



# Direct numerical simulations of turbulent flows using high-order asynchrony-tolerant schemes: Accuracy and performance

Komal Kumari, Diego A. Donzis\*

Department of Aerospace Engineering, Texas A&M University, College Station, TX 77843, United States



## ARTICLE INFO

### Article history:

Received 9 November 2019

Accepted 31 May 2020

Available online 8 July 2020

### Keywords:

Direct numerical simulations

Turbulent flows

Asynchronous simulations

## ABSTRACT

Direct numerical simulations (DNS) are an indispensable tool for understanding the fundamental physics of turbulent flows. Because of their steep increase in computational cost with Reynolds number ( $Re_\lambda$ ), well-resolved DNS are realizable only on massively parallel supercomputers, even at moderate  $Re_\lambda$ . However, at extreme scales, the communications and synchronizations between processing elements (PEs) involved in current approaches become exceedingly expensive and are expected to be a major bottleneck to scalability. In order to overcome this challenge, we developed algorithms using the so-called Asynchrony-Tolerant (AT) schemes that relax communication and synchronization constraints at a mathematical level, to perform DNS of decaying and solenoidally forced compressible turbulence. Asynchrony is introduced using two approaches, one that avoids synchronizations and the other that avoids communications. These result in periodic and random delays, respectively, at PE boundaries. We show that both asynchronous algorithms accurately resolve the large-scale and small-scale motions of turbulence, including instantaneous and intermittent fields. We also show that in asynchronous simulations the communication time is a relatively smaller fraction of the total computation time, especially at large processor count, compared to standard synchronous simulations. As a consequence, we observe improved parallel scalability up to 262144 processors for both asynchronous algorithms.

© 2020 Elsevier Inc. All rights reserved.

## 1. Introduction

Turbulence is the most common fluid state of motion and is inherent in a large number of natural and physical phenomena. The physics of turbulence can be modeled mathematically by the Navier-Stokes (NS) equations under the continuum limit. These equations are highly non-local and non-linear and have resisted analytical analyses except for overly simplified canonical problems. Numerical simulations, thus, serve as an indispensable tool for understanding turbulence. One of the most important characteristics of turbulence is the inherent wide range of spatial and temporal scales. This range of scales increases with the Reynolds Number ( $Re_\lambda$ ), the ratio between the inertial and viscous forces, which is typically very high in applications. In accurate numerical simulations the computational domain has to be large enough to accommodate the largest scales of motion in the flow and the grid spacing small enough to resolve the so-called Kolmogorov scale [1], the

\* Corresponding author.

E-mail address: donzis@tamu.edu (D.A. Donzis).

smallest dynamically relevant scale in a turbulent flow. Furthermore, the simulation time should be sufficiently long to capture the slow evolution of the largest scale while the time-step size should be small enough to capture the fast Kolmogorov time scale characteristic of the smallest scales. Simulations that follow these stringent constraints and consequently, accurately resolve the physics of all relevant scales are known as Direct Numerical Simulations (DNS) [2,3]. Using classical scaling relations based on Kolmogorov ideas [1] for grid spacing and a CFL condition for time-step size, the computational work grows steeply as  $Re_\lambda^6$ , though more recent work suggests  $Re_\lambda^8$  if all intermittent events are to be resolved [4]. Due to this steep power-law dependence, high fidelity DNS are computationally prohibitively expensive and even with highly scalable codes run on today's most powerful supercomputers, unachievable for conditions of practical relevance.

Several numerical methods have been used for DNS of the Navier-Stokes equations to study turbulence, depending upon the complexity of the domain and the nature of problem of interest. Spectral methods [5], known for accurate computation of derivatives, have been used extensively in incompressible simulations. However, these present challenges when extended to non-periodic boundary conditions. An alternative to these methods, that is more amenable to the choice of boundary conditions, is the compact difference schemes that have spectral like resolution [6]. These are widely used for simulations of multi-scale phenomena like turbulence [7–11]. However, computation of derivatives using compact schemes involves a system of linear equations. This imposes constraints on the computational domain since each processor must have entire range of data in the direction of computation of derivative. Such codes require multiple collective communication calls, which in turn can make communication time quite significant [12,13,9] for both compact and spectral implementations.

Explicit finite difference schemes have also been extensively used for approximation of derivatives in partial differential equations (PDEs) including in massive simulations of turbulent reacting flows [14,15]. For explicit schemes, the derivative at a grid point in the domain is approximated as a linear combination of the values at its neighboring points only. Because of this local dependence, different processors can work concurrently on different parts of the domain. However, at the processor boundaries, processors need to communicate to obtain data from the neighboring processors in order to compute the derivatives. Although these are local communications as opposed to the collective communications for compact or spectral schemes, processors still incur in overheads due to the need to communicate and synchronize at every time step to meet accuracy requirements. While simulations have been successfully done using hundreds of thousands of processors [9,16,17], the synchronizations and communication overheads, irrespective of the choice of numerical methods, pose a serious challenge to scalability at extreme scales [18]. In order to overcome this bottleneck, some work has focused on relaxing the synchronization requirements among the processors and perform so-called asynchronous numerical simulations. Early work in the literature dealt with asynchronous simulations but severely limited to lower orders of accuracy and restricted to certain class of PDEs [19–22]. A new and more generalized approach, extensible to arbitrarily high orders of accuracy, has been recently developed [23,24] to derive the so-called Asynchrony-Tolerant (AT) finite-difference schemes.

However, these studies investigated numerical accuracy and stability for simplified model problems in low dimensions. The ability of these schemes to accurately simulate realistic three-dimensional turbulent flows has not been done before. Without careful assessment of the numerical and parallel performance of these schemes it is unclear whether they can indeed provide a path towards exascale simulations in future massively parallel systems. This is the main thrust of this paper.

In particular, we use AT schemes to perform, a first of a kind, asynchronous simulation of three-dimensional compressible turbulence. Our focus is on the effect of asynchrony on important turbulent characteristics such as evolution of the turbulent kinetic energy, the spectra and PDFs of velocity gradients, enstrophy and dissipation. In order to conduct these simulations, in addition to the selection of appropriate AT schemes, one needs to consider how asynchrony is introduced which has implications in terms of both numerical and computational performance. We propose two approaches for introducing asynchrony: one that avoids synchronizations and the other that avoids communications. While the former leads to reduction in processor idling time and results in random delays at processor boundaries, the latter leads to periodic delays and reduction in the volume of communications. Since power consumption, especially for data movement is expected to be a major concern for the next generation exascale machines, the reduced frequency of communications in communication avoiding algorithm, makes it a viable energy efficient alternative to standard approaches.

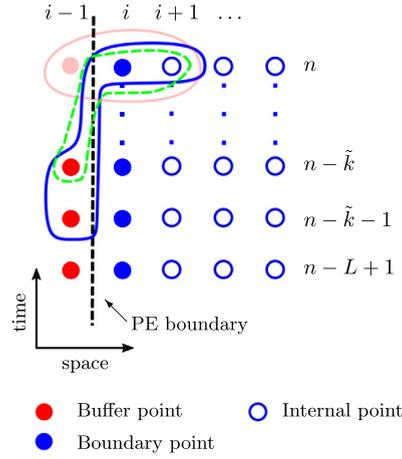
The rest of the paper is organized as follows. In section 2 we present the governing equations and the details of the spatial and temporal discretization schemes. In section 3 we discuss the implementation details and introduce the algorithms to allow for asynchrony along with stability analysis. In section 4 we present the numerical results for DNS of decaying and solenoidally forced turbulence, showing the excellent agreement between standard synchronous and asynchronous simulations for both large and small-scale characteristics. We conclude the section with the discussion on the computational performance of the asynchronous solver. Conclusions and scope of future work are discussed in section 4. The appendix lists the AT schemes used in the paper and detailed stability analysis of the schemes.

## 2. Governing equations and numerical schemes

The NS equations, which represent conservation of mass, momentum and energy can be written as,

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i} (\rho u_i) = 0, \quad (1)$$

$$\frac{\partial}{\partial t} (\rho u_i) + \frac{\partial}{\partial x_j} (\rho u_i u_j) = -\frac{\partial p}{\partial x_i} + \frac{\partial}{\partial x_j} (\sigma_{ij}) + \rho f_i, \quad (2)$$



**Fig. 1.** Left boundary of a Processing Element (PE) with  $L$  time levels. The solid blue points are boundary points and hollow blue points are internal points. The solid red points are the buffer points communicated from the neighboring PE. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

$$\frac{\partial}{\partial t}(\rho e) + \frac{\partial}{\partial x_i}(\rho e u_i) = -p \frac{\partial u_i}{\partial x_i} + \frac{\partial}{\partial x_i} \left( k \frac{\partial T}{\partial x_i} \right) + \sigma_{ij} S_{ij}, \quad (3)$$

with  $\rho$  being the density,  $u_i$  the  $i$ th component of velocity,  $e$  the internal energy per unit mass which depends upon temperature ( $T$ ) according to the perfect gas law,  $k$  the coefficient of thermal conductivity,  $p$  the pressure, and  $f_i$  the external forcing. The viscous stress and the strain rate tensors are given, respectively, by,

$$\sigma_{ij} = \mu \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \frac{2}{3} \delta_{ij} \frac{\partial u_k}{\partial x_k} \right), \quad (4)$$

$$S_{ij} = \frac{1}{2} \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right), \quad (5)$$

where the dynamic viscosity,  $\mu$ , follows Sutherland viscosity law.

The above equations are solved numerically using finite difference approximations for the spatial derivatives. In order to do so, the physical domain is discretized into  $N$  grid points in each direction and this discretized domain is then decomposed into  $P$  sub-domains, where  $P$  is the number of processing elements (PEs). Fig. 1 shows the left boundary (dashed black line) of one such PE in 1D, with internal points in hollow blue, boundary points in solid blue and the points communicated from the neighboring PE, known as the buffer points in solid red. For standard finite difference schemes, the derivative at the  $i$ th grid point is a weighted average of the values at the neighboring points. Mathematically, this approximation for second derivative, at time level  $n$ , is given by

$$\frac{\partial^2 u}{\partial x^2} \Big|_i^n \approx \frac{1}{\Delta x^2} \sum_{m=-M}^M c_m u_{i+m}^n + \mathcal{O}(\Delta x^p), \quad (6)$$

where  $M$  is the stencil size in each direction and the  $c_m$ 's are the weights that are computed using Taylor expansion of  $u_{i+m}$  in space such that the order  $p$  in the truncation error term is the highest. For example, a standard synchronous implementation of a second order scheme ( $M = 1$ ) is outlined in faded red in Fig. 1 and given by

$$\frac{\partial^2 u}{\partial x^2} \Big|_i^n = \frac{u_{i-1}^n - 2u_i^n + u_{i+1}^n}{\Delta x^2}, \quad (7)$$

where all points are at time level  $n$ . Computation of derivatives at the boundary points requires updated values (*i.e.* at time level  $n$ ) at the buffer points which are communicated from the neighboring PE. This forces the communications across the PEs to synchronize and leads to additional overheads due to processor idling.

To avoid this, one can instead allow computations to proceed asynchronously, that is, without waiting for the most updated value at the PE boundaries. This results in delayed values at buffer points. Explicitly, the derivative is then computed as

$$\frac{\partial^2 u}{\partial x^2} \Big|_i^n = \frac{u_{i-1}^{n-\tilde{k}} - 2u_i^n + u_{i+1}^n}{\Delta x^2}, \quad (8)$$

where  $n - \tilde{k}$  is the latest available time level written in terms of the delay  $\tilde{k}$ . This scheme is schematically shown as a dashed green curve in Fig. 1. Because of the delay ( $\tilde{k}$ ), the accuracy of the standard scheme degrades severely. In fact it can be shown [23] that the resulting scheme is zeroth-order. This prevents the use of the standard finite difference schemes asynchronously and necessitates the need for numerical methods that are resilient to asynchrony. Such family of schemes has been put forth in [24]. These so-called Asynchrony-Tolerant (AT) schemes preserve the order of accuracy, despite asynchrony and are described next.

### 2.1. Spatial AT schemes

AT schemes can be seen as a generalization of standard finite differences, where the computation of spatial derivatives uses function values of neighboring points in both space and time. For example, the second derivative of a spatially and temporally varying function,  $u(x, t)$ , at grid point  $i$  and time level  $n$  can then be written as,

$$\left. \frac{\partial^2 u}{\partial x^2} \right|_i^n \approx \frac{1}{\Delta x^2} \sum_{l=0}^L \sum_{m=-M}^M c_m^l u_{i+m}^{n-l} + \mathcal{O}(\Delta x^p). \quad (9)$$

Here the weights  $c_m^l$ 's are computed by solving a system of linear equations constructed by imposing order of accuracy constraints on the Taylor series expansion of  $u_{i+m}^{n-l}$  in space and time. The choice of stencil and the general methodology for the derivation of these AT schemes has been explained in detail in [24]. As an example, a second order AT scheme at the left boundary with stencil  $M = 1$  in space, and a delay of  $\tilde{k}$  can be written as,

$$\left. \frac{\partial^2 u}{\partial x^2} \right|_i^n = \frac{-\tilde{k}u_{i-1}^{n-\tilde{k}-1} + (\tilde{k} + 1)u_{i-1}^{n-\tilde{k}} - 2u_i^n + u_{i+1}^n}{\Delta x^2}, \quad (10)$$

where we have used a diffusive CFL relation of the form  $\Delta t \sim \Delta x^2$  to relate spatial and temporal resolutions. This scheme is shown schematically in Fig. 1 with a solid blue curve. These schemes use multiple consecutive time levels on the delayed side, depending upon the order of accuracy. An interesting feature of this kind of AT schemes is that the coefficients are a function of delay  $\tilde{k}$  and they reduce to standard coefficients in the absence of delays ( $\tilde{k} = 0$ ). Note that these delays depend upon the machine characteristics such as clock rate, network latency, bandwidth and topology. Because of this, the delays and, consequently, the coefficients are random and computed dynamically at runtime.

In this work we use fourth-order AT schemes at processor boundaries for spatial derivatives in each direction which require communication across six faces of each PE in a 3D domain. At the internal points we use standard fourth-order finite differences for spatial derivatives. Computation of mixed derivatives is challenging as they require communication across more neighboring PEs or communication of additional quantities such as gradients. Both of these are detrimental to parallel performance. As an alternative, we limit our communications per PE to six by computing mixed derivatives at the boundary points in three steps. For example, for  $(\partial(\partial u/\partial y)/\partial x)$ , we first compute  $\partial u/\partial y$  and  $\partial u/\partial x$  using AT schemes at the boundaries. Next we compute  $(\partial(\partial u/\partial y)/\partial x)$  and  $(\partial(\partial u/\partial x)/\partial y)$  using standard one sided finite difference schemes in  $x$  and  $y$  direction, respectively. Since  $(\partial(\partial u/\partial y)/\partial x) = (\partial(\partial u/\partial x)/\partial y)$ , we take the average of  $(\partial(\partial u/\partial y)/\partial x)$  and  $(\partial(\partial u/\partial x)/\partial y)$  to minimize errors and use this averaged value as the final approximation of the corresponding mixed derivatives.

Because the NS equations represent conservation laws, it is important that the numerical discretization of these laws also satisfies the global conservation. For example, in a 1D form of conservation law,

$$\frac{\partial u}{\partial t} + \frac{\partial f}{\partial x} = 0, \quad (11)$$

where  $f(x, t)$  is the flux, the total variation of  $u(x, t)$  over a domain  $[0, 1]$  depends only upon the flux through the boundaries. This can be expressed more precisely by integrating Eq. (11) over the domain,

$$\frac{d}{dt} \int_0^1 u(x, t) dx = - \int_0^1 \left( \frac{\partial f}{\partial x} \right) dx = f(0, t) - f(1, t), \quad (12)$$

showing explicit dependence of variation in  $u(x, t)$  only on the flux at the boundaries. For periodic boundary condition i.e.  $f(0, t) = f(1, t)$ , this flux is equal to zero. When the derivatives are approximated numerically, it is desirable that the discrete form of the above conservation law is also satisfied to a given accuracy. Consider a generalized spatial discretization given by Eq. (9), for  $N$  grid points and time level  $n$ , to yield

$$\int_0^1 \left. \frac{\partial f}{\partial x} \right|_i^n dx = \sum_{i=1}^N \left( \frac{1}{\Delta x} \sum_{l=0}^L \sum_{m=-M}^M c_m^l f_{i+m}^{n-l} \right). \quad (13)$$

For  $M = 1$ , corresponding to an AT scheme with leading truncation error term of order  $\mathcal{O}(\Delta x^a)$  where  $a = 2$  when  $\Delta t \sim \Delta x^2$  [24], and a domain decomposed into 2 PEs such that PE<sup>(1)</sup> holds gridpoints  $i \in [1, N/2]$  and PE<sup>(2)</sup> holds gridpoints  $i \in [N/2 + 1, N]$  and satisfies periodic boundary conditions, we can write Eq. (13) as

$$\int_0^1 \left. \frac{\partial f}{\partial x} \right|^n dx = \frac{f_2^n - (\tilde{k}_l^{(1)} + 1)f_N^{n-\tilde{k}_l^{(1)}} + \tilde{k}_l^{(1)} f_N^{n-\tilde{k}_l^{(1)}-1}}{2\Delta x} + \sum_{i=2}^{N/2-1} \left( \frac{f_{i+1}^n - f_{i-1}^n}{2\Delta x} \right) + \frac{(\tilde{k}_r^{(1)} + 1)f_{N/2+1}^{n-\tilde{k}_r^{(1)}} - \tilde{k}_r^{(1)} f_{N/2+1}^{n-\tilde{k}_r^{(1)}-1} - f_{N/2-1}^n}{2\Delta x} + \frac{f_{N/2+2}^n - (\tilde{k}_l^{(2)} + 1)f_{N/2}^{n-\tilde{k}_l^{(2)}} + \tilde{k}_l^{(2)} f_{N/2}^{n-\tilde{k}_l^{(2)}-1}}{2\Delta x} + \sum_{i=N/2+2}^{N-1} \left( \frac{f_{i+1}^n - f_{i-1}^n}{2\Delta x} \right) + \frac{(\tilde{k}_r^{(2)} + 1)f_1^{n-\tilde{k}_r^{(2)}} - \tilde{k}_r^{(2)} f_1^{n-\tilde{k}_r^{(2)}-1} - f_{N-1}^n}{2\Delta x}, \tag{14}$$

where  $\tilde{k}_l^{(1)}$  and  $\tilde{k}_r^{(1)}$  are the delays on left and right boundary for PE<sup>(1)</sup> and  $\tilde{k}_l^{(2)}$  and  $\tilde{k}_r^{(2)}$  are the delays on left and right boundary for PE<sup>(2)</sup> and periodic boundary conditions are used. Because of the telescoping effect, the above expression can be simplified to

$$\int_0^1 \left. \frac{\partial f}{\partial x} \right|^n dx = \frac{-(\tilde{k}_l^{(1)} + 1)f_N^{n-\tilde{k}_l^{(1)}} + \tilde{k}_l^{(1)} f_N^{n-\tilde{k}_l^{(1)}-1}}{2\Delta x} + \left( \frac{f_{N/2}^n - f_1^n}{2\Delta x} \right) + \frac{(\tilde{k}_r^{(1)} + 1)f_{N/2+1}^{n-\tilde{k}_r^{(1)}} - \tilde{k}_r^{(1)} f_{N/2+1}^{n-\tilde{k}_r^{(1)}-1}}{2\Delta x} + \frac{-(\tilde{k}_l^{(2)} + 1)f_{N/2}^{n-\tilde{k}_l^{(2)}} + \tilde{k}_l^{(2)} f_{N/2}^{n-\tilde{k}_l^{(2)}-1}}{2\Delta x} + \left( \frac{f_N^n - f_{N/2+1}^n}{2\Delta x} \right) + \frac{(\tilde{k}_r^{(2)} + 1)f_1^{n-\tilde{k}_r^{(2)}} - \tilde{k}_r^{(2)} f_1^{n-\tilde{k}_r^{(2)}-1}}{2\Delta x}. \tag{15}$$

For the standard synchronous case,  $\tilde{k}_l^{(1)} = \tilde{k}_r^{(1)} = \tilde{k}_l^{(2)} = \tilde{k}_r^{(2)} = 0$ , that is, when delays are absent, all terms on the right-hand side of Eq. (15) cancel each other and the conservative property is trivially satisfied. In the presence of delays, on the other hand, this is not immediately obvious from Eq. (15). Further simplification of this equation can be done using a Taylor series expansion in time which leads to similar cancellation of all low-order terms yielding a residual of the order of  $\Delta x^3$ . More generally, for larger  $M$ , that is, for AT schemes of order  $a = 2M$  and  $\Delta t \sim \Delta x^2$ , the residual is found to be

$$\int_0^1 \left. \frac{\partial f}{\partial x} \right|^n dx = \mathcal{O}(\Delta x^{a+1}). \tag{16}$$

Thus, we conclude that the AT schemes retain the conservative property up to an order higher than the order of the scheme. For the fourth-order schemes used below, conservation is satisfied up to  $\mathcal{O}(\Delta x^5)$ .

### 2.2. Temporal schemes

For the evolution of a system of PDEs in time, spatial schemes need to be coupled with a temporal scheme of appropriate order. High order explicit temporal methods including multistage Runge-Kutta (RK) schemes and multistep Adams-Bashforth schemes, are very common choices of temporal discretizations. While RK schemes are known for their good stability characteristics, the computation of a stage of RK requires communication of previous stages across all neighbors at all times. In a 3D domain this is equivalent to  $26 \times s$  communications and  $s$  computations of the right hand side of the PDE, per PE, per time-step, for an  $s$ -stage RK scheme. Consequently, RK schemes are computation, communication and synchronization intensive. On the other hand, multi-step Adams-Bashforth (AB) schemes offer more flexibility in terms of implementation and require less communications. A general AB scheme with  $T$  steps for an equation of the form  $\partial u/\partial t = f$ , can be written as,

$$u_i^{n+1} = u_i^n + \Delta t \sum_{m=0}^{T-1} \beta_m f_i^{n-m}, \tag{17}$$

where the coefficients  $\beta_m$  depend upon the desired order of accuracy [25]. Not only can AB be efficiently implemented with only six communications per PE per time-step, it only requires computation of  $f^n$  every time-step since  $f^{n-m}, m > 0$  is used from previous steps. Furthermore, the computation of  $f^{n-m}$  using AT schemes does not alter the order of accuracy of AB schemes [24]. Thus, here we use second-order AB schemes for the temporal evolution in both synchronous and asynchronous simulations.

The CFL, relating time-step size to the grid spacing, can be used to determine the leading order error term of a fully discretized PDE in order to ensure that global order of accuracy is preserved. For example, for a convective CFL ( $r_c$ ), the time-step  $\Delta t$  is computed as,

$$\Delta t = \frac{r_c \Delta x}{u_{max}} \quad (18)$$

where  $u_{max}$  is the global maximum velocity. Since this maximum is computed across all PEs, it requires a collective blocking communication call at every time step and leads to more synchronization overheads. To avoid this, instead of a CFL condition, one can use a fixed  $\Delta t$  [15,26]. This is the approach we adopt here. For consistency, synchronous simulations are also done at the same fixed  $\Delta t$ . In summary, we use fourth-order AT schemes in space for boundary points, fourth-order finite difference schemes at internal points and AB schemes in time with a fixed  $\Delta t$ .

### 3. Implementation

The compressible flow solver is parallelized using a 3D domain decomposition and each PE is responsible for computations in a piece of this 3D domain. Communications between PEs are localized to the nearest neighbors only. In Fig. 1 we see a simple 1D domain decomposition where every PE has  $N_T$  grid points and two neighbors. The number of internal points ( $N_I$ ) i.e. the number of points that use standard (synchronous) finite differences, with  $M$  points in each direction is equal to  $N_T - 2M$ . The total number of boundary points ( $N_B$ ) and buffer points ( $N_{bf}$ ) in this case are equal to  $2M$ . Clearly  $N_B \cup N_I = N_T$  and  $N_B \cap N_I = \emptyset$ .

Extending this idea to a 3D computational topology, such that each PE has a total of twenty six neighboring PEs, we can compute the total number of internal and boundary points. Consider a general 3D domain with  $N_x$ ,  $N_y$ , and  $N_z$  grid points and  $P_x$ ,  $P_y$ , and  $P_z$  processors in the  $x$ ,  $y$ , and  $z$  directions, respectively. Then the total number of grid points ( $N_T$ ) per PE is

$$N_T = \frac{N_x N_y N_z}{P_x P_y P_z}. \quad (19)$$

Using spatial schemes which require  $M$  points on each side for all three directions, it is easy to show that the number internal points ( $N_I$ ) is

$$N_I = \left( \frac{N_x}{P_x} - 2M \right) \left( \frac{N_y}{P_y} - 2M \right) \left( \frac{N_z}{P_z} - 2M \right). \quad (20)$$

Since communications are done across all six faces of a PE, Eq. (19) and Eq. (20) gives us the exact number of boundary points ( $N_B$ ) or the points that use AT schemes for the computation of spatial derivatives,

$$N_B = \frac{N_x N_y N_z}{P_x P_y P_z} - \left( \frac{N_x}{P_x} - 2M \right) \left( \frac{N_y}{P_y} - 2M \right) \left( \frac{N_z}{P_z} - 2M \right). \quad (21)$$

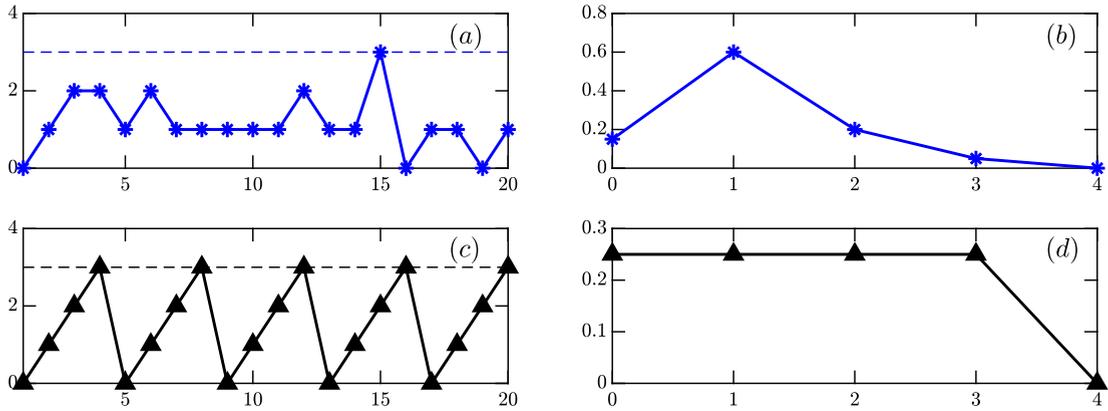
We can then compute the percentage of points that use AT schemes,

$$N_B(\%) = 100 \left( \frac{N_B}{N_T} \right), \quad (22)$$

which can be used as a metric of the extent in space in which asynchrony affects the computations of derivatives directly.

#### 3.1. Algorithm

We solve the NS equation for five variables ( $\rho, \rho u_1, \rho u_2, \rho u_3, \rho e$ ) at every time step. Since data at older time levels is used for AT schemes, each PE stores  $5 \times (N_I + N_B + N_{bf}) \times (L + \tau_\ell + 1)$  data points, where  $L$  is the maximum allowed delay that can also be used as a control parameter for error and stability as we show below and  $\tau_\ell$  is the number of consecutive time-levels required for the computation of derivatives by AT schemes. We use two-sided non-blocking MPI calls ( $MPI\_Isend, MPI\_Irecv$ ) for asynchronous communications between the PEs across the six faces of the 3D computational domain. In each direction, these communications are limited to immediate neighbors only. The status of these non-blocking communications is checked using  $MPI\_test$  and it is utilized to compute delay at each PE boundary. To control the manner in which asynchrony appears we utilize two control parameters  $c_r$  and  $L$ . The *communication rate*  $c_r$  specifies the frequency of communication in each direction, that is to say, PEs initiate communication calls every  $c_r$  consecutive time steps. The second parameter is the maximum allowed delay,  $L$ . PEs impose explicit synchronization by invoking  $MPI\_Wait$  whenever instantaneous delays at PE boundaries cross this threshold  $L$ . This synchronization is imposed only in the direction in which the delay is larger than  $L$  and is thus local in nature. These two parameters determine the nature of delays. For example, if  $c_r > 1, L > 0$  delays are periodic and if  $c_r = 1, L > 0$  then delays at PE boundaries are random. In both the cases the delays are however bounded by  $L$ . Synchronous simulations are realized when  $c_r = 1$  and  $L = 0$ . Irrespective of  $c_r$  and  $L$ , global communications and synchronizations involving all PEs are done only for I/O.



**Fig. 2.** Simulated time series of delays for (a) SAA with  $c_r = 1$ ,  $L = 3$  and (c) CAA with  $c_r = 4$ ,  $L = 3$ . PDF of these simulated delays for (b) SAA and (d) CAA.

### 3.1.1. Synchronization avoiding algorithm (SAA): random delays

For  $c_r = 1$  and  $L > 1$ , we have what we call a synchronization avoiding algorithm: a local synchronization is applied *if and only if* the delay ( $\tilde{k}$ ) at a PE boundary is greater than the maximum allowed delay  $L$ . We use circular send ( $U_{send}$ ) and receive ( $U_{recv}$ ) buffers in each direction for communicating to and from the neighboring PEs, respectively. At each time step, PEs send data at only one time level across the boundaries. A generalized SAA is listed in Algorithm 1, where computations can proceed without waiting for updated values whenever  $\tilde{k} \leq L$ . When synchronization is not imposed, communications can complete in the background, facilitating overlap between communications and computations. Despite the reduction of synchronization overheads and PE idling, this method does require communication at every step. The delay observed at PE boundaries is a function of machine characteristics, such as, network performance, processor and memory speeds etc., and is therefore a random variable, with a different value at each of the six PE boundaries. Since the delays and, consequently, coefficients of AT schemes may be different for each PE boundary, some additional numerical errors can be introduced due to random nature of these delays. Numerical simulations show that this effect is negligible for values of  $L$  that satisfy stability (section 3.2).

A typical time series of random delays for SAA is shown in Fig. 2(a) along with its PDF in Fig. 2(b). In this example, the delay is bounded by  $L = 3$  as shown by the dashed blue line in Fig. 2(a). The statistical moments of the distribution of delays have an effect on the accuracy of the solution [24]. Since statistical characteristics of the delays can be controlled by forced synchronization,  $L$  becomes a parameter for error control.

---

**Algorithm 1:** Synchronization Avoiding Algorithm (SAA). Here  $U^n$  is the variable array at time level  $n$ ,  $U_{send}$  is the send buffer,  $U_{recv}$  is the receive buffer,  $U_{buffer}$  is the data at buffer points,  $U_{boundary}^n$  is the data at boundary points and  $f$  evaluates the discretized NS equation using AT schemes in space and AB schemes in time for number of time steps equal to  $st$ .

---

Synchronous Loop: Initialize  $L + \tau_\ell$  levels of  $U$ ,  $U_{send} \leftarrow U_{boundary}$ ,  $U_{buffer} \leftarrow U_{recv}$

Asynchronous Loop:

**for**  $n = L + \tau_\ell + 1, \dots, st$  **do**

$U^{n+1} = f(U^n, U^{n-1}, \dots, U^{n-T+1})$

$U_{send} \leftarrow U_{boundary}^{n+1}$

    Send data across 6 faces:  $MPI\_Isend$

**for**  $face = 1 : 6$  **do**

        Check communication status:  $MPI\_Test$

        Compute delay ( $\tilde{k}$ )

**if** delay ( $\tilde{k}$ )  $\leq L$  **then**

$U_{buffer} \leftarrow U_{recv}$

**else**

            Force synchronization:  $MPI\_Wait$

            Update delay ( $\tilde{k}$ )

$U_{buffer} \leftarrow U_{recv}$

**end**

        Compute coefficients of the AT schemes (Appendix B)

**end**

**end**

---

### 3.1.2. Communication avoiding algorithm (CAA): periodic delays

As alternative to communicating at every time step, we propose the so-called communication avoiding algorithm, in which the PEs communicate periodically every  $c_r > 1$  steps. As a result, the delay changes periodically from 0 (no delay) to a maximum allowed delay  $L$  which satisfies  $L = c_r - 1$ . Because of this periodicity, the delay across all the PE boundaries is the same in every direction. Since PEs communicate every  $c_r$  time step, the send and receive buffers now have data at  $\min(\tau_\ell, c_r)$  time levels. This multiple time level data is required for computation of derivatives using AT schemes at the communication avoiding time steps. We have listed a generalized implementation of CAA in Algorithm 2, where the delay is incremented by one when PEs do not communicate. A typical time series of delays bounded by  $L = 3$  (dashed black line) is shown in Fig. 2(c) for CAA with  $c_r = 4$ . The delay in this case is deterministic and the PDF shown in Fig. 2(d) has a uniform distribution. Both delay and its PDF are independent of the machine characteristics and depend only upon the control parameters, contrary to SAA where the delay is random and its PDF is machine specific. CAA reduce the total latency time by a factor of  $c_r$  in comparison to synchronous avoiding or standard synchronous algorithms and are therefore particularly effective in latency-dominated machines. Furthermore, because of the reduction in frequency of communications, the energy consumption for these algorithms is also expected to be reduced. One drawback of the communication avoiding algorithms is the larger size of send and receive buffers that could adversely affect performance for bandwidth-dominated machines.

---

**Algorithm 2:** Communication avoiding algorithm. Here  $U^n$  is the variable array at time level  $n$ ,  $U_{send}$  is the send buffer,  $U_{recv}$  is the receive buffer,  $U_{buffer}$  is the data at buffer points,  $U_{boundary}^n$  is the data at boundary points and  $f$  evaluates the discretized NS equation using AT schemes in space and AB schemes in time for number of time steps equal to  $st$ .

---

Synchronous Loop: Initialize  $L + \tau_\ell$  levels of  $U$ ,  $U_{send} \leftarrow U_{boundary}$ ,  $U_{buffer} \leftarrow U_{recv}$

Compute  $\ell = \min(\tau_\ell, c_r)$

Asynchronous Loop:

```

for  $n = L + 1, \dots, st$  do
   $U^{n+1} = f(U^n, U^{n-1}, \dots, U^{n-T+1})$ 
  for  $face = 1 : 6$  do
    if  $(\text{mod}(n, c_r) == 0)$  then
       $U_{send} \leftarrow U_{boundary}^{n+1}, \dots, U_{boundary}^{n-\ell+2}$ 
       $MPI\_Isend$  and  $MPI\_Irecv$ 
       $U_{buffer} \leftarrow U_{recv}$ 
       $\tilde{k} \leftarrow 0$ 
    else
      Update delay:  $\tilde{k} \leftarrow k + 1$ 
    end
    Compute coefficients of the AT schemes (Appendix B)
  end
end

```

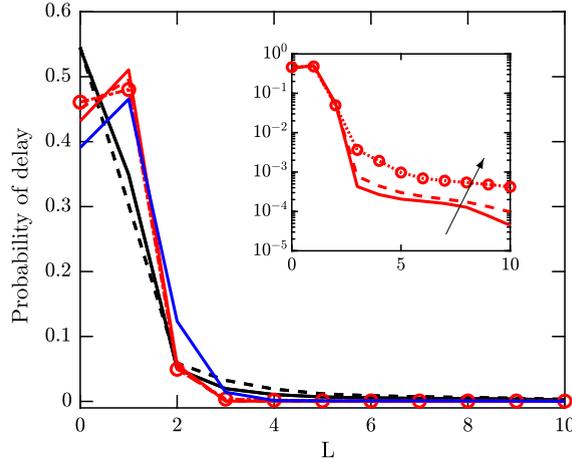
---

### 3.2. Maximum delay $L$ and stability

The maximum allowed delay  $L$  is an important control parameter as it determines the error and stability of the AT schemes as well as the computational performance of the solver. As shown in [24], the error due to asynchrony in AT schemes is a function of statistical moments of delays which depend upon the architecture of the machine, communication links and patterns, latency, bandwidth and clock speed. Since the asynchronous error grows with  $L$  [24], very large values of  $L$  can affect the accuracy of simulations. Furthermore, the memory requirement of all stored variables, the size of send and receive buffers and the rate of synchronizations and communications are also directly affected by the choice of parameter  $L$ . It is therefore critical that  $L$  be chosen judiciously in simulations and this choice can be based on two main factors that are described next.

First,  $L$  has implications in terms of the computational implementation of the solver. Increasing  $L$  increases the number of times levels that need to be stored which increases memory requirement. At the same time, if  $L$  is too small then synchronization will be forced more often than required and asynchrony will not be leveraged efficiently. In practice, one can run a short simulation with a very large  $L$  and obtain the PDF of the delays ( $\tilde{k}$ ). From this data, one can calculate an appropriate  $L$  by requiring  $P(\tilde{k} > L) \lesssim c$ , that is to say, one would expect forced synchronization  $c\%$  of the time. Thus,  $c$  exposes tradeoff between performance and accuracy through the degree of asynchrony. For example, at  $c = 0$  the simulation is completely asynchronous, i.e., synchronization is never imposed, which is detrimental to accuracy if  $L$  is large.

For illustration purposes, in Fig. 3 we show PDF of delays ( $k$ ) on three large systems at Texas Advanced Computing Center (TACC), namely, Stampede2, Frontera and Lonestar5 for  $L = 10$  and different processor counts. From the black lines, we can clearly see that the probability of delays decreases with increasing delay and  $P(\tilde{k} > 3) \lesssim 0.05$  on Stampede2. This implies that for a simulation with  $L = 3$ , synchronizations will be forcefully imposed less than 5% of the time. The trend



**Fig. 3.** PDF of delays on Stampede2 (black), Frontera (red) and Lonestar5 (blue) with maximum allowed delay of  $L = 10$ . Different lines are  $P = 4096$  (solid red, solid blue),  $P = 8192$  (solid black, dashed red) and  $P = 16384$  (dashed black) and  $P = 262144$  (dotted-circles). Inset is PDF of delays on Frontera in linear-log scale with arrow denoting increasing  $P$ .

is consistent even if we double the number of processors from  $P = 8192$  (solid black) to  $P = 16384$  (dashed black). For Frontera (red lines in Fig. 3) we see that the probability of  $\tilde{k} = 1$  is higher than the probability of  $\tilde{k} = 0$  for all the three cases. This points to a slow network that is expected to adversely affect the scaling for standard synchronous simulations. We see similar behavior for Lonestar5 (blue), with probability of  $\tilde{k} = 1$  being the maximum. For both Frontera and Lonestar5,  $P(\tilde{k} > 3) \lesssim 0.05$ , for all the processor counts shown in Fig. 3. Thus,  $L = 3$  is a reasonable choice for these three machines. Note that for CAA this is equivalent to a reduction in the volume of communications by a factor of four. This reduction will be particularly critical when the PE count is high as envisioned in the exascale machines.

It is also worth noting the non-vanishing probability of delays as large as  $L = 10$  in Fig. 3 (inset) for Frontera, indicating that at least some fraction of communications were synchronized. While this is not of much consequence at low processor counts, for an increasingly large number of PEs, even a small probability of large delays can account for severe overheads. For example, for a seemingly low probability of  $P(\tilde{k} = 10) \sim \mathcal{O}(10^{-4})$ , at a processor count of  $\mathcal{O}(10^5)$ , at least  $\mathcal{O}(10^1)$  processors see a delay of  $L = 10$  at the boundaries and are forced to synchronize at every time step. Considering that the probability of large delays increases with increasing number of PEs as seen in Fig. 3 (inset), a much larger fraction of processors would see large delays in the next generation exascale machines where the number of PEs is expected to be of  $\mathcal{O}(10^6)$ - $\mathcal{O}(10^9)$  with increased architectural inhomogeneity. Thus, even with a large value for the maximum allowed delay, a significant number of PEs would be subject to forced synchronization at extreme scales. However, these synchronizations would still be extremely small in comparison to the standard synchronous algorithms that require all PEs to synchronize at all times.

The second and equally important factor to be considered in the choice of  $L$  is numerical stability. Asynchrony and the associated random nature of delays and coefficients introduce random numerical errors. These errors can trigger instabilities, especially if the delay ( $\tilde{k}$ ), bounded by  $L$ , is very large. We will discuss this effect for a simple 1D diffusion equation,

$$\frac{\partial u}{\partial t} = \alpha \frac{\partial^2 u}{\partial x^2} \quad (23)$$

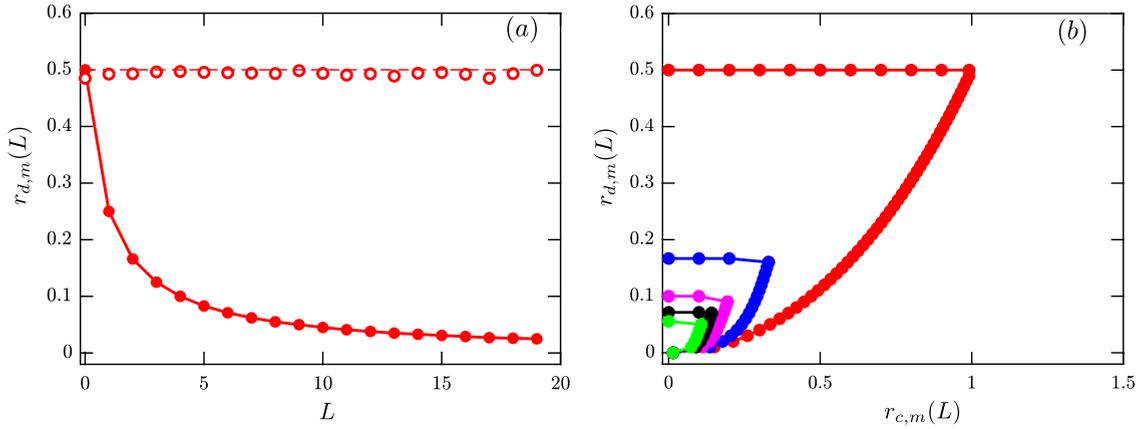
where  $\alpha$  is the diffusivity constant and  $u(x, t)$  is the velocity field. This equation is discretized using a second order AT scheme in space and forward Euler in time. Following [24], we can discretize Eq. (23) at the  $i$ -th grid point with delay  $\tilde{k}_l$  at the left boundary and  $\tilde{k}_r$  at the right boundary as,

$$u_i^{n+1} = u_i^n + \frac{\alpha \Delta t}{\Delta x^2} \left( (\tilde{k}_l + 1)u_{i-1}^{n-\tilde{k}_l} - \tilde{k}_l u_{i-1}^{n-\tilde{k}_l-1} - 2u_i^n + (\tilde{k}_r + 1)u_{i+1}^{n-\tilde{k}_r} - \tilde{k}_r u_{i+1}^{n-\tilde{k}_r-1} \right). \quad (24)$$

For the above discretization we have considered an extreme case scenario where  $P = N$  and  $N_T = 1$  i.e. every PE has only one grid point. It can be shown that Eq. (24) preserves the order of accuracy despite delays on both boundaries. Next we define  $U^n := [u_1^n, \dots, u_N^n]$  and  $V^n := [U^n, \dots, U^{n-\tilde{k}-1}]^T$ , where  $\tilde{k} = \max(\tilde{k}_l, \tilde{k}_r)$ . Using these definitions, we can write the matrix form of the evolution equation as,

$$V^{n+1} = \mathbf{A}(\tilde{k}_l, \tilde{k}_r) V^n \quad (25)$$

where the coefficient matrix is



**Fig. 4.** (a) Variation of stability limit  $r_d$  (solid) and  $\tilde{r}_d = (L + 1)r_d$  (hollow) with  $L$  for diffusion equation. (b) Stability limit in  $r_c$ - $r_d$  plane for advection-diffusion equation for  $L = 0$  (red),  $L = 2$  (blue),  $L = 4$  (magenta),  $L = 6$  (black) and  $L = 8$  (green).

$$\mathbf{A}(\tilde{k}_l, \tilde{k}_r) = \begin{bmatrix} \mathbf{A}_0 & \mathbf{A}_1 & \dots & \mathbf{A}_{\tilde{k}} & \mathbf{A}_{\tilde{k}+1} \\ \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{I} & \mathbf{0} \end{bmatrix}. \tag{26}$$

While this equation is very general, we specialize this system to same delay on both sides ( $\tilde{k} = \tilde{k}_l = \tilde{k}_r$ ) for all processors, which can be thought as a worst case scenario. Defining  $r_d = \alpha \Delta t / \Delta x^2$  as the diffusive CFL, we can then write,

$$\mathbf{A}_0(\tilde{k}) = \begin{bmatrix} 1 - 2r_d & \mathcal{L}_1^0 r_d & 0 & \dots & 0 & \mathcal{L}_1^0 r_d \\ \mathcal{L}_1^0 r_d & 1 - 2r_d & \mathcal{L}_1^0 r_d & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathcal{L}_1^0 r_d & 0 & 0 & \dots & \mathcal{L}_1^0 r_d & 1 - 2r_d \end{bmatrix}, \quad \mathbf{A}_{\tilde{k}}(\tilde{k}) = \begin{bmatrix} 0 & \mathcal{J} & 0 & \dots & 0 & \mathcal{J} \\ \mathcal{J} & 0 & \mathcal{J} & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathcal{J} & 0 & 0 & \dots & \mathcal{J} & 0 \end{bmatrix}$$

with  $\mathcal{J}$  defined as,

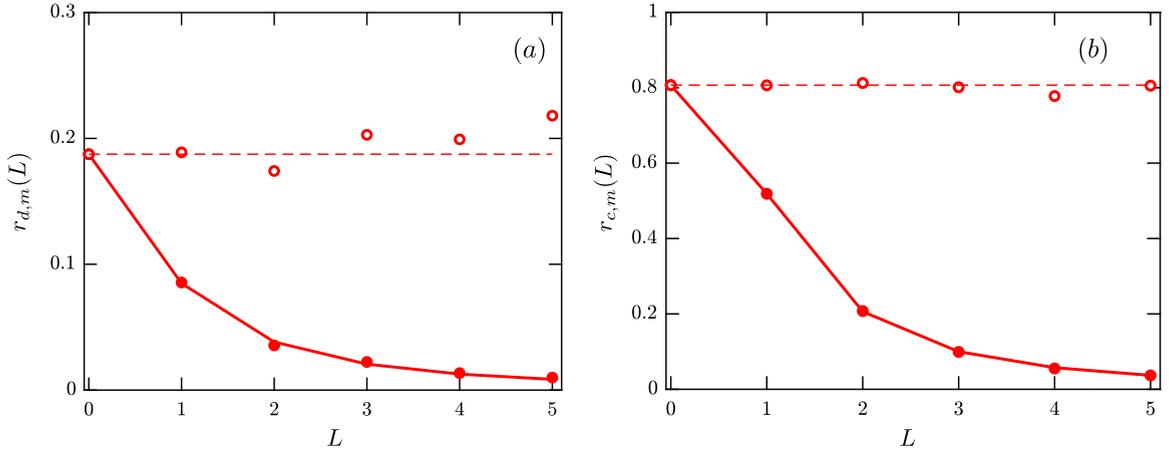
$$\mathcal{J} = \mathcal{L}_1^m r_d (\tilde{k} + 1) - \mathcal{L}_2^m r_d \tilde{k} \tag{27}$$

which is used to set the coefficient as  $r_d(\tilde{k} + 1)$  for  $\mathbf{U}^{n-\tilde{k}}$  and  $-r_d \tilde{k}$  for  $\mathbf{U}^{n-\tilde{k}-1}$ . For this we use  $\mathcal{L}^m$ , which is the Lagrange polynomial of order  $L$ ,

$$\mathcal{L}_1^m(\tilde{k}) = \prod_{l \neq m} \frac{\tilde{k} - l}{m - l} \quad \mathcal{L}_2^m(\tilde{k}) = \prod_{l \neq m} \frac{\tilde{k} + 1 - l}{m - l}. \tag{28}$$

By definition,  $\mathcal{L}_1^m(\tilde{k})$  takes value 1 if  $m = \tilde{k}$  and zero otherwise. Similarly  $\mathcal{L}_2^m(\tilde{k})$  is 1 if  $m = \tilde{k} + 1$  and 0 for other values of  $m$ . The number of Lagrange polynomials is equal to the number of time levels in the AT scheme, which for the second-order scheme used here is equal to two. In the absence of delays we have,  $\mathbf{A} = \mathbf{A}_0$ , which is equivalent to the standard second-order finite difference system.

For stability, the spectral radius of  $\mathbf{A}(\tilde{k})$  should be bounded by unity to ensure that the numerical perturbations do not grow unboundedly in time. Because of the complexity of the system, the spectrum has to be computed numerically. Again as a worst case scenario [24], we assume a Dirac delta distribution of delays, such that,  $\tilde{k} = L$  at all points. For a given  $L$ , we compute the maximum or critical  $r_d$  for which all the eigenvalues of the evolution matrix are less than unity. This is the largest value for which the numerical scheme is stable, and is denoted by  $r_{d,m}(L)$ . The results of this analysis are shown in Fig. 4. In the synchronous limit ( $L = 0$ ), we obtain the well known stability limit for a second order central difference scheme in space with forward Euler in time,  $r_{d,m}(0) = 0.5$  [27]. As we increase  $L$ , this stability limit decreases as can be seen from the solid red circles in Fig. 4(a). Similar analysis was also done for the advection-diffusion equation which has both first and second derivatives and thus both convective ( $r_c = c \Delta t / \Delta x$ ) and diffusive CFLs are used to determine stability. Here again we fix the delay  $\tilde{k} = L$  and compute the stability limit in the  $r_c$ - $r_d$  plane. The procedure was repeated for different values of  $L$ . The result is plotted in Fig. 4(b). For  $L = 0$ , we get the well known stability bound,  $r_{c,m}(0) \leq 2(r_{d,m}(0))^2 \leq 1$  [27]. As  $L$  is increased we see that both  $r_{c,m}(L)$  and  $r_{d,m}(L)$  decrease resulting in a smaller stability region. Thus, for stability a time-step smaller than that for  $L = 0$  is required whenever  $L > 0$ . We do point out that these stability bounds are based on worst case scenario assumptions and are thus strict. In more realistic scenarios (Appendix A), the effect of asynchrony on stability is relatively weaker.



**Fig. 5.** Variation of stability limit (a)  $r_d$  (solid) and  $\tilde{r}_d = (0.74L^2 + 0.47L + 1)r_d$  (hollow) with  $L$  for diffusion equation and (b)  $r_c$  (solid) and  $\tilde{r}_c = (0.90L^2 - 0.35L + 1)r_c$  (hollow) for NS equation, using fourth-order AT scheme in space and AB2 in time.

In order to characterize the reduction in stability limits, it is of interest to obtain the stability limit in asynchronous simulations from the known stability limit of a synchronous implementation. This can be formulated as

$$r_{d,m}(L) = r_{d,m}(0)/f(L), \quad (29)$$

where the yet unknown function  $f(L)$  characterizes the effect of delays. Clearly,  $f(0) = 1$ . Some guidance on a plausible functional form for  $f(L)$  can be obtained by a careful examination of Eq. (24) where we observe that, in the presence of delays,  $r_d$  at the boundary points always appears in conjunction with functions of delays that are also the coefficients of the AT scheme. In the present case, from Eq. (27) we have  $r_d(\tilde{k} + 1)$  and  $-r_d\tilde{k}$  in the evolution matrix  $\mathbf{A}(\tilde{k})$ . Since both terms are linear in the delay, it is natural to expect that, for  $\tilde{k} = L$ , stability, and thus  $f(L)$ , would be a linear function of  $L$ . In fact, a best fit approximation for  $r_{d,m}(L) = r_{d,m}(0)/f(L)$  does yield a linear relation  $f(L) \approx L + 1$ . Both  $r_{d,m}(0)/(L + 1)$  (solid line) and  $r_{d,m}(L)$  (solid circles) are plotted in Fig. 4(a) and are in excellent agreement with each other. Furthermore, we can re-arrange  $r_{d,m}(L) = r_{d,m}(0)/f(L)$  to read as  $r_{d,m}^a = r_{d,m}(L) \times f(L) = r_{d,m}(0)$ . This implies that with a correct approximation for  $f(L)$ , we can express stability in terms of an *effective asynchronous* CFL ( $r_{d,m}^a$ ), which is independent of delay  $L$  and essentially equal to the synchronous stability limit ( $r_{d,m}(0)$ ). The numerical data do support this argument as can be seen from Fig. 4(a) where  $r_{d,m}^a$  (hollow circles) are constant for all  $L$  and close to  $r_{d,m}(0) = 0.5$  (dashed line).

We also computed the stability limit for the schemes used for the turbulence simulations in this work, namely, fourth-order AT schemes coupled with AB2 in time. This is shown in Fig. 5(a) for the diffusion equation. In this case,  $r_d$  appears multiplied by the coefficients in this fourth-order AT scheme (Appendix B) in the discrete equation which are seen to be quadratic in  $L$ . Then, based on the argument above, we expect  $f(L)$  also to be quadratic in  $L$ . From Fig. 5(a) we can see that  $r_{d,m}(L)$  (solid circles) decreases with  $L$  and is in good agreement with  $r_{d,m}(0)/f(0.74L^2 + 0.47L + 1)$  (solid line). Moreover,  $r_{d,m}^a$  (hollow circles) is close to  $r_{d,m}(0) \approx 0.18$  (dashed line) for all  $L$ . This again supports the proposed rescaling in Eq. (29).

One can understand this effect more intuitively as follows. When there is a delay at the PE boundaries, data from multiple delayed time levels is used at these points for computation of derivatives. As a result the effective time-step, as seen by the numerical scheme, increases. This effective time-step is essentially equivalent to  $\Delta t_L = \Delta t \times f(L)$  and is apparent when  $r_{d,m}^a$  is written as

$$r_{d,m}^a = r_{d,m}(L) \times f(L) = \frac{\alpha(\Delta t \times f(L))}{\Delta x^2} = \frac{\alpha(\Delta t_L)}{\Delta x^2}. \quad (30)$$

For fixed grid spacing  $\Delta x$ , this increase in time-step is compensated by a decrease in  $r_{d,m}(L)$  to ensure stability. On the other hand,  $r_{d,m}^a$  which is already expressed in terms of  $\Delta t_L$ , remains approximately constant with  $L$  and is equal to  $r_{d,m}(0)$ .

For a complex system of equations, such as the Navier-Stokes equations, an analytical stability analysis is difficult. However, stability limits can be computed numerically either by gradually increasing the CFL until the system becomes unstable or by using the bisection method. We obtained the stability limit for decaying turbulence at  $Re_\lambda \approx 35$ , by imposing a fixed delay  $L$  at all the six faces at every time step. Since both diffusive and convective terms are present in the NS equations, the time-step is determined by the smallest physical time scale, which for the simulations presented is always the latter. Thus, the stability limit is obtained in terms of a convective CFL ( $r_c$ ) and is shown in Fig. 5(b) with  $r_{c,m}(L)$  (solid circles) decreasing with  $L$ . As before, this effect is accurately captured by  $r_{c,m}(L) = r_{c,m}(0)/f(L)$  (solid line), where  $f(L) \approx 0.90L^2 - 0.35L + 1$ . Here again,  $r_{c,m}^a = r_{c,m}(L) \times f(L)$  (hollow circles) is seen to be a constant consistent with the synchronous limit  $r_{c,m}(0) \approx 0.8$  for all  $L$ .

In general, this analysis shows that the stability limit for the AT schemes decreases with  $L$  as  $r_{d,m}(0)/f(L)$ . This dependency can also be expressed using the *effective asynchronous* CFL ( $r_{c,m}^a$  or  $r_{d,m}^a$ ) which satisfies the same limit as the

**Table 1**

DNS parameters: number of grid points  $N^3$ , percentage of boundary points  $N_B\%$ , Taylor Reynolds number  $Re_\lambda$ , resolution  $\eta/\Delta x$  and  $\kappa_{max}\eta$  and normalized simulation time  $t/T_e$ . Normalization is done using the initial values ( $Re_\lambda(0)$ ,  $\eta(0)$ ,  $T_e(0)$ ) for the decay cases and using average computed over stationary state for the forced case.

Decaying					
$N^3$	$N_B(\%)$	$Re_\lambda(0)$	$\eta(0)/\Delta x$	$\kappa_{max}\eta(0)$	$t/T_e(0)$
256 <sup>3</sup>	57.8	100	0.5	1.4	24
512 <sup>3</sup>	50.8	145	0.5	1.5	24
Forced					
$N^3$	$N_B(\%)$	$Re_\lambda$	$\eta/\Delta x$	$\kappa_{max}\eta$	$t/T_e$
64 <sup>3</sup>	57.8	35	0.5	1.6	10
256 <sup>3</sup>	57.8	100	0.5	1.8	19

synchronous case ( $r_{c,m}(0)$  or  $r_{d,m}(0)$ ) and uses an effective time-step ( $\Delta t_L = f(L) \times \Delta t$ ). Here  $f(L)$ , which is of the same order in  $L$  as the coefficients in the corresponding AT scheme, gives a quantitative measure of the effect of delays on the stability limit. For example, for a large value of  $f(L)$ , in order to keep  $r_{c,m}^a = f(L) \times r_{c,m}(L)$  constant,  $r_{c,m}(L)$  needs to be small. This implies that a small  $\Delta t$  is required for the asynchronous simulation to be stable, which in turn can increase the computational cost. However, we note that while simulations of turbulent flows at  $r_c = 1$  are prevalent in literature, recent studies have shown that for adequate temporal resolution, a much smaller  $r_c$  should be used [28]. Thus, the CFL (or  $\Delta t$ ) dictated by those resolution requirements, could be much smaller than the reduced stability limit discussed above.

Summarizing the results from this section, the maximum allowable delay ( $L$ ) is chosen such that the PEs incur in minimal overheads because of forced synchronization and communications, without additional computational cost to ensure stability.

## 4. Numerical results

We have implemented the synchronous and asynchronous numerical methods and algorithms described in the previous sections to perform DNS of decaying and forced isotropic turbulence at different Reynolds numbers to assess the effect (or lack thereof) of asynchrony. The resolution used for both synchronous and asynchronous implementations is  $\eta/\Delta x \approx 0.5$  or  $\kappa_{max}\eta \approx 1.5$ , where  $\eta = (\nu^3/\langle\epsilon\rangle)^{1/4}$  is the Kolmogorov length scale,  $\nu$  is the kinematic viscosity and  $\kappa_{max} = \sqrt{2}N/3$  is the highest resolvable wave number for commonly used pseudospectral simulations in a cubic domain of length  $2\pi$  on each side and  $N^3$  points [5,28]. This resolution has been shown to lead to well-resolved simulations for the conditions and quantities of interest presented here [29,30]. As discussed in section 2, the time-step size  $\Delta t$  is fixed at a value that yields an initial CFL of  $\mathcal{O}(0.1)$  consistent with the recommendation in [28]. To facilitate comparisons both synchronous and asynchronous simulations use the same time step. We use periodic boundary condition in all directions. The initial velocity field is a stationary state obtained by forcing the large scales of motion as done in [31,29] and is same for both synchronous and asynchronous simulations. The important simulation parameters including resolution, percentage of points directly affected by asynchrony ( $N_B\%$ ),  $Re_\lambda$ , and simulation time in terms of eddy turnover time  $T_e = \mathcal{L}/u_{rms}$ , where  $\mathcal{L}$  is the integral length scale and  $u_{rms}$  is the root mean square of velocity fluctuations, are tabulated in Table 1. The level of compressibility is commonly defined in terms of the turbulent Mach number  $M_t = \langle u_i u_i \rangle^{1/2}/c$ , where  $c$  is the mean speed of sound,  $u_i$  is the velocity fluctuation,  $\langle \cdot \rangle$  is the average computed across the entire domain and summation convention is used. For the simulations in this paper  $M_t \approx 0.3$  which represents a case where dilatational effects start becoming important [29].

For the rest of this section, we will refer the synchronous simulations using standard finite differences as SFD. The asynchronous simulations using AT schemes with random delays will be referred to as SAA and that with periodic delays will be referred to as CAA. We also have tenth-order compact schemes (C10) with third order RK scheme in time for one of the cases for comparison purposes to highlight that our finite difference simulations are comparable to the most well resolved simulations of compressible turbulence in literature [32,9,31,30].

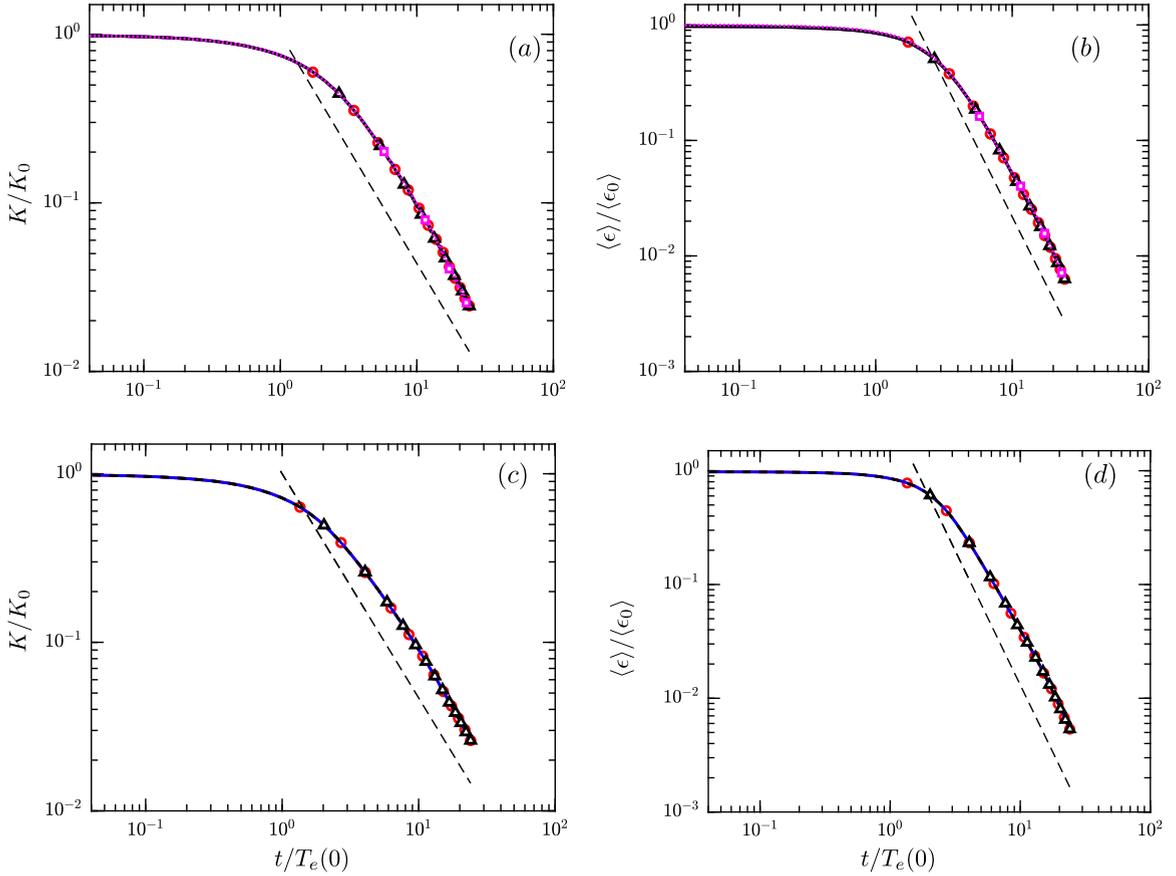
### 4.1. Decaying turbulence

#### 4.1.1. Low order statistics in physical space

It is important for any numerical scheme to accurately capture the large scale behavior of the system. An important and widely studied [33,34] large scale quantity in fluid turbulence is the mean turbulent kinetic energy per unit mass defined as,

$$K = \frac{1}{2} \langle \rho u_i u_i \rangle. \quad (31)$$

In the absence of energy input to the system,  $K$  decays in time as shown in Fig. 6(a), (c), where  $K$  is normalized by its initial value  $K_0$  and time is normalized by initial eddy turnover time,  $T_e(0) = \mathcal{L}/u_{rms}$ . After an initial transient, the decay obeys a power-law in time observed as a straight line on a log-log scale in Fig. 6(a), (c). The decay exponent is seen to be



**Fig. 6.** Evolution of space averaged turbulent kinetic energy normalized by the initial turbulent kinetic energy  $K_0$  (left) and evolution of space averaged dissipation rate normalized by the initial dissipation rate  $\epsilon_0$  (right) for  $Re_\lambda(0) \approx 100$  (a), (b) and  $Re_\lambda(0) \approx 145$  (c), (d). Different lines are: SFD (red-circle), CAA (black-triangle) and SAA (blue) both with  $L = 3$ . The black-dashed line corresponds to  $K/K_0 \propto (t/T_e(0))^{-1.4}$  in (a) and  $\langle \epsilon \rangle / \langle \epsilon_0 \rangle \propto (t/T_e(0))^{-2.4}$  in (b). Magenta line in (a), (b) is C10.

consistent with that found in the literature for similar conditions [34,35]. The excellent agreement between SFD, CAA and SAA in Fig. 6(a), (c) at all times shows that asynchronous implementations have accuracy comparable to SFD.

The rate at which kinetic energy is dissipated is given by  $\langle \epsilon \rangle = 2 \langle \sigma_{ij} S_{ij} \rangle$ . Because most of the contribution to dissipation comes from small scales (or high wavenumbers) it is therefore sensitive to how accurately high wavenumbers are resolved by the numerical methods. The decay of  $\langle \epsilon \rangle$  (normalized by its initial value) is shown in Fig. 6(b), (d) for SFD, CAA and SAA with no observable differences. Thus, we find that the asynchronous implementations are able to capture the evolution of low-order large and small scale quantities with accuracy comparable to the standard finite differences. Also shown in Fig. 6(a), (b) is the evolution obtained for C10 (magenta line), which is identical to the evolution obtained for both asynchronous and synchronous finite difference.

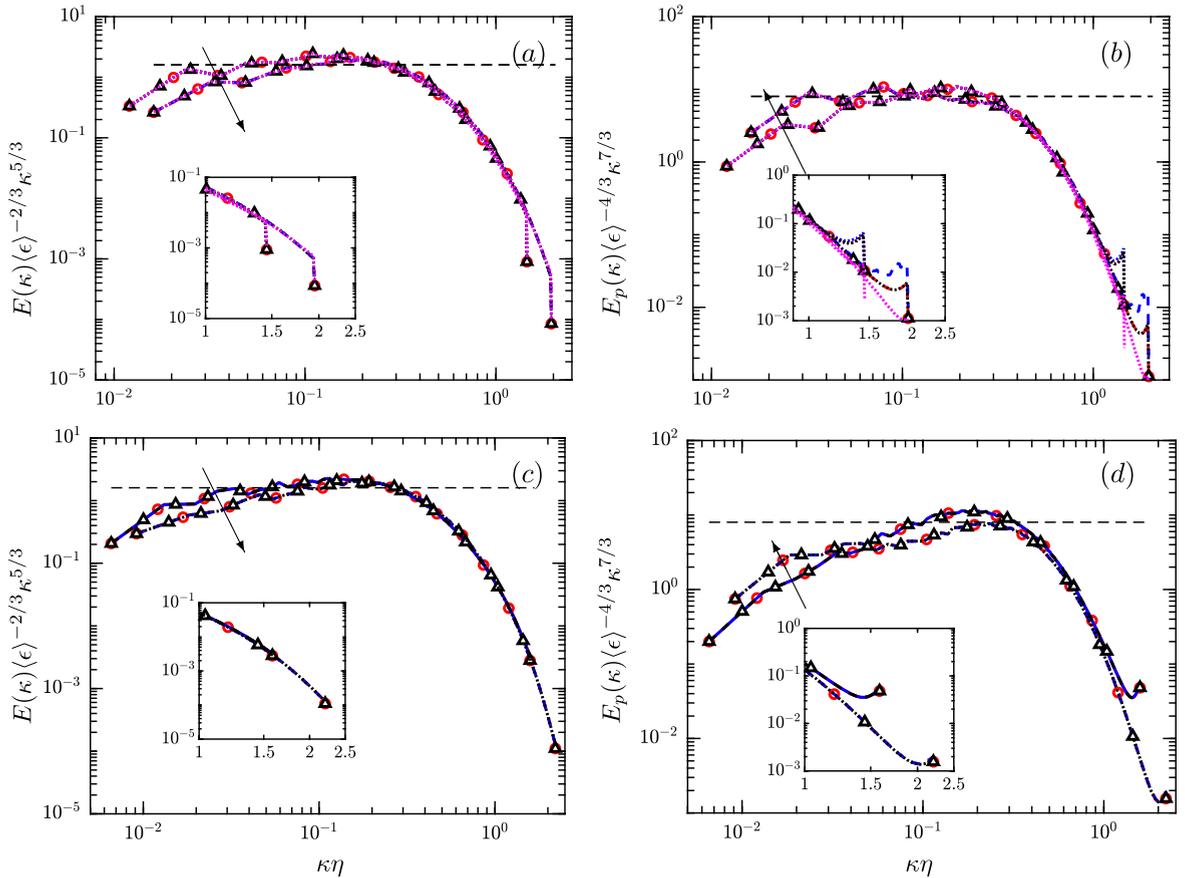
#### 4.1.2. Low order statistics in spectral space

Fluid turbulence comprises a wide range of interacting scales [36]. The energy distribution across these scales is characterized by the energy spectrum, which according to Kolmogorov self-similarity hypothesis (K41) [1] is given by,

$$E(\kappa) = C \langle \epsilon \rangle^{2/3} \kappa^{-5/3} f(\kappa \eta), \quad (32)$$

where  $C$  is the Kolmogorov constant,  $\kappa$  is the wavenumber and  $\eta = (v^3 / \langle \epsilon \rangle)^{1/4}$  is the Kolmogorov length scale [1] and  $f(\kappa \eta)$  is a universal function. This has been compared against simulations and experiments extensively and shown to be a good representation of the spectrum across different flows and Reynolds numbers for incompressible [37] and compressible flows [31,38,33] at low  $M_t$ . In the so-called inertial range ( $1/\mathcal{L} \ll k \ll 1/\eta$ ),  $f(k\eta) = 1$  and the classical 5/3 scaling for the energy spectrum [39,37,40] can be seen as a flat region in the compensated energy spectrum,

$$\frac{E(\kappa)}{\langle \epsilon \rangle^{2/3} \kappa^{-5/3}} = C, \quad (33)$$



**Fig. 7.** Compensated energy spectrum (left) and compensated pressure spectrum (right) for  $Re_\lambda(0) \approx 100$  (a), (b) and  $Re_\lambda(0) \approx 145$  (c), (d) at  $t/T_e(0) \approx 1$  and 4. Different lines are: SFD (red-circle), CAA (black-triangle) and SAA (blue) with  $L = 3$ . The arrow denotes increasing time. Magenta line in (a), (b) is C10.

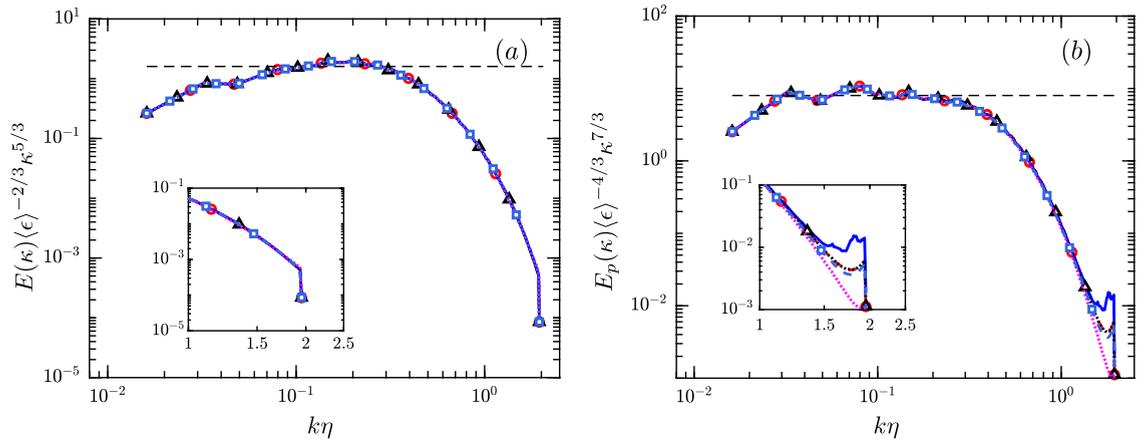
which becomes wider with an increase in Reynolds number. The height of this flat region gives the Kolmogorov constant which has been estimated to be  $C = 1.6$  from simulations and experiments in incompressible turbulence [41–43]. This value has been shown to be consistent for compressible simulations [31]. At high wavenumbers,  $f(k\eta)$  is a decaying exponential [44–47] which may retain a weak Reynolds number effect at very high wavenumbers [48].

In Fig. 7(a), (c) we show the compensated energy spectrum at  $t/T_e(0) \approx 1$  and 4 for  $Re_\lambda(0) \approx 100$  and 145 for SFD, CAA and SAA implementations. A plateau in this normalization corresponding to the inertial range can be seen at short times over a narrow range of scales. Because of the decrease in  $Re_\lambda$  with time due to the decay, the inertial range becomes less prominent at later times. We also see that the high wavenumbers are universal as expected from Eq. (32). Both SAA and CAA retain the universality at small scales and accurately capture the evolution of inertial and large scales. We see a virtually perfect agreement even at the smallest scales (inset in Fig. 7(a), (c)) for CAA as well as SAA with SFD. Moreover, the energy spectrum is also identical to the one obtained with C10 (magenta line in Fig. 7(a)) from some of the most well-resolved simulation of compressible turbulence [31,29].

Similar to the energy spectrum, K41 also predicts a scaling in the inertial range for pressure fluctuations [49,50] which reads,

$$E_p(\kappa) = C_p \langle \epsilon \rangle^{4/3} \kappa^{-7/3}. \quad (34)$$

The inertial range can be identified as the plateau in the compensated pressure spectrum plot, if  $Re_\lambda$  is high enough. Since  $M_t \approx 0.3$  for our simulation is fairly low, the pressure spectrum should be similar to the incompressible spectrum [31,50]. This is indeed observed in Fig. 7(b), (d) for  $Re_\lambda(0) \approx 100$  and 145 at  $t/T_e(0) \approx 1$  and 4 for the universal part of the spectrum. A horizontal dashed line at  $C_p = 8$  is also included for reference obtained from incompressible flows [50]. These spectra are consistent with those in the literature at similar conditions [31] with a collapse at the high wave-numbers similar to the energy spectrum. The data for CAA and SAA agree closely with that for SFD at both times for both  $Re_\lambda$ . However, for  $k\eta \geq 1.5$ , SAA spectrum has a small pileup at the high wavenumbers. This difference in the spectrum for CAA and SAA can be attributed to the difference in the nature of delays which for the former is deterministic and random for the latter. The randomness associated with SAA can lead to numerical errors that are absent in CAA and can cause a



**Fig. 8.** (a) Compensated energy spectrum and (b) compensated pressure spectrum for  $Re_\lambda(0) \approx 100$  at  $t/T_e(0) \approx 4$ . Faded dashed-blue line with squares is the sixth-order asynchronous scheme with random delays and solid blue line is fourth-order AT scheme with random delays. Rest of the lines are same as in Fig. 7. Insets zoom in on high wavenumbers.

small pileup of energy at the high wavenumbers as seen in the pressure spectrum in Fig. 7(b) for SAA. The differences in Fig. 7(b) are magnified because of the prefactor  $k^{7/3}$  but they are concentrated only in a few wavenumbers and represent an extremely small contribution to e.g., pressure variance. We have also performed simulations at higher  $M_t \approx 0.6$  and found that this small pileup disappears. Thus, this seems to be a low- $M_t$  effect which can be explained by noting that as  $M_t$  increases, there is stronger interaction between the so-called solenoidal and dilatational velocity components [51] which can help mix these already small perturbations at PE boundaries for SAA. For CAA, no pileup is observed at any  $M_t$ .

A general conclusion one can draw from both energy and pressure spectrum plots, is that the dynamics of the flow at the scales of interest is accurately captured despite asynchrony even though there are some very small deviations at the high wavenumbers in the pressure spectrum for SAA. Furthermore, we see from Fig. 7(b) that the pressure spectrum for SFD itself is not identical to the spectrum obtained for C10 at higher wavenumbers. Thus, it is not unexpected that asynchronous schemes present a different behavior at high wavenumbers. The errors in SAA, though already very small, can be mitigated if higher order schemes or higher resolution is used. As an example, in Fig. 8(b), the compensated pressure spectrum is shown for  $Re_\lambda(0) \approx 100$  at  $t/T_e(0) \approx 4$  using fourth and sixth order AT scheme (Appendix B) for SAA. While the SAA with fourth-order AT scheme (solid blue) peels off at  $\kappa\eta \approx 1.5$ , SAA with sixth-order AT scheme (faded-blue square), follows the SFD spectrum till the highest  $\kappa\eta$ .

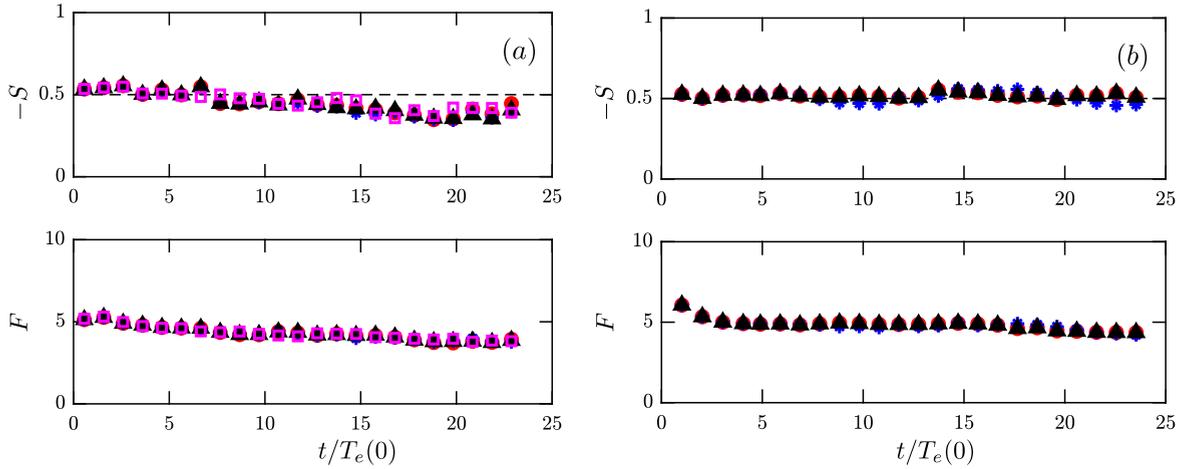
#### 4.1.3. Statistics of velocity gradients

An important feature of 3D turbulence is the generation of vortical motions, often quantified with the so-called enstrophy ( $\Omega = \langle \omega_i \omega_i \rangle$ , where  $\boldsymbol{\omega} = \nabla \times \mathbf{u}$  is the vorticity vector). A normalized metric for the production of enstrophy, which is also representative of the non-linear transfer of energy from large scales to small scales, is the skewness of the longitudinal velocity gradient,  $S = \langle (\partial u_1 / \partial x_1)^3 \rangle / (\langle (\partial u_1 / \partial x_1)^2 \rangle)^{3/2}$  [49,37,52]. The negative of the skewness ( $-S$ ) is constant at about  $\sim 0.5$  as long as the Reynolds number is not too low. This has been extensively documented in experiments and numerical simulations [53,37,54,40]. In Fig. 9(a), (c) we show the time evolution of  $-S$  for initial  $Re_\lambda$  of 100 and 145, respectively. We see that  $-S$  is close to 0.5 and this is consistent for SFD, CAA and SAA, with some small differences at later times. Despite odd-order moments being more sensitive to resolution [55] and susceptible to numerical errors, we see that the asynchronous algorithms capture skewness very well and close to the skewness computed using C10.

Another intrinsic characteristic of turbulent flows is the phenomena of intermittency which is a tendency to have localized events of fluctuations that are orders of magnitude larger than the mean [45,37,52,56–58]. These events add to the complexity of the turbulent flows, specifically at the smallest scales. One way to quantify this phenomenon is through the moments of velocity gradients as most of their contribution stems from the small scales and it is thus an excellent quantity to check small scale resolution. These moments transition from Gaussian to anomalous as Reynolds number increases [59–61]. In Fig. 9(b), (d) we show the normalized fourth-order moment or flatness ( $F = \langle (\partial u_1 / \partial x_1)^4 \rangle / (\langle (\partial u_1 / \partial x_1)^2 \rangle)^2$ ) of the longitudinal velocity gradient. The flatness is close to 6 [53,37] at initial times and tends to decrease because of decrease in Reynolds number for decaying turbulence. We see an excellent agreement between synchronous and both the asynchronous simulations with no observable differences from C10. Even though the computation of the gradient  $\partial u_1 / \partial x_1$  is directly affected by asynchrony, the higher order moments of the same exhibit trends similar to SFD and C10.

#### 4.1.4. Instantaneous enstrophy field

The average quantities discussed in the sections show good agreement between the asynchronous and synchronous simulations. A stricter test of accuracy would comprise the instantaneous flow fields which can potentially show some differences because of different truncation errors for different schemes in the computation of derivatives at the boundaries.



**Fig. 9.** Negative of skewness (top row) and flatness (bottom row) of the longitudinal velocity gradient vs. normalized time for (a)  $Re_\lambda(0) \approx 100$  and (b)  $Re_\lambda(0) \approx 145$ . Different symbols are: SFD (red-circle), CAA (black-triangle) and SAA (blue asterisk) with  $L = 3$ . The dashed black line indicates skewness of 0.5 and magenta squares in (a), (c) are C10.

As argued above, enstrophy is known to be sensitive to small scale resolution and is highly intermittent [55,62] and thus provides a stringent test of the numerical performance of schemes. In Fig. 10 we show the contours of the enstrophy normalized by its mean ( $\Omega/\langle\Omega\rangle$ ) in the  $yz$  plane at  $x = \pi$ . Qualitatively, all the large and small structures look identical for SFD, CAA and SAA. In particular, a concern with asynchronous schemes is the behavior close to the processor boundaries. If we closely look along these PE boundaries (faded lines in Fig. 10) there are no perceptible differences between enstrophy contours for SAA, CAA and SFD. Moreover, even complex structures spanning across multiple PE boundaries, for example, inside black circle in Fig. 10, is consistent for all the three cases. Besides some very small but not apparent localized differences in the intensity of enstrophy for SAA, the asynchronous algorithms accurately resolve the highly intermittent instantaneous enstrophy field. The instantaneous dissipation field (not shown here) exhibits similar behavior and is captured accurately.

#### 4.2. Forced turbulence

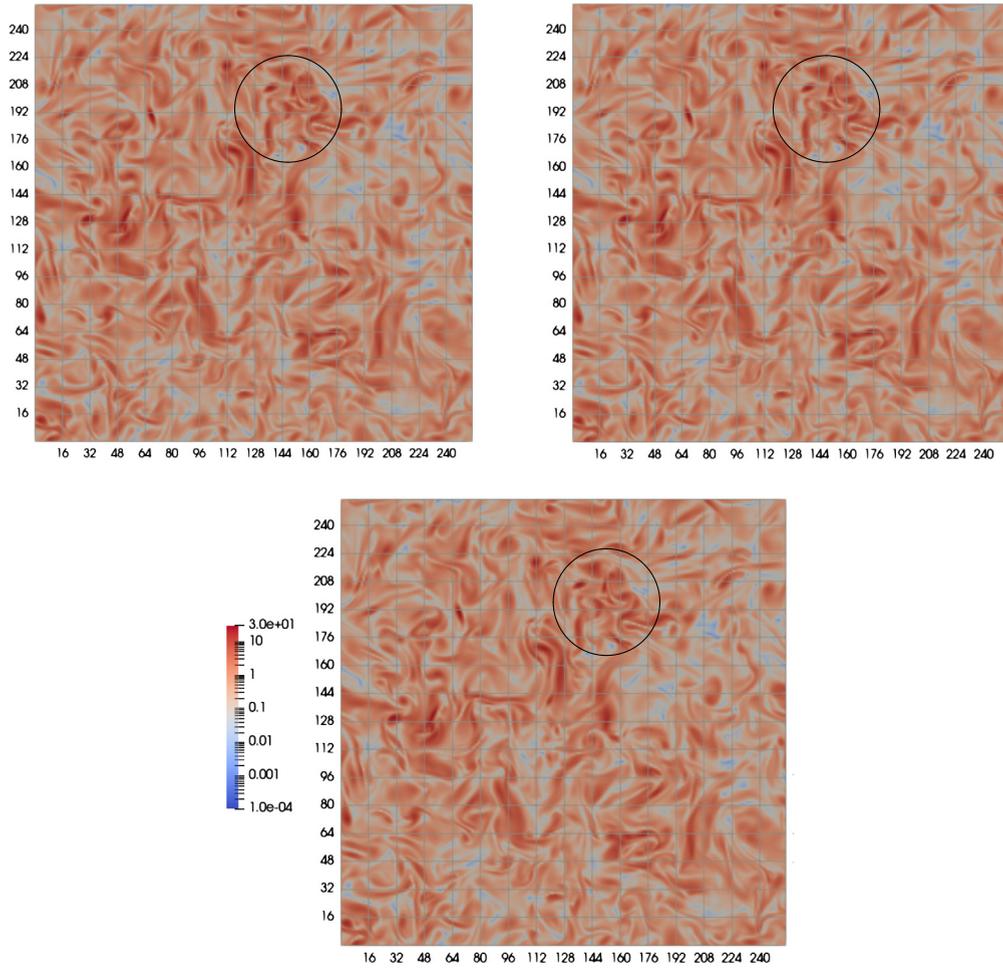
In the preceding section we focused on the DNS of decay of stationary state initial velocity field and observed a close agreement between the synchronous and asynchronous numerical simulations. In this section we discuss the effect of asynchrony on forced turbulence. Here, energy is injected at the large scales, or wavenumbers ( $\kappa$ ) in a spherical shell of radius  $\kappa_f$ , where  $\kappa \leq \kappa_f$ , ( $\kappa_f = 3$ ), through the term  $f$  in the momentum equation (Eq. (2)). The details of the stochastic forcing implemented can be found in [63] and have been extensively used in [8,31,29,64] for compressible turbulence. Through the non-linear interactions this injected energy cascades down to the inertial and small scales, where it is dissipated into internal energy by the action of viscosity. One can derive the evolution equation of the mean turbulent kinetic energy ( $K$ ) by multiplying Eq. (2) by  $u_i$  and taking the mean, which reads as

$$\frac{dK}{dt} = \langle p'\theta' \rangle - \langle \epsilon \rangle + \langle f_i u_i \rangle \quad (35)$$

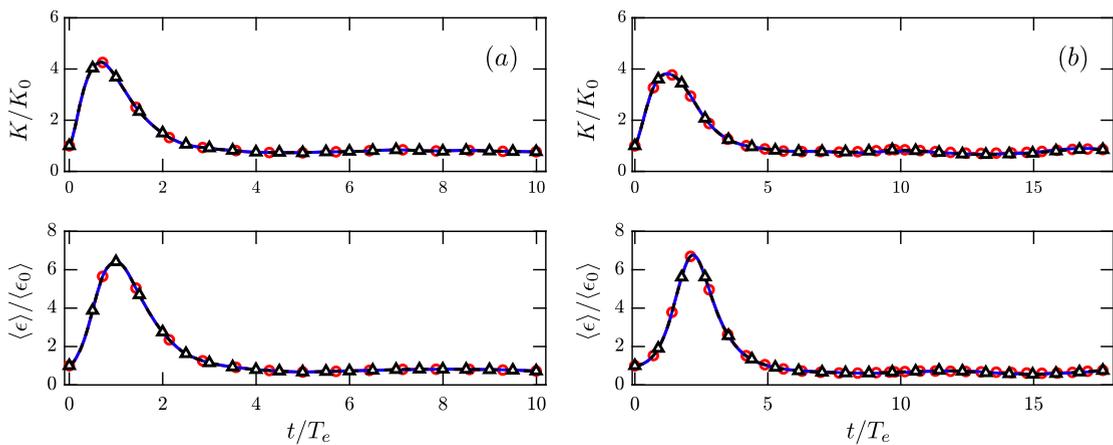
where  $\theta = \partial u_i / \partial x_i$  is the dilatation,  $\langle p'\theta' \rangle$  is the mean pressure-dilatation correlation and the mean dissipation  $\langle \epsilon \rangle$ . The external forcing  $f$  acts against the dissipative effect of viscosity to sustain turbulent fluctuations. We can also write the equation of the mean internal energy ( $\langle \rho e \rangle$ ) from Eq. (3) as

$$\frac{d\langle \rho e \rangle}{dt} = -\langle p'\theta' \rangle + \langle \epsilon \rangle. \quad (36)$$

The pressure-dilatation and viscous dissipation are responsible for the exchange between kinetic and internal energy. While the former is a bi-directional exchange depending upon the value of turbulent Mach number,  $M_t$  [31,29], the latter converts kinetic energy into internal energy irreversibly. Since no external sink is added to the energy equation, the internal energy of the system always increases. The time evolution of  $K$  and  $\langle \epsilon \rangle$ , normalized by their initial values, is plotted in Fig. 11. We can see that  $K$  increases initially, because of the input of energy due to forcing at large scales. Once the cascade develops and transfers energy to the smallest dissipative scales, the mean kinetic energy starts to decrease. At the same time, dissipation also increases initially, after an initial lag, until it reaches an equilibrium. At this point the rate of energy input is equal to rate of dissipation and a quasi-stationary state is reached [38]. In Fig. 11, this state is achieved at  $t/T_e \approx 5$  for  $Re_\lambda \approx 35$  (a), and  $t/T_e \approx 6$  for  $Re_\lambda \approx 100$  (b), where  $T_e$  is the average eddy turnover time. The average eddy turnover

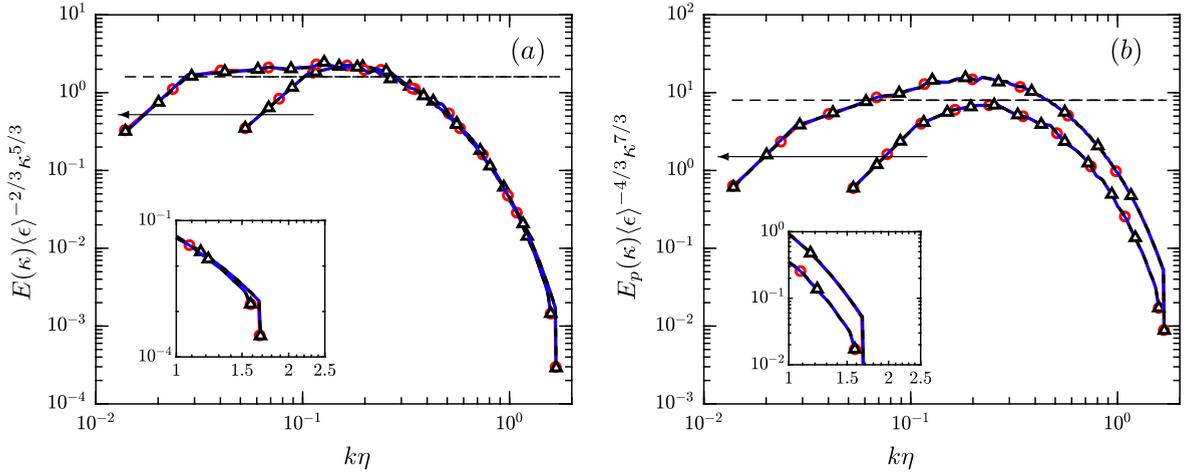


**Fig. 10.** Normalized instantaneous enstrophy ( $\Omega/\langle\Omega\rangle$ ) field at  $t/T_e(0) \approx 4$  for (a) SFD (b) CAA ( $L = 3$ ) and (c) SAA ( $L = 3$ ) in the  $yz$  plane at  $x = \pi$  for  $Re_\lambda(0) \approx 100$ . The faded lines represent processor boundaries.

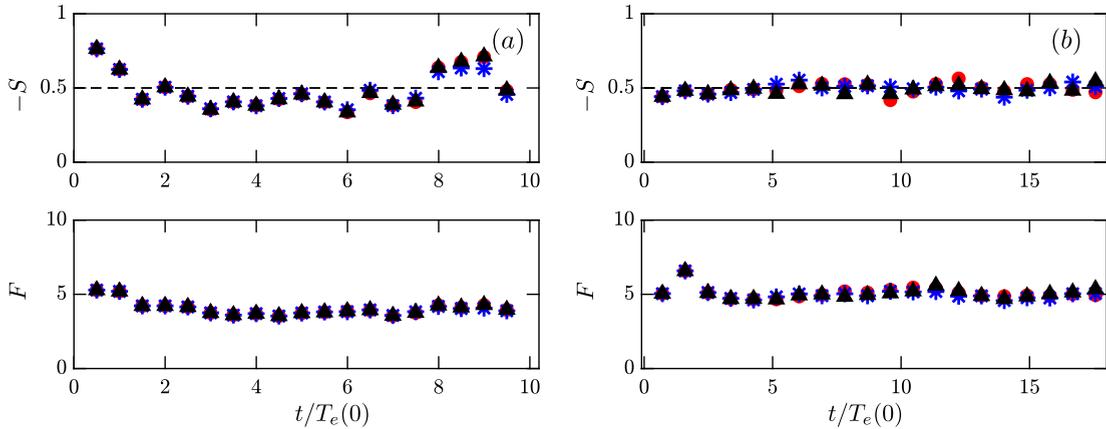


**Fig. 11.** Evolution of space averaged turbulent kinetic energy normalized by the initial turbulent kinetic energy  $K_0$  (top row) and space averaged dissipation rate normalized by the initial dissipation rate  $\epsilon_0$  (bottom row) for (a)  $Re_\lambda \approx 35$  and (b)  $Re_\lambda \approx 100$ . Different lines are: SFD (red-circle), CAA (black-triangle) and SAA (blue) with maximum allowed delay level of  $L = 3$ . Time is normalized by the average eddy turnover time ( $T_e$ ).

time is computed from the average taken at ten checkpoints from  $t/T_e \geq 5$  for  $Re_\lambda \approx 35$  and at fifteen checkpoints from  $t/T_e \geq 6$  for  $Re_\lambda \approx 100$ . The net increase in the total energy is, at this point, equal to the increase in the internal energy.



**Fig. 12.** (a) Compensated energy spectrum and (b) compensated pressure spectrum for  $Re_\lambda \approx 35$  and  $Re_\lambda \approx 100$ . Different lines are: SFD (red-circle), ATP (black-triangle) with  $L = 3$  and ATR (blue) with  $L = 2$ . Arrow indicated increasing  $Re_\lambda$ . Dashed line is Kolmogorov constant  $C = 1.6$  in (a) and  $C_p = 8$  in (b).

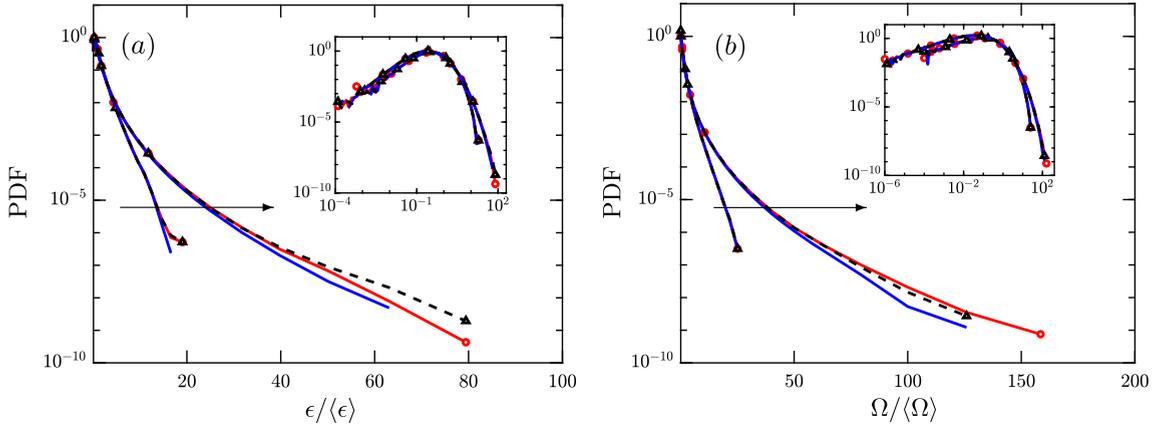


**Fig. 13.** Negative of skewness (top row) and flatness (bottom row) of the longitudinal velocity gradient vs. normalized time for (a)  $Re_\lambda \approx 35$  and (b)  $Re_\lambda \approx 100$ . Different symbols are: SFD (red-circle), ATP (black-triangle) with  $L = 2$  and ATR (blue asterisk) with  $L = 2$ . The dashed black line indicates skewness of 0.5.

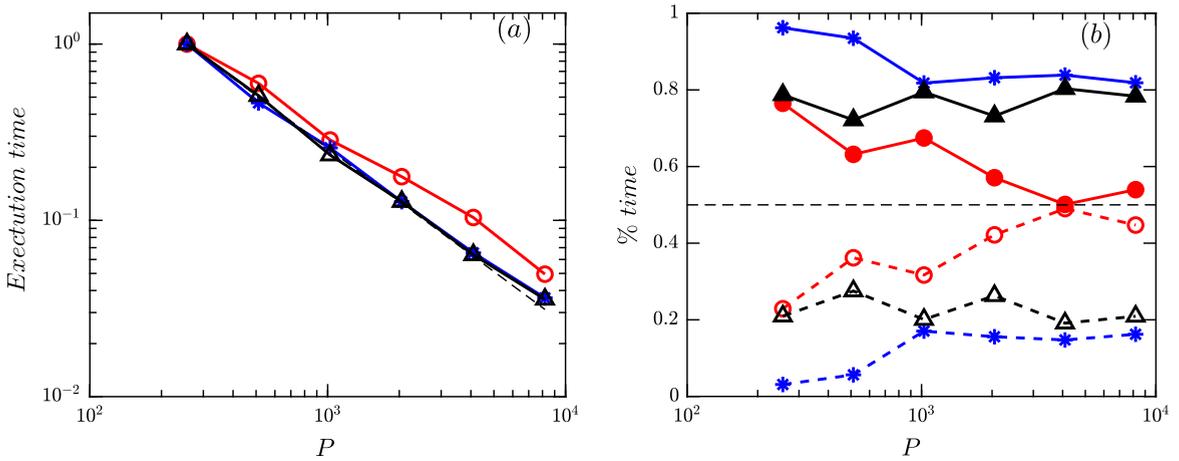
As in the case of decaying turbulence, we see a good agreement between the synchronous and asynchronous simulations in Fig. 11 for both high and low  $Re_\lambda$ .

We are also interested in the energy and pressure spectrum, which are plotted in Fig. 12. These spectra are the average taken for ten and fifteen checkpoints, respectively for  $Re_\lambda \approx 35$  and 100, after the quasi-stationary state is reached. These energy spectra are shown in Fig. 12(a) where we see that both CAA and SAA simulations are accurately resolved, with good collapse at all wavenumbers. For the pressure spectrum in Fig. 12(b), the CAA and SAA agree equally well with the SFD, unlike the decaying case where small errors were seen at the large wavenumbers for SAA. These spectra are also consistent with [31] at similar conditions. The higher order moments of the longitudinal velocity gradients are also plotted in Fig. 13. We see that  $-S$  fluctuates around 0.5 [61] and the values are fairly consistent for SAA, CAA and SFD. Even better agreement is seen for  $F$  in Fig. 13(c), (d), with value close to 6 for  $Re_\lambda \approx 100$  and smaller for  $Re_\lambda \approx 35$  [61].

Finally we look at the PDF of enstrophy density ( $\Omega$ ) and dissipation rate ( $\epsilon$ ) [55,65,62]. Both dissipation and enstrophy are crucial in the understanding of the small-scale motions [37] and are highly intermittent. Because of extreme events in  $\epsilon$  and  $\omega$ , the corresponding PDFs of the normalized quantities,  $\epsilon/\langle\epsilon\rangle$  and  $\omega/\langle\omega\rangle$ , are characterized by wide tails. The PDF of  $\epsilon/\langle\epsilon\rangle$  and  $\Omega/\langle\Omega\rangle$ , averaged over checkpoints as in case of averaged spectrum, are plotted in Fig. 14(a), (b). We can clearly see the tails of both the PDFs become wider as Reynolds number is increased from 38 to 100. This suggests that the propensity of events that are an order of magnitude more intense than the mean, increases with the Reynolds number [55,62]. Furthermore, we also observe that the tails for the PDF of  $\Omega/\langle\Omega\rangle$  in Fig. 14(b) are wider than the tails for PDF of  $\epsilon/\langle\epsilon\rangle$  in Fig. 14(a). This implies that enstrophy is more intermittent than dissipation and this has been consistently established in several past studies [53,66,37,67,55]. These features of the PDF are captured well by both the asynchronous



**Fig. 14.** PDF of (a) normalized dissipation rate ( $\epsilon/\langle\epsilon\rangle$ ) and (b) normalized enstrophy ( $\Omega/\langle\Omega\rangle$ ) in log-linear scale. The insets are the same PDFs in log-log scale. Different lines are: SFD (solid red with circle), CAA (dashed black with triangle) and SAA (solid blue with triangle) with  $L = 3$ . The arrow indicates increasing  $Re_\lambda$ .



**Fig. 15.** Strong scaling for  $N = 128$ . (a): Total execution time normalized by the execution time for  $P = 256$ . (b): Computation time and communication time as a percentage of the total execution time. Different lines are: SFD (red), CAA (black) and SAA (blue), dotted black in (a) is ideal scaling and in (b) is 50% of total time. Dashed lines with hollow symbols in (b) is communication time and solid lines with solid symbols is computation time.

algorithms with very small differences at the far tails. Thus, the AT schemes accurately resolve even the finest scales of turbulence including very highly intermittent events in dissipation and enstrophy.

### 4.3. Computational performance

The preceding sections demonstrated the ability of asynchronous algorithms in resolving important physical characteristics of turbulent flows including instantaneous field and high order statistics. Now we show that the asynchronous simulations are computationally more efficient than their synchronous counterpart. To study this we look at so-called strong and weak scaling of the solver. In the former the problem size remains fixed, while in the latter the computational work is kept constant. Ideally, for a fixed problem size, the computation time should decrease linearly on increasing the processor count. However, with increasing number of processors, the necessary communications and synchronizations increase the communication time until it eventually dominates the total execution time. This is essentially the communication bottleneck and is expected to be a major challenge to scalability [18,9,23,24]. In Fig. 15(a), we have plotted the total execution time for synchronous and asynchronous implementations for our compressible flow solver. These times are an average of five runs of 6000 steps each and a maximum allowed delay of  $L = 4$  for both SAA and CAA. For reference we have also plotted ideal scaling as a dashed black line. In Fig. 15(a) clear departures from ideal scaling are seen at  $P = 512$  for SFD. This, as is evident from Fig. 15(b), happens because the percentage of communication time (dashed red) grows with processor count ( $P$ ) until it becomes comparable to the computation time. On the other hand, both CAA and SAA (black and blue lines) are close to the ideal scaling in Fig. 15(a) for a much larger processor count of  $P = 8192$ . The improved scaling is attributed to the fact that only a small percentage ( $\sim 20\%$ ) of the overall time is spent on communications. This percentage (Fig. 15(b)) remains fairly constant on increasing the number of processors for the asynchronous implementations, whereas grows to larger than 50% for the synchronous case.

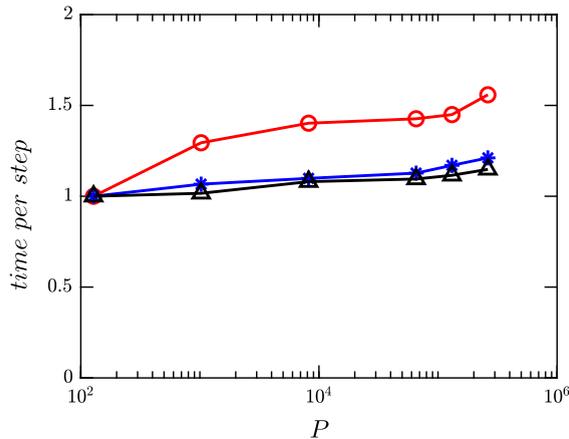


Fig. 16. Weak scaling: time per step for  $N^3/P = 2048$  normalized by time per step for  $P = 128$ . Different lines are: SFD (red), CAA (black) and SAA (blue).

Next we look at the weak scaling, where ideally because of fixed computational work, the time per step should remain constant on increasing the processor count. The time per step for a computational load of  $N^3/P = 2048$  is plotted in Fig. 16. For the synchronous case, this time per step scaling grows by a factor of 60% because of increase in communication and synchronization overheads at large core count ( $P = 262144$ ). This can only be expected to get worse at much higher levels of parallelism expected in exascale machines. On the other hand, the asynchronous algorithms show improved scaling, with a much smaller 21% increase in time per step for SAA and only 14% increase for CAA on increasing the number of processors from  $P = 128$  to  $P = 262144$ . This also implies that reduction in the overall volume of communication (CAA) at extreme scales provides more improvement in scaling than reducing forced synchronizations (SAA).

Both weak and strong scaling analysis lead us to the same conclusion that the asynchronous algorithms remove synchronization and communication overheads, leading to an effective overlap between communications and computations and, consequently, an improvement in scaling.

## 5. Conclusions and future implementation considerations

Numerical simulations of PDEs, governing complex natural and engineering phenomena, using standard numerical methods on parallel supercomputers, require PEs to communicate and synchronize frequently to ensure accuracy. This synchronization and communication cost and the resulting PE idling grows with increasing levels of parallelism and presents a major challenge to scalability to exascale computing. In order to mitigate this bottleneck, these constraints were relaxed at a mathematical level to derive the so-called Asynchrony-Tolerant (AT) of arbitrary order of accuracy in [24]. By allowing for asynchrony, these AT schemes can be used to allow computations to proceed in a PE without having to wait for updated values at the boundaries, thus removing synchronizations.

In this work we presented, first of a kind, asynchronous simulations of compressible turbulence using high-order Asynchrony-Tolerant (AT) schemes to study the effect of asynchrony on the physics of turbulence at different scales and on the computational performance of the solver. We show analytically that these schemes preserve the conservative property of standard finite differences up to an order higher than the order of the scheme. Stability analysis of these schemes shows that their stability limit is smaller than their synchronous counterpart. This reduction in the stability limit can be expressed in terms of the synchronous stability limit and a function of delays that gives a quantitative measure of the effect of delays. Numerical data also suggests that these two can be rearranged to obtain an *effective asynchronous* CFL that is essentially equal to the known synchronous stability limit and independent of  $L$ . We introduced two ways to allow for asynchrony, namely, communication avoiding and synchronization avoiding algorithms (CAA and SAA, respectively). While the former leads to deterministic delays with a uniform probability distribution, the latter leads to random delays with a machine specific delay distribution.

The aforementioned asynchronous algorithms are used for the simulation of decaying and solenoidally forced turbulence. Important low and high order statistics obtained for the asynchronous algorithms are compared with that for the standard synchronous finite differences (SFD) at the same resolution and order and also with high-order compact difference schemes (C10). We found excellent agreement between SFD and CAA for the time evolution of turbulent kinetic energy and dissipation for both decaying and forced turbulence, including the transients for the latter. The distribution of energy at different scales as shown by the velocity and the pressure spectrum is resolved by CAA with same level of accuracy as SFD and C10, even at the largest wavenumbers. Higher-order moments of longitudinal velocity gradient, including skewness and flatness, also showed excellent agreement between SFD, C10 and CAA. No observable differences are seen in the complex distribution of the contours of instantaneous enstrophy field. The PDF of highly intermittent quantities such as dissipation and enstrophy, that are also very sensitive to the accuracy of numerical schemes and small scale resolution, are also captured well by CAA, with some statistical differences at extreme tails.

For SAA as well, the evolution of turbulent kinetic energy and dissipation for decaying and forced turbulence (including transients) is in excellent agreement with SFD and C10. While no differences were seen for the energy spectrum, the pressure spectrum showed some small differences at high wavenumbers. However, these differences do not affect the dynamics of the scales of interest and, as we show, are easily mitigated if higher order AT schemes are used. Similar to CAA, the instantaneous enstrophy field, the flatness and skewness of longitudinal velocity gradient and the PDF of dissipation and enstrophy is shown to be in excellent agreement with the synchronous simulations.

Taken together, the results obtained for both CAA and SAA and their comparison with synchronous simulations (SFD and C10), clearly show that the physics of turbulence even at the finest scales is resolved accurately by the asynchronous algorithms, even though more than 50% of total gridpoints are affected directly by asynchrony.

We also presented the effect of asynchrony on computational performance. In particular, both strong and weak scaling results showed a near ideal scaling for the asynchronous algorithms and significant departures from the same for synchronous case. This improved scaling can be traced back to a significant reduction in communications (CAA) and synchronizations (SAA), resulting in an overall lower fraction of communication compared to computation for both CAA and SAA. We also observed that at very high processor count ( $P = 262144$ ), the reduction in overall volume of communications (CAA) is more effective in improving the scaling than relaxing explicit synchronization (SAA). This improvement in scaling is expected to be more consequential as we increase the problem size and processor count to levels anticipated on exascale machines.

In conclusion, asynchronous simulations can accurately resolve the physics at all scales and provide better parallel performance as problem size increases. Thus, asynchronous computing presents an effective alternative to standard computing approaches for simulation of turbulence and other complex phenomena at unprecedented levels of physical realism on the next generation exascale machines.

We close by mentioning some important future extensions of the current work. First, in this work CAA and SAA were presented as two separate algorithms, however, a combination of the two can also be used. This will potentially lead to further reduction in overheads associated with the communication and synchronization. Second, there are generalizations that can be introduced where the maximum delay level ( $L$ ) is different across different regions in the domain, depending upon the level of accuracy required. This does require critical analysis of load balancing to ensure that the processors synchronizing more often (smaller  $L$ ), have less computational work, so that the synchronization time in these processors does not affect the total execution time. Third, in order to further reduce the communication time in CAA, new AT schemes can be derived which use only one time level information from the buffer points and multiple delayed levels at internal points instead. The size of message in CAA for these new AT schemes is the same as that for the algorithms which communicate at every time step. The effect of these schemes on the performance and accuracy is part of our ongoing research. Lastly, from the performance analysis, we observed that while the asynchronous algorithms showed an improved scaling compared to the standard synchronous algorithm, the cache miss rate for the former was found to be higher. Though this miss rate reduces as the processor count is increased, optimization in implementation will help further push the limits of scaling and reduce the overall computation time. This will be discussed elsewhere.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

The authors acknowledge funding from The National Science Foundation (Grant No. 1439145), and XSEDE and Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing high-performance computing resources.

### Appendix A. Von-Neumann analysis of AT schemes

Stability analysis in the frequency domain, more commonly known as the *Von Neumann method*, has been widely used for linear problems with constant coefficients [27]. Since this method also requires all the points in the domain to use the same numerical scheme, we assume that each processor has only one grid point ( $P = N$ ) and the AT scheme uses same delay ( $\tilde{k}$ ) on both sides for all PEs. To proceed further, consider a diffusion equation discretized using a second-order AT scheme in space and forward Euler in time,

$$u_i^{n+1} = u_i^n + \frac{\alpha \Delta t}{\Delta x^2} \left( (\tilde{k} + 1)u_{i-1}^{n-\tilde{k}} - \tilde{k}u_{i-1}^{n-\tilde{k}-1} - 2u_i^n + (\tilde{k} + 1)u_{i+1}^{n-\tilde{k}} - \tilde{k}u_{i+1}^{n-\tilde{k}-1} \right), \quad (\text{A.1})$$

where  $\tilde{k}$  is the delay at both left and right boundary. Using a Fourier decomposition,  $u_i^n = v^n e^{Ii\phi}$  where  $I = \sqrt{-1}$  and  $\phi = \kappa \Delta x$ , we can simplify Eq. (A.1) as

$$v^{n+1} = r_d \left( v^{n-L}(L+1) - v^{n-L-1}L \right) e^{-I\phi} + (1 - 2r_d)v^n + r_d \left( v^{n-L}(L+1) - v^{n-L-1}L \right) e^{I\phi}, \quad (\text{A.2})$$

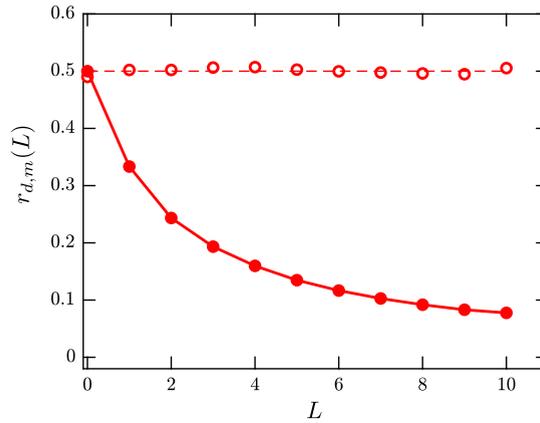


Fig. A.17. Variation of stability limit  $r_{d,m}(L)$  (solid) and  $r_{d,m}^a$  (hollow) with  $L$  for diffusion equation with delays on only one side.

where  $r_d = \alpha \Delta t / \Delta x^2$  is the diffusive CFL and delay is equal to the maximum allowed delay i.e.  $\tilde{k} = L$ . On taking the Z-transform of Eq. (A.2) with

$$v^n = v(z); \quad v^{n-L} = z^{-L} v(z) \tag{A.3}$$

we get a polynomial of order  $L + 2$  in  $z$  which reads as,

$$z^{L+2} - z^{L+1}(1 - 2r_d) - r_d(e^{-L\phi} - e^{L\phi})(L + 1)z - r_d(e^{-L\phi} - e^{L\phi})L = 0 \tag{A.4}$$

For stability we require the amplitude all the harmonics to not grow in time. This is possible only if all the roots of Eq. (A.4) are within a unit disc in the complex  $Z$ -plane and the root  $z = 1$  has multiplicity one [27]. The roots of the above equation can be computed numerically for different  $L$  and these can be used to obtain the maximum  $r_d$  such that all roots lie within a unit disc. This is the largest limit for which the numerical scheme is stable and as before, we denote it as  $r_{d,m}(L)$ . Instead of computing the roots, the sufficient conditions that guarantee that all the roots lie within the disc of radius one can also be computed using the *Schur-Cohn* criteria [68]. As an example, we use *Schur-Cohn* criteria to get these conditions for  $L = 1$ . Since for  $L = 1$ , Eq. (A.4) reduces to a cubic polynomial, we get a total of three conditions,

$$4r_d^2 \cos^2(\phi) - 1 < 0, \tag{A.5}$$

$$16r_d^4 \cos^4(\phi) - 4r_d^2(4r_d(rd + 1) + 3) \cos^2(\phi) + 1 > 0, \tag{A.6}$$

$$8r_d \sin^2\left(\frac{\phi}{2}\right) \left(-2r_d^2 + 2r_d(2r_d \cos(\phi) - r_d \cos(2\phi) + \cos(\phi)) + 1\right)^2 \left(3r_d^3 \cos(3\phi) + 3(3r_d^2 - 1)r_d \cos(\phi) + (r_d - 1)(2r_d^2 \cos(2\phi) + 2r_d^2 - 1)\right) > 0. \tag{A.7}$$

These equations have only two parameters which can be varied to determine  $r_{d,m}(L)$  that satisfies all the three conditions, for all values of  $\phi \in [0, \pi]$ . This gives us  $r_{d,m}(1) \approx 0.25$  which is consistent with the stability limit obtained from the matrix stability analysis presented in section 3. For larger values of  $r_d$ , at least one of the above conditions is violated and thus instabilities can be triggered. We can see that this value is less than the traditional stability limit for a synchronous scheme ( $r_{d,m}(0) = 0.5$ ). It is worthwhile to note that similar constraints can also be written for  $L = 0$  or the standard synchronous scheme which give us the expected limit of  $r_{d,m}(0) = 0.5$ . This idea can be extended to get the constraints for the stability for higher  $L$  as well. [68] discusses techniques to extend the *Schur-cohn* criteria to higher order polynomials efficiently.

The above analysis assumes a worst case scenario with two-sided delays at all points. A more practical scenario would be to have delays on either the left or the right PE boundary. Since the schemes are symmetric, the stability limit for an AT scheme with delays on left boundary, is also applicable to AT scheme with delays at the right boundary. Repeating the above procedure, but now considering delay  $\tilde{k} = L$  on the left boundary and  $\tilde{k} = 0$  on the right, gives us following three conditions for all the roots of the  $z$  equation to lie with a unit disc,

$$r_d^2 - 1 < 0, \tag{A.8}$$

$$1 - r_d^2(3 + 4r_d(1 + r_d)) + r_d(2r_d^2 + 4r_d^3) \cos(\phi) > 0, \tag{A.9}$$

$$8r_d(-1 + r_d^2) \sin^2\left(\frac{\phi}{2}\right) \left(-1 + r_d + 4r_d^2 + r_d^3 - 4r_d^4 + r_d(-1 + 2r_d)(1 + 2r_d)^2 \cos(\phi) + r_d^2(1 - 4r_d(1 + r_d)) \cos(2\phi) + r_d^3 \cos(3\phi)\right) > 0. \tag{A.10}$$

**Table B.2**  
Asynchrony-tolerant (AT) schemes for left and right boundary used in numerical simulations (in section 4) for first and second derivative.

(Derivative, order)	Boundary	Scheme
(2,4)	Left	$\frac{1}{2}(\tilde{k}^2 + 3\tilde{k} + 2) \left( -u_{i+2}^n + 16u_{i+1}^n - 30u_i^n + 16u_{i-1}^{n-\tilde{k}} - u_{i-2}^{n-\tilde{k}} \right) / 12\Delta x^2$ $- (\tilde{k}^2 + 2\tilde{k}) \left( -u_{i+2}^n + 16u_{i+1}^n - 30u_i^n + 16u_{i-1}^{n-\tilde{k}-1} - u_{i-2}^{n-\tilde{k}-1} \right) / 12\Delta x^2$ $+ \frac{1}{2}(\tilde{k}^2 + \tilde{k}) \left( -u_{i+2}^n + 16u_{i+1}^n - 30u_i^n + 16u_{i-1}^{n-\tilde{k}-2} - u_{i-2}^{n-\tilde{k}-2} \right) / 12\Delta x^2$
(2,4)	Right	$\frac{1}{2}(\tilde{k}^2 + 3\tilde{k} + 2) \left( -u_{i+2}^{n-\tilde{k}} + 16u_{i+1}^{n-\tilde{k}} - 30u_i^n + 16u_{i-1}^n - u_{i-2}^n \right) / 12\Delta x^2$ $- (\tilde{k}^2 + 2\tilde{k}) \left( -u_{i+2}^{n-\tilde{k}-1} + 16u_{i+1}^{n-\tilde{k}-1} - 30u_i^n + 16u_{i-1}^n - u_{i-2}^n \right) / 12\Delta x^2$ $+ \frac{1}{2}(\tilde{k}^2 + \tilde{k}) \left( -u_{i+2}^{n-\tilde{k}-2} + 16u_{i+1}^{n-\tilde{k}-2} - 30u_i^n + 16u_{i-1}^n - u_{i-2}^n \right) / 12\Delta x^2$
(1,4)	Left	$\frac{1}{2}(\tilde{k}^2 + 3\tilde{k} + 2) \left( -u_{i+2}^n + 8u_{i+1}^n - 8u_{i-1}^{n-\tilde{k}} + u_{i-2}^{n-\tilde{k}} \right) / 12\Delta x$ $- (\tilde{k}^2 + 2\tilde{k}) \left( -u_{i+2}^n + 8u_{i+1}^n - 8u_{i-1}^{n-\tilde{k}-1} + u_{i-2}^{n-\tilde{k}-1} \right) / 12\Delta x$ $+ \frac{1}{2}(\tilde{k}^2 + \tilde{k}) \left( -u_{i+2}^n + 8u_{i+1}^n - 8u_{i-1}^{n-\tilde{k}-2} + u_{i-2}^{n-\tilde{k}-2} \right) / 12\Delta x$
(1,4)	Right	$\frac{1}{2}(\tilde{k}^2 + 3\tilde{k} + 2) \left( -u_{i+2}^{n-\tilde{k}} + 8u_{i+1}^{n-\tilde{k}} - 8u_{i-1}^n + u_{i-2}^n \right) / 12\Delta x$ $- (\tilde{k}^2 + 2\tilde{k}) \left( -u_{i+2}^{n-\tilde{k}-1} + 8u_{i+1}^{n-\tilde{k}-1} - 8u_{i-1}^n + u_{i-2}^n \right) / 12\Delta x$ $+ \frac{1}{2}(\tilde{k}^2 + \tilde{k}) \left( -u_{i+2}^{n-\tilde{k}-2} + 8u_{i+1}^{n-\tilde{k}-2} - 8u_{i-1}^n + u_{i-2}^n \right) / 12\Delta x$
(2,6)	Left	$\frac{1}{6}(\tilde{k}^3 + 6\tilde{k}^2 + 11\tilde{k} + 6) \left( 2u_{i+3}^n - 27u_{i+2}^n + 270u_{i+1}^n - 490u_i^n + 270u_{i-1}^{n-\tilde{k}} - 27u_{i-2}^{n-\tilde{k}} + 2u_{i-3}^{n-\tilde{k}} \right) / 180\Delta x^2$ $- \frac{1}{2}(\tilde{k}^3 + 5\tilde{k}^2 + 6\tilde{k}) \left( 2u_{i+3}^n - 27u_{i+2}^n + 270u_{i+1}^n - 490u_i^n + 270u_{i-1}^{n-\tilde{k}-1} - 27u_{i-2}^{n-\tilde{k}-1} + 2u_{i-3}^{n-\tilde{k}-1} \right) / 180\Delta x^2$ $+ \frac{1}{2}(\tilde{k}^3 + 4\tilde{k}^2 + 3\tilde{k}) \left( 2u_{i+3}^n - 27u_{i+2}^n + 270u_{i+1}^n - 490u_i^n + 270u_{i-1}^{n-\tilde{k}-2} - 27u_{i-2}^{n-\tilde{k}-2} + 2u_{i-3}^{n-\tilde{k}-2} \right) / 180\Delta x^2$ $- \frac{1}{6}(\tilde{k}^3 + 3\tilde{k}^2 + 2\tilde{k}) \left( 2u_{i+3}^n - 27u_{i+2}^n + 270u_{i+1}^n - 490u_i^n + 270u_{i-1}^{n-\tilde{k}-3} - 27u_{i-2}^{n-\tilde{k}-3} + 2u_{i-3}^{n-\tilde{k}-3} \right) / 180\Delta x^2$
(2,6)	Right	$\frac{1}{6}(\tilde{k}^3 + 6\tilde{k}^2 + 11\tilde{k} + 6) \left( 2u_{i+3}^{n-\tilde{k}} - 27u_{i+2}^{n-\tilde{k}} + 270u_{i+1}^{n-\tilde{k}} - 490u_i^n + 270u_{i-1}^n - 27u_{i-2}^n + 2u_{i-3}^n \right) / 180\Delta x^2$ $- \frac{1}{2}(\tilde{k}^3 + 5\tilde{k}^2 + 6\tilde{k}) \left( 2u_{i+3}^{n-\tilde{k}-1} - 27u_{i+2}^{n-\tilde{k}-1} + 270u_{i+1}^{n-\tilde{k}-1} - 490u_i^n + 270u_{i-1}^n - 27u_{i-2}^n + 2u_{i-3}^n \right) / 180\Delta x^2$ $+ \frac{1}{2}(\tilde{k}^3 + 4\tilde{k}^2 + 3\tilde{k}) \left( 2u_{i+3}^{n-\tilde{k}-2} - 27u_{i+2}^{n-\tilde{k}-2} + 270u_{i+1}^{n-\tilde{k}-2} - 490u_i^n + 270u_{i-1}^n - 27u_{i-2}^n + 2u_{i-3}^n \right) / 180\Delta x^2$ $- \frac{1}{6}(\tilde{k}^3 + 3\tilde{k}^2 + 2\tilde{k}) \left( 2u_{i+3}^{n-\tilde{k}-3} - 27u_{i+2}^{n-\tilde{k}-3} + 270u_{i+1}^{n-\tilde{k}-3} - 490u_i^n + 270u_{i-1}^n - 27u_{i-2}^n + 2u_{i-3}^n \right) / 180\Delta x^2$
(1,6)	Left	$\frac{1}{6}(\tilde{k}^3 + 6\tilde{k}^2 + 11\tilde{k} + 6) \left( u_{i+3}^n - 9u_{i+2}^n + 45u_{i+1}^n - 45u_{i-1}^{n-\tilde{k}} + 9u_{i-2}^{n-\tilde{k}} - u_{i-3}^{n-\tilde{k}} \right) / 60\Delta x$ $- \frac{1}{2}(\tilde{k}^3 + 5\tilde{k}^2 + 6\tilde{k}) \left( u_{i+3}^n - 9u_{i+2}^n + 45u_{i+1}^n - 45u_{i-1}^{n-\tilde{k}-1} + 9u_{i-2}^{n-\tilde{k}-1} - u_{i-3}^{n-\tilde{k}-1} \right) / 60\Delta x$ $+ \frac{1}{2}(\tilde{k}^3 + 4\tilde{k}^2 + 3\tilde{k}) \left( u_{i+3}^n - 9u_{i+2}^n + 45u_{i+1}^n - 45u_{i-1}^{n-\tilde{k}-2} + 9u_{i-2}^{n-\tilde{k}-2} - u_{i-3}^{n-\tilde{k}-2} \right) / 60\Delta x$ $- \frac{1}{6}(\tilde{k}^3 + 3\tilde{k}^2 + 2\tilde{k}) \left( u_{i+3}^n - 9u_{i+2}^n + 45u_{i+1}^n - 45u_{i-1}^{n-\tilde{k}-3} + 9u_{i-2}^{n-\tilde{k}-3} - u_{i-3}^{n-\tilde{k}-3} \right) / 60\Delta x$
(1,6)	Right	$\frac{1}{6}(\tilde{k}^3 + 6\tilde{k}^2 + 11\tilde{k} + 6) \left( u_{i+3}^{n-\tilde{k}} - 9u_{i+2}^{n-\tilde{k}} + 45u_{i+1}^{n-\tilde{k}} - 45u_{i-1}^n + 9u_{i-2}^n - u_{i-3}^n \right) / 60\Delta x$ $- \frac{1}{2}(\tilde{k}^3 + 5\tilde{k}^2 + 6\tilde{k}) \left( u_{i+3}^{n-\tilde{k}-1} - 9u_{i+2}^{n-\tilde{k}-1} + 45u_{i+1}^{n-\tilde{k}-1} - 45u_{i-1}^n + 9u_{i-2}^n - u_{i-3}^n \right) / 60\Delta x$ $+ \frac{1}{2}(\tilde{k}^3 + 4\tilde{k}^2 + 3\tilde{k}) \left( u_{i+3}^{n-\tilde{k}-2} - 9u_{i+2}^{n-\tilde{k}-2} + 45u_{i+1}^{n-\tilde{k}-2} - 45u_{i-1}^n + 9u_{i-2}^n - u_{i-3}^n \right) / 60\Delta x$ $- \frac{1}{6}(\tilde{k}^3 + 3\tilde{k}^2 + 2\tilde{k}) \left( u_{i+3}^{n-\tilde{k}-3} - 9u_{i+2}^{n-\tilde{k}-3} + 45u_{i+1}^{n-\tilde{k}-3} - 45u_{i-1}^n + 9u_{i-2}^n - u_{i-3}^n \right) / 60\Delta x$

These new conditions are satisfied for  $r_{d,m}(1) \approx 0.33$  for all  $\phi \in [0, \pi]$ , which is larger than the previously computed bound of  $r_{d,m}(1) \approx 0.25$ . In Fig. A.17 we show the  $r_{d,m}(L)$  (solid circles) obtained for varying  $L$  for one-sided delays. Also plotted is  $r_{d,m}(0)/f(L)$  (solid line) defined in Eq. (29), where  $f(L) \approx 0.56L + 0.9$ , and  $r_{d,m}^0 = r_{d,m}(0) \times f(L)$  in hollow circles with a dashed line corresponding to the synchronous limit  $r_{d,m}(0) = 0.5$ .

We see that while  $r_{d,m}(L)$  decreases as delay increases, the *effective asynchronous CFL* ( $r_{d,m}^0$ ) is close to 0.5 for all  $L$  and supports the argument presented in section 3.

### Appendix B. Asynchrony-tolerant schemes

See Table B.2.

### References

- [1] A.N. Kolmogorov, Local structure of turbulence in an incompressible fluid for very large Reynolds numbers, Dokl. Akad. Nauk SSSR 30 (1941) 299–303.
- [2] P. Moin, K. Mahesh, Direct numerical simulation: a tool in turbulence research, Annu. Rev. Fluid Mech. 30 (1998) 539–578.
- [3] T. Ishihara, T. Gotoh, Y. Kaneda, Study of high-Reynolds number isotropic turbulence by direct numerical simulation, Annu. Rev. Fluid Mech. 41 (2009) 165–180.

- [4] V. Yakhot, K.R. Sreenivasan, Anomalous scaling of structure functions and dynamic constraints on turbulence simulations, *J. Stat. Phys.* 121 (2005) 823–841.
- [5] C. Canuto (Ed.), *Spectral Methods in Fluid Dynamics*, corr. 3rd print edition, Springer Series in Computational Physics, Springer-Verlag, Berlin; New York, 1988.
- [6] S.K. Lele, Compact finite difference schemes with spectral-like resolution, *J. Comput. Phys.* 103 (1992) 16–42.
- [7] K. Mahesh, S. Lee, S.K. Lele, P. Moin, The interaction of an isotropic field of acoustic waves with a shock wave, *J. Fluid Mech.* 300 (1995) 383–407.
- [8] M.R. Petersen, D. Livescu, Forcing for statistically stationary compressible isotropic turbulence, *Phys. Fluids* 22 (2010) 116101.
- [9] S. Jagannathan, D.A. Donzis, Massively parallel direct numerical simulations of forced compressible turbulence: a hybrid MPI/OpenMP approach, in: *Proceedings of the 1st Conference of the Extreme Science and Engineering Discovery Environment: Bridging from the eXtreme to the Campus and Beyond*, XSEDE '12, ACM, New York, NY, USA, 2012, pp. 23:1–23:8.
- [10] C.H. Chen, D.A. Donzis, Shock-turbulence interactions at high turbulence intensities, *J. Fluid Mech.* 870 (2019) 813–847.
- [11] S. Khurshid, D.A. Donzis, Decaying compressible turbulence with thermal non-equilibrium, *Phys. Fluids* 31 (2019) 015103.
- [12] A.W. Cook, W.H. Cabot, P.L. Williams, B.J. Miller, B.R.d. Supinski, R.K. Yates, M.L. Welcome, Tera-scalable algorithms for variable-density elliptic hydrodynamics with spectral accuracy, in: *Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*, SC'05, IEEE Computer Society, Washington, DC, USA, 2005, p. 60, No. 1.
- [13] D.A. Donzis, P.K. Yeung, D. Pekurovsky, Turbulence simulations on  $O(10^4)$  processors, in: *TeraGrid 2008 Conference*, 2008.
- [14] J.B. Chen, H.G. Im, Stretch effects on the burning velocity of turbulent premixed hydrogen/air flames, *Proc. Combust. Inst.* 28 (2000) 211–218.
- [15] J.H. Chen, A. Choudhary, B. de Supinski, M. DeVries, E.R. Hawkes, S. Klasky, W.K. Liao, K.L. Ma, J. Mellor-Crummey, N. Podhorszki, R. Sankaran, S. Shende, C.S. Yoo, Terascale direct numerical simulations of turbulent combustion using S3D, *Comput. Sci. Discov.* 2 (2009) 015001.
- [16] M. Lee, N. Malaya, R.D. Moser, Petascale direct numerical simulation of turbulent channel flow on up to 786K cores, in: *SC '13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, 2013, pp. 1–11.
- [17] Y. Minamoto, J.H. Chen, DNS of a turbulent lifted DME jet flame, *Combust. Flame* 169 (2016) 38–50.
- [18] J. Dongarra, P. Beckman, T. Moore, P. Aerts, G. Aloisio, J.-C. Andre, D. Barkai, J.-Y. Berthou, T. Boku, B. Braunschweig, F. Cappello, B. Chapman, Xuebin Chi, A. Choudhary, S. Dossanjh, T. Dunning, S. Fiore, A. Geist, B. Gropp, R. Harrison, M. Hereld, M. Heroux, A. Hoisie, K. Hotta, Zhong Jin, Y. Ishikawa, F. Johnson, S. Kale, R. Kenway, D. Keyes, B. Kramer, J. Labarta, A. Lichnewsky, T. Lippert, B. Lucas, B. Maccabe, S. Matsuoka, P. Messina, P. Michielse, B. Mohr, M.S. Mueller, W.E. Nagel, H. Nakashima, M.E. Papka, D. Reed, M. Sato, E. Seidel, J. Shalf, D. Skinner, M. Snir, T. Sterling, R. Stevens, F. Streit, B. Sugar, S. Sumimoto, W. Tang, J. Taylor, R. Thakur, A. Trefethen, M. Valero, A. van der Steen, J. Vetter, P. Williams, R. Wisniewski, K. Yelick, The international exascale software project roadmap, *Int. J. High Perform. Comput. Appl.* 25 (2011) 3–60.
- [19] D. Amitai, A. Averbuch, S. Itzikowitz, M. Israeli, Parallel adaptive and time-stabilizing schemes for constant-coefficient parabolic PDEs, *Comput. Math. Appl.* 24 (1992) 33–53.
- [20] D. Amitai, A. Averbuch, S. Itzikowitz, E. Turkel, Asynchronous and corrected-asynchronous finite difference solutions of PDES on MIMD multiprocessors, *Numer. Algorithms* 6 (1994) 275–296.
- [21] D. Mudigere, S.D. Sherlekar, S. Ansumali, Delayed difference scheme for large scale scientific simulations, *Phys. Rev. Lett.* 113 (2014) 218701.
- [22] A. Mittal, S. Girmaji, Proxy-equation paradigm: a strategy for massively parallel asynchronous computations, *Phys. Rev. E* 96 (2017) 033304.
- [23] D.A. Donzis, K. Aditya, Asynchronous finite-difference schemes for partial differential equations, *J. Comput. Phys.* 274 (2014) 370–392.
- [24] K. Aditya, D.A. Donzis, High-order asynchrony-tolerant finite difference schemes for partial differential equations, *J. Comput. Phys.* 350 (2017) 550–572.
- [25] J. Stoer, R. Bulirsch, *Introduction to Numerical Analysis*, vol. 12, Springer Science & Business Media, 2013.
- [26] A. Gruber, E.S. Richardson, K. Aditya, J.H. Chen, Direct numerical simulations of premixed and stratified flame propagation in turbulent channel flow, *Phys. Rev. Fluids* 3 (2018) 110507.
- [27] C. Hirsch, *Numerical Computation of Internal and External Flows*, vol. 1, Wiley, New York, 1994.
- [28] P.K. Yeung, K.R. Sreenivasan, S.B. Pope, Effects of finite spatial and temporal resolution in direct numerical simulations of incompressible isotropic turbulence, *Phys. Rev. Fluids* 3 (2018) 064603.
- [29] S. Jagannathan, D.A. Donzis, Reynolds and Mach number scaling in solenoidally-forced compressible turbulence using high-resolution direct numerical simulations, *J. Fluid Mech.* 789 (2016) 669–707.
- [30] J. Wang, T. Gotoh, T. Watanabe, Spectra and statistics in compressible isotropic turbulence, *Phys. Rev. Fluids* 2 (2017) 013403.
- [31] D.A. Donzis, S. Jagannathan, Fluctuations of thermodynamic variables in stationary compressible turbulence, *J. Fluid Mech.* 733 (2013) 221–244.
- [32] J. Wang, L.-P. Wang, Z. Xiao, Y. Shi, S. Chen, A hybrid numerical simulation of isotropic compressible turbulence, *J. Comput. Phys.* 229 (2010) 5257–5279.
- [33] S. Kida, S.A. Orszag, Energy and spectral dynamics in decaying compressible turbulence, *J. Sci. Comput.* 7 (1992) 1–34.
- [34] R. Samtaney, D.I. Pullin, B. Kosović, Direct numerical simulation of decaying compressible turbulence and shocklet statistics, *Phys. Fluids* 13 (2001) 1415–1430.
- [35] S.R. Yoffe, W.D. McComb, Onset criteria for freely decaying isotropic turbulence, *Phys. Rev. Fluids* 3 (2018) 104605.
- [36] S.B. Pope, *Turbulent Flows*, Cambridge University Press, 2000.
- [37] K.R. Sreenivasan, R.A. Antonia, The phenomenology of small-scale turbulence, *Annu. Rev. Fluid Mech.* 29 (1997) 435–472.
- [38] S. Kida, S.A. Orszag, Energy and spectral dynamics in forced compressible turbulence, *J. Sci. Comput.* 5 (1990) 85–125.
- [39] S.K. Lele, Compressibility effects on turbulence, *Annu. Rev. Fluid Mech.* 26 (1994) 211–254.
- [40] T. Ishihara, T. Gotoh, Y. Kaneda, Study of high-Reynolds number isotropic turbulence by direct numerical simulation, *Annu. Rev. Fluid Mech.* 41 (2009) 165–180.
- [41] K.R. Sreenivasan, On the universality of the Kolmogorov constant, *Phys. Fluids* 7 (1995) 2778–2784.
- [42] P.K. Yeung, Y. Zhou, Universality of the Kolmogorov constant in numerical simulations of turbulence, *Phys. Rev. E* 56 (1997) 1746–1752.
- [43] D.A. Donzis, K.R. Sreenivasan, The bottleneck effect and the Kolmogorov constant in isotropic turbulence, *J. Fluid Mech.* 657 (2010) 171–188.
- [44] R.H. Kraichnan, The structure of isotropic turbulence at very high Reynolds numbers, *J. Fluid Mech.* 5 (1959) 497.
- [45] R.H. Kraichnan, Intermittency in the very small scales of turbulence, *Phys. Fluids* 10 (1967) 2080.
- [46] L. Sirovich, L. Smith, V. Yakhot, Energy spectrum of homogeneous and isotropic turbulence in far dissipation range, *Phys. Rev. Lett.* 72 (1994) 344–347.
- [47] C. Foias, O. Manley, L. Sirovich, Empirical and Stokes eigenfunctions and the far-dissipative turbulent spectrum, *Phys. Fluids A, Fluid Dyn.* 2 (1990) 464–467.
- [48] S. Khurshid, D.A. Donzis, K.R. Sreenivasan, Energy spectrum in the dissipation range, *Phys. Rev. Fluids* 3 (2018) 082601.
- [49] A.S. Monin, A.M. Yaglom, *Statistical Fluid Mechanics*, vol. II, MIT, 1975.
- [50] T. Gotoh, D. Fukayama, Pressure spectrum in homogeneous turbulence, *Phys. Rev. Lett.* 86 (2001) 3775–3778.
- [51] D.A. Donzis, J. Panickacheril John, Universality and scaling in compressible turbulence, arXiv:1907.07871, 2019.
- [52] P. Davidson, *Turbulence: An Introduction for Scientists and Engineers*, vol. 2, Oxford University Press, 2015.
- [53] R.M. Kerr, Higher-order derivative correlations and the alignment of small-scale structures in isotropic numerical turbulence, *J. Fluid Mech.* 153 (1985) 31–58.
- [54] T. Gotoh, D. Fukayama, T. Nakano, Velocity field statistics in homogeneous steady turbulence obtained using a high-resolution direct numerical simulation, *Phys. Fluids* 14 (2002) 1065–1081.

- [55] D.A. Donzis, P.K. Yeung, K.R. Sreenivasan, Dissipation and enstrophy in isotropic turbulence: resolution effects and scaling in direct numerical simulations, *Phys. Fluids* 20 (2008) 045108.
- [56] D.A. Donzis, S. Jagannathan, On the relation between small-scale intermittency and shocks in turbulent flows, *Proc. IUTAM* 9 (2013) 3–15.
- [57] D. Porter, A. Pouquet, P. Woodward, Intermittency in compressible flows, in: R. Moreau, U. Frisch (Eds.), *Advances in Turbulence VII*, vol. 46, Springer Netherlands, Dordrecht, 1998, pp. 255–258.
- [58] L. Pan, P. Padoan, A.G. Kritsuk, Dissipative structures in supersonic turbulence, *Phys. Rev. Lett.* 102 (2009) 034501.
- [59] V. Yakhot, D.A. Donzis, Anomalous exponents in strong turbulence, *Phys. D: Nonlinear Phenom.* 384–385 (2018) 12–17.
- [60] V. Yakhot, D. Donzis, Emergence of multiscaling in a random-force stirred fluid, *Phys. Rev. Lett.* 119 (2017) 044501.
- [61] J. Schumacher, J.D. Scheel, D. Krasnov, D.A. Donzis, V. Yakhot, K.R. Sreenivasan, Small-scale universality in fluid turbulence, *Proc. Natl. Acad. Sci.* 111 (2014) 10961–10965.
- [62] P.K. Yeung, X.M. Zhai, K.R. Sreenivasan, Extreme events in computational turbulence, *Proc. Natl. Acad. Sci. USA* 112 (2015) 12633–12638.
- [63] V. Eswaran, S. Pope, An examination of forcing in direct numerical simulations of turbulence, *Comput. Fluids* 16 (1988) 257–278.
- [64] D.A. Donzis, A.F. Maqui, Statistically steady states of forced isotropic turbulence in thermal equilibrium and non-equilibrium, *J. Fluid Mech.* 797 (2016) 181–200.
- [65] P.K. Yeung, D.A. Donzis, K.R. Sreenivasan, Dissipation, enstrophy and pressure statistics in turbulence simulations at high Reynolds numbers, *J. Fluid Mech.* 700 (2012) 5–15.
- [66] E.D. Siggia, Numerical study of small-scale intermittency in three-dimensional turbulence, *J. Fluid Mech.* 107 (1981) 375.
- [67] T. Zhou, R.A. Antonia, Reynolds number dependence of the small-scale structure of grid turbulence, *J. Fluid Mech.* 406 (2000) 81–107.
- [68] J. Miller, On the location of zeros of certain classes of polynomials with applications to numerical analysis, *J. Inst. Math. Appl.* 8 (1971).