Progress in Petascale Computations of Turbulence at high Reynolds number

P. K. Yeung
Georgia Tech, pk.yeung@ae.gatech.edu

Blue Waters Symposium
NCSA, Urbana-Champaign, May 2014

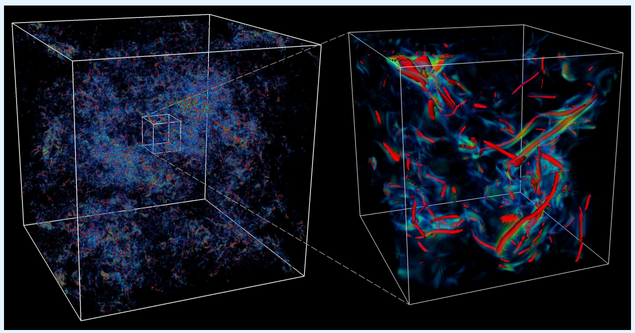

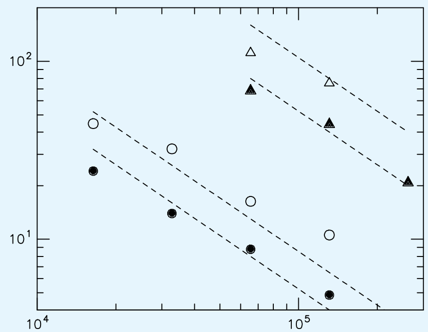

Thanks...

NSF: PRAC and Fluid Dynamics Programs
Blue Waters Project & Cray staff
Co-PI's, collaborators, graduate students

- *Key Challenges*
  - Fluid turbulence at high Reynolds number
  - Towards 8192^3 simulation (over 0.5 trillion grid points)
- *Why It Matters*
  - Fundamental advances, useful in engineering and environment
  - Range of scales, and spatial resolution
- *Why Blue Waters*
  - Resources needed for the first production 8192^3 simulation
  - Excellent, dedicated support via PRAC program
- *Accomplishments and Science Goals*
  - Intermittency, including extreme dissipation and enstrophy
  - Towards studies of turbulent mixing and dispersion

Turbulence at High Reynolds Number

Fluctuations: Unsteady, 3D, stochastic, wide range of scales, nonlinear

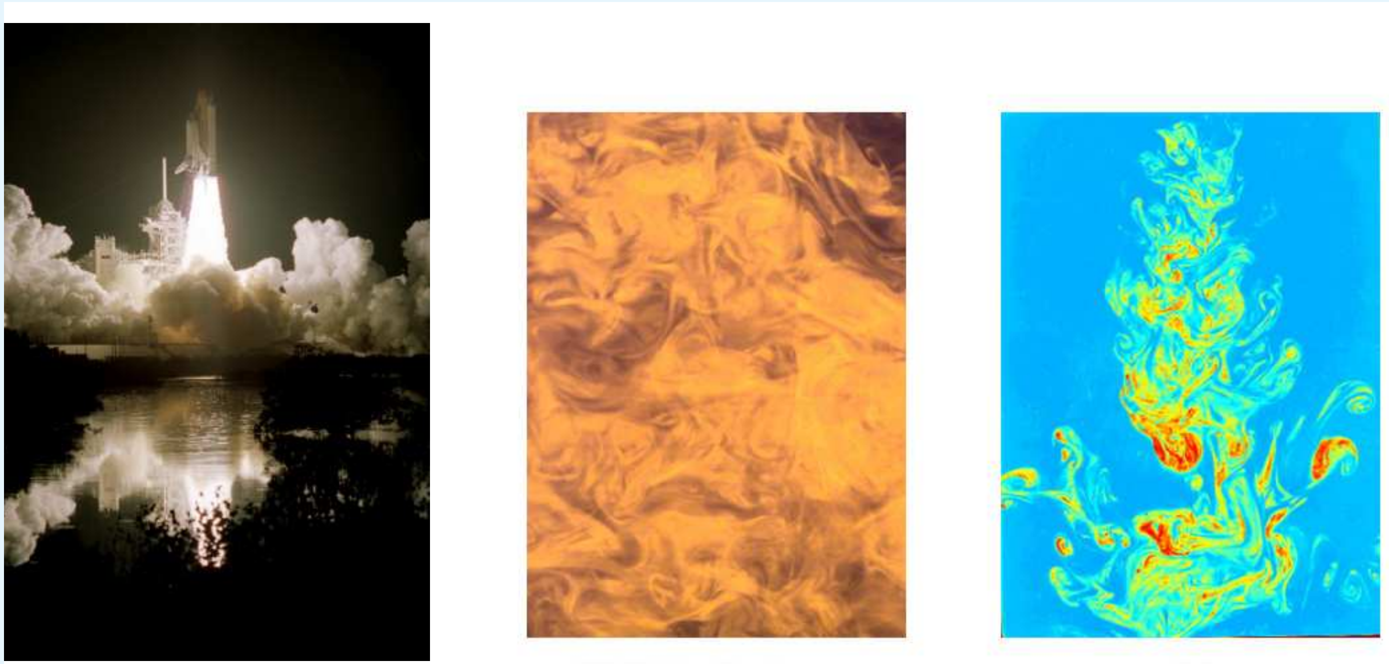

Space Shuttle | **Grid turbulence** | **Jet**

Understanding turbulence is the key to improved engineering devices and prediction/management of natural phenomena
— astrophysics, combustion, meteorology, marine ecology, etc.

About large-scale DNS

- Direct numerical simulation: solve exact equations of motion, massive detail helpful for understanding and modeling
- Yokokawa *et al.* 2002 performed the first 4096^3 DNS, on Earth Simulator (see also Ishihara, Gotoh & Kaneda, ARFM 2009)
- Yet, path to the next frontier has not been easy:
  - codes both computation and communication-intensive
  - requirements increase rapidly with problem size
  - large, multi-year allocation needed on multi-Petaflop machine
- Several interrelated components:
  - algorithmic development and performance enhancement
  - data generation: $O(10^5)$ time steps, w/ suitable parameters
  - data analysis: quantitative and qualitative

2D Domain Decomposition

- Partition a cube along two directions, into “pencils” of data

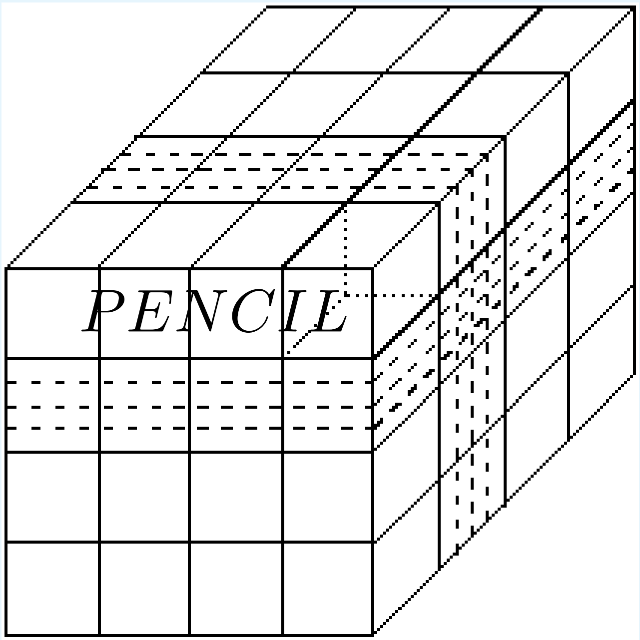

- MPI: 2-D processor grid, $M_1(\text{rows}) \times M_2(\text{cols})$

3D FFT, starting with pencils in x:

- Transform in x
- Transpose to pencils in z
- Transform in z
- Transpose to pencils in y
- Transform in y

Transposes by message-passing, collective communication, with non-contiguous data

Performance: Communication

Main bottleneck for scalability at high core counts is in collective (alltoall) communication, especially if using pure MPI

- OpenMP multithreading (share memory on multicored processors)
  - fewer MPI processes, lower communication overhead
  - favorable at very large core counts, subject to memory access
- Remote memory addressing (Fortran Co-Arrays, Cray Compiler)
  - declare major buffers as co-arrays, accessible to other processes
  - one-sided “get” operation for pairwise exchange
  - copy of data between regular and co-arrays
- Best performance is obtained when running on a reserved partition designed to minimize contention from network traffic

Performance (Reserved Partitions)

- 262144 MPI tasks, Fortran co-arrays, single-prec, RK2

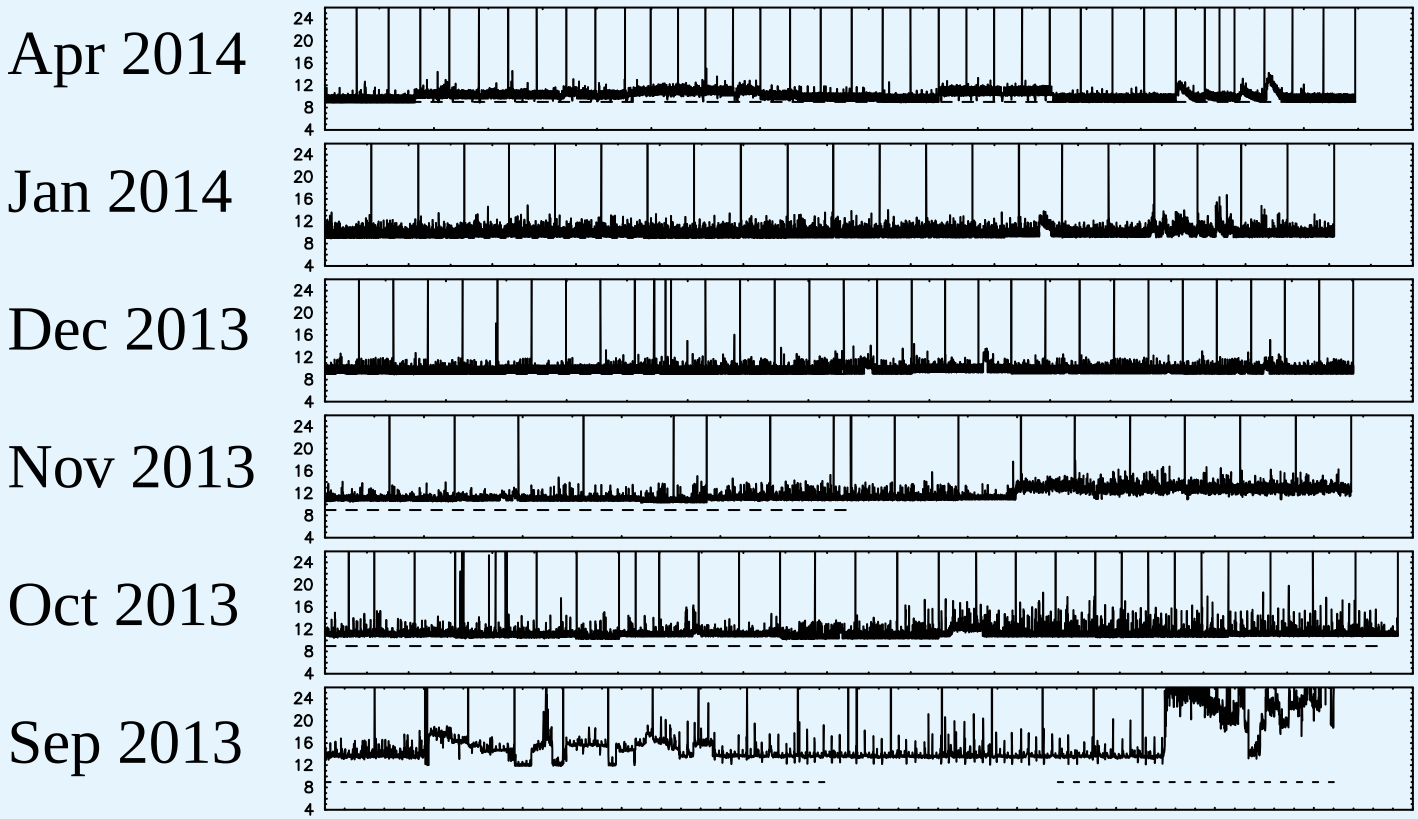

- Best timing was 8.897 secs/step; with other traffic minimized
- I/O on Blue Waters is good: 40 secs to write 8192^3 checkpoint

8192^3 DNS

- A 8192^3 DNS (although arduous) is now in progress (up to $8192 \times 32 = 262144$ Cray XE cores, on Blue Waters)
- Forced turbulence in stationary state, many snapshots saved
- Cannot have it all: careful planning is required — targeting $R_\lambda \approx 1300$, $\Delta x/\eta \approx 1.5$

What are some of the science questions that cannot be settled w/o good data at both high Reynolds number and high resolution?

- Fine-scale intermittency
- Turbulent mixing (Eulerian, fixed reference frame)
- Turbulent dispersion (Lagrangian, moving reference frame)
- ....

Energy Spectrum

- $E(k)\langle\epsilon\rangle^{-2/3}k^{5/3}$: relatively wide inertial range ($1/L_1 \ll k \ll 1/\eta$)

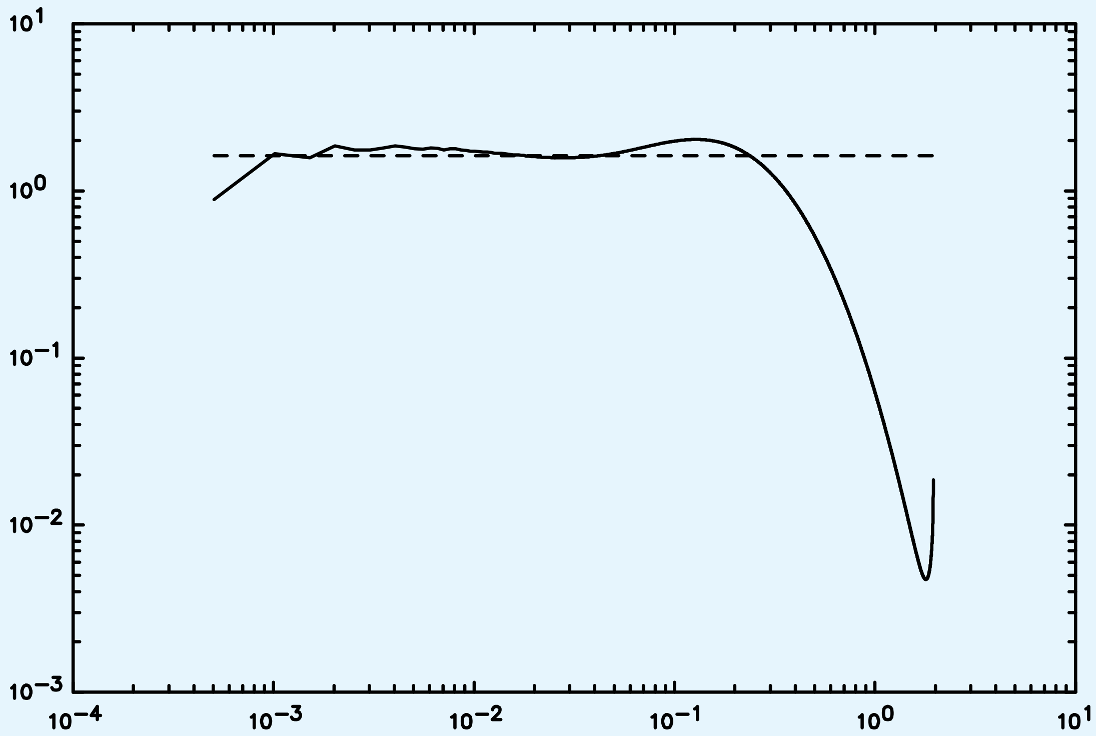

$k\eta$

- Consistent with Kolmogorov constant ≈ 1.62, "bottleneck" at $k\eta \approx 0.2$, and small intermittency correction in exponent

Dissipation and Enstrophy

- Dissipation: $\epsilon = 2\nu s_{ij} s_{ij}$ (strain rates squared)
  Enstrophy: $\Omega = (\nu)\omega_i\omega_i$ (rotation rates squared)
- Intense strain rate can break a flame surface and lead to combustion instabilities
- In relative dispersion, straining pulls particle pairs apart but rotation makes them move around together
- Same mean values in homogeneous turbulence. Do extreme fluctuations in each scale similarly at high Reynolds no.?
- Joint distribution of ϵ and Ω determines statistics of

$$\nabla^2(p/\rho) = (\Omega - \epsilon/\nu)/2$$

3D Visualization

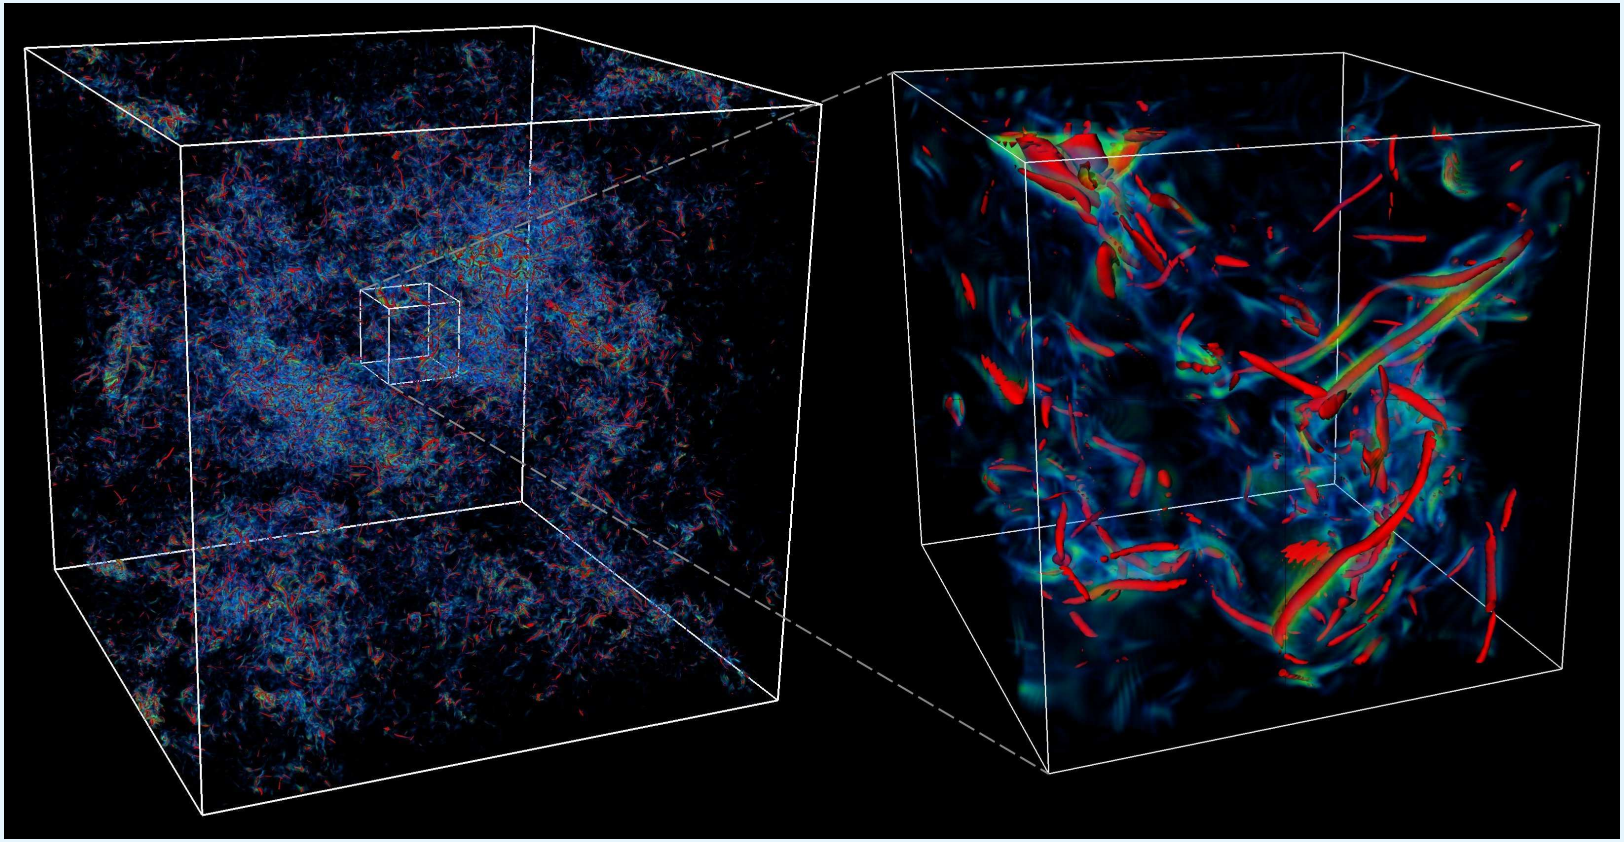

[TACC visualization staff] 4096^3, $R_\lambda \approx 650$: intense enstrophy (red) has worm-like structure, while dissipation (blue) is more diffuse.

Visualization of 8192^3 data on BW

1024^3 sample sub-cube, thanks to D. Semeraro; parallel VISIT

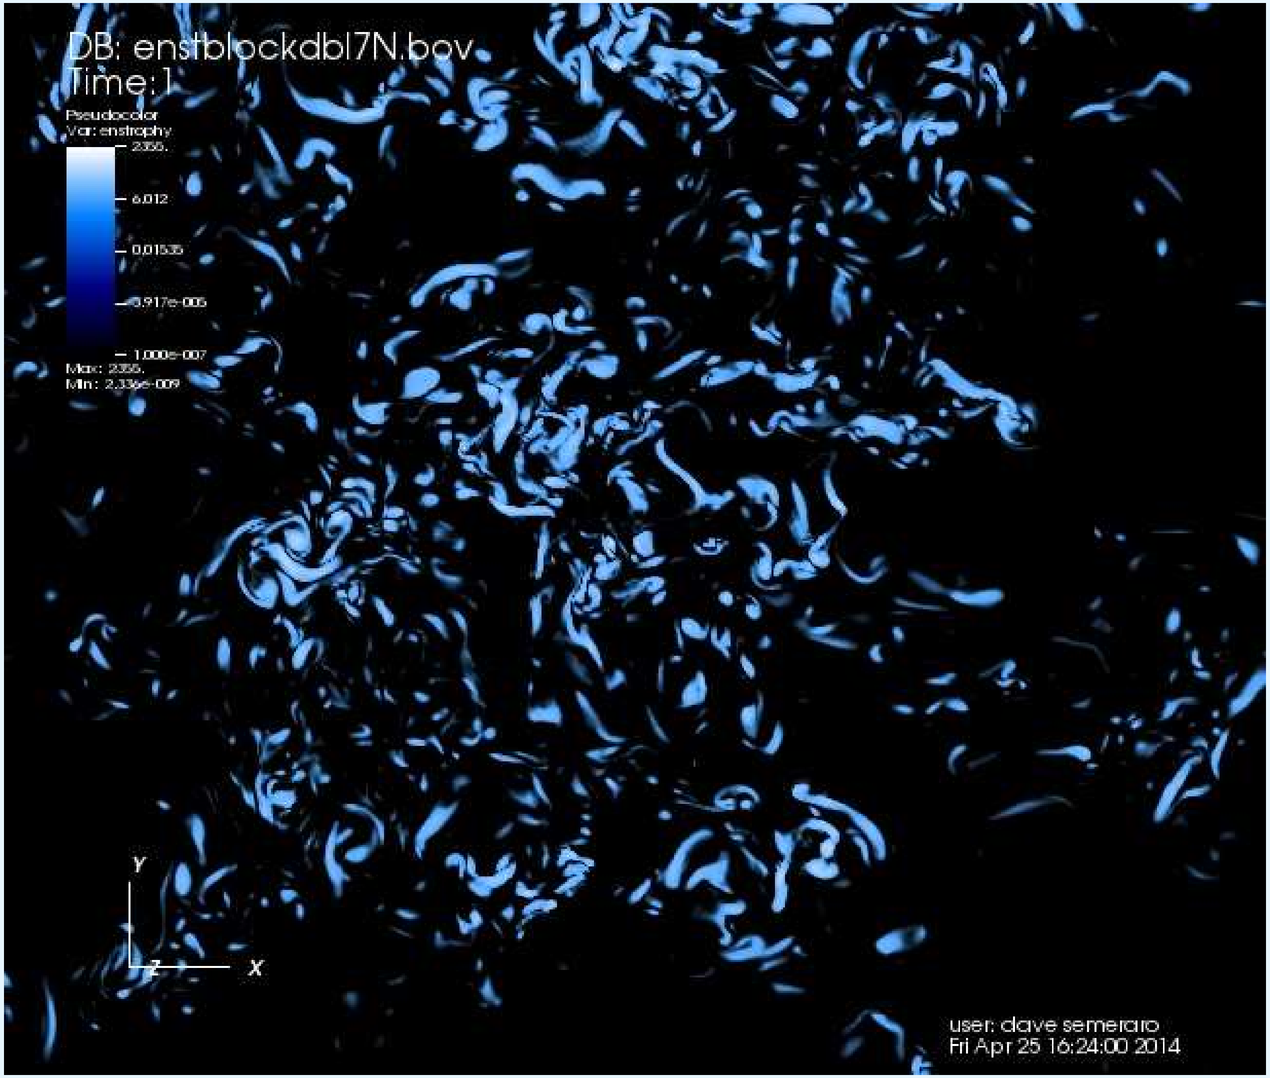

PDFs of Dissipation and Enstrophy

- Yeung, Donzis & Sreenivasan JFM June 2012 (Vol 700)
- Highest Re and best-resolved at moderate Re (both 4096^3)

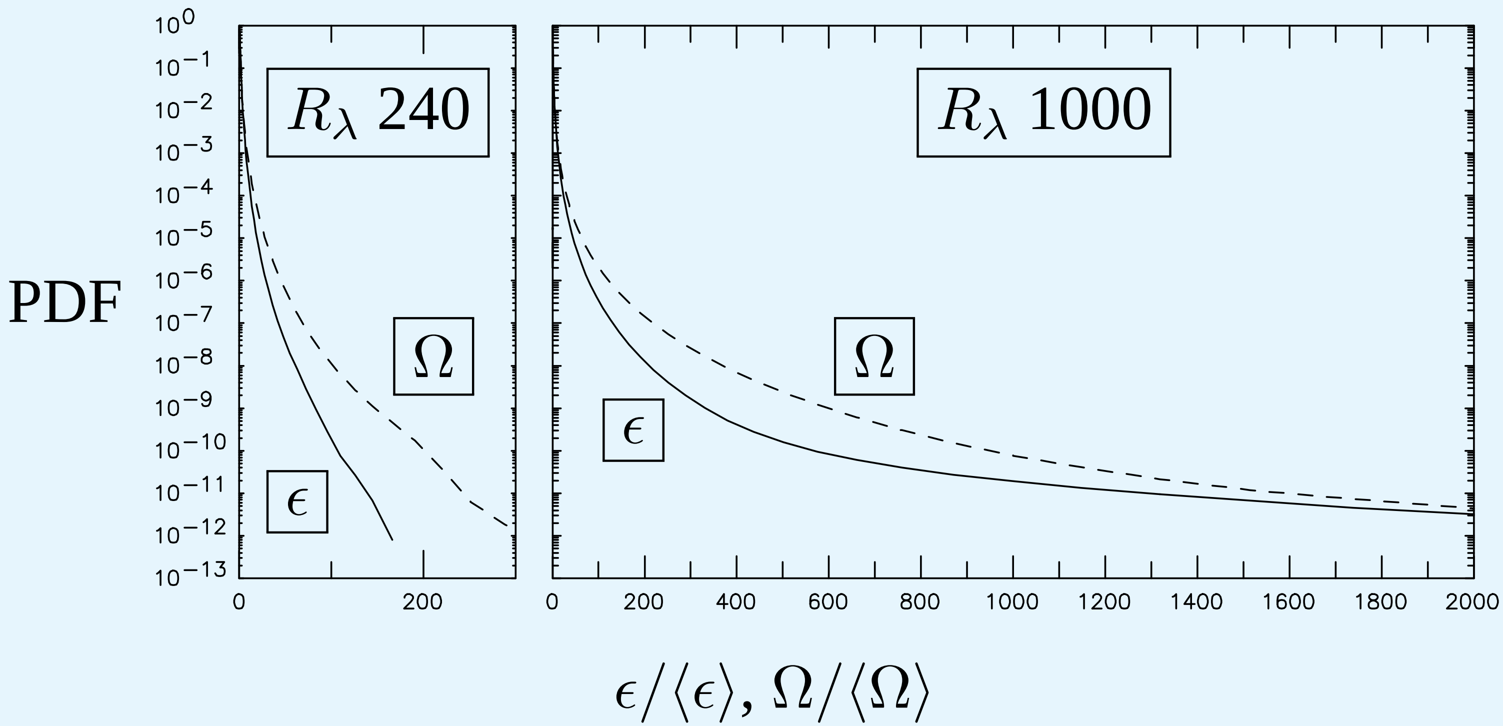

- High Re: most intense events in both scale similarly
- Study trends from data at different Reynolds numbers

Preliminary results at 8192^3

- Dissipation and enstrophy PDFs, 14 single-time snapshots:

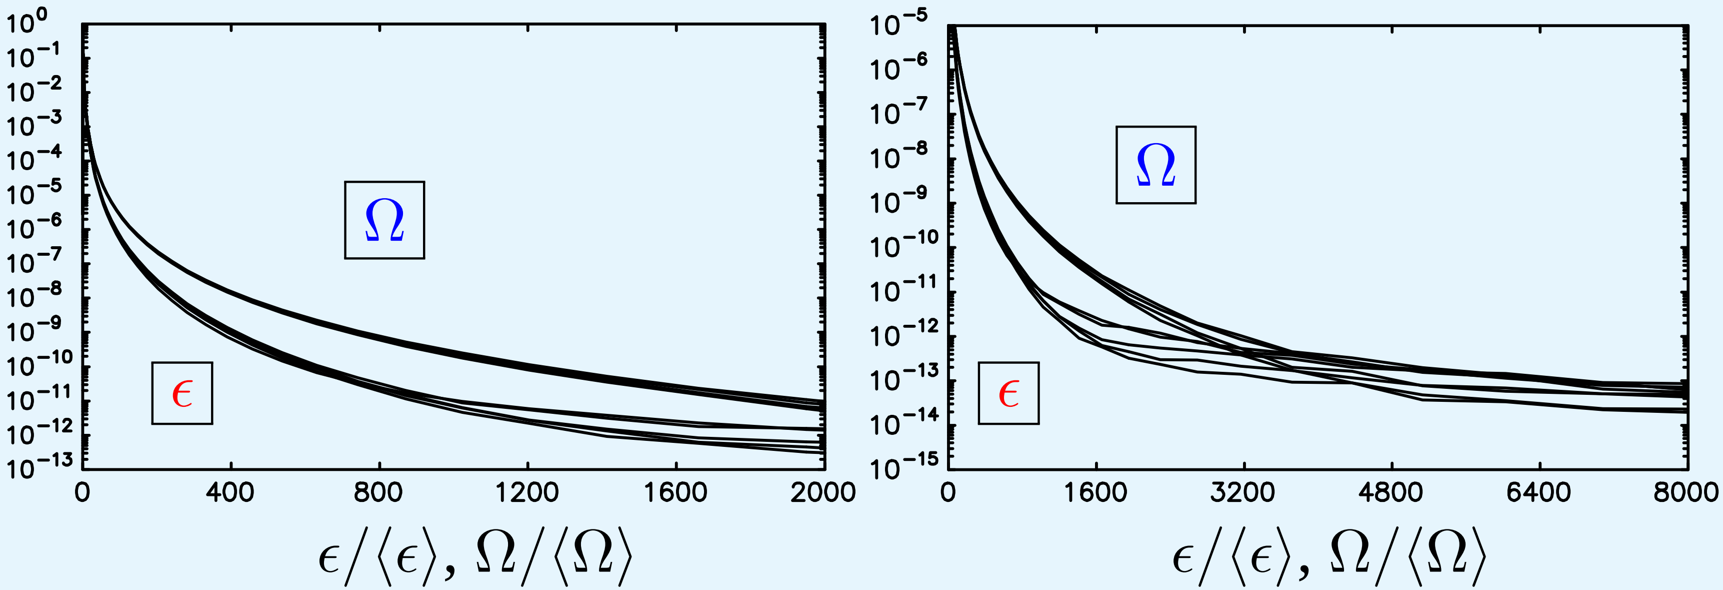

- Two distinct groups of lines, merging at ~ 4000 times the mean (*qualitatively* consistent with Yeung *et al.* JFM 2012, but slightly higher Reynolds no. and better resolved)

Statistics of the Acceleration

- Lagrangian view: rate of change of fluid particle velocity
- Eulerian view: material derivative of velocity, as in N-S equations

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

- Acceleration PDF has very wide tails
  (probability of extreme fluctuations is significant)
- A fluid particle in a region of large velocity gradients is expected to experience large acceleration (Yeung *et al.* JFM 2007)
- Modeling of acceleration is crucial in capturing Reynolds number dependence in stochastic modeling (Sawford 1991)

Acceleration PDF

- Compare properties of $D\mathbf{u}/Dt$ with $\partial\mathbf{u}/\partial t$ and $\mathbf{u}\cdot\nabla\mathbf{u}$

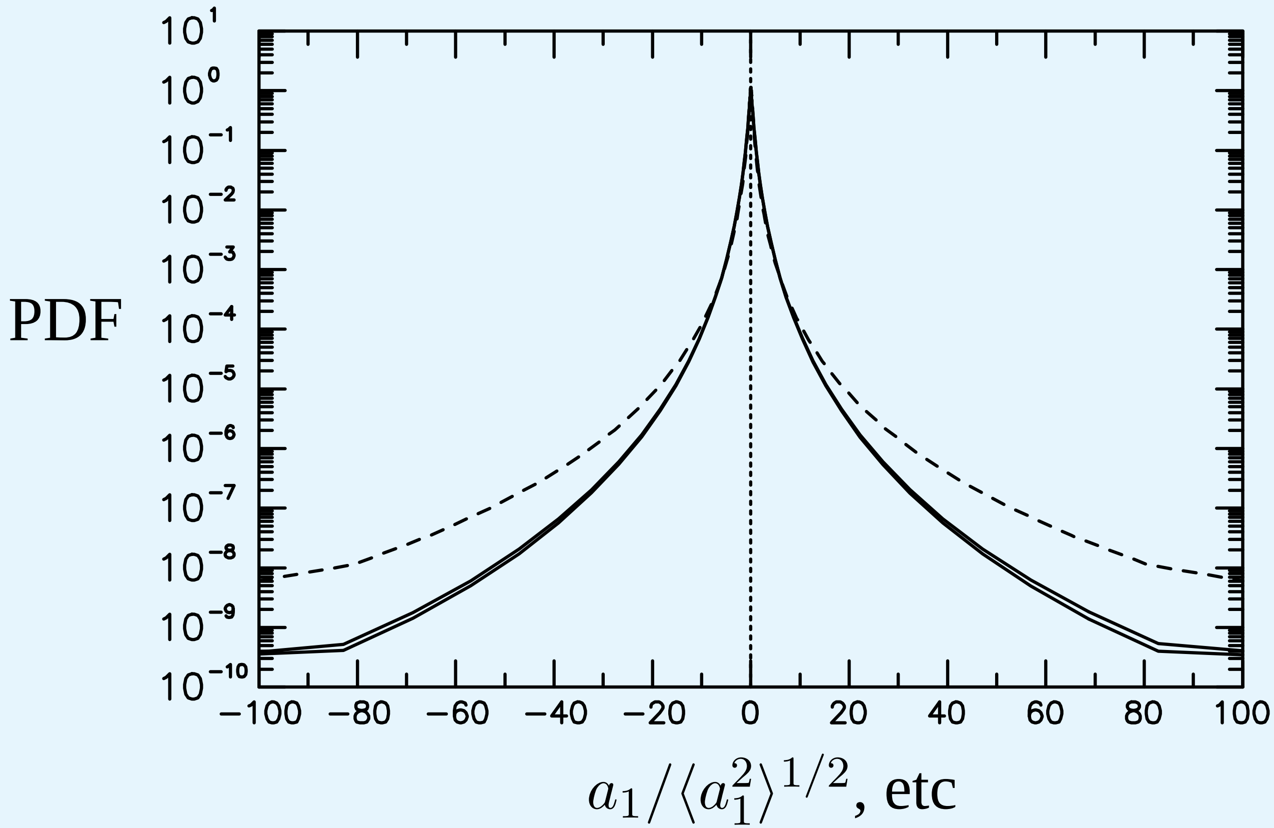

- Total acceleration more intermittent than both local and convective (strong mutual cancellation between the two contributions)

Turbulent Mixing: Scaling Regimes

- Scaling ranges depend on $Sc \ll 1$, $O(1)$, or $\gg 1$:

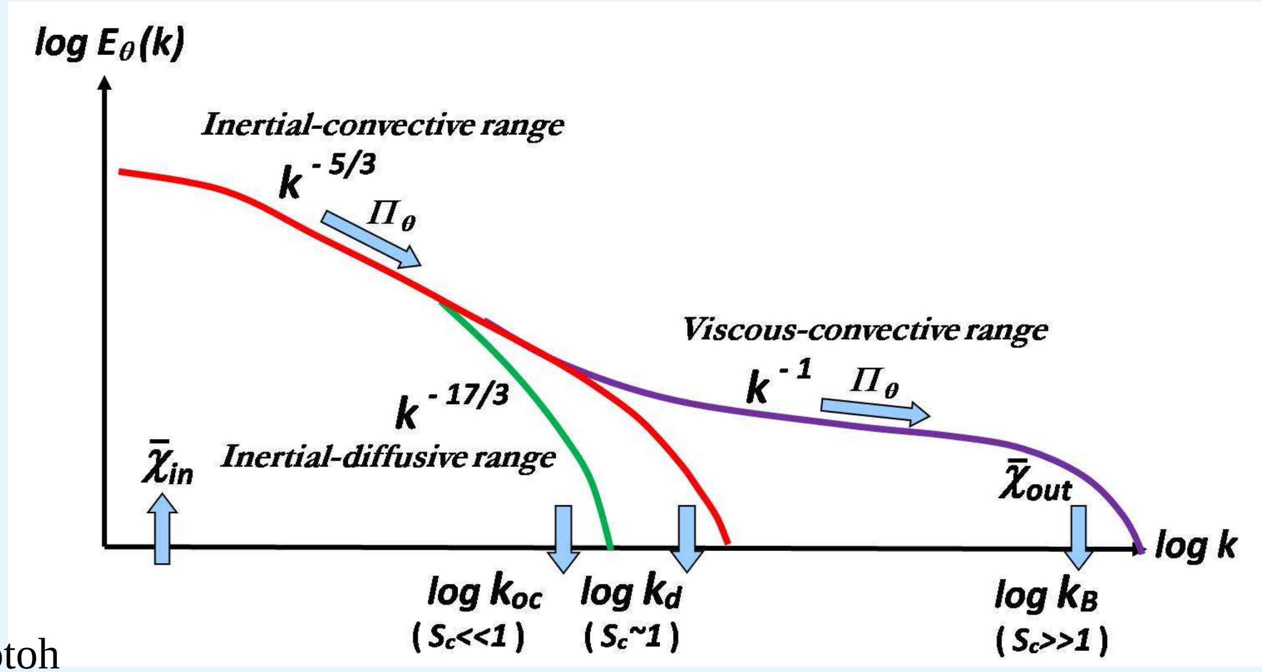

Courtesy T. Gotoh

- Inertial convective range, at high Reynolds number:
  - but scalar field needs better resolution
  - study anomalous scaling (Gotoh & Yeung 2013)

Turbulent Dispersion (Lagrangian)

- Lagrangian viewpoint of fluid mechanics:
  - continuous fluid medium represented by a large number of infinitesimal fluid elements (fluid particles)
  - fluid particle velocity by interpolation from Eulerian grid points at instantaneous particle position
- Monin & Yaglom 1975: Kolmogorov similarity for Lagrangian statistics (dependent on a wide range of time scales)
  - will need very high Reynolds numbers, but some statistics scale more readily (Sawford & Yeung PoF 2011)
- Multiparticle clusters stretched and deformed by turbulence can be related to statistics of contaminant concentration from localized sources (e.g. in air-quality modeling)

Other Work and Future Plans

Onwards with 8192^3 DNS: track fluid particles and solve for transported scalar fields

Challenges ahead: visualization, overlapping computation w/ communication, data analyses

Flows with more complex physics: solid-body rotation, stratification, magnetohydrodynamics