Present and future computing needs in high-speed turbulent flows

Sanjiva Lele and Johan Larsson
Stanford University

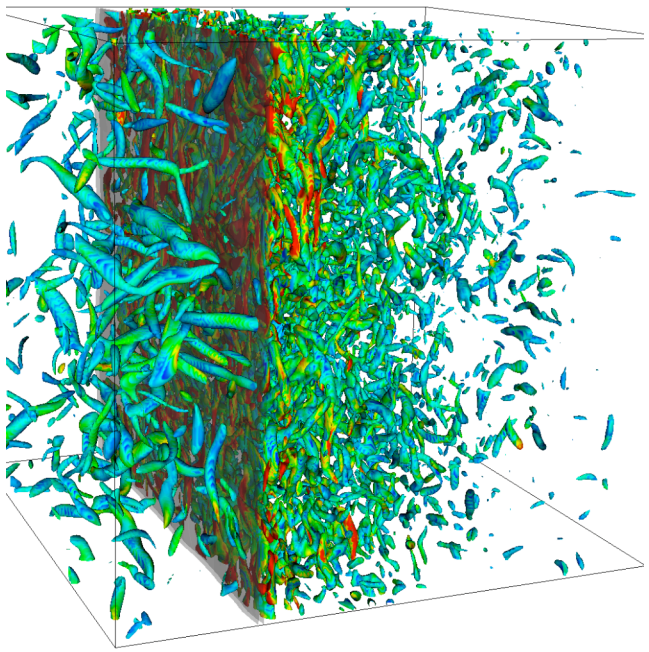

DOE ASCR workshop
Jan 5-6, 2011

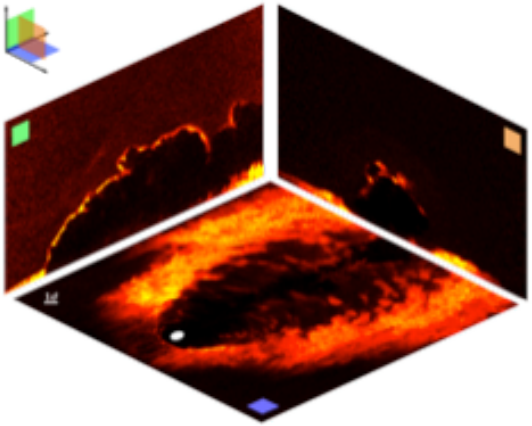

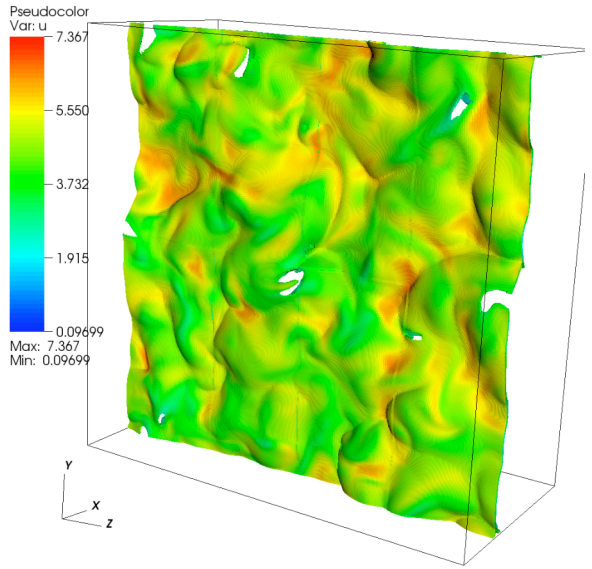

Overview

SciDAC: Sanjiva Lele (PI); Moin (Stanford); Cook, Cabot, Sjogreen (LLNL); Zhong (UCLA); Yee (NASA)

Other projects: PSAAP, NASA

- Objective: Increase fundamental understanding and improve predictive/modeling capability of general compressible turbulence problems, especially shock/turbulence interaction, shock/material-interface interaction.

- Present focus is on canonical problems (simple geometries), gradual move towards realistic problems (complex geometries) underway

- In the next 3 years:
  - A few even larger runs than present (2.5B->10B grid points)
  - Move to unstructured code, both big (few) and small (many) of runs
  - Include chemical reactions (potentially memory consuming)

Current HPC methods

- *Hybrid* (structured grid) and *Charles* (unstructured, more capabilities) codes
- Both codes:
  - Compressible Navier-Stokes (5 coupled nonlinear PDEs on weak form)
  - Fully explicit in space and time, no large matrix inversions
  - Solution-adaptive, different numerics around discontinuous features (potential load-imbalance)
  - MPI with domain-decomposition
  - Memory footprint: 60 doubles/grid point (*Hybrid*), hundreds of doubles/grid point (*Charles*); additional 1Gb or more per core for chemistry table (if reactive flow)
- Examples of computations in 2010 (mainly *Hybrid* code):
  - Big (3 runs): 2.5B grid points, 18 hrs on 65,536 BG/P cores or 48 hrs on 12,288 Cray XT-4 cores
  - Small (hundreds of runs): 2-50M grid points, 64-512 cores, at NERSC or Stanford
  - Post-processing: sequentially read in 100 snapshots, compute statistics, filtering, dump silo-files for VisIt, etc. Memory-limited, ideally “interactive”.

Current HPC requirements

- Machines: BG/P (ALCF), Cray XT (NERSC), Dell (Stanford)
- Hours/year: 12M (BG/P), 4.6M (NERSC), ? (Stanford)
- File I/O:
  - Dumping 100 files during run (1-2 Tb total), use these for post-processing
  - MPI-I/O, typically only few % of total time spent on I/O
- Memory usage:
  - Current runs: 0.5 Gb/core or less, clearly not a bottleneck
  - Future runs with chemistry tables: need 1 Gb table on each process
  - Post-processing: completely memory-limited, choosing #cores based on memory

- Anecdotes:
  - Different scratch file-systems at NERSC => have moved large files from Franklin to Euclid for post-processing (not efficient)
  - Post-processing should ideally be “interactive” => queue time is disruptive of the thought-process => users submit in debug queues (ALCF) or resort to transferring files to other (including local) machines with shorter queue times

HPC usage and methods for the next 3-5 years

- More use of unstructured *Charles* code, but otherwise basically the same with simply normal “growth” (of job sizes etc)

- Potential big change: Chemistry-modeling through tables
  - Currently modeling-fidelity is limited by memory on present cores (3-dimensional table since this requires 1 Gb)
    - Highly unsatisfactory to make physical modeling decisions based on memory considerations!
  - Some experimentation with hybrid programming for a Stanford cluster (single 15 Gb table on 24-core node with OpenMP/MPI)

- Strategy for many-core machines: Interactions with Liszt (domain-specific language for CPU/GPU/???) developers at Stanford
  - Implementation of parts of production code imminent

Summary

- *What new science results might be afforded by improvements in NERSC computing hardware?*
  - Dynamic visualization of shock-turbulence interactions: *mechanics of shock-holes*
  - *Improved sub-grid models* for shock-turbulence interaction and multi-material mixing
- *Recommendations on NERSC architecture, system configuration?*
  - We view NERSC as complimentary to ALCF -- this is good!
  - Trustworthiness of systems is vital -- waiting in queue OK if the run will go through and complete
- *What significant scientific progress could you achieve over the next 3 years with access to 50X NERSC resources?*
  - Improved capturing and modeling of inter-scale energy transfer in non-equilibrium compressible turbulence
- *What "expanded HPC resources" are important for your project?*
  - More hours!!
  - Help with visualization
  - Post-processing (beyond viz) is still a bottleneck -- more software than hardware issue
- *General discussion*
  - NERSC resources prepared us for ALCF/INCITE
  - We returned to NERSC for further science exploration