Abstract Submitted
for the DFD19 Meeting of
The American Physical Society

In situ data compression for large-scale computational fluid dynamics simulations via interpolative decomposition methods[1] HEATHER PACELLA, Stanford University, ALEC DUNTON, ALIREZA DOOSTAN, University of Colorado, Boulder, GIANLUCA IACCARINO, Stanford University — Over the next decade, exascale supercomputers will provide a thousand-fold increase in floating-point performance, with memory increase of only a factor of one hundred. As a result, data generation from computational fluid dynamics simulations will easily surpass available memory capacity. Additional data generation from ensemble simulations for uncertainty quantification, inference, and optimization will also contribute to this discrepancy. To address this, we implement a lossy in situ compressive algorithm, interpolative decomposition (ID), within the solvers themselves, which allows us to store simulation results at a fraction of memory cost. Because ID algorithms operate independently on subregions of the fluid domain, they are a natural fit for the flexibility that task-based parallelism programming systems provide. Legion is one such programming system; it allows for implicitly extracted parallelism, easy performance tuning, and porting to various heterogeneous architectures. We will discuss the implementation of both the sub-sampled and single-pass ID algorithms in a high-order Navier-Stokes solver written in Regent, as well as performance, scalability, and cost of both ID algorithms in a task-parallel environment.

[1]This work was funded by the US Department of Energy's National Nuclear Security Administration under the Predictive Science Academic Alliance Program II at Stanford University, Grant495 DE-NA-0002373.

Heather Pacella
Stanford University

Date submitted: 31 Jul 2019 Electronic form version 1.4