



Small-scale properties from exascale computations of turbulence on a $32\,768^3$ periodic cube

P.K. Yeung^{1,2} , Kiran Ravikumar¹ , Rohini Uma-Vaideswaran,¹
Daniel L. Dotson,¹ Katepalli R. Sreenivasan³ , Stephen B. Pope,⁴
Charles Meneveau⁵  and Stephen Nichols⁶

¹School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

²School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

³Tandon School of Engineering, New York University, New York, NY 10012, USA

⁴Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA

⁵Department of Mechanical Engineering & IDES, Johns Hopkins University, Baltimore, MD 21205, USA

⁶Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

Corresponding author: P.K. Yeung, pk.yeung@ae.gatech.edu

(Received 12 June 2025; revised 18 July 2025; accepted 18 July 2025)

To study the physics of small-scale properties of homogeneous isotropic turbulence at increasingly high Reynolds numbers, direct numerical simulation results have been obtained for forced isotropic turbulence at Taylor-scale Reynolds number $R_\lambda = 2500$ on a $32\,768^3$ three-dimensional periodic domain using a GPU pseudo-spectral code on a 1.1 exaflop GPU supercomputer (Frontier). These simulations employ the multi-resolution independent simulation (MRIS) technique (Yeung & Ravikumar 2020, *Phys. Rev. Fluids*, vol. 5, 110517) where ensemble averaging is performed over multiple short segments initiated from velocity fields at modest resolution, and subsequently taken to higher resolution in both space and time. Reynolds numbers are increased by reducing the viscosity with the large-scale forcing parameters unchanged. Although MRIS segments at the highest resolution for each Reynolds number last for only a few Kolmogorov time scales, small-scale physics in the dissipation range is well captured – for instance, in the probability density functions and higher moments of the dissipation rate and enstrophy density, which appear to show monotonic trends persisting well beyond the Reynolds number range in prior works in the literature. Attainment of range of length and time scales consistent with classical scaling also reinforces the potential utility of the present high-resolution data for studies of short-time-scale turbulence physics at high Reynolds

numbers where full-length simulations spanning many large-eddy time scales are still not accessible. A single snapshot of the $32\,768^3$ data is publicly available for further analyses via the Johns Hopkins Turbulence Database.

Key words: turbulence simulation, intermittency, isotropic turbulence

1. Introduction

It is well known (e.g. Kim & Leonard 2024) that direct numerical simulations (DNS) are a powerful tool for studying the fundamental physics of turbulence in canonical settings. The first DNS were performed for a three-dimensional (3-D) periodic domain more than 50 years ago by Orszag & Patterson (1972), with 32^3 grid points at Taylor-scale Reynolds number (R_λ) 64. Since then, with computer power growing approximately a million times every 25 years, the drive towards higher Reynolds numbers has been a prime motivation for turbulence researchers (Yokokawa *et al.* 2002; Lee, Malaya & Moser 2013; Ravikumar, Appelhans & Yeung 2019) seeking to harness resources that have marked the arrival of the terascale, petascale and more recently exascale eras (Atchley *et al.* 2023). At the same time, it has become clear (Yakhot & Sreenivasan 2005) that scales finer than the Kolmogorov scales appear more copiously as the Reynolds numbers increase, which implies a need to resolve the small scales better than was often practised in the past, and limits the Reynolds number achievable at the required resolution. This consideration also makes the pursuit of extreme-scale simulation capabilities ever more important.

An important science question requiring access to 3-D data at increasing Reynolds numbers and improved spatial and temporal resolution at the smallest scales of turbulence is whether scaling and other trends observed at moderate Reynolds numbers continue at ever-increasing Reynolds number, or whether qualitative transitions occur at some point. An important example of such trends concerns enstrophy having been found to be more important than dissipation in most known data sources. Past literature (Nelkin 1999) included arguments that since dissipation and enstrophy are both small-scale quantities belonging to the same symmetry group, the tails of their probability density functions (PDFs) must have the same scaling characteristics (i.e. their moments could differ by no more than a Reynolds-number-independent constant). Despite recent progress (Yeung, Sreenivasan & Pope 2018) that addressed the impact of finite resolution, high-fidelity data at significantly higher Reynolds numbers are still required.

The largest Fourier pseudo-spectral DNS dataset that can fit into the memory of the world's first exaflop leadership class supercomputer (named Frontier) is $40\,960^3$, but we focus on $32\,768^3$, which is more realistic for our present purposes. However, the viability of a specific simulation still depends on both the cost per step (at least proportional to N^3 on an N^3 grid) and the number of time steps for a desired physical time span (as N increases). As reported in Yeung *et al.* (2025), the best timing for a $32\,768^3$ grid using 44 % of the machine is 13 s per step. At $R_\lambda \approx 2500$, each Kolmogorov time scale (τ_η) takes approximately 1000 time steps, while the ratio of large-eddy turnover time (T_E) to τ_η is approximately 300. At this scale, a long simulation spanning $10T_E$ enabling long-time averaging would involve two years or more of virtually non-stop computing. It is clear that a different strategy is necessary.

This paper is focused on recently available DNS datasets at high resolution and high Reynolds numbers that may be used to address fundamental questions pertaining to properties of the small scales of turbulence, such as on comparisons between the

asymptotic scaling of tails of dissipation and enstrophy, at increasing Reynolds numbers. The data are obtained using the multi-resolution independent simulations (MRIS) approach introduced by Yeung & Ravikumar (2020), here extended to facilitate attainment of higher Reynolds numbers. The technique will be described in more detail in § 2, but it suffices here to say that this method includes ensemble averaging over short segments (referred to as MRIS segments) with statistically equivalent initial conditions. The short duration of the simulations calls for a separate assessment of inertial range properties, including some tests of statistical convergence, but dissipative scales that are spatially and temporally finer appear to be adequately captured. We present these results in §§ 3 and 4, including fundamental properties of the energy spectrum, the scaling of the mean energy dissipation rate, and new insights into the PDFs of energy dissipation and squared vorticity. We conclude the paper in § 5, where we also note additional simulation capabilities developed on Frontier.

2. Technical approach

In this section, we first briefly review our basic DNS algorithm and its high-performing implementation on the leadership-class graphics processing unit (GPU) platform on Frontier. The use of the MRIS protocol for both improving resolution in time and space and for achieving progressively higher Reynolds number is then discussed in some detail.

2.1. *The GPU-enabled extreme-scale turbulence simulations on Frontier*

The GPU-enabled extreme-scale turbulence simulations code (Yeung *et al.* 2025) on Frontier performs DNS of incompressible isotropic turbulence on a 3-D periodic domain. Numerical integration in time is carried out in wavenumber space, using the Fourier pseudo-spectral formulation of Rogallo (1981), with solenoidal numerical forcing at the largest scales (in the lowest few wavenumber shells). The principal mathematical operation is a discrete 3-D fast Fourier transform (FFT), which is invoked multiple times between physical space and wavenumber space at each Runge–Kutta substep. Distributed-memory parallelism requires a solution domain partitioned in one or two directions simultaneously, while FFTs are taken one direction at a time. To complete a 3-D FFT, the domain must be re-divided in different directions, which is achieved by exchanging data among multiple parallel processes via message-passing communication.

Data movements in communication calls often add substantially to the overall simulation cost, especially at large problem sizes or large node counts, ultimately limiting the size of problem accessible on a given machine. However, the arrival of exascale computing defined by a theoretical peak of 10^{18} floating point operations per second has led to improved prospects, by enabling much faster computation, while increased memory per node allows using fewer parallel processes, thus reducing the communication overhead. In addition, while early GPU platforms often had limited memory, parity in memory between central processing unit (CPU) and GPU on Frontier allows the GPUs to perform almost all operations, including computation and communication, thus reducing data movement between CPUs and GPUs to a minimum. Further algorithmic details and per-step timing data are given in Yeung *et al.* (2025).

2.2. *Successive grid refinement and increasing the Reynolds number*

We next consider the physical and numerical issues that determine the number of time steps required for specific purposes, and how best to push the limits of available computing power to obtain high-Reynolds-number results for specific purposes.

The MRIS approach is built on the premise that while long simulations at high resolution may not be attainable, long simulations at low resolution are much more affordable, and that high-resolution results can be obtained from the latter by refining resolution in space and/or time, and then running at high resolution for only a short period of time. In pseudo-spectral methods, grid refinement results in an additional set of high-wavenumber modes, which soon start to receive energy by spectral transfer from wavenumbers close to the maximum represented on the coarser grid. Some small-scale statistics will adjust quickly because their own time scales are typically short. Indeed, there is evidence (Yeung *et al.* 2018; Yeung & Ravikumar 2020) that when resolution is improved by halving the time step size or the grid spacing, it takes only approximately $2\tau_\eta$ for some equilibration to occur of small scales on the order of η of the coarser grid. The core MRIS strategy thus consists of a three-phased sequence: (1) running a reasonably long simulation at temporal and spatial resolution adequate for scales larger than η ; (2) continuing but halving the time step and running for a few τ_η ; and finally (3) halving the grid spacing and running at the highest resolution also for a small number of τ_η . For smoother transitions in the numerics, increased resolution can be implemented as a series of small increments: e.g. by starting from N_1 grid points in each direction, and first refining to $3N_1/2$ before ultimately reaching $N_2 = 2N_1$. We are also able to improve statistical sampling by repeating this sequence using different instantaneous velocity field snapshots for phase 1, and taking ensemble averages over multiple MRIS segments. Generally, phase 2 in each MRIS segment involves reducing the Courant number $C = [|u| + |v| + |w|]_{\max} \Delta t / \Delta x$ (where the velocity maximum is taken over all grid points) from 0.6 to 0.3, while phase 3 involves increasing $k_{\max} \eta = (\sqrt{2}N/3)\eta \approx 2.96/(\Delta x/\eta)$ from 1.4 to 2.8. The value of η corresponds to the highest resolution, but it does not change significantly with resolution.

Table 1 shows the main parameters in the computations reported in this paper. Although the number of phase 3 simulations available is not large, there is still appreciable benefit from ensemble averaging over the initial conditions (which is rarely discussed in the literature). We have used the forcing scheme of Donzis & Yeung (2010), where the total modal energy in each of the first three wavenumber shells ($k < k_F$, where $k_F = 2.5$ on a $(2\pi)^3$ domain) is kept constant at values derived from long-time averages from stochastic forcing (Eswaran & Pope 1988). A test of a single segment of $R_\lambda = 1000$ at 8192^3 and spanning $27\tau_\eta$ suggests that the results presented are insensitive to the values of T_3/τ_η in the table, presumably because the energy of the large scales is kept frozen by the forcing scheme. For consistency, we have used MRIS results at all Reynolds numbers except the lowest, even if longer simulations at lower Reynolds numbers are available from prior work. The MRIS procedure also has some conceptual similarities with scale enrichment in large-eddy simulations (Ghate & Lele 2020).

It should be noted that all of the simulations reported used single precision arithmetic. Numerical tests in Yeung *et al.* (2018) focused on the peak values of dissipation and enstrophy showed that effects of machine precision were small compared to those associated with spatial and temporal resolution as well as statistical sampling. Recently, there has been increasing interest in DNS regimes where round-off errors may accumulate over time (Okamoto *et al.* 2025a), while a rigorous computer science approach can provide insights into the effects of different levels of machine precision on quantities of different types (Karp *et al.* 2025). While rigorous testing is not readily available, based on the above, and since our MRIS segments are only approximately $2\tau_\eta$ long, it seems unlikely that machine precision would have strongly affected the conclusions drawn in the present work.

To reach higher Reynolds numbers, we start from developed velocity fields at, say, $R_{\lambda 1}$ on an N_1^3 grid, and reduce the viscosity while keeping both the forcing details and the

R_λ	N^3	ν	M	T_1/τ_η	T_2/τ_η	T_3/τ_η
390	2048 ³	4.370×10^{-4}	4	> 60	5.00	3.74
650	4096 ³	1.732×10^{-4}	3	—	—	4.06
1000	8192 ³	6.883×10^{-5}	3	83	2.63	2.56
1600	16 384 ³	2.7319×10^{-5}	3	107	3.01	1.84
2500	32 768 ³	1.10555×10^{-5}	3	43	1.93	1.93

Table 1. Basic parameters in this work: for each approximate R_λ , highest grid resolution (N^3) used, kinematic viscosity (ν), number of MRIS segments (M), and (see text in § 2.2) time spans T_1 , T_2 , T_3 in τ_η for phases 1–3, with $\{k_{max}\eta, C\} \approx \{1.4, 0.6\}$, $\{1.4, 0.3\}$ and $\{2.8, 0.3\}$, respectively. Here, T_1 , T_2 and T_3 are averaged over different segments since they have varied somewhat. Phases 1 and 2 were not needed for the $R_\lambda = 650$ case, which started from a $C = 0.15$ stationary state in Yeung *et al.* (2018). Early portions of data from phase 3 are not used in the time averaging.




Figure 1. (a) Evolution of kinetic energy (black), mean dissipation (red), S_ϵ (green) and R_λ (blue), all normalised by initial values, during transition from $R_\lambda = 2080$ on a $12\,288^3$ grid to $R_\lambda = 2500$ on $16\,384^3$, by reducing the viscosity. Time t is normalised by the final value of τ_η . (b) Plot of $D(k)$ compared with data at $R_\lambda = 2080$ data (gold), at $t/\tau_\eta \approx 0.01, 5.7, 12.8, 20.5, 28.8, 37.8$ (red, green, blue, black, cyan, magenta, respectively).

resolution parameter $k_{max}\eta$ unchanged. We refine the grid to N_2^3 with $\eta \propto 1/N$. Since the mean dissipation rate is nearly independent of viscosity in the stationary state, the definition $\eta = (\nu^3/\langle\epsilon\rangle)^{1/4}$ implies $\nu \propto (1/N)^{4/3}$. In addition, with the root mean square velocity u' also unchanged, the standard local isotropy relation $\langle\epsilon\rangle = 15\nu(u'/\lambda)^2$ predicts $R_\lambda \propto N^{2/3}$. This implies that at a higher resolution, the expected new Reynolds number can be estimated by $R_{\lambda 2} \approx R_{\lambda 1}(N_2/N_1)^{2/3}$, which amounts to a 58 % increase in the case $N_2 = 2N_1$.

Figure 1 shows the manner in which quantities dominated by different scale sizes evolve during transition to a higher Reynolds number, when subjected to the procedures described above. We have chosen an example where the velocity field evolves from $R_\lambda = 2080$, which served as an intermediate step between $R_\lambda = 1600$ and $R_\lambda = 2500$. In figure 1(a), it can be seen that the kinetic energy varies little, while mean dissipation recovers from the sudden reduction in viscosity within approximately $20\tau_\eta$. For a diagnostic more strongly linked to the small scales, we consider the ‘dissipation skewness’ defined from (despite its name) a fourth-moment integral of the energy spectrum, i.e. $S_\epsilon = (4/35)(15\nu/\langle\epsilon\rangle)^{3/2} \int_0^\infty \nu k^4 E(k) dk$, which is equal to the negative of the velocity gradient skewness in stationary isotropic turbulence. As expected, S_ϵ recovers sooner.




Figure 2. Scaling behaviours with respect to R_λ of (a) integral-scale Reynolds number, (b) length scale ratio (triangles) and time-scale ratio (circles), and (c) the normalised energy dissipation rate. Dotted lines in (a) and (b) show expected power laws. Horizontal dashed lines in (c) indicate the range in various simulations.

When grid refinement begins, R_λ rises artificially from its stable value $R_{\lambda 1}$ on a $12\,288^3$ grid to $R_\lambda(0) = R_{\lambda 1}/0.6814$, then decreases to approximately $0.84R_{\lambda 0} = 1.23R_{\lambda 1}$, which agrees well with the expectation of a 21.1 % increase (corresponding to $N_2/N_1 = 4/3$).

Unlike the quest for resolution improvement alone in Yeung & Ravikumar (2020), here a reduction of viscosity directly affects almost all Fourier modes. Figure 1(b) shows, without normalisation, that the dissipation spectrum settles very quickly onto a new and approximately stationary form at higher Reynolds number. Results in this figure correspond to ‘phase 1’ in our MRIS protocol, for a case of grid refinement by ratio $N_2/N_1 = 4/3$. In most cases, we operate with $N_2/N_1 = 2$, which requires a longer transition period, with the Reynolds number dropping further before increasing towards the new value expected on the finer grid.

One of the three $32\,768^3$ snapshots (3-D fields of three velocity components and pressure, encompassing half a petabyte of data) has been ingested into a public database system (the Johns Hopkins Turbulence Database) where the data are accessible using web-services-enabled virtual sensors.

3. Basic scaling behaviours versus the Reynolds number

As seen in table 1 the DNS data reported here cover more than six-fold spread of R_λ , with the values at the upper end exceeding those in most of the DNS literature. It is thus useful to compare the new data with some basic classical scaling relations, to confirm the fidelity of the new results and to check if classical scaling continues to apply at these higher Reynolds numbers without qualitative change. Figure 2(a) shows the large-scale Reynolds number (R_ℓ , where the length scale ℓ is taken to be the longitudinal integral length scale L_1 of the velocity field) versus the Taylor-scale Reynolds number R_λ . The dashed line corresponds to $C_\epsilon R_\lambda^2/15$, by assuming that $C_\epsilon = \langle \epsilon \rangle L_1/u'^3$ is 0.45 (figure 2c). Figure 2(b) shows the length scale ratio L_1/η and the time scale ratio $(L_1/u')/\tau_\eta$. The lines show the expected $3/2$ power and the linear relations, again taking $C_\epsilon = 0.45$. Figure 2(c) shows the variation of the normalised mean energy dissipation rate $C_\epsilon = \langle \epsilon \rangle L_1/u'^3$, vis-à-vis the dissipative anomaly or the zeroth law of turbulence (Sreenivasan 1984, 1998). The data show a weak trend up to $R_\lambda \approx 1500$, beyond which they stabilise to a constant, slightly larger than 0.4. Continued validity of the zeroth law of turbulence to R_λ over 2500 has important implications for turbulence theory and modelling practices.




Figure 3. (a) The 3-D energy spectrum $E(k)$ at $R_\lambda \approx 390$ (red), 650 (green), 1000 (blue), 1600 (black), 2500 (magenta), after averaging in time and over available MRIS segments. The dotted line has slope $-5/3$. (b) Compensated 3-D spectrum $\langle \epsilon \rangle^{-2/3} k^{5/3} E(k)$ (upper curve) and the one-dimensional longitudinal spectrum $E_{11}(k)$ in the Kolmogorov variable. The dashed horizontal lines are the corresponding (generally accepted) values 1.62 and $(1.62)(18/55) = 0.53$.

A most investigated but still elusive topic in fundamental turbulence theory is the form of the energy spectrum in the inertial range. Figure 3(a) suggests a range $E(k) \propto k^{-5/3}$, which widens with an increase of R_λ . Yet a closer look at figure 3(b) shows that the compensated quantity $E(k)\langle \epsilon \rangle^{-2/3} k^{5/3}$ does not display a plateau. Instead, the data support a power law for $E(k)$ with a slope steeper than $-5/3$, and a bump, called the bottleneck, which is strongest at $k\eta$ between 0.1 and 0.2. The first observation is consistent with the phenomenology of intermittency (Kolmogorov 1962) that predicts a correction (steepening) of the slope and has motivated extensive research, including analyses of experimental data on one-dimensional spectra at higher Reynolds numbers (see Sreenivasan & Antonia 1997). The second issue has also spawned past papers (Donzis & Sreenivasan 2010) and proposed spectral fitting functions that include both intermittency correction and bottleneck (Meyers & Meneveau 2008). It is understood that because of the $5/3$ exponent, the peak of the bottleneck is at a wavenumber slightly below that for the dissipation spectrum $D(k) \propto k^2 E(k)$, hence viscous effects certainly play a role. Figure 3(b) also shows similar features in the one-dimensional spectrum $E_{11}(k)$. We have verified consistency with isotropy relations such as $E(k) = k^3 d[(1/k)d E_{11}(k)]/dk$, and noted that the bottleneck in $E_{11}(k)$ is weaker than $E(k)$. Similarly to the work by Ishihara *et al.* (2016), the high Reynolds number in the current data is expected to be useful for further work on this topic.

4. The PDFs of dissipation and enstrophy density

The PDFs of dissipation rate ϵ (formed from the square of strain rate) and enstrophy density $\Omega \equiv \omega^2$ (formed from the square of vorticity – henceforth enstrophy, for brevity) are prominent indicators of small-scale turbulence. Earlier measurements have suggested that the PDF tails of both quantities are well represented by stretched exponentials, with enstrophy being the more intermittent variable.

To firmly understand how the two variables behave, high Reynolds numbers, high resolution and large datasets are necessary. In this work, we have extracted the PDFs of dissipation and enstrophy from a large number of instants of time even though the MRIS segments themselves are short. Although one may question whether these PDFs have fully equilibrated from forcing over the short time interval covered by the MRIS segments, we can mitigate such concerns by plotting the two PDFs against each other,




Figure 4. Base-10 logarithms of PDFs of enstrophy versus the same for dissipation at (a) $R_\lambda = 390$ and (b) $R_\lambda = 2500$. Symbols in red denote data points where ϵ and Ω are both below their mean values; blue indicates samples above the mean values. Black dashed lines of slopes $1/3$ (upper) or 1 (lower) have been added for some comparisons (in the text).

This is done in figure 4, where different colours represent the left tails (red, for small ϵ or Ω) and right tails (blue, for large ϵ or Ω). Deviations from the dashed line of slope 1 indicate that the two PDFs do not have identical form. We will qualify this statement below. Both very low and very high values of enstrophy are more probable than those of dissipation. The behaviour in the left tail can be explained by noting (Yeung, Donzis & Srenivasan 2012; Gotoh & Yeung 2025) that zones of low strain rate and low enstrophy are nearly Gaussian, hence the PDFs behave like a chi-square distribution of order n , which holds for the sum of squares of n independent and identically distributed Gaussian variables. For small enstrophy $n = 3$, while for dissipation $n = 5$, since there are (due to the incompressibility constraint) five independent strain-rate components. Consequently (for small x), the chi-square result is $f(x) \propto x^{n/2-1}$, which then gives $f_\epsilon(x) \propto x^{3/2}$, while $f_\Omega(x) \propto x^{1/2}$. Combining these results, we obtain

$$f_\Omega(x) \propto [(f_\epsilon(x)^{2/3})^{1/2} \propto [f_\epsilon(x)]^{1/3}, \quad (4.1)$$

which is confirmed by a dashed line of slope $1/3$ in figure 4. This result is kinematic and holds equally well at both Reynolds numbers of this figure.

The right tails of $f_\epsilon(\cdot)$ and $f_\Omega(\cdot)$ can be represented quite well by expressions of the form

$$f_\epsilon(x) \propto \exp(-a_\epsilon x^{b_\epsilon}), \quad f_\Omega(x) \propto \exp(-a_\Omega x^{b_\Omega}). \quad (4.2)$$

The data are reasonably represented by these expressions with $b_\Omega = b_\epsilon$, in which case extremely high Ω , being more probable, implies $a_\Omega < a_\epsilon$. The right tail portion in figure 4(b) appears close to linear, with slope close to 0.9. Thus we can conclude that differences in high-order moments of the two variables occur only via the coefficients a_Ω and a_ϵ .

The stretched exponential expressions above are relevant in the ranges where $\epsilon/\langle\epsilon\rangle$ and $\Omega/\langle\Omega\rangle$ far exceed unity. For a closer examination of the parameters involved, figure 5 shows, for three Reynolds numbers, the quantity $\ln(-\ln[f_\epsilon(x)/f_\epsilon(1)])$ versus $\epsilon/\langle\epsilon\rangle$ for $x > 1$ only, and likewise for the enstrophy. Perfect fits give straight lines of slope equal to the exponent and intercepts on the y-axis equal to the prefactor. As R_λ increases, non-zero samples are found further to the right, while the lines for dissipation (red) and enstrophy (blue) become more nearly parallel, with the latter having a smaller intercept on the y-axis. These observations are consistent with the discussion of figure 4 above,




Figure 5. Test of the stretched exponential form of the right tails of the PDFs of $\epsilon/\langle\epsilon\rangle$ (red) and $\Omega/\langle\Omega\rangle$ at approximately (a) 390, (b) 1000, (c) 2500. Here, $f_X(\cdot)$ denotes the PDF of either $\epsilon/\langle\epsilon\rangle$ or $\Omega/\langle\Omega\rangle$. The dotted lines are at slope 0.518, which corresponds to 0.225 reported in Gotoh & Yeung (2025) using the common log instead of the natural log as here.




Figure 6. Reynolds number dependence of (a) normalised moments of ϵ , orders $n = 2, 3, 4$ (bottom to top), (b) normalised moments of Ω , and (c) ratio of normalised moments of Ω to those of ϵ . Red triangles are from integration of time-averaged PDFs. Blue circles are from averaging in physical space over several instantaneous snapshots.

while also demonstrating the value of high Reynolds number in establishing trends more conclusively than possible before.

The PDFs obtained here are smooth enough, extending to large values of the respective variable, that they can be used to compute high-order moments as measures of intermittency. Having checked for satisfactory convergence of the moment integrands associated with dissipation and enstrophy PDFs, in figure 6 we show the behaviour of second, third, and fourth moments of $\epsilon/\langle\epsilon\rangle$ and $\Omega/\langle\Omega\rangle$ at all five Reynolds numbers of table 1. As expected from the contrast in their PDFs, the moments of the normalised enstrophy for any given order $n > 1$ are higher than those of normalised dissipation. Their ratios in figure 6(c) are nearly independent of the Reynolds number, which is consistent with arguments by Nelkin (1999) that while enstrophy is more intermittent than dissipation, their scaling with Reynolds number must be the same asymptotically.

We note in passing that the skewness and the flatness factor of the longitudinal velocity gradient agree well with the data compiled by Sreenivasan & Antonia (1997).

5. Concluding remarks

Three-dimensional turbulence governed by the Navier–Stokes equations continues to be a grand challenge in computational science, even as leadership-class computing power

has recently advanced past the exascale. Direct numerical simulations (DNS) at $32\,768^3$ resolution are now possible on the world's first exascale computer (named Frontier). However, as the range of time scales also grows, full-length simulations spanning multiple large-eddy time scales at this high resolution remain prohibitively expensive, especially if rigorous resolution requirements (Yakhot & Sreenivasan 2005; Yeung *et al.* 2018) are to be met. In this paper we have not fully addressed the implications of these rigorous requirements for all aspects of DNS, but we have shown that a protocol named MRIS (for multi-resolution independent simulations, introduced in Yeung & Ravikumar 2020), consisting of ensemble averaging over results from a number of short simulation segments at high resolution, can produce reliable results on small-scale turbulence.

This work is focused on statistically stationary isotropic turbulence with a forcing scheme, which produces very stable energetics at the large scales. Reynolds numbers are increased by reducing the viscosity at modest resolution, followed by short simulations that improve accuracy by increasing both temporal and spatial resolution. The highest Taylor-scale Reynolds number reached is approximately 2500. Even though the MRIS segments for resolution improvement are sometimes as short as approximately two Kolmogorov scales, many fundamental scaling properties, including the property of dissipative anomaly, and effects of intermittency and a bottleneck in the energy spectrum, are reproduced well. This study also successfully confirms and extends prior knowledge of the scaling properties of the probability densities and moments of the energy dissipation rate and enstrophy, where extreme events are of great interest. Our results show that these PDFs behave like chi-square distributions at low amplitudes, but as stretched exponentials at high amplitudes and asymptotics, consistent with the arguments provided by Nelkin (1999).

Turbulence simulations have been advancing steadily in size, with the most prominent drivers reaching higher Reynolds number and achieving better resolution at the small scales. As Reynolds numbers increase, new results may either lead to confirmation of prior trends, thus removing doubts that might otherwise blunt the path forwards, or suggest new physics not apparent from results at lower Reynolds number or insufficient resolution. Examples in this work include figures 2(c) and 6(a,b), where conclusions would differ if data points at $R_\lambda \approx 1600$ and 2500 were absent. Both types of contributions continue to provide motivation for the meticulous work (as documented in Yeung *et al.* 2025) required to extract the best performance from the best machines available. (The magnitude of the computational effort exerted in obtaining the present results can be inferred from the triangles in figure 2(b), where a 2.5-fold increase in R_λ from 1000 to 2500 (third data point to the last) requires a 4-fold increase in the number of grid points in each direction – which translates to at least a ($4^3 = 64$)-fold increase in elapsed wall time per step, and a ($4^4 = 256$)-fold increase for a given physical time span, on a machine with at least 64 times more memory.)

We note again that this paper provides primarily dissipation range results, while a discussion of inertial range quantities is postponed to a later paper. The MRIS technique may not be as directly applicable when one wishes to study phenomena that are allowed to evolve on time scales associated with motions in the inertial or energy-containing range. Examples include non-stationary turbulence or flows subjected to forcing schemes that supply a highly variable rate of energy input. It would seem that longer MRIS segments may help us to understand the dominant phenomena of interest that occur over intervals of several Kolmogorov time scales. Continuing work includes extensions of MRIS to study the small scales in passive scalar mixing, and short-time Lagrangian properties such as statistics of the fluid particle acceleration or rapidly separating particle pairs in relative dispersion.

Note added in proof. The authors have become aware of a new manuscript (Okamoto *et al.* 2025b) which provides some new insights relevant to figure 3(b) in this paper.

Acknowledgements. We thank especially Dr B. Messer, Director of Science at OLCF, for his constant support and encouragement. The first author also wishes to thank Professors T. Gotoh, Y. Kaneda and S.K. Lele for helpful recent discussions.

This manuscript has been authored in part by U T-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The publisher acknowledges the US government licence to provide public access under the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

Funding. This research used resources allocated through the INCITE program at the Oak Ridge Leadership Computing Facility, which is a US Department of Energy Office of Science User Facility supported under contract DE-AC05-00OR22725. The work at Georgia Tech is further supported by a subaward from NSF via CSSI grant 2103874 to Johns Hopkins University. The research of K.R.S. was funded by New York University.

Declaration of interests. The authors report no conflict of interest.

Data availability statement. A single instantaneous snapshot of the $32\,768^3$ velocity field (including pressure fluctuations derived therefrom) obtained in this work is publicly available from the Johns Hopkins Turbulence Database. See <https://turbulence.idies.jhu.edu/datasets/homogeneousTurbulence/isotropic32768>.

REFERENCES

- ATCHLEY, S., *et al.* 2023 Frontier: exploring exascale. In *SC '23: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, Article 52, pp. 1–16. ACM.
- DONZIS, D.A. & SREENIVASAN, K.R. 2010 The bottleneck effect and the Kolmogorov constant in isotropic turbulence. *J. Fluid Mech.* **657**, 171–188.
- DONZIS, D.A. & YEUNG, P.K. 2010 Resolution effects and scaling in numerical simulations of passive scalar mixing in turbulence. *Physica D: Nonlinear Phenom.* **239**, 1278–1287.
- ESWARAN, V. & POPE, S.B. 1988 An examination of forcing in direct numerical simulations of turbulence. *Comput. Fluids* **16**, 257–278.
- GHATE, A.S. & LELE, S.K. 2020 Gabor mode enrichment in large eddy simulations of turbulent flow. *J. Fluid Mech.* **903**, A13.
- GOTOH, T. & YEUNG, P.K. 2025 On one-dimensional surrogacy of probability density function of enstrophy in isotropic turbulence. *Phys. Fluids* **37**, 025178.
- ISHIHARA, T., MORISHITA, K., YOKOKAWA, M., UNO, A. & KANEDA, Y. 2016 Energy spectrum in high-resolution direct numerical simulation of turbulence. *Phys. Rev. Fluids* **1**, 082403.
- KARP, M., *et al.* 2025 Effects of lower floating-point precision on scale-resolving numerical simulations of turbulence. *J. Comput. Phys.* (under review).
- KIM, J. & LEONARD, A. 2024 The early days and rise of turbulence simulation. *Annu. Rev. Fluid Mech.* **56**, 21–44.
- KOLMOGOROV, A.N. 1962 A refinement of previous hypotheses concerning the local structure of a viscous incompressible fluid. *J. Fluid Mech.* **13**, 82–85.
- LEE, M., MALAYA, N. & MOSER, R.D. 2013 Petascale direct numerical simulation of turbulent channel flow on up to 786K cores. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pp. 61:1–61:11. ACM.
- MEYERS, J. & MENEVEAU, C. 2008 A functional form for the energy spectrum parametrizing bottleneck and intermittency effects. *Phys. Fluids* **20**, 065109.
- NELKIN, M. 1999 Enstrophy and dissipation must have the same scaling exponent in the high Reynolds number limit of fluid turbulence. *Phys. Fluids* **11**, 2202–2204.
- OKAMOTO, N., ISHIHARA, T., YOKOKAWA, M. & KANEDA, Y. 2025a Effects of finite arithmetic precision on large-scale direct numerical simulation of box turbulence by spectral method. *Phys. Rev. Fluids* **10**, 064603.
- OKAMOTO, N., ISHIHARA, T., YOKOKAWA, M. and KANEDA, Y. 2025b Statistics of energy dissipation rate and enstrophy in high resolution direct numerical simulation of turbulence in a periodic box. *Phys. Rev. Fluids* (under review).
- ORSZAG, S.A. & PATTERSON, G.S. 1972 Numerical simulation of three-dimensional homogeneous isotropic turbulence. *Phys. Rev. Lett.* **28**, 76–79.

- RAVIKUMAR, K., APPELHANS, D. & YEUNG, P.K. 2019 GPU acceleration of extreme scale pseudo-spectral simulations of turbulence using asynchronism. In *Proceedings of The International Conference for High Performance Computing, Networking and Storage Analysis (SC'19)*, Article 8, pp. 1–22. ACM. <https://doi.org/10.1145/3295500.3356209>
- ROGALLO, R.S. 1981 Numerical experiments in homogeneous turbulence. *Tech. Memo.* 81315. NASA Ames Research Center.
- SREENIVASAN, K.R. 1984 On the scaling of the turbulence energy-dissipation rate. *Phys. Fluids* **27**, 1048–1051.
- SREENIVASAN, K.R. 1998 An update on the energy-dissipation rate in isotropic turbulence. *Phys. Fluids* **10**, 528–529.
- SREENIVASAN, K.R. & ANTONIA, R.A. 1997 The phenomenology of small-scale turbulence. *Annu. Rev. Fluid Mech.* **29**, 435–472.
- YAKHOT, V. & SREENIVASAN, K.R. 2005 Anomalous scaling of structure functions and dynamic constraints on turbulence simulations. *J. Stat. Phys.* **121**, 823–841.
- YEUNG, P.K., DONZIS, D.A. & SREENIVASAN, K.R. 2012 Dissipation, enstrophy and pressure statistics in turbulence simulations at high Reynolds numbers. *J. Fluid Mech.* **700**, 5–15.
- YEUNG, P.K. & RAVIKUMAR, K. 2020 Advancing understanding of turbulence through extreme-scale computation: intermittency and simulations at large problem sizes. *Phys. Rev. Fluids* **5**, 110517.
- YEUNG, P.K., RAVIKUMAR, K., NICHOLS, S. & UMA-VAIDESWARAN, R. 2025 GPU-enabled extreme-scale turbulence simulations: Fourier pseudo-spectral algorithms at the exascale using OpenMP offloading. *Comput. Phys. Commun.* **306**, 109364.
- YEUNG, P.K., SREENIVASAN, K.R. & POPE, S.B. 2018 Effects of finite spatial and temporal resolution on extreme events in direct numerical simulations of incompressible isotropic turbulence. *Phys. Rev. Fluids* **3**, 064603.
- YOKOKAWA, M., ITAKURA, T., UNO, A., ISHIHARA, T. & KANEDA, Y. 2002 16.4-Tflops direct numerical simulation of turbulence by a Fourier spectral method on the Earth Simulator. In *Proceedings of the Supercomputing Conference*, pp. 22. IEEE.