Extreme events, resolution and a new parallel algorithm for turbulent mixing on Blue Waters

P.K. Yeung (PI)
Matthew P. Clay (PhD student)

Georgia Tech (Schools of AE and ME)
E-mail: pk.yeung@ae.gatech.edu

NSF:	PRAC (1036170, 1640771) and Fluid Dynamics Programs
BW Team, Cray:	Scaling, Reservations, Help Requests, Storage, Visualization
Collaborators:	T. Gotoh, S.B. Pope, B.L. Sawford, K.R. Sreenivasan

Blue Waters Symposium, May 16-19, 2017

Introduction: Turbulence and Reynolds Number

Disorderly fluctuations: unsteady, 3D, multiscale, nonlinear

— prevalent in many fields of science and engineering

— effective mixing (coupling with molecular diffusion)

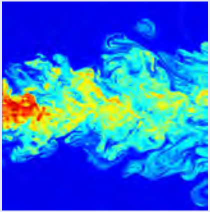

(Images taken from Wikipedia page on Turbulence)

Higher Reynolds no (UL/ν): wider range of scales, more uncertainty, larger number of degrees of freedom; $\Longrightarrow$ more CPU power needed

A Two-Part Presentation

1. Extreme events and resolution effects in simulation of high Reynolds no. turbulence, on world-leading scale
    - Fundamentals of wide implication, despite idealized geometry
    - Fourier pseudo-spectral, benefits from favorable network topology
    - 8192^3 and 16384^3 grid resolution
    - Publication: *Proc. Nat Acad Sci*, Yeung *et al* 2015
2. Turbulent mixing at low molecular diffusivity, with novel algorithm that achieves 6% of peak on BW
    - A multi-resolution problem, suggesting hybrid approaches
    - Velocity on coarser grid, passive scalar on finer grid
    - Compact finite difference, nested OpenMP parallelism
    - *Comput. Phys. Comm.* 2017 (in revision, acceptance likely)

Extreme Events and Turbulence

- High intensity, rare, localized in space and/or short-lived in time
- Fluid elements experiencing extreme local deformation
  - rate of strain (change in shape)
  - rate of rotation (change in orientation)
- Search for self-similarity: scaling exponents of dissipation rate are central in turbulence theory addressing fine-scale intermittency
- Very sensitive to Reynolds number, and more:
  - small-scale resolution and sampling are both important
- On Blue Waters: first 8192^3 simulation of homogeneous isotropic turbulence on a periodic domain, focus on fundamental issues
- A short simulation at 16384^3 has also been performed

The Computational Approach

- Direct numerical simulation (DNS): use exact equations of motion (Navier-Stokes; $\nabla\cdot\mathbf{u} = 0$ for constant density)

$$\partial\mathbf{u}/\partial t + \mathbf{u}\cdot\nabla\mathbf{u} = -\nabla(p/\rho) + \nu\nabla^2\mathbf{u} + \mathbf{f}$$

- Fourier pseudo-spectral: high accuracy, but communication- intensive
  - massive parallelism: 2D (pencils) domain decomposition
- BW: MPI, Co-Array Fortran, 8192^3 w/ favorable topology:
  - 8.897 secs/step on 262,144 cores; 30 secs on 65,536
  - I/O is usually fast: 4 TB in a minute or less
  - postprocessing and on-the-fly processing
  - VISIT for 3D scientific visualization
- To span several large-eddy time scales: $O(10^5)$ time steps

Intermittency and Local Averaging

- Dissipation rate and enstrophy as quadratic measures of local strain and rotation rates (vorticity):

$$\epsilon \equiv 2\nu s_{ij} s_{ij} \quad ; \quad \Omega \equiv \omega_i \omega_i$$

- Kolmogorov (1941): averaged $\langle \epsilon \rangle$ represents rate of energy transfer (cascade) from large scales to small scales
- Kolmogorov Refined Similarity (1962): average locally over a volume of space of linear dimension r at inertial (intermediate) scales

$$\epsilon_r(\mathbf{x}, t) = \frac{1}{Vol} \int_{Vol} \epsilon(\mathbf{x} + \mathbf{r}', t) \; d\mathbf{r}'$$

- Although 3D averages are important, they are not often reported:
  - averaging along a line (1D) is *much* easier
  - 1D surrogate $(\partial u/\partial x)^2$ often used in experiments
  - DNS: also, nontrivial due to domain decomposition

Higher-order Moments of the Local Averages

- At smallest scales, very large higher-order moments, such as $\langle(\epsilon_r/\langle\epsilon\rangle)^m\rangle$ $(m > 1)$ are expected, as a result of intermittency and extreme events (Yeung *et al.*JFM 2012, PNAS 2015):
- Data from high *Re* 8192^3 simulation: ϵ (solid) and Ω (dashed)

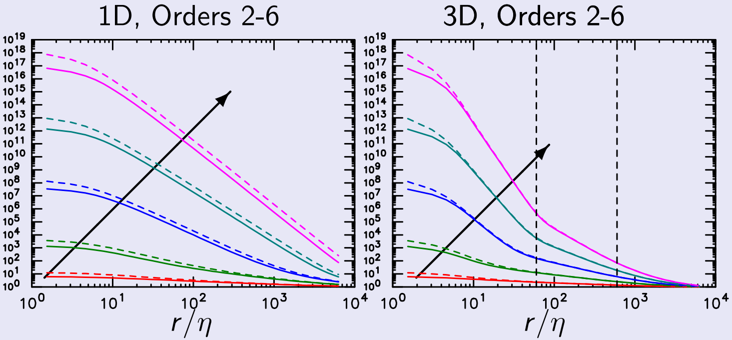

- The importance of 3D averages: strong indication of an “inertial” scaling range around $60 \le r/\eta \le 600$; both ϵ and Ω

Higher-order Scaling Exponents (Preliminary)

- Inertial: $\langle \epsilon_r^q \rangle / \langle \epsilon \rangle^q \propto (r/\eta)^{-\tau_q}$
- Find τ_q: look for best fit of flat region for $\langle \epsilon_r^q \rangle / \langle \epsilon \rangle^q (r/\eta)^{\tau_q}$.
- Orders 2-6, 8192^3 datasets
  0.23, 0.70, 1.40, 2.35, 3.40
  Close to Log-normal theory:
  $$\tau_q = \mu q(q-1)/2$$
  0.23, 0.69, 1.38, 2.30, 3.45
- To compare with past 1D results (Sreenivasan & Antonia 1997)
- Same exponents for ϵ_r (solid) and Ω_r (dashed)

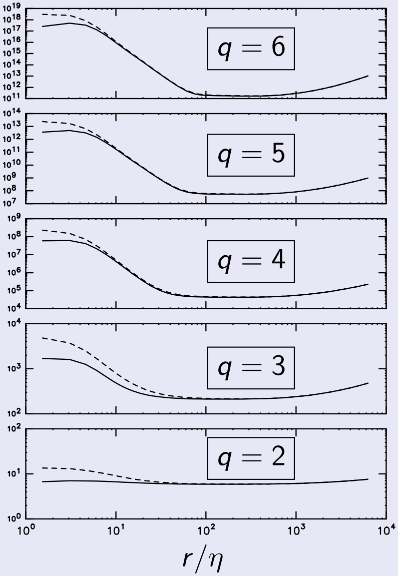

Are these results accurate and reliable?

8192^3 DNS: a decent match in Re for many experiments, but showing much more "extreme" fluctuations than in past literature

How do we assess adequacy of small-scale resolution?

- Refine the grid spacing, run again with same physical parameters
    - perform a yet-larger simulation: expensive, may be unfeasible
    - comparisons contaminated by statistical variability
- Take existing dataset, coarsen grid spacing, compare the results
    - if discrepancies are small, then solution is accurate enough
    - can quantify, e.g. what fraction of extreme events would be missed if resolution were degraded
    - a post-processing task at modest cost, that allows us to isolate effects of truncation error from statistical sampling

Tests of accuracy, up to 16384^3

Compensated plots using same exponents as found for
(a) 8192^3 ensemble-averaged scaling of 3D averages, $k_{max}\eta \approx 2$
(b) single 16384^3 snapshot, by grid refinement, $k_{max}\eta \approx 3.8$
(c) filter from (b) to 8192^3 resolution, $k_{max}\eta \approx 1.9$

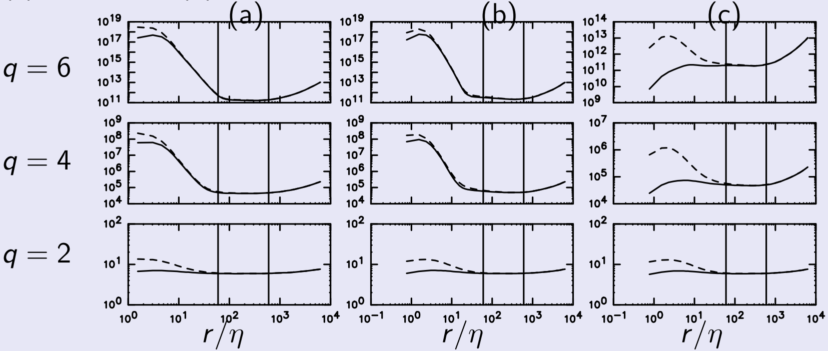

- Great variability at small r (dominated by extreme events)
- But relatively robust scaling in inertial range

Why Blue Waters

- Turbulence as a Grand Challenge in Science:
  - unsteady, 3D, nonlinear, stochastic, wide range of scales
  - smaller simulations often compromised in physics or accuracy needed for applications where turbulence is the critical process
- Turbulence as a Grand Challenge in Computing:
  - first 4096^3 simulation was performed in Japan (2002)
  - on BW: the first production 8192^3 (16X more expensive)
- Would be impossible if not for BW:
  - very large resource allocation on multi-Pflop machine
  - dedicated and expert staff assistance (even late nights!)
  - generous storage capacity (2 PB)

A New Algorithm for Turbulent Mixing at Low Diffusivity

Temperature or concentration fields in a turbulent flow

- Dynamically passive scalars governed by advection-diffusion equation
- The Schmidt number ($Sc = \nu/D$) varies over a wide range
  - *Sc*: $\mathcal{O}(0.01)$ liquid metals, $\mathcal{O}(1)$ gas-phase, $\mathcal{O}(1000)$ salinity in ocean

Low diffusivity is more difficult in both experiment and DNS

- Fluctuations arise at scales smaller than those in velocity field
- Fundamental differences in shape of spectrum, intermittency, etc.

A dual-resolution, dual-numerical-scheme code

- Velocity on coarser grid, scalar on finer grid
- Compact finite differences for scalar (Gotoh *et al.* JCP 2012)
- How do we design a parallel algorithm for best efficiency?

Broad range of scales in scalar field at high Sc (low diffusivity)

- Small scales for velocity field on the order of the Kolmogorov scale η
- Small scales in scalar field given by the Batchelor scale $\eta_B = \eta Sc^{-1/2}$

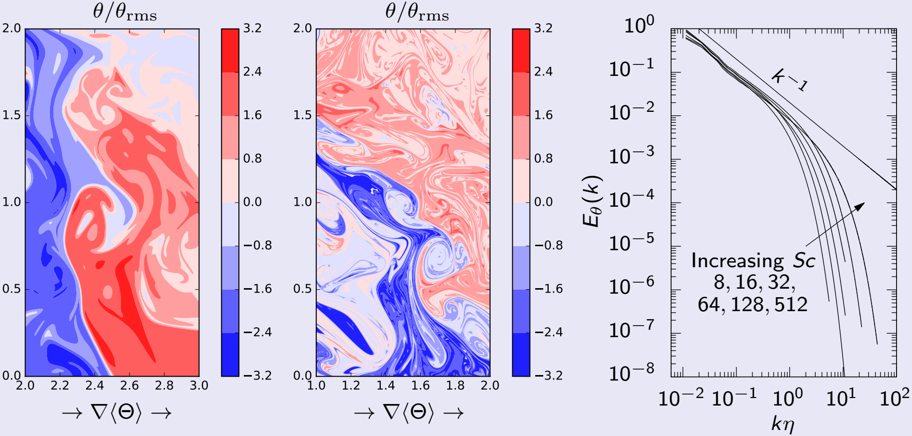

Figure: Scalar (left, 1024^3) $Sc = 8$ and (middle, 8192^3) $Sc = 512$ at $R_\lambda = 140$.

Equations and Dual Numerical Scheme

- Velocity field on coarser grid: Navier-Stokes equations, via usual Fourier pseudo-spectral method (3D FFTs)
- Scalar fluctuations on finer grid (Gotoh *et al.*, JCP 2012), with uniform mean scalar gradient:

$$\partial\theta/\partial t + \mathbf{u}\cdot\nabla\theta = D\nabla^2\theta - \mathbf{u}\cdot\nabla\langle\Theta\rangle$$

  - Eighth-order combined compact finite differences (Mahesh, JCP 1998)
  - Computes first and second derivatives in all 3 directions
  - Much less communication than typical FPS codes

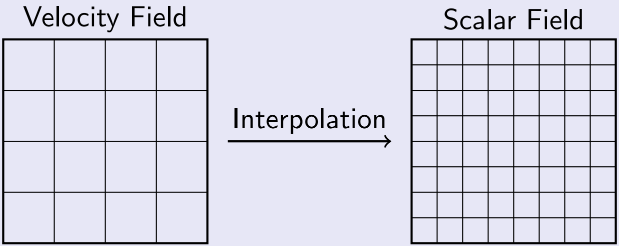

Algorithm for $Sc \gg 1$ on Blue Waters

Want high Sc, while ensuring accuracy at moderate Reynolds no.

- Velocity: $R_\lambda = 140$, $N_v = 1024$, $k_{\max,v}\eta = 5.6$ (512 cores)
- Scalar: $Sc = 512$, $N_\theta = 8192$, $k_{\max,\theta}\eta_B = 2.0$ (262,144 cores)

Use disjoint groups of processors for velocity and scalar fields

- To form advective term, send well-resolved velocity field to scalar communicator, and perform tricubic interpolation
- Overlap inter-communicator transfer with CCD operations on scalar

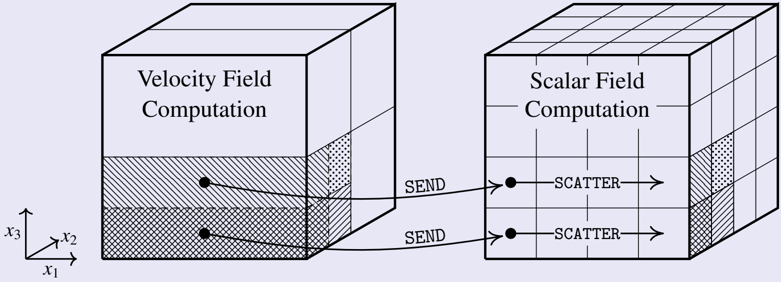

CCD Scheme & Opportunities to Improve Scalability

- Scheme is implicit: all points along grid line coupled. Must solve linear system $A\mathbf{x} = \mathbf{b}$ for each grid line in all three coordinate directions.
- To avoid memory transposes, adopt a static 3D domain decomposition
  - Implies that no processor ever has data in-core to solve CCD system
  - Adopt parallel algorithm (Nihei *et al.*) to solve system

Op.	Operation Summary
A	Fill ghost layers for scalar field with `SEND` and `RECV` operations
B	Form right-hand-side of linear system and obtain solution
C	Pack and distribute data for reduced system with `MPI_ALLTOALL`
D	Unpack data and solve reduced linear system
E	Pack and distribute data for final solution with `MPI_ALLTOALL`
F	Unpack data and finalize solution of CCD linear system

- Operations for three coordinate directions are independent
  - Try to overlap communication with computation

Overlapping Communication with Computation

Achieve overlap by interleaving communication and computation operations for all three coordinate directions in one subroutine

1. Use non-blocking communication calls from MPI
   - Post communication call for next coordinate direction (e.g., x_2) before proceeding with computations for current direction (e.g., x_1)
   - Use `MPI_WAIT` to ensure results are ready, when needed
2. Using dedicated communication threads in a MPI/OpenMP approach
   - Goal: one thread per NUMA domain to communicate, while others compute
   - Thread synchronization: use OpenMP locks
     - Use one lock for each coordinate direction
     - Thread must obtain lock for a given coordinate direction before doing work
   - Work-sharing the computations: use nested OpenMP parallelism
     - Initial comput. thread spawns nested parallel region to use rest of threads
     - Loops cannot be partitioned evenly: explore `GUIDED` and `DYNAMIC` scheduling

Using Dedicated Communication Threads

tid=0

```
CALL OMP_SET_LOCK(x2)
CALL OMP_SET_LOCK(x3)

! Spin until the x1 lock is set.
test=.TRUE.
DO WHILE(test)
   test=OMP_TEST_LOCK(x1)
   IF (test) CALL OMP_UNSET_LOCK(x1)
END DO

COMMUNICATE x2 [A2]
CALL OMP_UNSET_LOCK(x2)

COMMUNICATE x3 [A3]
CALL OMP_UNSET_LOCK(x3)

CALL OMP_SET_LOCK(x1)
COMMUNICATE x1 [C1]
```

tid=1

```
CALL OMP_SET_NUM_THREADS(nth-1)
!$OMP PARALLEL

!$OMP MASTER
CALL OMP_SET_LOCK(x1)
! Spin until the x2 lock is set.
! Spin until the x3 lock is set.
!$OMP END MASTER
!$OMP BARRIER

!$OMP DO; COMPUTE x1; !$OMP END DO

!$OMP MASTER
CALL OMP_SET_LOCK(x2)
CALL OMP_UNSET_LOCK(x1)
!$OMP END MASTER
!$OMP BARRIER
```

Scalability of New Hybrid PSDNS-CCD3D Code

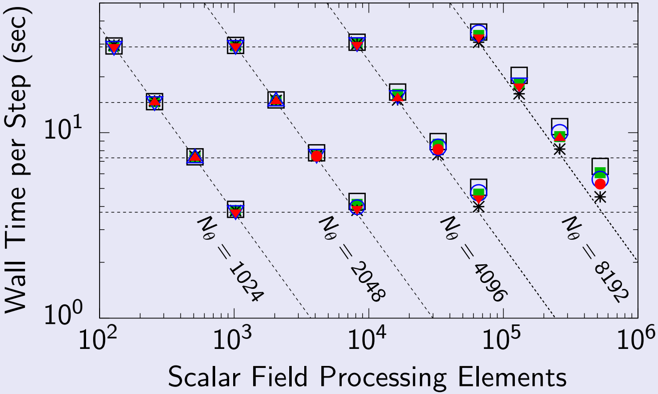

Figure: Scalability of scalar field computation using different versions of the CCD routines: □ single-threaded, blocking; △▽○ (2,4,8 threads) multi-threaded, blocking; ■ single-threaded, overlapped; ▲▼● multi-threaded, overlapped; ∗ one dedicated communication thread per NUMA domain.

Spectrum of Passive Scalar at High Schmidt Number

- Theory predicts k^{-1} in the viscous-convective range $(1/\eta \ll k \ll 1/\eta_B)$
- Kraichnan: exponential, not Gaussian (Batchelor), in diffusive range

$$E_\theta(k) = C_B \langle \chi \rangle (\langle \epsilon \rangle / \nu)^{-1/2} k^{-1} (1 + \sqrt{6 C_B} k \eta_B) \exp(-\sqrt{6 C_B} k \eta_B)$$

- Considerable scatter in C_B (Donzis *et al.* FTC 2010, Gotoh *et al.* 2014)

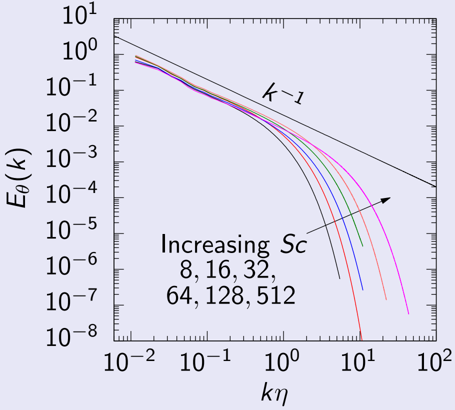

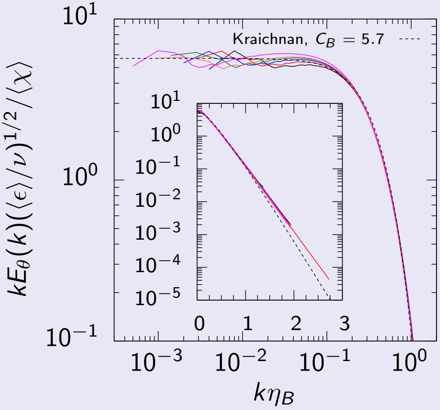

Concluding Remarks

- Simulations of turbulence at 8192^3 grid resolution conducted using PRAC allocation of BW resources:
  - isotropic turbulence at high Reynolds number w/ good scale resolution
  - study of turbulent dispersion under the same conditions (reported at BW Symposium last year)
  - turbulent mixing at high Schmidt number, using a newly-developed advanced parallel algorithm (Clay *et al.* CPC 2017)
- Ongoing and future work (re: PRAC renewal award)
  - A penultimate Petascale computational turbulence laboratory
  - Magnetohydrodynamic turbulence (at low magnetic Reynolds no.)