Towards Petascale Simulations of Turbulence and Turbulent Mixing: Recent results and Future Opportunities

Diego A. Donzis

donzis@tamu.edu

Department of Aerospace Engineering, Texas A&M University

SciDAC 2010—Chattanooga, TN
July 14, 2010

Acknowledgments

- Collaborators:
  - Science: R.D. Moser (UT Austin), J.J. Riley (Wash. U.), K.R. Sreenivasan (New York U.), P.K. Yeung (GaTech)
  - Computing: D. Pekurovsky, A. Majumdar (SDSC)
- Supercomputing centers
  - *Oak Ridge National Lab*: D.M. Frederick, D.B. Kothe, H. Marsha, M.A. Matheson, W. Renaud,...
  - *Texas Advanced Computing Center:* W.L. Barth, J. Boisseau, K. Gaither, C. Hempel, T. Minyard, K. Schulz,...
  - *National Institute for Computational Sciences:* P. Andrews, P. Kovatch, B. Loftis, K. Wong,...
- **INCITE** and TeraGrid allocations (TACC, NICS)

Main themes

- Direct numerical simulation (DNS) of turbulence
- Resolution requirements, HPC, trends
- Turbulent mixing

Direct Numerical Simulations

- Turbulence: most common state of fluid motion
  - random fluctuations over a wide range of non-linearly interacting scales
  - range of scales increases with Reynolds number R_λ (very high in applications)

Direct Numerical Simulations

- Turbulence: most common state of fluid motion
  - random fluctuations over a wide range of non-linearly interacting scales
  - range of scales increases with Reynolds number $\mathbf{R_\lambda}$ (very high in applications)
- Navier-Stokes eqns.: conservation of mass and momentum

$$\begin{aligned} \nabla \cdot \mathbf{u} &= 0 \\ \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} &= -\frac{1}{\rho}\nabla p + \nu \nabla^2 \mathbf{u} \end{aligned}$$

- **Direct Numerical Simulations** (DNS): resolve all "dynamically relevant" length and time scales
  - tremendous detail (beyond experiments), systematic variation of parameters
  - computationally intensive in both time and memory

Numerical Method

- Fourier pseudo-spectral in space (Rogallo 1981)
  - accurate and efficient: Fast Fourier Transforms with operations count proportional to $N^3 \log_2 N$
  - optimized FFT libaries (FFTW, IBM-ESSL)
  - modes formally decoupled in wavenumber space, hence readily parallelizable
- 2nd and 4th order Runge-Kutta in time
- Most time-consuming task: 3D FFT
  - domain decomposition: critical

Three-dimensional FFTs

- Decomposition in "**pencils**": 2D processor grid $(M_1 \times M_2)$

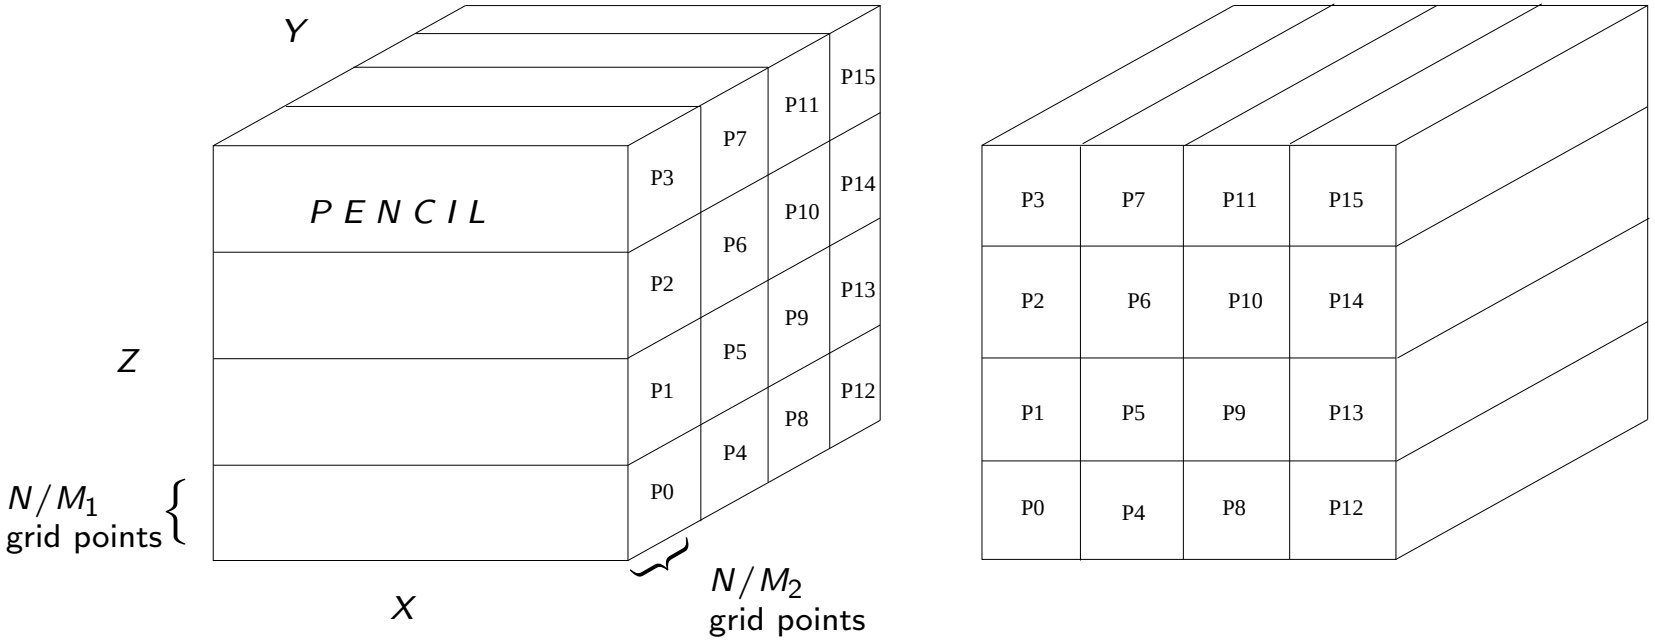

- FFT(X)$\rightarrow$ transpose$\rightarrow$ FFT(Z) $\rightarrow$transpose$\rightarrow$ FFT(Y)
- Two collective communications among $\sim \sqrt{M}$ procs.
- $\boxed{M \text{ up to } N^2}$ (Very general: allows $M_1 \neq M_2$ or not integer factors of N)

Strong scaling

- For a given problem size N^3: $\boxed{time/step \propto M^{-1}}$ if perfect

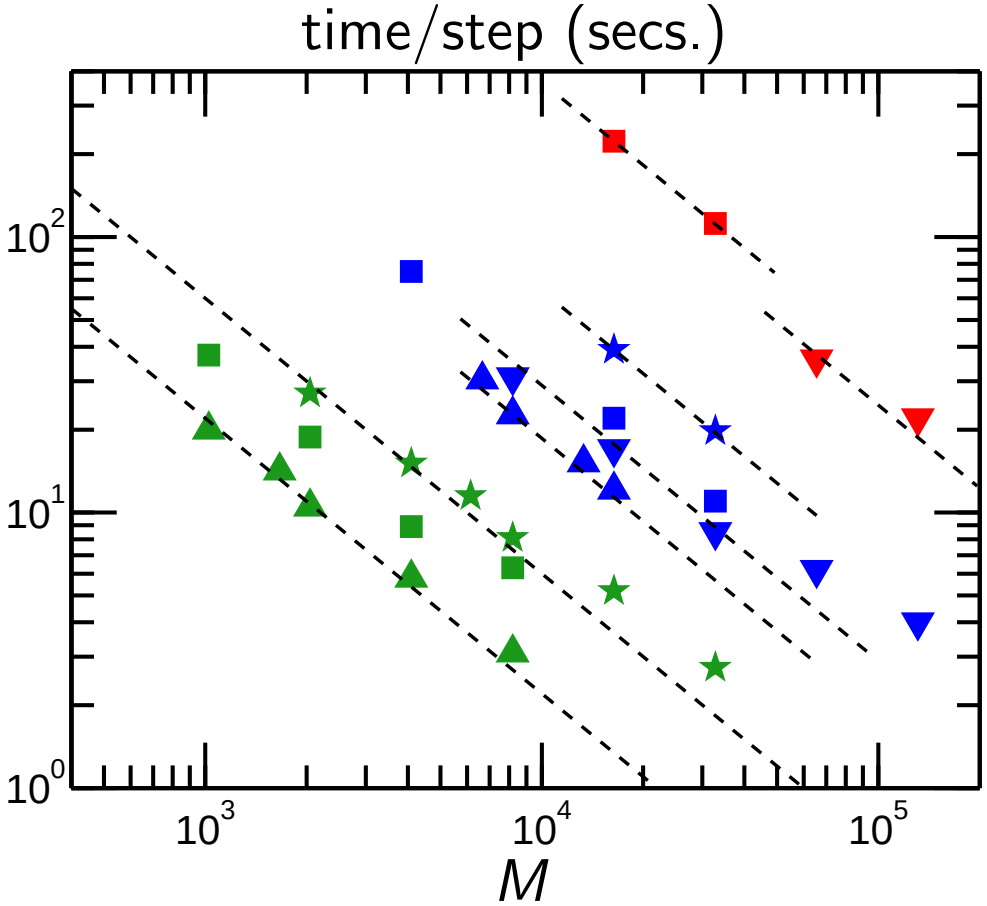

$N^3 = 2048^3, 4096^3, 8192^3$
Squares: Ranger (TACC)
Stars: Blue Gene
Up triangles: Cray XT4
Down triangles: Cray XT5

- Excellent scaling e.g.: $> 95\%$ from $M = 16k$ to $32k$
- Sensitive to proc. grid $M_1 \times M_2$ (depends on network topology)
- Bottleneck: communication

Strong scaling

- For a given problem size N^3: $\boxed{time/step \propto M^{-1}}$ if perfect

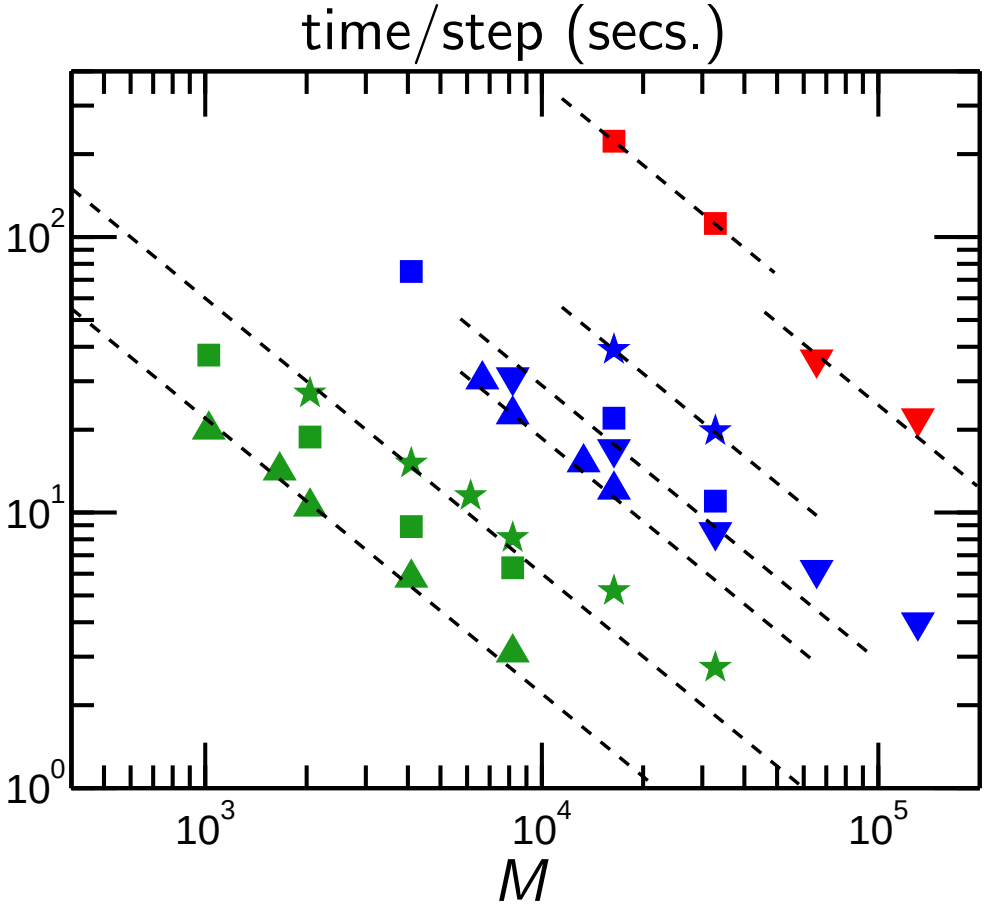

$N^3 = 2048^3, 4096^3, 8192^3$
Squares: Ranger (TACC)
Stars: Blue Gene
Up triangles: Cray XT4
Down triangles: Cray XT5

- Excellent scaling e.g.: $> 95\%$ from $M = 16k$ to $32k$
- Sensitive to proc. grid $M_1 \times M_2$ (depends on network topology)
- Bottleneck: communication

- 4096^3 in production, matching world-record in published work
- Largest (and best resolved) simulation of turbulent mixing

The "dynamically relevant" scales

Full range of scales in space and time

Size of domain	$L_0 > L$	(integral length scale)
Grid spacing	$\Delta x \lesssim \eta$	(Kolmogorov length scale)
Length of simulation	$T \gg T_E$	(eddy-turnover time)
Time step	$\Delta t \ll \tau_\eta$	(Kolmogorov time scale)

The "dynamically relevant" scales

Full range of scales in space and time

Size of domain	$L_0 > L$	(integral length scale)
Grid spacing	$\Delta x \lesssim \eta$	(Kolmogorov length scale)
Length of simulation	$T \gg T_E$	(eddy-turnover time)
Time step	$\Delta t \ll \tau_\eta$	(Kolmogorov time scale)

- Classical textbook scale estimates (K41) give:

$$\mathbf{N^3 \sim (L/\eta)^3 \sim R_\lambda{}^{9/2}} \quad \text{and} \quad \mathbf{T_E/\tau_\eta \sim R_\lambda}$$

Computational work: $\mathbf{W \sim R_\lambda{}^6}$

- Moore's law: R_λ grows exponentially in time

HPC and DNS evolution

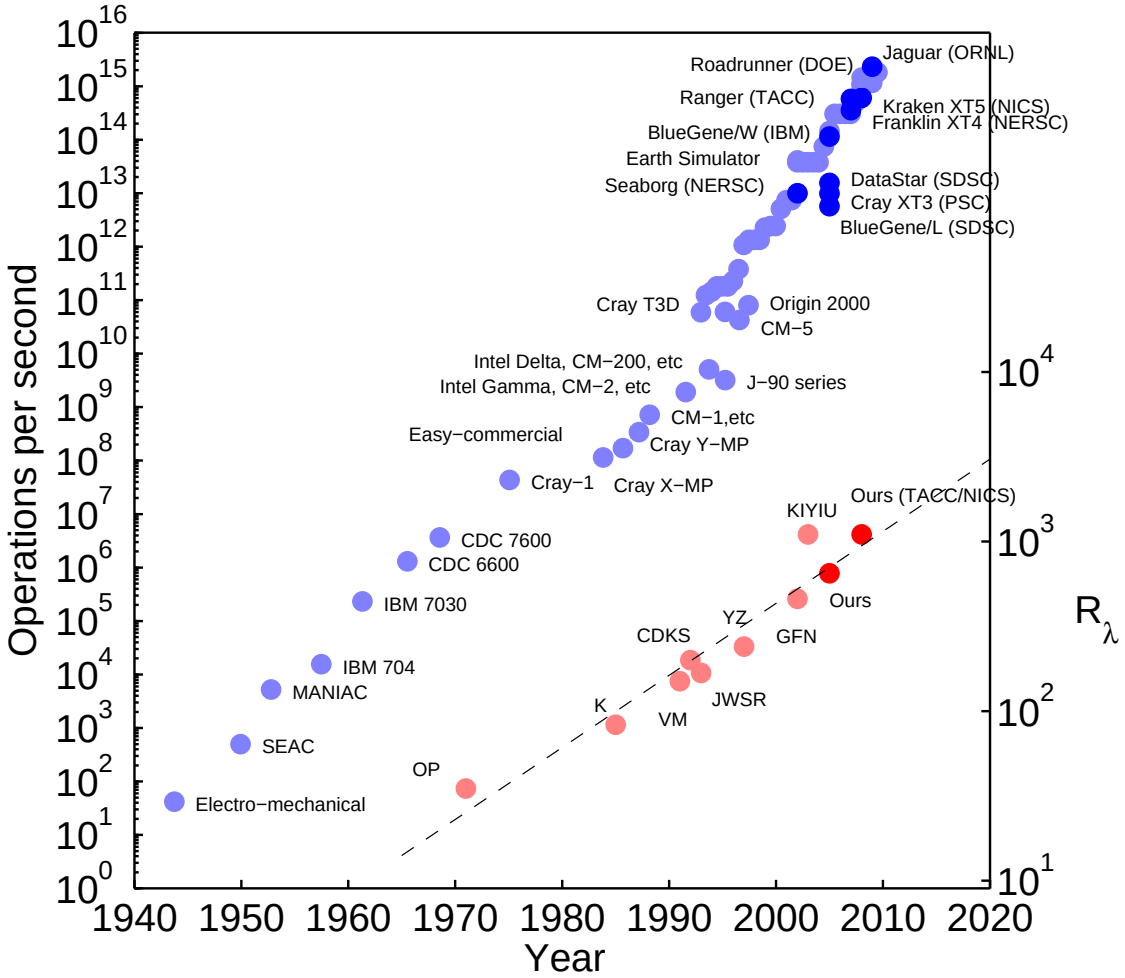

- $R_\lambda \approx 10^{0.042 * year - 82}$

HPC and DNS evolution

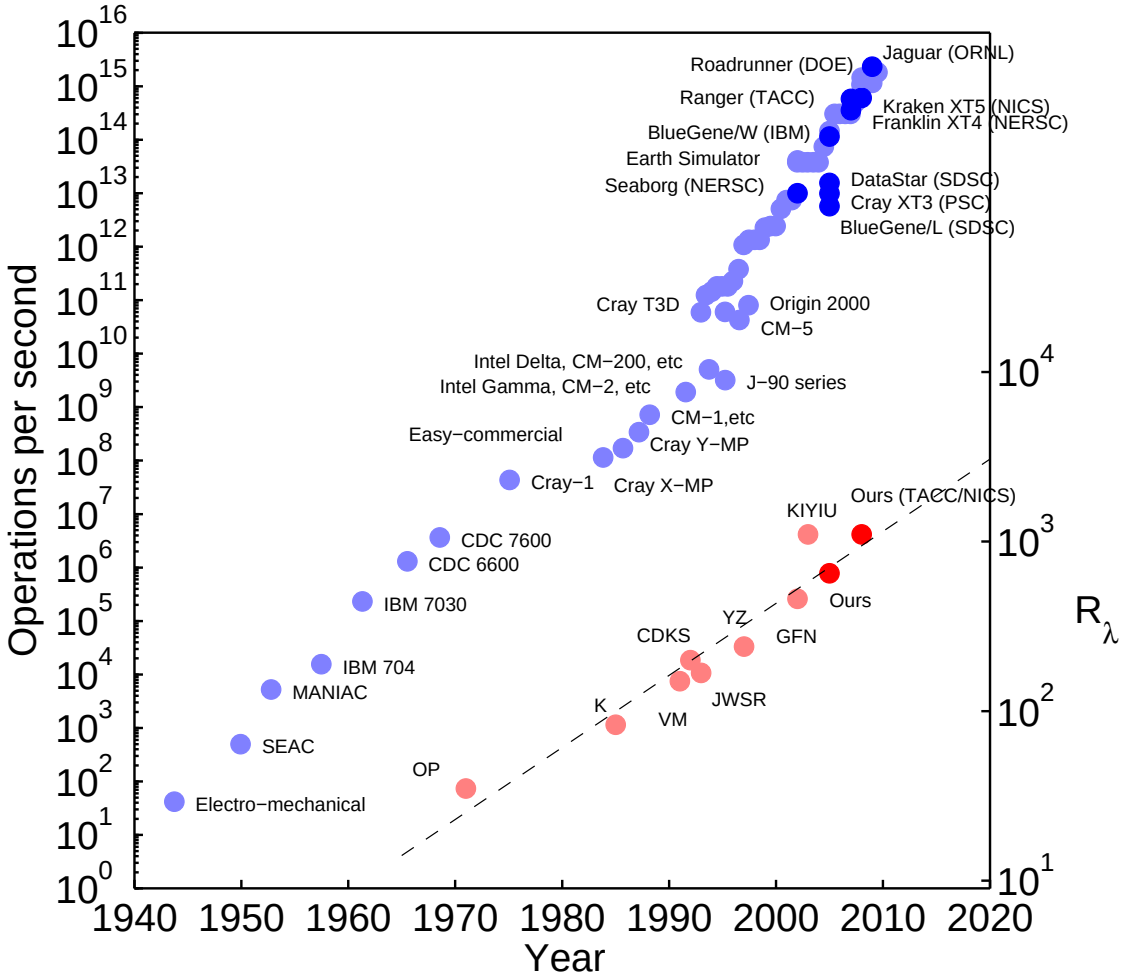

- $\boxed{R_\lambda \approx 10^{0.042 * year - 82}}$
- In 2020 (Exascale?): $R_\lambda \approx 3000$
- But is $W \sim {R_\lambda}^6$ too optimistic?
- Yakhot & Sreenivasan 2005: $W \sim {R_\lambda}^8$!
  In 2020: $R_\lambda \approx 500$!

Small-scale intermittency

Strong, localized fluctuations in space/time, especially at high R_λ

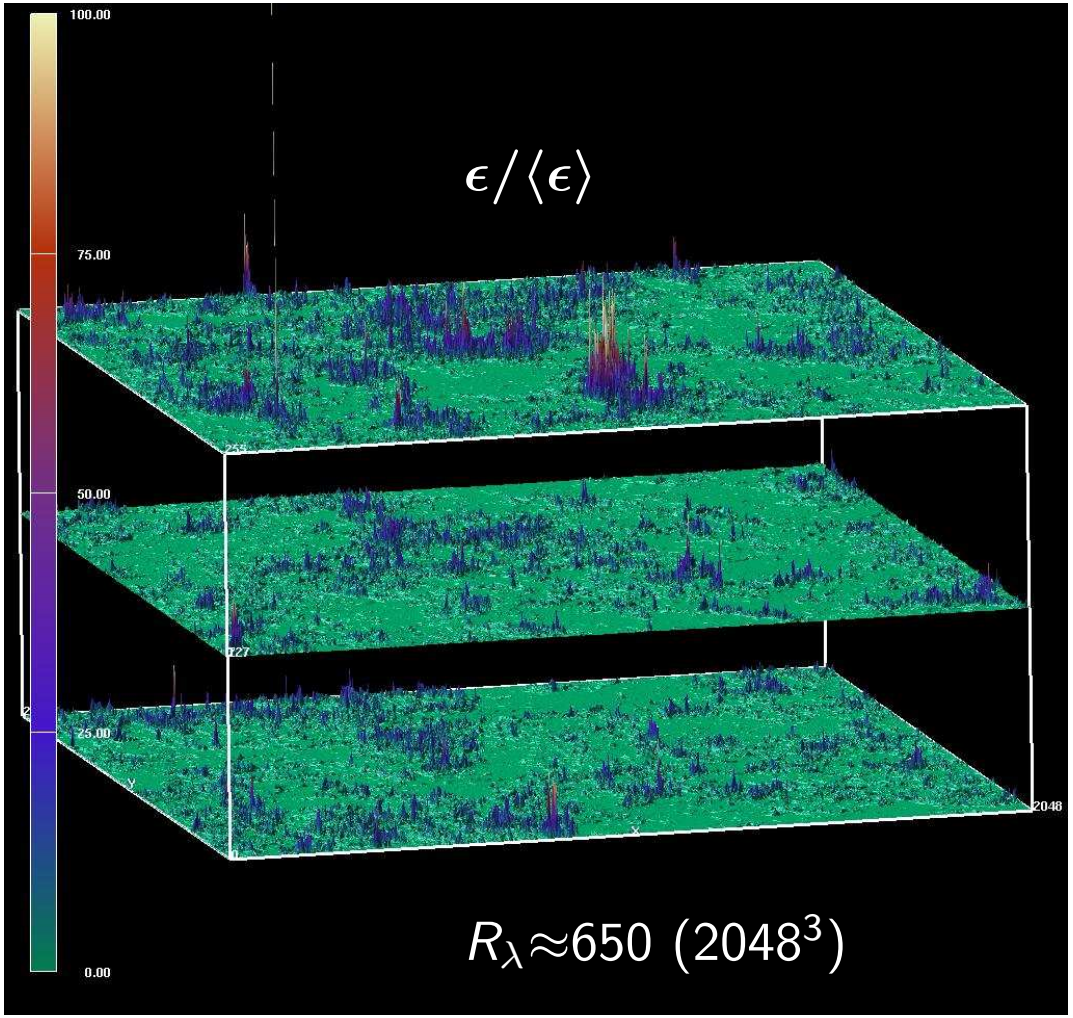

- Instantaneous energy dissipation rate
$$\epsilon = 2\nu s_{ij} s_{ij}$$
- What is the smallest scale we need to resolve?

Small-scale resolution

- For spectral codes, resolution criterion typically based on

$$\mathbf{k_{max}\eta \approx \frac{2.98}{\Delta x/\eta}}$$

where $k_{max} = \sqrt{2}N/3$

- DNS aimed at high R_λ: $k_{max}\eta \approx 1-2$ $\quad (\Delta x/\eta \approx 1.5-2)$

DNS data

- Small scale activity: **dissipation rate ϵ** and **enstrophy Ω**
- Database at $R_\lambda \approx 140$ and $\mathbf{k_{max}\eta}$ from **1.5** to **11** (up to 2048^3)

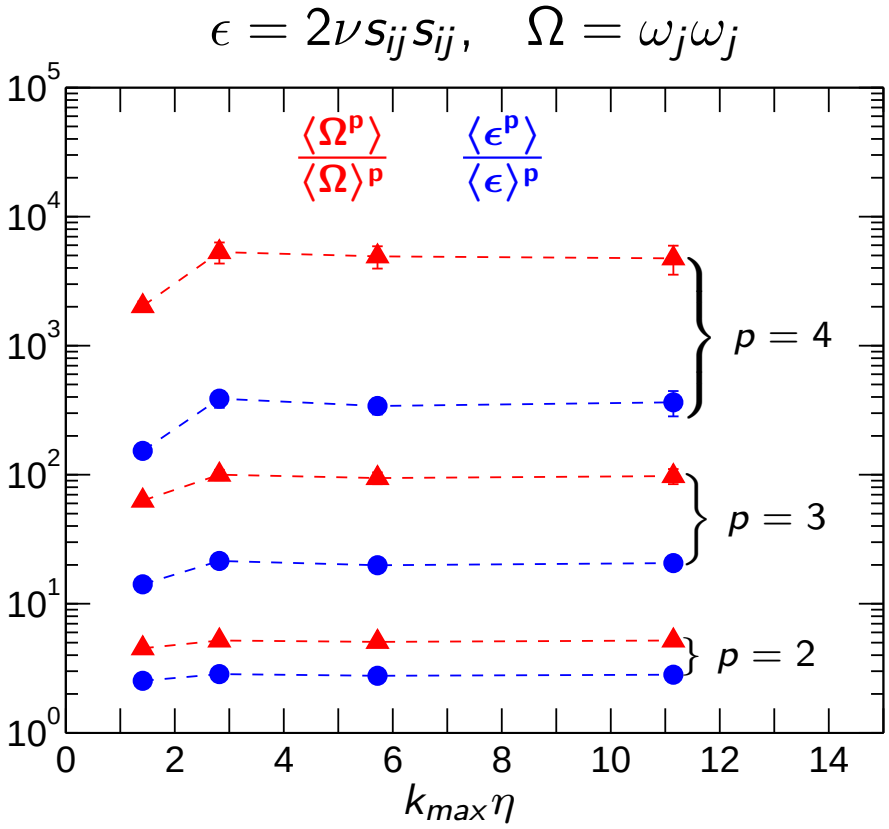

Donzis, Yeung & Sreenivasan PoF '08

- Standard $k_{max}\eta \approx 1.5$ underestimates moments
- $k_{max}\eta \approx 3$ enough for $p \lesssim 4$
- Error increases with order

DNS data

- Small scale activity: **dissipation rate ϵ** and **enstrophy Ω**
- Database at $R_\lambda \approx 140$ and $\mathbf{k_{max}\eta}$ from **1.5** to **11** (up to 2048^3)

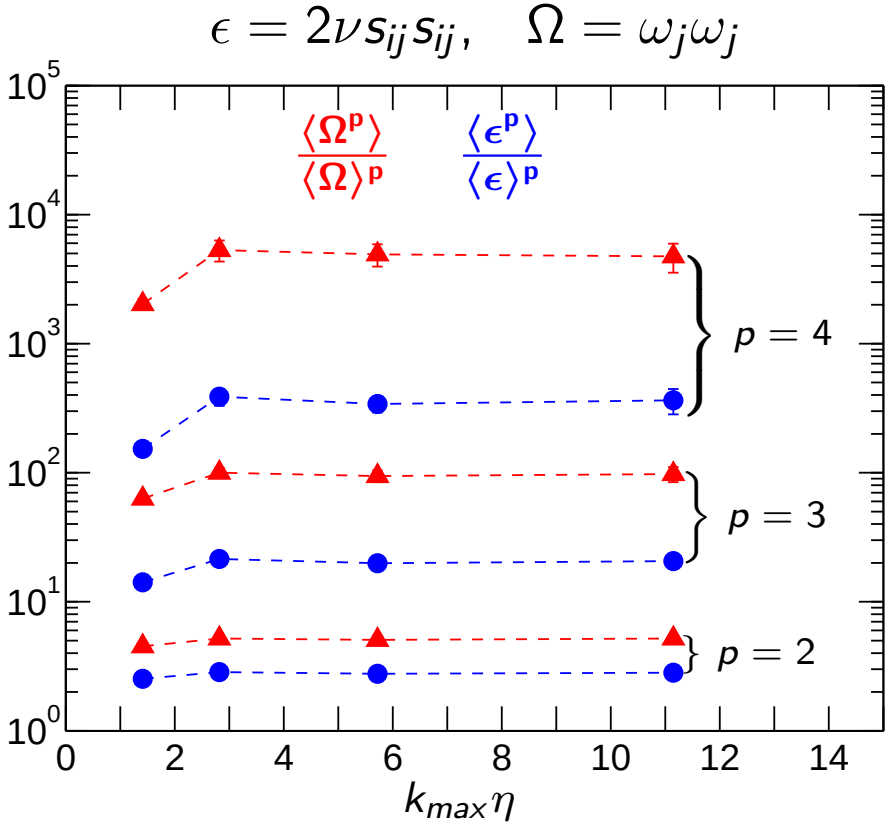

Donzis, Yeung & Sreenivasan PoF '08

- Standard $k_{max}\eta \approx 1.5$ underestimates moments
- $k_{max}\eta \approx 3$ enough for $p \lesssim 4$
- Error increases with order
- Resolution requirement depends on the quantity of interest!

$$W \sim p^4 R_\lambda^{6.16}$$

- So, what do we want to compute in 2020?

Mixing of passive scalars: parameters and scales

- Turbulence can mix very efficiently; orders of magnitude faster than molecular processes
- Scalar fluctuations: $\boxed{\partial\phi/\partial t + \mathbf{u}\cdot\nabla\phi = -\mathbf{u}\cdot\nabla\Phi + D\nabla^2\phi}$
- Characterized by Schmidt number $\mathbf{Sc} \equiv \nu/\mathbf{D}$
  - $O(10^{-2})$ liquid metals, ionized plasmas
  - $O(1)$ temperature in air
  - $O(10^3)$ dye or saline solution in water, organic liquids
  - $O(10^6)$ smoke, clouds

Mixing of passive scalars: parameters and scales

- Turbulence can mix very efficiently; orders of magnitude faster than molecular processes
- Scalar fluctuations: $\boxed{\partial\phi/\partial t + \mathbf{u}\cdot\nabla\phi = -\mathbf{u}\cdot\nabla\Phi + D\nabla^2\phi}$
- Characterized by Schmidt number $\mathbf{Sc} \equiv \nu/\mathbf{D}$
  - $O(10^{-2})$ liquid metals, ionized plasmas
  - $O(1)$ temperature in air
  - $O(10^3)$ dye or saline solution in water, organic liquids
  - $O(10^6)$ smoke, clouds
- If Sc is large:
  - smallest scale (Batchelor): $\eta_\mathbf{B} = \eta\mathbf{Sc}^{-1/2}$ $(\eta_B \ll \eta)$
  - difficult in experiments and simulations
  - scaling less understood

Additional constraints

- Classical phenomenology: $W \sim {R_\lambda}^6 Sc^2$ if η_B is to be resolved (subject to CFL condition)

Additional constraints

- Classical phenomenology: $W \sim {R_\lambda}^6 Sc^2$ if η_B is to be resolved (subject to CFL condition)
  - **But numerically unstable at high R_λ!** (unexpected)

Additional constraints

- Classical phenomenology: $W \sim {R_\lambda}^6 Sc^2$ if η_B is to be resolved (subject to CFL condition)
  - **But numerically unstable at high R_λ!** (unexpected)
- Stability analysis for 1D scalar equation:
  - stability limits:
    $CFL \sim Sc^{-2/3}$ (RK2) and $CFL \sim Sc^{-1/2}$ (RK4)
  - consistent with 3D tests
- So,

$$\boxed{W \sim {R_\lambda}^{6.16} Sc^{5/2}} \qquad \boxed{R_\lambda Sc^{2/5} \approx 10^{0.041 * year - 80}}$$

- In 2020: $R_\lambda Sc^{2/5} \approx 2500$

The physics of mixing: local isotropy

- Isotropy at small scales, regardless of large-scale geometrical features (e.g. mean gradient $\nabla\Phi$)
- Implication: odd-order moments of gradients should vanish
  - separate Sc and R_λ effects

The physics of mixing: local isotropy

- Isotropy at small scales, regardless of large-scale geometrical features (e.g. mean gradient $\nabla\Phi$)
- Implication: odd-order moments of gradients should vanish
  - separate Sc and R_λ effects

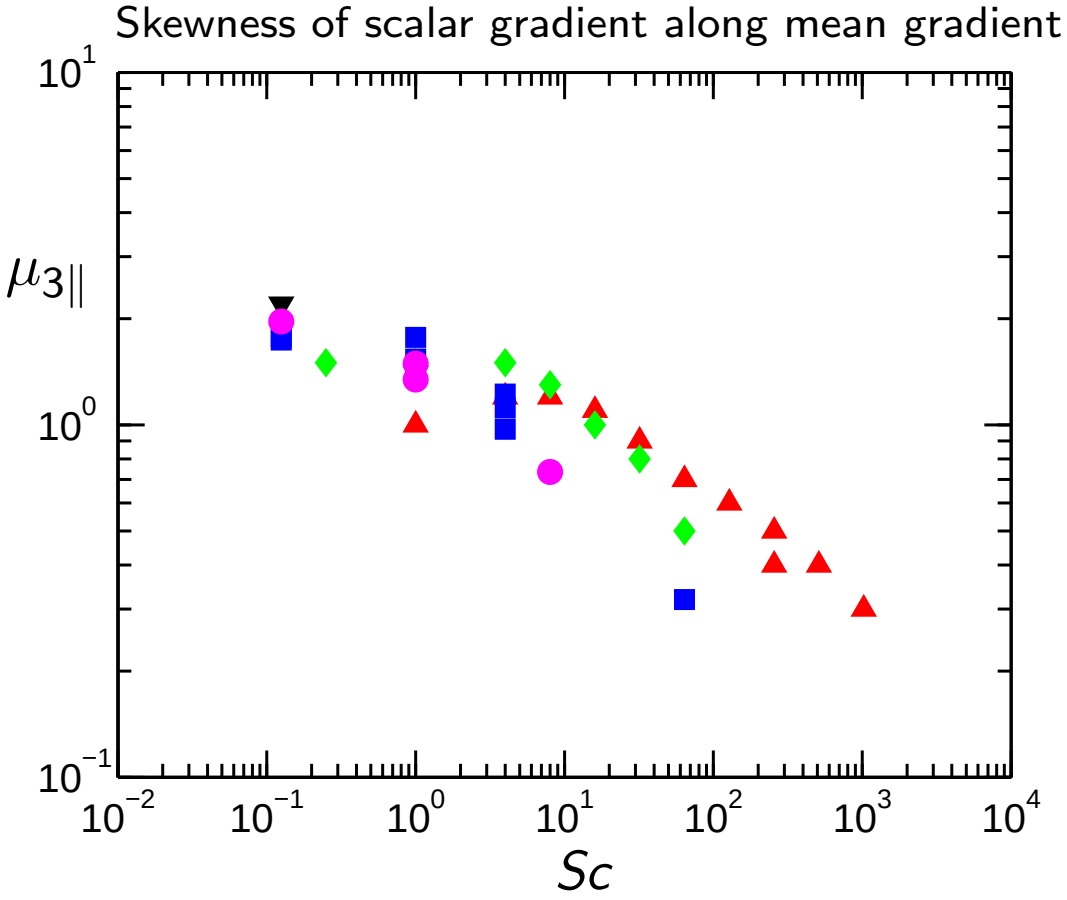

$R_\lambda \approx$ 8, 38, 140, 240, 650

- Skewness of $\nabla_\| \phi$ decreases with Sc
- Isotropy recovered at $Sc \to \infty$
- Trend depends on R_λ

Local isotropy: R_λ effects

- Literature: persistent anisotropies with R_λ
- Problematic from theoretical and practical points of view (Warhaft 2000)

R_λ	$k_{max}\eta$	$\mu_{3\parallel}$	$\mu_{4\parallel}/\mu_{4\perp}$
140	1.4	1.39	1.138
240	1.4	1.34	1.057
400	1.4	1.34	1.092
650	1.4	1.38	1.129

Local isotropy: R_λ effects

- Literature: persistent anisotropies with R_λ
- Problematic from theoretical and practical points of view (Warhaft 2000)

R_λ	$k_{max}\eta$	$\mu_{3\parallel}$	$\mu_{4\parallel}/\mu_{4\perp}$
140	1.4	1.39	1.138
240	1.4	1.34	1.057
400	1.4	1.34	1.092
650	1.4	1.38	1.129
140	**5.5**	**1.77**	**1.272**
240	**5.1**	**1.49**	**1.099**
650	**2.7**	**1.33**	**1.043**

- Low-resolution data mask isotropy trend! (Donzis & Yeung, Physica D 2010)

Mixing at sub-Kolmogorov scales

- What drives scalar fluctuations at $k\eta > 1$?
- Batchelor 1962: most compressive principal strain $\langle\gamma\rangle$
  - From data $\langle\gamma\rangle = 1/C_B\tau_\eta, \quad \tau_\eta = \sqrt{\nu/\langle\epsilon\rangle} \quad C_B \approx 2$
  - $E_\phi(k) = C_B\langle\chi\rangle\tau_\eta^{-1}k^{-1}f(k\eta_B)$

Mixing at sub-Kolmogorov scales

- What drives scalar fluctuations at $k\eta > 1$?
- Batchelor 1962: most compressive principal strain $\langle\gamma\rangle$
  - From data $\langle\gamma\rangle = 1/C_B\tau_\eta, \quad \tau_\eta = \sqrt{\nu/\langle\epsilon\rangle} \quad C_B \approx 2$
  - $E_\phi(k) = C_B\langle\chi\rangle\tau_\eta^{-1}k^{-1}f(k\eta_B)$

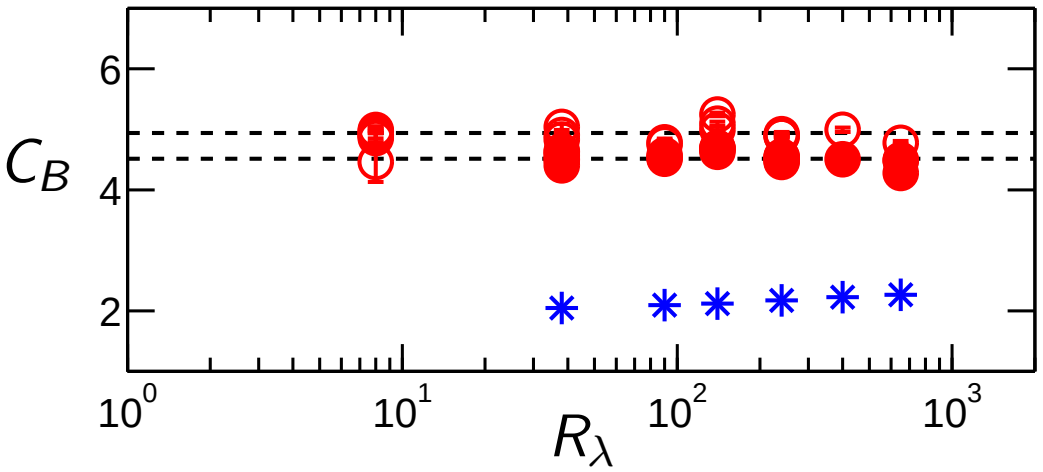

- Database (up to 4096^3):
  - $R_\lambda \approx 8 - 650$
  - $Sc = 1/8 - 1024$
- Batchelor estimate $C_B \approx 2$ too small

- $C_B \approx 5$ implies mixing is dominated by weaker compression (not the mean $\langle\gamma\rangle$)

Compressive strain

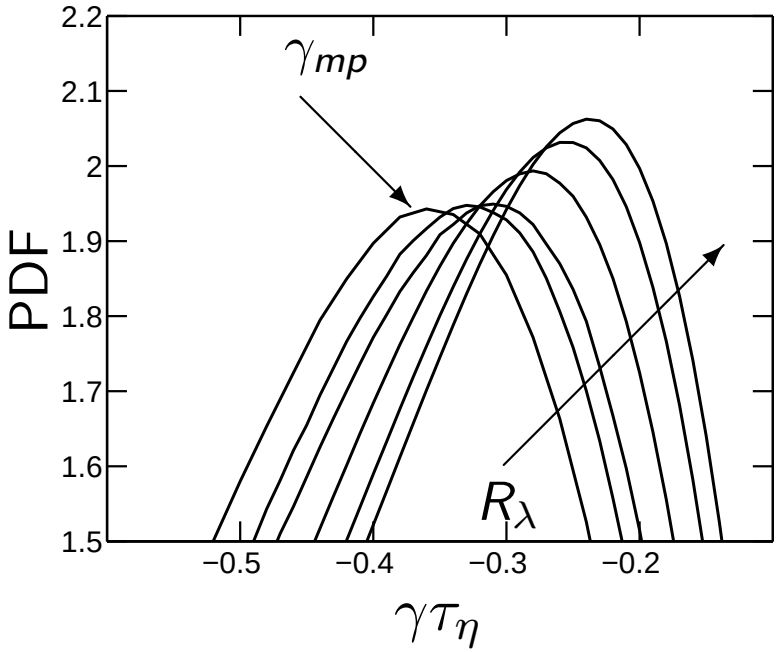

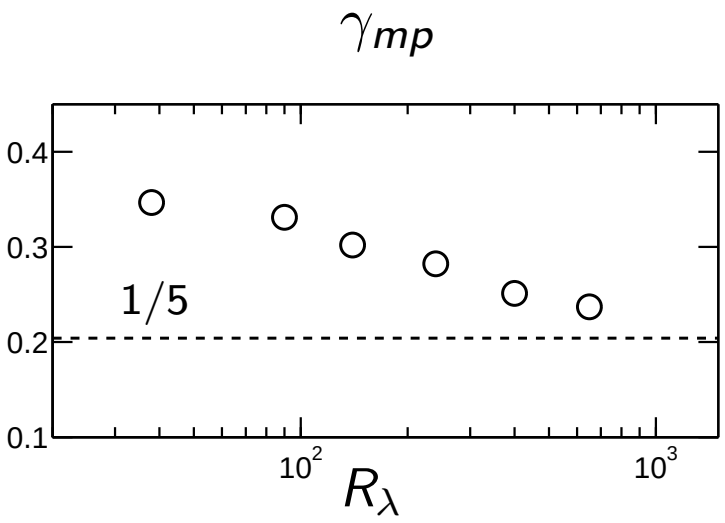

- Regions of weakly compressive strain rate become more probable
- Consistency with Batchelor's picture with γ_{mp} at high R_λ

So, what happens in real high-R_λ flows? Need Petascale simulations!

Concluding remarks: turbulence and beyond

- Computing power = more realistic simulations (obviously)
- But... how to use that power?
  - intermittency: order-dependent resolution criterion
  - scalars: further constraints (space and time)
  - achievable R_λ, lower that previously thought
- Well resolved simulations of turbulent mixing
  - return to isotropy, masked at lower resolutions
  - consistency with Batchelor theory at high R_λ

Concluding remarks: computing turbulence

- Towards Peta- and Exa-scale computing:
  - communication is the bottleneck: network contention as more cores share network resources
  - use fewer cores to communicate: threaded FFTWs
  - hybrid MPI/OpenMP (Blue Waters, BG/Q,...)
  - more control over communications: one-sided, pack-send-unpack (more memory)
  - overlap communication with computation

Concluding remarks: computing turbulence

- Towards Peta- and Exa-scale computing:
  - communication is the bottleneck: network contention as more cores share network resources
  - use fewer cores to communicate: threaded FFTWs
  - hybrid MPI/OpenMP (Blue Waters, BG/Q,...)
  - more control over communications: one-sided, pack-send-unpack (more memory)
  - overlap communication with computation
- The challenge of massive data:
  - a snapshot of the velocity and scalar fields: 1.25TB
  - archival? transfer the data? allow remote post-processing?
- Close collaboration with CS essential: new tools, techniques, programming models, ...