their total I/O throughput. Accordingly, as we advance toward exascale, fewer bits/FLOP can be written to disk. This possibly prevents simulations from saving enough data and could reduce the simulation's fidelity and sensitivity. Moreover, accessing large datasets from heavily subscribed file systems can be time consuming and suffer from significant performance variability. Users could benefit from the option of reading smaller, less precise versions of the data for better performance, supplemented with error-bounds to understand the uncertainties in their analyses. A team of ORNL-led researchers has offered preliminary solutions and articulated strategies for further improvements. Lossy compression plays an important role, because it can reduce the volume of information to be written. Compression is used in a way to achieve progressive refinement – multiple-quality versions of the same data, available in case of file system variability. Heavy emphasis is given to careful monitoring, to ensure salient features are not lost in reduction and that the compression is performant enough to warrant its usage. Increased mathematical rigor is also stressed to better understand the compression's impact on the scientific analysis, and a mathematically-principled approach for lossless decimation is presented and compared to other methods.