
Three examples of reduced precision with sparsity in high-fidelity simulations

Hans Johansen^{*1}, Zhe Bai¹, and Rochi Chowdhury¹

¹LBNL – United States

Abstract

High-fidelity computational simulations typically require double precision operations in time integrators, geometric calculations, discretized operators, and solvers. With the goal of improving memory bandwidth without compromising accuracy, we explore applying sparsity and reduced precision operations to these key areas. In time integrators for hyperbolic PDE's, we apply machine learning to estimate entire time series; it is relatively accurate and can be used as an effective preconditioner for parallel-in-time (PinT) algorithms. For geometric calculations and spatial discretizations, we explore the impact of precision on operator spectrum and accuracy, highlighting a "sparse-sparse" mixed precision representation of operators to reduce memory footprints. And finally, we use combinations of approaches to speed up implicit time integrators with better preconditioners. We show that, in total, this can produce significant speed-ups with a reduction in memory for a high-fidelity PDE simulation.

^{*}Speaker