Abstract Submitted
for the DFD17 Meeting of
The American Physical Society

A multithreaded and GPU-optimized compact finite difference algorithm for turbulent mixing at high Schmidt number using petascale computing[1] M. P. CLAY, P. K. YEUNG, Georgia Tech, D. BUARIA, Max Planck Institute for Dynamics and Self-Organization, Germany, T. GOTOH, Nagoya Inst Tech, Japan — Turbulent mixing at high Schmidt number is a multiscale problem which places demanding requirements on direct numerical simulations to resolve fluctuations down the to Batchelor scale. We use a dual-grid, dual-scheme and dual-communicator approach where velocity and scalar fields are computed by separate groups of parallel processes, the latter using a combined compact finite difference (CCD) scheme on finer grid with a static 3-D domain decomposition free of the communication overhead of memory transposes. A high degree of scalability is achieved for a 8192^3 scalar field at Schmidt number 512 in turbulence with a modest inertial range, by overlapping communication with computation whenever possible. On the Cray XE6 partition of Blue Waters, use of a dedicated thread for communication combined with OpenMP locks and nested parallelism reduces CCD timings by 34% compared to an MPI baseline. The code has been further optimized for the 27-petaflops Cray XK7 machine Titan using GPUs as accelerators with the latest OpenMP 4.5 directives, giving 2.7X speedup compared to CPU-only execution at the largest problem size.

[1]Supported by NSF Grant ACI-1036170, the NCSA Blue Waters Project with sub-award via UIUC, and a DOE INCITE allocation at ORNL.

Matthew Clay
Georgia Institute of Technology

Date submitted: 01 Aug 2017 Electronic form version 1.4