Using ZFP Lossy Compression in HPC: Theoretical Support for the New Floating-Point Representation

Alyson Fox, Lawrence Livermore National Laboratory

The amount of data generated and gathered in scientific simulations is continuously growing, putting mounting pressure on storage and bandwidth limitations. New creative ways of handling the data movement bottleneck problem have been proposed. In this talk, we review our ongoing efforts to provide theoretical support for adopting ZFP lossy compression as a new data-type for scientific simulations. ZFP is a lossy compression algorithm, proposed by Peter Lindstrom, that was explicitly designed for floating-point data arrays. It decomposes the data set into blocks, which are compressed and decompressed independently. ZFP itself is quite effective for I/O operations and in-memory storage of static data. However, there is potential for significant performance gains by using compressed ZFP data arrays within a simulation to store the solution state. As ZFP arrays are most efficient using lossy compression, it is essential to understand how the error from lossy compression impacts the accuracy of the solution through repeated cycles of compression and decompression. We discuss the error analysis work for ZFP as well as present potential new flexible precision options for ZFP, further extending the compression capabilities for HPC applications.