



Delft University of Technology

## **Fake it till you make it**

### **Synthetic turbulence to achieve swift converged turbulence statistics in a pressure-driven channel flow**

Patil, Akshay; García-Sánchez, Clara

#### **DOI**

[10.1016/j.compfluid.2025.106733](https://doi.org/10.1016/j.compfluid.2025.106733)

#### **Publication date**

2025

#### **Document Version**

Final published version

#### **Published in**

Computers and Fluids

#### **Citation (APA)**

Patil, A., & García-Sánchez, C. (2025). Fake it till you make it: Synthetic turbulence to achieve swift converged turbulence statistics in a pressure-driven channel flow. *Computers and Fluids*, 301, Article 106733. <https://doi.org/10.1016/j.compfluid.2025.106733>

#### **Important note**


To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

#### **Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

#### **Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.



# Fake it till you make it: Synthetic turbulence to achieve swift converged turbulence statistics in a pressure-driven channel flow

Akshay Patil<sup>1</sup>\*, Clara García-Sánchez<sup>1</sup>

<sup>1</sup>3D Geoinformation Research Group, Delft University of Technology, Delft, 2628BL, Zuid Holland, The Netherlands

## ARTICLE INFO

Dataset link: <https://github.com/AkshayPatil1994/Synthetic-Eddy-Method-KCX2013>, [https://github.com/AkshayPatil1994/transition\\_data](https://github.com/AkshayPatil1994/transition_data)

### Keywords:

Direct numerical simulations  
Turbulence  
Transition  
Channel Flow  
CaNS

## ABSTRACT

In this study, we introduced a simple yet innovative application: the isotropic synthetic turbulence field generator (iSTFG), based on the synthetic turbulent inflow generator. The iSTFG leverages the homogeneity in the streamwise direction for channel flows and triggers turbulence to achieve statistically stationary flow conditions faster than standard community-used strategies. We compare this new method with two other well-established methods: linear and log-law profiles superposed with random noise and descending counter-rotating vortices. We find that iSTFG provides a computationally cheap and effective way to reduce simulation spin-up costs/time/emissions to achieve statistically stationary flow conditions when a precursor turbulent initial condition is unavailable. At a one-time cost between 1-10 Central Processing Unit (CPU) hour(s) to generate the synthetic turbulent initial condition based on the target friction Reynolds numbers (1 CPU hour -  $Re_\tau = 500$ , 7 CPU hours -  $Re_\tau = 2000$ ), the flow achieves statistically stationary turbulent flow (SSTF) state within three eddy turnovers for all the parameters of interest in wall-bounded pressure-driven channel flow simulations when compared to other alternatives that can take more than ten eddy turnovers resulting in substantial savings in the computational cost. We also demonstrate that the transition and convergence to an SSTF state using conventional methods are sensitive to the computational domain size, while iSTFG is agnostic to the domain size. Furthermore, we explored the sensitivity of the iSTFG method to the non-dimensional integral length scale parameter and mismatch in reference and target input data to find iSTFG robust in such scenarios.

## 1. Introduction

Turbulent channel flows have significantly improved our understanding of wall-bounded turbulence as detailed in the extensive body of literature [1–8, to list a few]. While such studies provide a fundamentally novel insight into flow physics, the computational resources required to simulate such flows can become increasingly demanding as the flow Reynolds number increases. Recent estimates by Horwitz [9] suggest that a turbulent channel flow simulated by Vela-Martín et al. [10] using 512 Graphics Processing Units (GPU) can use up to  $5.98 \times 10^5$  kilo-Watt-hours (kWh) of energy and emit an equivalent of 6894 kg of CO<sub>2</sub> when running over a course of 2734 wall-clock hours. The simulation by Vela-Martín et al. [10] solved for the three-dimensional flow field at a friction Reynolds number ( $Re_\tau \equiv u_\tau H/\nu$ ) of 5303, where  $u_\tau$  is the friction velocity,  $H$  is the channel half-height, and  $\nu$  is the kinematic viscosity of the fluid. As illustrated by this simple estimate, this can pose a significant hurdle when statistically stationary flow conditions are to be achieved, which requires long simulation spin-up times before flow parameters can be averaged to gain generalisable insights.

Without appropriate initial conditions to begin the channel flow simulations, most simulation frameworks generate three-dimensional flow fields initialised using simple analytical profiles for the streamwise velocity components superposed with white noise [11,12]. For established datasets, such as the Johns Hopkins turbulence data set maintained by Graham et al. [13] also mentions a similar simulation spin-up where the flow is run with a constant bulk velocity forcing, i.e., momentum forcing, similar to Nelson and Fringer [11] to obtain the initial conditions. In some studies, a pair of counter-rotating vortices [14] is added to the linear-log-law profile to trigger the transition to turbulence and accelerate momentum mixing in the vertical direction, thus reducing the CPU time spent arriving at statistically stationary flow conditions [12]. However, as discussed by Nelson and Fringer [11] and Costa [12], the simulations require over 10 turnover ( $T_e \equiv H/u_\tau$ ) periods to reach stationary flow conditions. This computational cost can scale to a drastically large amount with increasing flow Reynolds number, which can be otherwise spent on collecting valuable statistics to support the requisite inferences. To bridge this

\* Corresponding author.

E-mail address: [a.l.patil@tudelft.nl](mailto:a.l.patil@tudelft.nl) (A. Patil).

<https://doi.org/10.1016/j.compfluid.2025.106733>

Received 15 January 2025; Received in revised form 4 May 2025; Accepted 15 June 2025

Available online 29 July 2025

0045-7930/© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

gap, we developed a scalable isotropic synthetic turbulence field generator (iSTFG) that can synthesise initial conditions for arbitrary channel flows with pre-existing user input data. In this study, we used the iSTFG initial conditions for two  $Re_\tau$  in a pressure-driven channel flow to understand how the spin-up time compares with existing methods to achieve stationary flow conditions for prototypical channel domain sizes. Additionally, we also compared the impact on convergence of flow statistics as a function of the computational domain size, model parameter, i.e., the isotropic integral length scale, and input data mismatch. The focus was to understand the effect of time-marching a given initial condition without any additional forcing apart from the constant pressure gradient that drives the flow. The following sections discuss the governing equations and numerical methods used to solve the flow equations and generate the initial conditions. This is followed by a detailed discussion of the results obtained from the various methods by comparing the quality of the statistics obtained against standard datasets and a brief discussion on the limitations of iSTFG, ending with some concluding remarks.

## 2. Problem formulation

In this study, we consider the non-dimensional incompressible Navier–Stokes momentum equations given by

$$\partial_t^* u_i^* + \partial_j^* u_j^* u_i^* = -\partial_i^* p^* + Re_\tau^{-1} \partial_j^* \partial_j^* u_i^* + \Pi_c \delta_{i1}, \quad (1)$$

subject to the incompressible continuity equation given by

$$\partial_i^* u_i^* = 0, \quad (2)$$

where  $t$  is time,  $u_i$  is the velocity vector,  $p$  is pressure,  $Re_\tau \equiv u_\tau H/\nu$ ,  $\Pi_c$  is the driving pressure gradient, and  $\delta_{ij}$  is the Kronecker delta function. The  $(\cdot)^*$  notation represents the non-dimensional parameters obtained using  $u_\tau$  and  $H$ . Using this choice of non-dimensionalisation, the driving pressure gradient is exactly unity. The governing equations are numerically integrated using the open-source massively parallel CaNS solver developed by Costa [12]. CaNS solves the governing equations using a second-order accurate spatial discretisation and a third-order accurate temporal discretisation using the low-storage Runge–Kutta 3-step method using the fractional step algorithm [15]. The flow variables are arranged on a staggered grid where scalars are placed at the cell centre while vector components are located at the cell faces [16]. CaNS has been extensively validated for channel flow simulations, and further details can be found in Costa [12] and will not be discussed for brevity.

For the medium-size domain, the channel has dimensions  $L_{x1} \times L_{x2} \times L_{x3} \equiv 4\pi H \times 2\pi H \times H$ , where  $x_i$  corresponds to the coordinate axes in streamwise, spanwise, and vertical directions, respectively. The flow is driven by a constant pressure gradient  $\Pi_c = 1.0$  subject to periodic boundary conditions in the streamwise and spanwise directions, and no-slip boundary condition at  $x_3 = 0$  and free slip boundary where  $\partial_3 u_1 = \partial_3 u_2 = 0$  and  $u_3 = 0$  at  $x_3 = H$ . The flow field is initialised using three different initial conditions: namely, inverse linear profile superposed with white noise and a pair of counter-rotating vortices (hereafter termed linear profile), linear-log-law profile superposed with white noise and a pair of counter-rotating vortices (hereafter termed log profile), and synthetically generated three-dimensional flow field using the iSTFG based on the work by Kim et al. [17] (hereafter termed synthetic profile). The linear profile initial condition is inspired by Nelson and Fringer [11], and is given by

$$u_i = 2U_o \left(1.0 - \frac{x_3}{H}\right) \delta_{i1} + \mathcal{U}(-\alpha U_b, \alpha U_b), \quad (3)$$

where  $\mathcal{U}$  is sampled from a uniform random distribution,  $\alpha$  is the white noise amplitude parameter set to a value of 0.7 as recommended by Nelson and Fringer [11] for all cases. The scaled bulk velocity for the linear profile is given by Nelson and Fringer [11]

$$U_o = \frac{u_\tau}{\kappa} \left[ \log \left( \frac{H}{z_o} \right) + \frac{z_o}{H} - 1 \right], \quad (4)$$

where  $z_o = \nu/(9u_\tau)$ . An inverse linear profile is used to effectively trigger the transition to turbulence as the shear stress at the bottom wall is large during the first eddy turnover. The plane-averaged shear stress is known a-priori for the analytical cases and is presented in detail in Appendix A.

The log velocity profile is only applied to the streamwise velocity component given by

$$u_1 = x_3 + \mathcal{U}(-\alpha U_b, \alpha U_b) \quad \forall \frac{x_3 u_\tau}{\nu} \leq 11.6, \quad (5)$$

$$= \left( \frac{u_\tau}{\kappa} \log \left( \frac{x_3 u_\tau}{\nu} \right) + 5.5 \right) + \mathcal{U}(-\alpha U_b, \alpha U_b) \quad \forall \frac{x_3 u_\tau}{\nu} > 11.6,$$

while the pair of counter-rotating vortices is initialised by prescribing the spanwise and vertical velocity components following Henningson and Kim [14]. The primary utility of adding this descending pair of vortices is to trigger a transition for the analytical profiles. As detailed in Nelson and Fringer [11], the linear and the log profiles may not always transition to a turbulent flow state. However, adding multiple pairs of vortices may not always lead to better performance, as it can introduce larger TKE levels during the initial transient that can take a long time to dissipate as the flow evolves, thus requiring a relatively longer spin-up. Consequently, for all the cases discussed in this work that make use of the pair of vortices, only a single pair was introduced. The bulk dimensional velocity  $U_b$  is derived from Pope [18] with an additional scaling factor of 0.5 to avoid overshoot given by

$$U_b = 0.5 \left[ \frac{\nu}{H} \right] \left[ \frac{Re_\tau}{0.09} \right]^{0.88}. \quad (6)$$

The choice for the two methods detailed in Eqs. (3) and (5) is motivated by Nelson and Fringer [11]; Costa [12].

While a wide range of synthetic turbulent inflow generators exists, as detailed in the extensive review by Wu [19], in this work, we focus on a relatively simple divergence-free method first proposed by Kim et al. [17] as detailed below. Much of the text below is based on the work by Kim et al. [17] and is repeated for reproducibility of the iSTFG code developed in this paper, with details on the minor differences in the implementation. The velocity field  $u_i$  is given by

$$u_i = \bar{U}_i + a_{ij} u_{*,j}, \quad (7)$$

where  $\bar{U}_i$  is the mean velocity profile known a-priori,  $a_{ij}$  is the amplitude tensor, and  $u_{*,j}$  is the unscaled fluctuations with zero mean, no correlation, and having unit variance. The amplitude tensor is obtained using the Cholesky decomposition of the Reynolds stress tensor ( $R_{ij}$ ) and has a form given by Lund et al. [20]


$$a_{ij} = \begin{pmatrix} \sqrt{R_{11}} & 0.0 & 0.0 \\ \frac{R_{21}}{a_{11}} & \sqrt{R_{22} - a_{21}^2} & 0.0 \\ \frac{R_{31}}{a_{11}} & \left( \frac{R_{32} - a_{21}a_{31}}{a_{22}} \right) & \sqrt{R_{33} - a_{31}^2 - a_{32}^2} \end{pmatrix}. \quad (8)$$

To generate the spatially correlated signal, a scalar field ( $\psi_m$ ) is generated using a digital filter method given by

$$\psi_m = \sum_{j=-N}^N b_j r_{m+j}, \quad (9)$$

where  $b_j$  is the model constant,  $m$  and  $j$  are position indices,  $N = 2n$ ,  $n = I_l/\Delta x_i$ ,  $I_l$  is the integral length scale, and  $\Delta x_i$  is the grid spacing in the coordinate directions. Unlike the work by Kim et al. [17], in this version of the iSTFG, the integral length scale is assumed to be isotropic in all directions for implementation simplicity. Here,  $\psi_m$  constitutes a one-dimensional array with zero mean, unit variance, and a spatially correlated signal. The model constant  $b_j$  is given by

$$b_j = \frac{b'_j}{\sqrt{\sum_{l=-N}^N b_l'^2}}, \quad (10)$$



**Fig. 1.** Snapshot of the initial condition generated using the iSTFG at various  $x_3^+$  locations above the wall for  $Re_\tau = 500$ . The rows from top to bottom correspond to the streamwise, spanwise, and vertical velocities, respectively. Two white lines at the bottom left corner of each panel provide a reference length scale of 400 wall units in the vertical and horizontal directions. The velocity marked using the colourmaps is non-dimensionalised using the friction velocity.

and  $b'_j = \exp\left(-\frac{\pi|j|}{2n}\right)$ . In most cases, the flow problem will be solved in three dimensions; thus, a two-dimensional extension of Eq. (9) can be formulated as

$$\psi_{m,l} = \sum_{j=-N}^N \sum_{k=-N}^N b_j b_k r_{m+j,l+k}. \quad (11)$$

Using the spatially correlated two-dimensional data as detailed in Eq. (11), the temporal correlations are implemented through the specification of  $u_{*,i}$  given by

$$u_{*,i}(t + \Delta t) = u_{*,i}(t) \exp\left(-\frac{C_{XC}\Delta t}{T}\right) + \psi_i(t) \left[1 - \exp\left(-\frac{2C_{XC}\Delta t}{T}\right)\right]^{\frac{1}{2}}, \quad (12)$$

where  $T = I_l/U_c$  is the Lagrangian time scale,  $\Delta t$  is the time step, and  $U_c$  is the convective velocity defined in this case as the vertically integrated bulk velocity ( $U_c = H^{-1} \int_{x_3'=0}^{x_3'=H} \bar{U}_i dx_3'$ ). A divergence-free form of the velocity field can be supplied by a constant correction of the mass flux for the synthetically generated velocity  $u_i$ .

For channel flows, the streamwise homogeneous direction allows one to exchange the spatial coordinate ( $x_1$ ) and time ( $t$ ) provided the proper convective velocity is applied based on the frozen turbulence hypothesis [21]. This work uses the vertically integrated mean input velocity as the convective velocity ( $U_c$ ). While the divergence-free form of the synthetic flow field is preferred, the first pressure-Poisson solution is sufficient to obtain such a divergence-free vector field; consequently, in this version of the iSTFG, only  $U_c$  is matched at every time step during the signal sampling phase. Enforcing the divergence-free condition is relayed to the flow solver during the first time step without significantly increasing the computational cost. The time step in the iSTFG is chosen based on the convection velocity and the grid spacing (known as a-priori) given by

$$\Delta t_{\text{iSTFG}} = \frac{\text{CFL} \Delta x_1}{U_c}, \quad (13)$$

where CFL is the Courant–Friedrichs–Lewy number set to a value of 0.95 unless specified otherwise, both in the iSTFG and the simulation, and  $\Delta x_1$  is the grid spacing in the streamwise direction. Here, the plane-averaged quantity is defined as

$$\langle g \rangle(x_3, t) = \frac{1}{L_{x_1} L_{x_2}} \int_0^{L_{x_1}} \int_0^{L_{x_2}} g(x_1, x_2, x_3, t) dx_1 dx_2, \quad (14)$$


where  $L_{x_1}$  and  $L_{x_2}$  are the domain lengths in the streamwise and spanwise directions, respectively, and  $g$  is any flow parameter of interest.

To simplify the digital filter implementation, a constant filter width is used in the spanwise and vertical directions despite the non-constant grid spacing used in the vertical direction. In the current implementation of the synthetic flow generator, three user parameters are required:

a. The target mean velocity profile  $\bar{U}_1(x_3)$ , b. The Reynolds stress tensor  $R_{ij}(x_3)$ , and c. Integral length scale  $I_l$ . A single scalar value for  $I_l^+ \equiv I_l u_\tau / \nu = 100.0$  is applied isotropically for the medium-size domain simulations, unlike the method used in Kim et al. [17]; this is primarily motivated by the need to keep the iSTFG implementation as simple as possible. Relatively more complex non-isotropic STFGs do exist, such as the one developed by Schau et al. [22]; however, in this paper, we chose a simpler approach to reduce complexity in the overall implementation of the code. Additionally, a relatively simple STFG also reduces the computational cost associated with the digital filter used to synthesise the flow field. The target mean velocity and Reynolds stress profiles act as input parameters, while the  $I_l^+$  can be viewed as a model parameter for the iSTFG. The serial and parallel implementation of the iSTFG can be accessed through the open-source, public repository: <https://github.com/AkshayPatil1994/Synthetic-Eddy-Method-KCX2013>.

Fig. 1 depicts the initial condition generated using the iSTFG, while Fig. 2 shows the plane-averaged profiles for the mean velocity and the root-mean-squared (rms) and  $\langle u'_1 u'_3 \rangle$  components compared against the data from Moser et al. [3]. The rescaling from  $Re_\tau = 590$  to  $Re_\tau = 500$  is done through a transformation between the non-dimensional profiles from the reference dataset. Specifically, knowing the values for  $v'$ ,  $H'$ , and  $u'_t$  for the target simulation, the dimensional profiles are first interpolated from the reference dataset onto the target simulation grid in the vertical direction. When interpolating the profiles on the target simulation grid, the vertical coordinate must be normalised as  $x_3^* = x_3/L_{ref}$  where  $x_3$  is the vertical coordinate and  $L_{ref}$  is the reference height used in the dataset which leads to a rescaling of the vertical coordinate over unit length. Following interpolation, the dimensional velocity profile  $\bar{U}_1$  is obtained using  $\bar{U}_1 = u'_t \bar{U}_1^+$ , the dimensional Reynolds stress profiles are obtained using  $\bar{u}'_i \bar{u}'_j = (u'_t)^2 (\bar{u}'_i \bar{u}'_j)^+$ , where,  $(\cdot)^+$  represents the non-dimensional profiles obtained from the reference dataset. The rescaling works best when the difference between the target and reference datasets is not large. However, for target simulations where a reference dataset is not available, the closest available dataset can provide a reasonable starting point for spinning up the simulation. As seen in Fig. 2, the plane-averaged profiles compare well as expected with the reference data and illustrate the suitability of this method to generate initial conditions for channel flow simulations. The novelty of this work does not arise from the iSTFG itself, but in the way it is applied to generate the initial conditions for channel flows.

Table 1 details the various simulations carried out in this paper. All the simulations were run on the DelftBlue (DB) high-performance computing centre at the Delft University of Technology. Individual compute node consists of 4th Generation Intel® Xeon® 2 × 32-core -



**Fig. 2.** (a) plane-averaged streamwise velocity profile comparison of the iSTFG data against linear, log-law profiles. (b) Root-mean-squared velocity and Reynolds stress profiles compared against the reference data from Moser et al. [3]. On both panels, data markers are shown for one in every four data points.

**Table 1**

Various simulations considered in this paper. Case naming convention first denotes the friction Reynolds number (e.g., Re500 for a friction Reynolds number of 500.0) followed by the type of initial condition used encoded by the first three letters (e.g., Syn for Synthetic Profile) followed by domain size denomination where S, M, and L are small, medium, and large domain sizes, respectively, and finally the integral length scale used in the iSTFG (e.g., I20 corresponds to  $I_1^+ = 20.0$ ). All cases except the large domain cases ending with the letter L have a spatial resolution  $\Delta x_1^+ = 4.18$ ,  $\Delta x_2^+ = 3.0$ ,  $\Delta x_{3,min}^+ = 0.4$ , while the large domain cases have a spatial resolution  $\Delta x_1^+ = 6.14$ ,  $\Delta x_2^+ = 4.18$ , and  $\Delta x_{3,min}^+ = 0.4$ , respectively. For cases with  $Re_\tau = 350$ ,  $\Delta x_{3,max}^+ = 5.4$  while for cases with  $Re_\tau = 500$ ,  $\Delta x_{3,max}^+ = 3.7$ . This resolution is sufficient for channel flows to resolve the requisite flow features of interest [23,24]. For both the friction Reynolds numbers, the synthetic profiles are obtained from the dataset by Moser et al. [3].

Case name	$Re_\tau$	Initial condition	Grid	Domain size	$I_1^+$	CPUs
Re350LinM	350	Linear Profile	$1048 \times 764 \times 128$	$4\pi H \times 2\pi H \times H$	–	64
Re350LogM	350	Log Profile	$1048 \times 764 \times 128$	$4\pi H \times 2\pi H \times H$	–	64
Re350SynM	350	Synthetic Profile	$1048 \times 764 \times 128$	$4\pi H \times 2\pi H \times H$	100	64
Re500LinM	500	Linear Profile	$1500 \times 1048 \times 256$	$4\pi H \times 2\pi H \times H$	–	128
Re500LogM	500	Log Profile	$1500 \times 1048 \times 256$	$4\pi H \times 2\pi H \times H$	–	128
Re500SynM	500	Synthetic Profile	$1500 \times 1048 \times 256$	$4\pi H \times 2\pi H \times H$	100	128
Re500LinS	500	Linear Profile	$750 \times 524 \times 256$	$2\pi H \times \pi H \times H$	–	64
Re500LogS	500	Log Profile	$750 \times 524 \times 256$	$2\pi H \times \pi H \times H$	–	64
Re500SynS	500	Synthetic Profile	$750 \times 524 \times 256$	$2\pi H \times \pi H \times H$	100	64
Re500LinL	500	Linear Profile	$2048 \times 1500 \times 256$	$8\pi H \times 4\pi H \times H$	–	256
Re500LogL	500	Log Profile	$2048 \times 1500 \times 256$	$8\pi H \times 4\pi H \times H$	–	256
Re500SynL	500	Synthetic Profile	$2048 \times 1500 \times 256$	$8\pi H \times 4\pi H \times H$	100	256
Re500SynMI20	500	Synthetic Profile	$1500 \times 1048 \times 256$	$4\pi H \times 2\pi H \times H$	20	128
Re500SynMI40	500	Synthetic Profile	$1500 \times 1048 \times 256$	$4\pi H \times 2\pi H \times H$	40	128
Re500SynMI60	500	Synthetic Profile	$1500 \times 1048 \times 256$	$4\pi H \times 2\pi H \times H$	60	128
Re500SynMI80	500	Synthetic Profile	$1500 \times 1048 \times 256$	$4\pi H \times 2\pi H \times H$	80	128

E5-6448Y 32C 2.1 GHz processor, totalling 64 cores per node. Each of the  $Re_\tau = 350$  cases required a total of 3392 CPU hours to simulate  $10 T_\epsilon$  using one full node. The small-size domain with  $Re_\tau = 500$  required a total 4200 CPU-hours to simulate  $10 T_\epsilon$  using one full node, while the medium-size domain with  $Re_\tau = 500$  cases required a total of 13 100 CPU-hours to simulate  $10 T_\epsilon$  using two full nodes, and the large-size domain with  $Re_\tau = 500$  required a total of 35 900 CPU-hours to simulate  $10 T_\epsilon$  using four full nodes. All simulations detailed in Table 1 were run for a total of  $10 T_\epsilon$  and the temporal evolution of the statistics was compared by averaging the statistics for the last  $5 T_\epsilon$  unless mentioned otherwise.

### 3. Results and discussion

#### 3.1. Medium domain size convergence behaviour

##### 3.1.1. Convergence history for the mean velocity and variances

First, we present the convergence of the shear stress as a function of time, as it represents the global balance between the imposed pressure gradient and shear stress at the wall. Thus, the shear/friction velocity ( $u_\tau$ ) can be deduced from the driving pressure gradient and the channel height given by

$$u_\tau^2 = \tau|_{x_3=0} = \nu \partial_3 \langle U_1 \rangle|_{x_3=0} = \Pi_c H, \quad (15)$$

here, it is assumed that the shear-stress has units of  $m^2/s^2$ , which is equivalent to setting  $\tau = \tau/\rho_0$ , where  $\rho_0$  is the density of the fluid.

Fig. 3 compares the time evolution of the normalised shear stress for the three initial conditions and two  $Re_\tau$  considered in this study. For cases with the log profile, the transition occurs identically for both values of  $Re_\tau$  such that initially, the shear stress is approximately half of the target value, followed by a sudden transition to elevated shear stress due to the downward convecting pair of vortices that trigger the flow to a turbulent state. Despite scaling the vortex pair and the initial condition by the bulk velocity associated with the target flow Reynolds number, the shear stress still experiences an overshoot that requires more than  $5 T_\epsilon$  to reach close to the  $\pm 5\%$  of the target value. Comparing the behaviour of the vortex pair with the linear profile results in the opposite trend, where the shear stress has a large magnitude that effectively triggers the flow to a turbulent state earlier when compared to the log profile. However, once the transition occurs, the shear stress is lower than observed for the log profile cases. This difference can mainly be attributed to the fact that in the linear profile, once the flow transitions to turbulence, the largest velocities close to the wall are suppressed due to momentum mixing in the vertical direction, thus resulting in smaller shear stress. In the log profile, the region close to the bottom wall exhibits a relatively higher velocity, thus resulting in an overall larger shear stress. As for the synthetically generated initial conditions, similar lower shear stress is observed within the first eddy turnover with a slight overshoot above the target value for case Re350SynM, which eventually asymptotes around the target value. Since the synthetically generated turbulence retains the scaled components of all the fluctuating quantities, the flow is much closer



to the target state and transitions to  $\pm 5\%$  of the target value in  $2T_e$ . For case Re500SynM, it is observed that the shear stress is consistently smaller than the target value in the transient phase while being within the  $\pm 5\%$  of the target value. However, despite this minor difference, the overall trend remains identical to case Re350SynM. Using a linear fit to the data after an initial transient of  $5T_e$ , the log profile is expected to converge at around  $21.5T_e$  and  $19.0T_e$ . In comparison, the linear profile is expected to converge at  $19.0T_e$  and  $15.0T_e$  for  $Re_\tau = 350$  and  $Re_\tau = 500$ , respectively. The same linear fit suggests that for the log profile to first enter the  $\pm 5\%$  range, it would take a total of  $\sim 11T_e$ , while for the linear profile, it would take  $\sim 12T_e$ . These linear fits are crude estimates based on the overall trend after  $5T_e$ . In addition to the three initial conditions discussed here, we also tested another standard technique to spin up the simulation, where a precursor simulation is run at a smaller Reynolds number or a coarse grid; typically  $Re_\tau^p = Re_\tau^t/2$ , where  $Re_\tau^p$  is the precursor friction Reynolds number and  $Re_\tau^t$  is the target friction Reynolds number. Typically, the precursor simulation is run for  $3T_e$  and then the velocity and pressure fields are mapped to the target Reynolds number simulation grid. While not compared extensively, Appendix B clearly details the shortcomings of this precursor method, which is typically used in channel flow simulations as it becomes increasingly expensive with increasing friction Reynolds number compared to the iSTFG method discussed in this paper. The precursor method requires large computational resources to obtain the initial conditions for the target simulation and a longer time to converge to the statistically stationary flow state. In our comparison, we only considered an initial spin-up of  $3T_e$  for the precursor method. It is important to note that this is potentially not long enough to ascertain sufficiently converged statistics before mapping the solution to the target simulation. However, this choice was made with the aim of allowing for a fair comparison between the iSTFG and the precursor method to obtain the initial conditions. Since the precursor method solves the governing equations while iSTFG merely samples a signal, the primary differentiator between the two is the algorithmic complexity. Thus, the computational cost for the precursor simulation is relatively higher and was limited to  $3T_e$ . Since the precursor method was not a focus of this study, the details have been listed in Appendix B and will not be discussed further. A visual comparison of the various initial conditions and the corresponding velocity profiles can be visualised in the supplementary movies included along with the manuscript. Movies 1 and 2 compare the streamwise velocity along the centre of the domain in conjunction with the plane-averaged streamwise velocity profile. Comparing the three initial conditions clearly illustrates the swift convergence to statistically stationary flow conditions when using the synthetic initial conditions (multimedia available online for both cases).

Fig. 4 compares the mean and rms velocity profiles for the three initial conditions detailed in the previous section. All profiles discussed below are averaged over the last  $5T_e$  as detailed in Fig. 3 with the right-pointing arrow. Since the mean velocity undergoes significant changes as a function of time during the transient phase, the time-averaged velocity changes as a function of the averaging window. Consequently, for a consistent comparison, the velocity is decomposed as

$$U_i(x_i, t) = \langle \bar{U}_i \rangle(x_i) + u'_i(x_i, t), \quad (16)$$

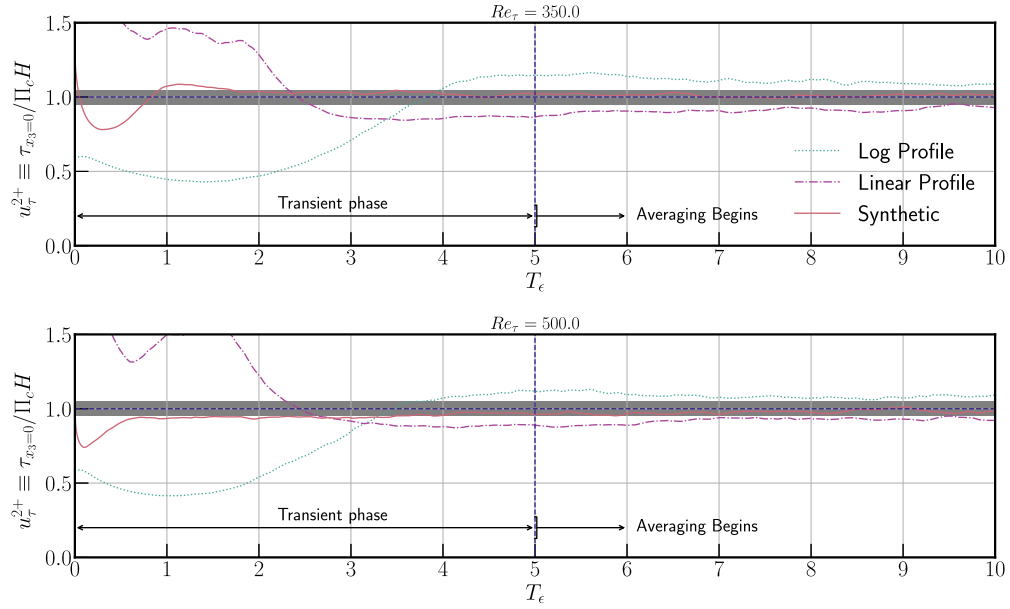
where,  $\langle \bar{U}_i \rangle$  is the plane- and time-average using the last  $5T_e$ . While the convergence of the shear stress provides a first indication of the stationary state of the flow, the velocity profiles and their variances also need to converge to carry out meaningful averages. As seen in the plane- and time-averaged velocity ( $\langle \bar{U}_i \rangle$ ), excellent agreement between the expected analytical log-law (black dashed line) and case Re350SynM is observed as expected, while case Re350LogM shows elevated velocity profile, and case Re350LinM shows slightly lower velocity profile. Case Re350SynM exhibits a slightly higher velocity magnitude in the inertial range, which can be primarily attributed to the relatively low Reynolds number and the averaging time used to compute the statistics

compared to standard datasets. Regardless, swift statistically stationary conditions are expected based on the convergence history of the shear stress detailed in Fig. 3. The synthetic profile compares well with the data from Moser et al. [3], further validating the convergence history observed in the shear stress. As the case Re350LogM shows elevated shear stress over the entire averaging period, the mean velocity profile suffers from an excess of total kinetic energy available that requires a longer time to be dissipated. For case Re350LinM, the opposite trend is observed and thus exhibits a relatively lower velocity magnitude. The rms velocity profiles and the Reynolds stress exhibit a similarly consistent trend, as observed in the mean velocity profile for all the cases with  $Re_\tau = 350$ .

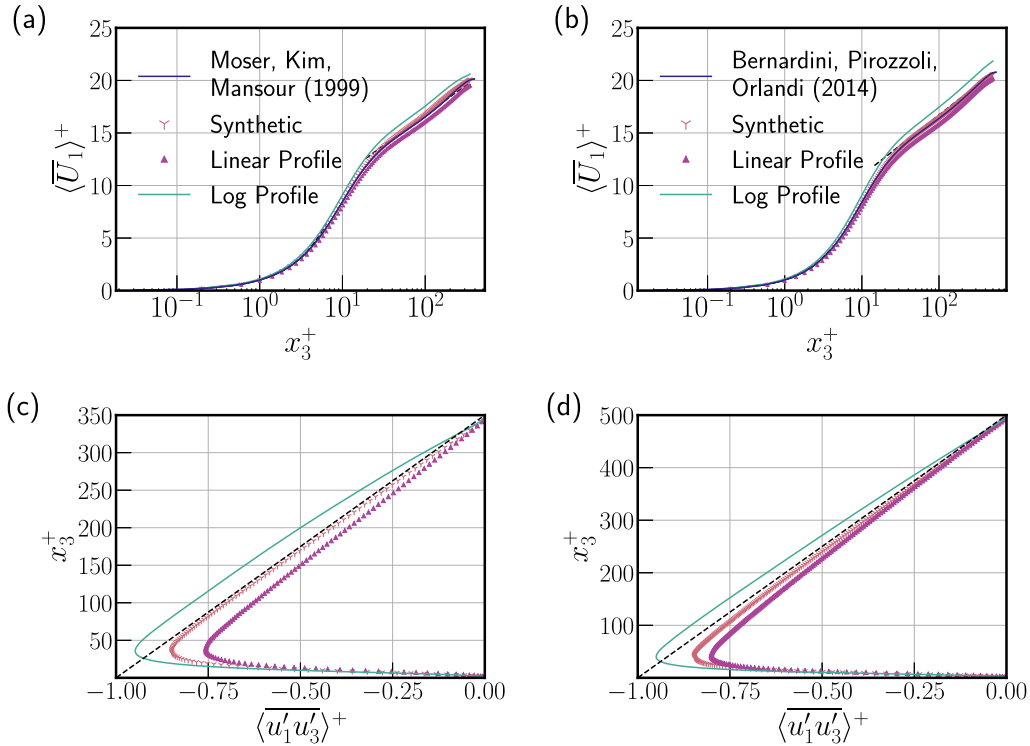
For cases with  $Re_\tau = 500$ , a similar trend is observed when the three initial conditions are compared against each other. Specifically, case Re500SynM shows an excellent agreement with the dataset from Bernardini et al. [24] for the mean velocity profile, while cases Re500LogM and Re500LinM exhibit a larger and smaller velocity magnitude when compared to the expected profiles, respectively. This trend between the three initial conditions is consistent across the rms velocity profiles and the Reynolds stress profiles, suggesting that the synthetic initial condition is expected to converge identically for higher friction Reynolds numbers. The rms velocity profiles and Reynolds stress are not compared against the dataset as the top wall boundary conditions are not identical, and minor differences can be observed when compared to a no-slip and free-slip channel. However, this does not affect the overall convergence trend presented in this work. Although not explicitly shown here, the absolute error between the synthetic profiles and the benchmark datasets across the entire velocity profile is less than 3%. These observations collectively suggest that the mean velocity profiles and the variances converge relatively quickly for the synthetic initial condition compared to the analytical profiles discussed in this paper. In the following section, the convergence will be assessed for the higher-order statistics, such as the spectral energy and the turbulent kinetic energy budgets.

### 3.1.2. Time evolution of the energy spectrum and turbulent kinetic energy budget

The time evolution of the energy spectrum complements the observations made in the previous sub-section. Specifically, in this case, we look at the plane-averaged spectrum for the velocity magnitude as detailed in Fig. 5. For the various wall-normal locations detailed in Fig. 5, a consistent trend between the three initial conditions is observed for both values of  $Re_\tau$ , where the linear profiles exhibit relatively smaller total energy at a given wavenumber compared to the synthetic and log-profiles, however, the difference between the three cases is not large. The time-evolution of the spanwise- and plane-averaged streamwise spectral energy for cases Re350LogM and Re500LogM has a consistent behaviour where the total energy during the first  $5T_e$  is relatively lower, thereby increasing to a consistent value (a trend also evident in the supplementary movies 1 and 2). The time evolution of the spectral energy for cases Re350LinM and Re500LinM also shows a consistent transition to turbulence. A key difference between the linear and the log profiles is that the excursion around the mean value for the linear profile is relatively smaller compared to the one observed for the log profile case. As for cases Re350SynM and Re500SynM, they do not seem to deviate much from their initial state, and most of the changes are observed at high wavenumber. In contrast, the small wavenumbers corresponding to the large-scale turbulent features of the flow are relatively converged. Since the synthetic initial condition preserves the form of the Reynolds stress tensor using the exponential kernel in space, the distribution of spectral energy is preserved at the respective wavenumbers, unlike the analytical profiles where uncorrelated white noise is added on top of a mean profile. This is clearly seen in the spectrum shape at  $T_e = 0$  compared to the averaged spectrum marked with black + symbols.



**Fig. 3.** Comparison of the convergence of  $u_\tau^2$  for the various initial conditions and  $Re_\tau$  discussed in this study. The grey region marks the  $\pm 5\%$  around the target value.



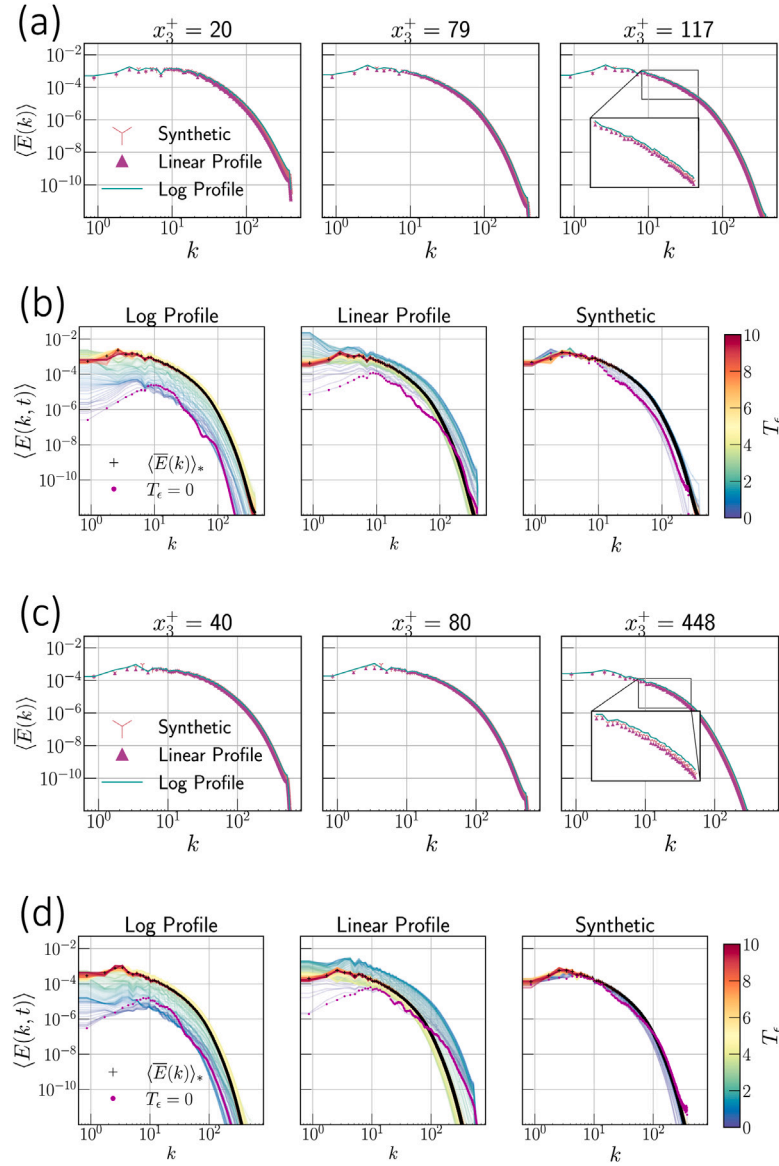
**Fig. 4.** (a) Comparison of plane- and time-averaged velocity profiles with  $Re_\tau = 350$  for the three initial conditions against the reference dataset by Moser et al. [3]. (b) Comparison of plane- and time-averaged velocity profiles with  $Re_\tau = 500$  for the three initial conditions against the reference dataset by Bernardini et al. [24]. The dashed line in panels (a) and (b) corresponds to the log-law detailed in Eq. (5). Panels (c) and (d) compare the Reynolds stress profiles for the three initial conditions considered in this paper, where panel (c) corresponds to  $Re_\tau = 350$  and panel (d) corresponds to  $Re_\tau = 500$ .

The turbulent kinetic energy (TKE) budget for a channel flow with a streamwise driving pressure gradient and homogeneous streamwise and spanwise directions is given by,

$$\partial_t \mathcal{K} = \mathcal{P}_k - \epsilon_k + \mathcal{V}_k - \mathcal{T}_k - \Pi_k, \quad (17)$$

where  $\mathcal{K} \equiv u'_i u'_i / 2$  is the TKE,  $\partial_t \mathcal{K}$  is the time rate of change of TKE and is zero for statistically stationary or steady-state conditions as will be assumed in this case,  $\mathcal{P}_k$  is the production of TKE via mean shear,  $\epsilon_k$  is the dissipation rate of TKE,  $\mathcal{V}_k$  is the viscous diffusion of TKE,

$\mathcal{T}_k$  is the turbulent transport of TKE, and  $\Pi_k$  is the pressure-diffusion of TKE, respectively. The last three terms are divergence terms and only transport the TKE within the domain without changing the net flux of TKE within the system. Fig. 6 compares the various terms in the TKE budget that are averaged in the homogeneous directions and over  $5T_\epsilon$  as indicated by the right arrow in Fig. 3. A consistent trend observed for the mean velocity and its variances is seen for the TKE budget terms; specifically, the TKE production, which is governed by the product of the Reynolds stress and the mean velocity gradient, has




**Fig. 5.** Comparison of spanwise- and time-averaged energy spectra for  $Re_\tau = 350$  (panel a) and  $Re_\tau = 500$  (panel c). Panels (b) and (d) compare the time evolution of the spectra at  $x_3^+ = 204$  and  $x_3^+ = 445$  for the three initial conditions discussed in this paper for  $Re_\tau = 350$  and  $Re_\tau = 500$ , respectively. The  $\langle \cdot \rangle_{*}$  operator signifies a plane-averaged quantity after the first  $5T_c$  as marked in Fig. 3. The x-axis on all panels marks the wavenumber ( $k$ ) while the y-axis denotes the energy as a function of the wavenumber ( $E(k)$ ).

a larger magnitude for the log profile cases and a smaller magnitude for the linear cases. The synthetic initial condition is observed to follow the baseline reference data quite accurately, with case Re350SynM exhibiting a relatively small overshoot mainly because of the relatively larger mean velocity estimates compared to the reference dataset of Moser et al. [3]. The TKE dissipation rate agrees with the baseline dataset for both Reynolds numbers quite accurately for the synthetic initial conditions. At the same time, the other two cases seem to differ quite substantially. The divergence terms in the TKE budget also compare well with the synthetic case despite the relatively short averaging and simulation window, suggesting that all the parameters of interest in the channel flow are adequately converged. Having demonstrated the relatively quick convergence to statistically stationary flow state using the iSTFG and medium-sized computational domain, in the following sub-section, we explore the sensitivity of iSTFG as a function of the computational domain size and the integral length scale (which acts as a model parameter) on the convergence using the  $Re_\tau = 500$  in the following sub-sections.


### 3.2. Effect of domain size and $I_l^+$ on convergence

To better understand the effect of domain size used to simulate the channel flow, we consider three different cases as detailed in Table 1 for cases with  $Re_\tau = 500$ . The time evolution of the plane-averaged shear stress is presented in Fig. 7, where it is clear that for the synthetic profile, the convergence to the target value is independent of the domain size. For cases with the linear profile as their initial condition, there is a delay observed in the time at which the transition occurs, and this delay is observed to increase with increasing domain size. A similar trend is observed for the cases with the log profile as their initial condition. As these two initial conditions rely on the vortex pair to trigger the turbulence in the channel, with increasing domain size, the flow requires relatively more time to reach the target state. Fig. 8 shows this sensitivity to the domain size, where all the synthetic profile cases (panels c, f, and i) are observed to have transitioned to a turbulent state. However, when comparing the linear profile cases, i.e., Figs. 8(a), (d), and (g), further elucidate the differences observed in the time evolution for this case as seen in Fig. 7. A similarly strong effect of the





**Fig. 6.** Comparison of the TKE budget terms where all the terms are non-dimensionalised using  $u_\tau^4/(\kappa\nu)$ . Panels (a)–(e) correspond to  $Re_\tau = 350$  while panels (f)–(j) correspond to  $Re_\tau = 500$ . Line style identical to Fig. 4.



**Fig. 7.** Comparison of the convergence history as a function of the initial condition and the domain size used to simulate the channel flow.

large domain is observed for the log profile cases (panels b, e, and h), where case Re500LogL exhibits a typical turbulent spot [25,26] due to the downward descending pair of vortices that eventually triggers the flow into the turbulent state. This strong dependence on the domain size for the log profile cases is expected, as only one pair of vortices is introduced in the domain as the initial condition. One could consider adding multiple pairs of such downward-descending vortices; however, as observed in the time evolution of the plane-averaged shear stress in Fig. 7, the combination of the log profile and the vortices leads to elevated levels of TKE within the flow domain, which can take a long time to dissipate [11]. The synthetic profile cases, on the other hand, are not affected by the variation of the domain size as the flow field is generated throughout the domain without any dependence on local transition or trigger mechanisms; this makes the iSTFG more appealing when large domain and  $Re_\tau$  are considered. This discussion illustrated that the synthetic method using the iSTFG is domain-size agnostic when comparing the convergence to a statistically stationary flow state.

Since the synthetic profile is independent of the domain size used to obtain the time evolution, the effect of the integral length scale used in the iSTFG as the only model parameter was also explored to quantify its sensitivity. For turbulent channel flows, a wide range of detailed studies have been carried out to better understand the transition to turbulent flow state [1,27,28, to list a few] detailing the various mechanics that lead to this transition. For a statistically stationary flow state, a fundamental length scale of interest is the integral length scale  $I_l^+$ , which for channel flows has been suggested to be  $I_l^+ \sim 100$  [1,18]; corresponding to the streamwise streaks responsible for near wall turbulence intermittency [1]. It has also been shown that this is the

relevant scale of interest when considering the non-linear regeneration cycle that sustains turbulent motion within the buffer layer of wall-bounded flows, thus proving to be a useful scale of interest when testing the sensitivity of iSTFG. To this end, we chose the medium-size computational domain to investigate the effect of varying values of  $I_l^+$  as detailed in Table 1. Since the iSTFG's computational cost scales with the width of the isotropic stencil used in generating the field, larger values of  $I_l^+$  for all other identical parameters result in a higher computational cost. For cases Re500SynMLI20, Re500SynMLI40, Re500SynMLI60, Re500SynMLI80, and Re500SynMLI100, the computational cost to generate the turbulent initial conditions are listed in Table 2. It can be readily concluded from the data presented in Table 2 that smaller values of  $I_l^+$  correspond to relatively lower computational costs. Thus, smaller values of  $I_l^+$  are preferred. Comparing the time-evolution of the plane-averaged shear stress shown in Fig. 9 suggests that most differences observed for varying values of  $I_l^+$  are during the early part of the transient phase. For case Re500MI20, the time evolution is similar to that of case Re500LogM primarily because the  $I_l^+$  sets the correlation length scale. Consequently, with decreasing values of  $I_l^+$ , the flow decorrelates relatively quickly and tends to the log profile superposed with white noise as seen in Fig. 10. The use of an isotropic  $I_l^+$  value for all the velocity components can be clearly seen in Fig. 10. During the initial part of the transient phase, there is a weak dependence observed as a function of the  $I_l^+$ . Smaller values of  $I_l^+$  seem to transition relatively later when compared to the larger values of  $I_l^+$ ; however, this trend asymptotes to the target value beyond  $3T_c$ . These observations suggest that the iSTFG initial condition does not impact the overall

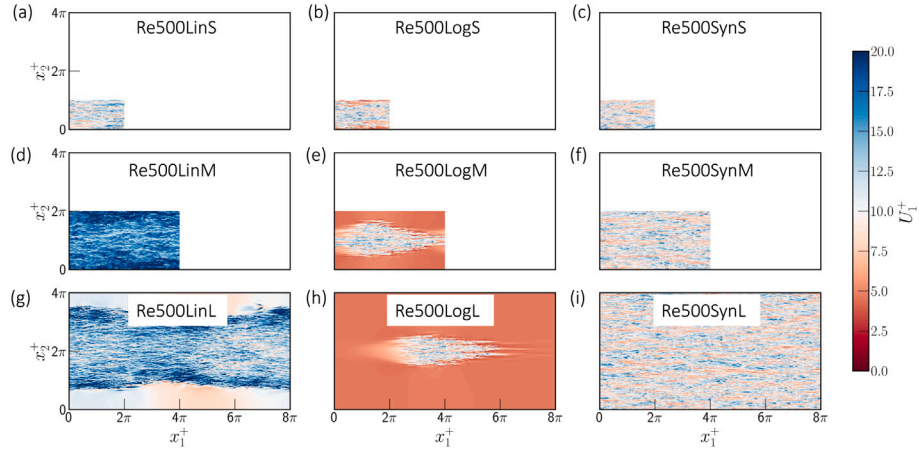


Fig. 8. Normalised streamwise velocity heatmap comparison at  $x_3 = H/8$  above the bottom wall at  $T_e = 2.0$  for the cases marked at the top centre on each panel of the figure.




Fig. 9. Time evolution of the plane-averaged shear stress as a function of varying values of  $I_l^+$  for the medium-sized computational domain. The inset shows a zoomed-in view of the first  $T_e = 2.5$  for the various cases. Data markers for the simulations are displayed once every 50 data points.

Table 2

Computational cost for the various cases used to investigate the impact of  $I_l^+$  on the convergence of flow statistics. The iSTFG was compiled using the OpenMPI mpi90 compiler [29] on a single socket 13th Generation Intel® Core™ i9-13900K CPU with 64 GiB of Random-Access-Memory (RAM) using 8 CPUs. For additional implementation details, the readers are referred to the public repository of the software: <https://github.com/AkshayPatil1994/Synthetic-Eddy-Method-KCX2013>.

Case name	Wall clock (s)	CPU-hours
Re500SynMI20	71	0.16
Re500SynMI40	117	0.26
Re500SynMI60	210	0.48
Re500SynMI80	295	0.66
Re500SynMI100	414	0.92

time evolution of the shear stress development and convergence to the statistically stationary flow state.


As discussed in the above sub-section, the time evolution of the plane-averaged shear stress is observed to be domain size and  $I_l^+$  agnostic for the initial conditions generated using iSTFG; thus, we will briefly discuss the absolute error for the plane-averaged streamwise velocity for the various  $I_l^+$  considered in this paper. The absolute error (in %) is calculated using

$$\epsilon_U = \text{abs}(U_{\text{ref}} - \langle U_1 \rangle) / (100 U_{\text{ref}}^{-1}), \quad (18)$$


where  $\epsilon_U$  is the per cent error,  $U_{\text{ref}}$  is the reference velocity profile, and  $\langle U_1 \rangle$  is the plane-averaged velocity profile at a given time instance. Fig. 11 compares the absolute error for each plane-averaged velocity profile at ten different eddy turnovers as a function of the distance from the bottom wall. Over the first three eddy turns, the error for

case Re500SynMI100 is above 5% close to the bottom wall, while above  $x_3^+ > 100.0$ , the error decreases rapidly. As for the other cases compared in Fig. 11, the error decreases significantly over the first three eddy turns in the near wall region and stays consistently lower than that of case Re500SynMI100. Far from the wall, the error increases for case Re500SynMI20 but falls below 1% at  $5T_e$ . Smaller values of  $I_l^+$  are observed to exhibit relatively lower absolute error during the transition phase; however, this trend is not as consistent and strong as the absolute error for case Re500SynMI80 is relatively larger when compared to Re500SynMI100 at  $5T_e$ . Overall, all the cases are observed to exhibit good agreement with the maximum absolute error at  $5T_e$  is  $\leq 3\%$  throughout the vertical extent.

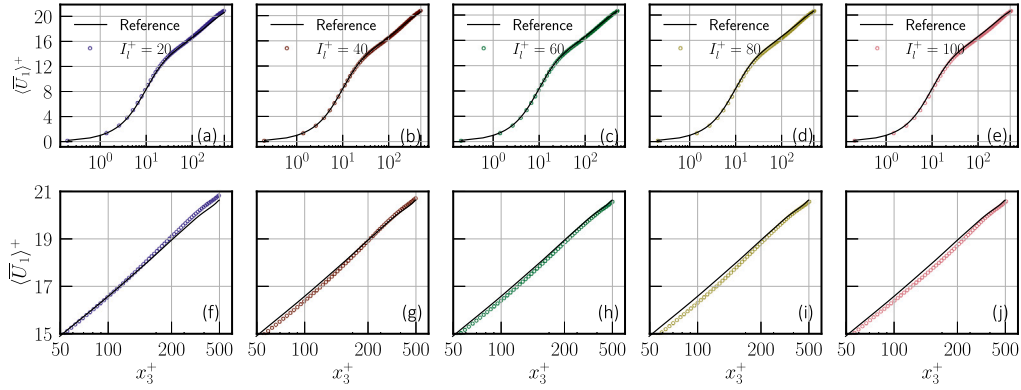
Fig. 12 compares the plane- and time-averaged velocity profiles against the reference velocity profile. Consistent with the shear stress convergence as a function of time, increasing values of  $I_l^+$  are observed to have relatively more significant errors when compared against the reference data throughout. This is largely supported by the previous discussion when the absolute error is compared in Fig. 11. Additionally, cases with relatively small values of  $I_l^+$  are observed to converge quickly in the lower portion of the log-law region. In contrast, closer to the top boundary, cases with relatively large values of  $I_l^+$  are in better agreement with the reference data as evidenced by Figs. 11 and 12. It is pertinent to recognise that the results presented in Fig. 12 are subject to improvement with increasing averaging time used to obtain the plane- and time-averaged results. These observations, combined with the computational cost associated with iSTFG as a function of  $I_l^+$  (see Table 2), suggest that of the five values for  $I_l^+$  discussed in this paper,  $I_l^+ = 20$  is the computationally efficient choice when all other parameters are identical. These results provide further support



**Fig. 10.** Auto-correlation of the streamwise ( $R_{11}$ ), spanwise ( $R_{22}$ ), and vertical ( $R_{33}$ ) velocity components for four distinct values of  $I_l^+$  as a function of the streamwise spacing. The autocorrelation data presented in this figure were generated for case Re500SynM, where  $I_l^+ = 1$  and  $I_l^+ = 5$  are only presented for qualitative comparison and were not simulated in this work.



**Fig. 11.** Absolute error (in %) as a function of the eddy turnover time for the medium-size domain with varying  $I_l^+$ . The reference velocity profile is obtained by running case Re500SynM for  $20T_ε$  and averaging over the last  $10T_ε$ . Data markers for the simulations are displayed once every 20 data points with the line colour and style identical to Fig. 9.



**Fig. 12.** Plane- and time-averaged velocity profile comparison for various values of the model parameter  $I_l^+$  simulated in this paper. The top row panels (a-e) show the entire velocity profile, while the bottom row panels (f-j) show the log-law region of the velocity profile. The velocity profiles are time-averaged after the initial five eddy turns. Data markers for the simulations are displayed for every third data point.

for the utility of the iSTFG in achieving fast and accurate stationary flow conditions without incurring a large computational cost.

### 3.3. Impact of friction Reynolds number mismatch


In the previous sections, a detailed discussion was presented to motivate the utility of the iSTFG to obtain converged statistics using reference data from precomputed channel flow data sets [13,24,30, to list a few]. However, it is easy to imagine a situation where such reference data is either unavailable or contains a friction Reynolds number mismatch. To make iSTFG functional in scenarios like this,

we carried out a sensitivity test for the iSTFG to achieve a target  $Re_τ = 500.0$  subject to an input data mismatch as detailed in Table 3. Specifically, we compared the convergence to an *a-priori* known target  $Re_τ^t = 500.0$  with 10% and 20% deviation in the initial conditions. For all cases detailed in Table 3, the model parameter  $I_l^+ = 20$  corresponds to the reference friction Reynolds number ( $Re_τ^t$ ). The CPU hours for the various cases differ as the underlying computational grid is designed to resolve the requisite scales of motion for  $Re_τ^t$ . Thus, for a relatively lower value of  $Re_τ^t$ , the stencil width for a fixed model parameter of  $I_l^+$  spans over relatively more grid points, resulting in a slightly higher computational cost and vice versa. For case Re500SynMI20-450, there is no reference dataset available from conventional channel flow

**Table 3**

Simulation details for cases used to verify the sensitivity of initial conditions with a  $Re_\tau$  mismatch between the reference and the target simulations. The case naming convention is identical to the definition detailed in Table 1 with the  $-(\text{num})$  indicating the  $Re_\tau^r$  from which the initial conditions are derived. Here,  $Re_\tau^r$  and  $Re_\tau^t$  correspond to the reference and the target friction Reynolds number, respectively, while the CPU-hours column denotes the computational cost associated with generating the initial conditions.

Case name	$Re_\tau^r$	$Re_\tau^t$	% mismatch	CPU-hours	Source
Re500SynMI20-395	500	395	~ 20% lower	0.165	Moser et al. [3]
Re500SynMI20-450	500	450	10% lower	0.153	Rescaled to $Re_\tau = 395$ [3]
Re500SynMI20-550	500	550	10% higher	0.140	Lee and Moser [30]
Re500SynMI20-590	500	590	~ 20% higher	0.132	Moser et al. [3]
Re500SynMI20-180rescaling	500	180	~ 64% lower	0.140	Rescaled to $Re_\tau = 500$ [3]
Re500SynMI20-1000rescaling	500	1000	100% higher	0.140	Rescaled to $Re_\tau = 500$ [30]



**Fig. 13.** Time evolution of the plane averaged shear stress as a function of varied  $Re_\tau^r$  mismatch against the target  $Re_\tau^t = 500.0$  for the cases detailed in Table 3. The solid blue line marks the time evolution of the plane-averaged shear stress for case Re500SynMI20, while the black dashed line marks the time-averaged value, with the grey shading denoting the  $\pm 5\%$  around this time-averaged value. Data markers for each of the cases are only shown for one in forty points.

simulations. As a result, the mean velocity and Reynolds stress tensor profiles are obtained using a friction Reynolds number re-scaling using the data from Moser et al. [3] as detailed in Section 2. For the last two cases in Table 3 (cases with  $<\text{num}>\text{rescaling}$  in their name), the time- and plane-averaged velocity can introduce a large mismatch due to the  $Re_\tau^r$  and  $Re_\tau^t$  difference. Since the logarithmic profile is assumed to be universal [31], the mean velocity input profile for these three cases is prescribed as per Eq. (5) except for the white noise terms.

Fig. 13 compares the effect of  $Re_\tau$  mismatch on the time evolution of the plane-averaged shear stress for the various cases detailed in Table 3. Since cases Re500SynMI20-395, Re500SynMI20-450, Re500SynMI20-550, and Re500SynMI20-590 are not scaled down to the target simulation friction Reynolds number, they clearly exhibit a sensitivity to initial conditions. Specifically, cases initialised with a relatively lower friction Reynolds number ( $Re_\tau^r < Re_\tau^t$ ) exhibit a slower convergence to the target value when compared to the cases initialised with a relatively higher friction Reynolds number ( $Re_\tau^r > Re_\tau^t$ ). This observation can be generalised when the time evolution of the vertically integrated bulk velocity is considered as shown in Fig. 14(a). Cases with a relatively larger  $Re_\tau$  mismatch take longer to reach the target value, with a systematic trend observed with varying  $Re_\tau$ . This observation can be further understood through the vertically integrated TKE and Reynolds stress time evolution shown in Fig. 14. For cases that are not rescaled and exhibit a mismatch (cases Re500SynMI20-395, Re500SynMI20-450, Re500SynMI20-550, and Re500SynMI20-590) show a wide range of variation in the time evolution of both the TKE and the Reynolds stress. For case Re500SynMI20-590 and Re500SynMI20-395, the differences are small during the transition but grow after the initial  $0.5T_\epsilon$  as seen in panels (b) and (c) of Fig. 14. More importantly, the relative amount of TKE between cases Re500SynMI20-395 and Re500SynMI20-590 is drastically different, leading to a slower convergence towards the target value.

As presented in Figs. 13 and 14, the time-evolution of the various parameters of interest suggest that iSTFG can reliably accelerate the convergence towards stationary turbulent state when the input data is correctly scaled to the target  $Re_\tau$ . This aims to solve two challenges associated with turbulent channel flow spinup: (a) transition to a turbulent flow state for all  $Re_\tau$  [11], (b) relatively fast convergence towards


a statistically stationary flow state when compared to conventional methods. Specifically, using iSTFG to generate initial conditions forced only with a constant pressure gradient is able to accelerate the transition to the statistically stationary flow state. As detailed extensively in the discussion above, iSTFG is not sensitive to the model parameter and mismatch in the available  $Re_\tau$  reference data as long as the mean velocity is specified through the classical log-law [31] and the Reynolds stress profiles are scaled to the target  $Re_\tau^t$ . As detailed in Figs. 13 and 14, when the transition to statistically stationary flow is considered, all cases using the correctly scaled  $Re_\tau^t$  converge to the target values ( $\pm 5\%$ ) within  $5T_\epsilon$  for all parameters of interest at a minimal computational cost.

### 3.4. Limitations of iSTFG

The previous section presented a detailed discussion on the performance of iSTFG to obtain stationary flow conditions. While iSTFG provides a computationally efficient method to synthesise initial conditions, it has limitations that can be further improved as detailed below.

- All of the tests presented in this work were carried out at a relatively modest friction Reynolds number ( $Re_\tau = 500$ ). This was mainly done to limit the steep increase in computational cost when carrying out the parameter sensitivity test. Although similar results have been obtained for relatively higher  $Re_\tau \sim 700$  (not presented in this work), additional numerical experiments are needed to generalise this approach for higher friction Reynolds numbers.
- The input data required for iSTFG must be carefully rescaled to the target friction Reynolds number as discussed in Section 3.3, as the time required to reach the statistically stationary flow state can be influenced by a Reynolds number mismatch. While this may seem like a severe limitation, appropriate rescaling results in a similar convergence trend for cases where exact mean velocity profiles and Reynolds stress tensor data may not be available.
- One of the strongest limitations of iSTFG in the way it is applied to turbulent channel flows is the intrinsic assumption of





**Fig. 14.** Time evolution of the plane-averaged and vertically integrated profiles (a) streamwise velocity, (b) Turbulent kinetic energy, and (c) Reynolds stress, respectively, as a function of time. The solid blue line on all the panels corresponds to case Re500SynMI20, while the black dashed line marks the time-averaged value, with the grey shading denoting the  $\pm 5\%$  around this time-averaged value. Data markers for each of the cases are only shown for one in forty points with line style and colour identical to Fig. 13.

streamwise and spanwise homogeneity. In its current form, iSTFG translates all the assumptions underlying classical synthetic eddy methods [17] and a generalised application beyond prototypical channel flows must be explored with caution. Specifically, iSTFG does not provide a generalised solution for problems where the spanwise and streamwise homogeneity is not preserved, thus limiting its application to these equally interesting turbulent flow problems.

#### 4. Conclusion

In this study, we evaluated three different methods to initialise flow in a pressure-driven channel, aiming to achieve statistically stationary flow conditions in a computationally efficient manner. Our results indicate that the synthetically generated three-dimensional turbulence (iSTFG) is the most computationally efficient and effective approach, compared to two analytically initialised velocity fields, for reducing simulation spin-up time. At a one-time computational cost of 1–10 CPU hour(s), the iSTFG achieved statistically stationary flow conditions within three eddy turnover times—a significant improvement over the alternative methods, which required more than 10 eddy turnovers, as demonstrated in our work for the first time. This reduction in spin-up time translates to substantial computational savings, making iSTFG particularly valuable when precursor turbulent initial conditions are unavailable for large computational domains and high friction Reynolds numbers.

We also found that iSTFG is domain-size and model-parameter agnostic, reaching a stationary turbulent flow state quickly—unlike the other two methods discussed in this paper. In view of the computational cost, we recommend using iSTFG with a model parameter of  $I_t^+ = 20$ . Additionally, we carried out systematic tests to better understand the sensitivity of iSTFG against input data mismatch. We found that iSTFG can reproduce a similar time evolution trend when the input data is rescaled to the target value, while a mismatch in input data can negatively impact the convergence rate. The synthetically generated turbulence method offers a robust and resource-efficient strategy for setting up initial conditions in wall-bounded, pressure-driven channel flow simulations when the appropriate convective velocity is chosen. By minimising the time to reach statistically stationary conditions, this approach enhances the efficiency of turbulence simulations, facilitating more rapid and cost-effective exploration of complex flow phenomena in such configurations.

#### CRediT authorship contribution statement

**Akshay Patil:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Clara García-Sánchez:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Funding acquisition, Formal analysis.

#### Acknowledgements

The authors acknowledge the use of computational resources of the DelftBlue supercomputer, provided by Delft High Performance Computing Centre (<https://www.tudelft.nl/dhpc>). AP would like to acknowledge the use of the colour-blind friendly Python library cblind, which was used to make all the figures in this manuscript. The authors would like to thank P. Costa for the support provided while using CaNS in this work. AP would also like to thank K. Schau for an insightful discussion during the preparation of this manuscript. AP and CGS would like to acknowledge the constructive and insightful comments provided by the reviewers that greatly improved the quality of the manuscript.

#### Funding statement

This research was carried out as a part of the EU-Project RefMAP. RefMAP has received funding from the Horizon Europe program under grant agreement No 101096698. The opinions expressed herein reflect the authors' views only. Under no circumstances shall the Horizon Europe program be responsible for any use that may be made of the information contained herein.

#### Declaration of competing interest

The authors report no conflict of interest.

#### Carbon footprint statement

This work made use of the DelftBlue supercomputer and had an estimated footprint of 2370 kg CO<sub>2</sub>-equivalent (at least if not higher) using the Green Algorithms (<http://calculator.green-algorithms.org/>). The input data used to arrive at these estimates were: Runtime — 3433 h, Number of cores — 64, Model Xeon X5660 (closest to the original hardware), and Memory available: 252 GiB located in the Netherlands. This is equivalent to taking 1 flight from New York, U.S.A., to Melbourne, Australia for one individual.






Fig. A.15. Comparison of the initial shear stress as a function of the wall-normal direction for the cases simulated in this paper.

### Appendix A. Mean shear profiles

Based on the analytical expressions for the initial conditions, it is possible to a-priori estimate the mean shear profiles used to set up the initial conditions. Since the channel is homogeneous in the streamwise and spanwise directions, the mean shear is non-zero for the vertical gradient of the streamwise component (i.e.,  $\partial_3 U_1$ ). For the linear profile, the shear stress that is provided as the initial condition is given by

$$\nu \partial_3 U_1 = \left[ \frac{-2u_\tau^2}{\kappa} \right] \left[ \frac{1}{Re_\tau} \right] \left[ \log \left( \frac{H}{z_o} \right) + \frac{z_o}{H} - 1 \right], \quad (\text{A.1})$$

while for the linear-log-law profile, the shear stress is given by

$$\begin{aligned} \nu \partial_3 U_1 &= \nu, & \forall \frac{x_3 u_\tau}{\nu} \leq 11.6 \\ &= \left[ \frac{u_\tau^2}{\kappa} \right] \left[ \frac{1}{Re_\tau^*} \right], & \forall \frac{x_3 u_\tau}{\nu} > 11.6 \end{aligned} \quad (\text{A.2})$$

where  $Re_\tau^* \equiv u_\tau x_3 / \nu$  is the local friction Reynolds number. Fig. A.15 compares the stress profiles for the three initial condition types and two friction Reynolds numbers. A constant stress magnitude is applied for the linear profile, while for the linear-log-law profile, the stress varies as a function of the distance away from the wall.

### Appendix B. Convergence history for precursor mapping and computational cost comparison

A typical precursor simulation uses a relatively low-friction Reynolds number simulation as a starting point for the large-friction Reynolds

number case using some form of interpolation and rescaling. In this test case, we use a precursor simulation with half the target friction Reynolds number to supply the initial conditions for the target simulation. The precursor simulation is run for  $3T_\epsilon$  and mapped to the target simulation using interpolation, as the precursor simulation does not have the same computational resolution/grid as the target simulation. The target simulation was only carried out for the  $Re_\tau = 500$  case to compare the computational cost against the synthetic case. Fig. B.16 shows a comparison of the shear stress time history marked in black solid lines with circle markers against the three methods used in this study. The time evolution of the shear stress clearly shows that while the transition to a turbulent flow condition was established relatively quickly for the precursor method, the shear stress does not converge to the desired value compared to the synthetic method. Additionally, the cost associated with the precursor simulation does not justify the viability of such a method for increasing friction Reynolds number. In the test case discussed here, the precursor simulation was run at  $Re_\tau^p = 250$  with a grid size  $750 \times 524 \times 128 \sim$  using 8 CPUs and required 20 h to simulate  $3T_\epsilon$ , which resulted in a total of 160 CPU-hours of computational cost. This computational cost for the case with a precursor friction Reynolds number of 250 is feasible; however, with increasing friction Reynolds number, this cost increases non-linearly as detailed in [9]. The iSTFG algorithm relies mainly on stochastic sampling, which is computationally less intensive compared to the precursor simulation that solves the governing momentum equations.

In contrast, despite the algorithmic difference between the precursor method and the iSTFG, the computational cost for the iSTFG scales linearly with the number of streamwise grid points, given a fixed computational cost for the number of spanwise and vertical grid points.




Fig. B.16. Comparison of the time history of the shear stress for the precursor mapping method against the three methods discussed in this paper.

**Table B.4**

Computational cost as a function of the grid size and the friction Reynolds number. The last column represents the CPU hours needed to obtain the initial conditions for the target simulation, where the precursor simulations are run at half the friction Reynolds number for each case. The precursor CPU hours for  $Re_\tau = 2000$  were estimated using the wall clock time averaged over the first 500 time steps for the precursor simulation with  $Re_\tau = 1000$  and not simulated in this paper.

$Re_\tau$	Grid	Total CPU-hours	Precursor CPU-hours
350	$1048 \times 764 \times 128$	0.22	NA
500	$1500 \times 1048 \times 256$	0.9	160
2000	$5120 \times 2500 \times 512$	7.4	$\approx 20000$

Table B.4 illustrates the computational cost comparison between the iSTFG and the precursor simulation to obtain the initial conditions. It is clear from computational cost and the convergence time history that the synthetically generated initial condition using the iSTFG provides the most optimal initial condition to spin up the simulation to a statistically stationary state beyond which time averages and higher order moments can be calculated.

### Appendix C. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.compfluid.2025.106733>.

### Data availability

The synthetic turbulence generator can be accessed via the public repository <https://github.com/AkshayPatil1994/Synthetic-Eddy-Method-KCX2013>. All of the data and figures detailed in this paper can be reproduced using the following open-source code repository [https://github.com/AkshayPatil1994/transition\\_data](https://github.com/AkshayPatil1994/transition_data).

### References

- [1] Jiménez J, Moin P. The minimal flow unit in near-wall turbulence. *J Fluid Mech* 1991;225:213–40.
- [2] Kasagi N, Tomita Y, Kuroda A. Direct numerical simulation of passive scalar field in a turbulent channel flow. *J Heat Transf* 1992;114:598–606.
- [3] Moser RD, Kim J, Mansour NN. Direct numerical simulation of turbulent channel flow up to  $Re_\tau=590$ . *Phys Fluids* 1999;11:943–5.
- [4] Marusic I, McKeon BJ, Monkewitz PA, Nagib HM, Smits AJ, Sreenivasan KR. Wall-bounded turbulent flows at high Reynolds numbers: Recent advances and key issues. *Phys Fluids* 2010;22:1–24.
- [5] Jiménez J. Cascades in wall-bounded turbulence. *Annu Rev Fluid Mech* 2012;44:27–45.
- [6] Smits AJ, Marusic I. Wall-bounded turbulence. *Phys Today* 2013;66:25–30.
- [7] Lozano-Durán A, Bae HJ, Encinar MP. Causality of energy-containing eddies in wall turbulence. *J Fluid Mech* 2020;882:A2.
- [8] Jie Y, Cui Z, Xu C, Zhao L. On the existence and formation of multi-scale particle streaks in turbulent channel flows. *J Fluid Mech* 2022;935:A18.
- [9] Horwitz JAK. Estimating the carbon footprint of computational fluid dynamics. *Phys Fluids* 2024;36.
- [10] Vela-Martín A, Encinar MP, García-Gutiérrez A, Jiménez J. A low-storage method consistent with second-order statistics for time-resolved databases of turbulent channel flow up to  $Re_\tau = 5300$ . *J Comput Sci* 2021;56:101476.
- [11] Nelson KS, Fringer OB. Reducing spin-up time for simulations of turbulent channel flow. *Phys Fluids* 2017;29.
- [12] Costa P. A FFT-based finite-difference solver for massively-parallel direct numerical simulations of turbulent flows. *Comput Math Appl* 2018;76:1853–62.
- [13] Graham J, Kanov K, Yang XIA, Lee M, Malaya N, Lalescu CC, et al. A web services accessible database of turbulent channel flow and its use for testing a new integral wall model for LES. *J Turbul* 2016;17(2):181–215. <http://dx.doi.org/10.1080/14685248.2015.1088656>.
- [14] Henningson DS, Kim J. On turbulent spots in plane poiseuille flow. *J Fluid Mech Digit Arch* 1991;228:183.
- [15] Kim J, Moin P. Application of a fractional-step method to incompressible Navier-Stokes equations. *J Comput Phys* 1985;59(2):308–23.
- [16] Ferziger JH, Perić M, Street RL. Computational methods for fluid dynamics. Springer; 2019.
- [17] Kim Y, Castro IP, Xie Z-T. Divergence-free turbulence inflow conditions for large-eddy simulations with incompressible flow solvers. *Comput & Fluids* 2013;84:56–68.
- [18] Pope SB. Turbulent flows. Cambridge University Press; 2000.
- [19] Wu X. Inflow turbulence generation methods. *Annu Rev Fluid Mech* 2017;49(1):23–49.
- [20] Lund TS, Wu X, Squires KD. Generation of turbulent inflow data for spatially-developing boundary layer simulations. *J Comput Phys* 1998;140:233–58.
- [21] Taylor GI. The spectrum of turbulence. *Proc R Soc Lond Ser A- Math Phys Sci* 1938;164(919):476–90.
- [22] Schau KA, Johnson C, Muller J, Oefelein JC. An ensemble synthetic eddy method for accurate treatment of inhomogeneous turbulence. *Comput & Fluids* 2022;248:105671.
- [23] Lozano-Durán A, Jiménez J. Effect of the computational domain on direct simulations of turbulent channels up to  $Re_\tau=4200$ . *Phys Fluids* 2014;26.
- [24] Bernardini M, Pirozzoli S, Orlandi P. Velocity statistics in turbulent channel flow up to. *J Fluid Mech* 2014;742:171–91.
- [25] Perry AE, Lim TT, Teh EW. A visual study of turbulent spots. *J Fluid Mech* 1981;104:387–405. <http://dx.doi.org/10.1017/S0022112081002966>.
- [26] Wu X. New insights into turbulent spots. *Annu Rev Fluid Mech Annu Rev Fluid Mech* 2023 2025;43:55. <http://dx.doi.org/10.1146/annurev-fluid-120720>.
- [27] Sandham N, Kleiser L. The late stages of transition to turbulence in channel flow. *J Fluid Mech* 1992;245:319–48.
- [28] Rabin S, Caulfield C, Kerswell R. Triggering turbulence efficiently in plane Couette flow. *J Fluid Mech* 2012;712:244–72.
- [29] Gabriel E, Fagg GE, Bosilca G, Angskun T, Dongarra JJ, Squyres JM, et al. Open MPI: Goals, concept, and design of a next generation MPI implementation. In: Recent advances in parallel virtual machine and message passing interface: 11th European PVM/MPI users' group meeting Budapest, Hungary, September 19–22, 2004. proceedings 11. Springer; 2004. p. 97–104.
- [30] Lee M, Moser RD. Direct numerical simulation of turbulent channel flow up to  $Re_\tau \approx 5200$ . *J Fluid Mech* 2015;774:395–415. <http://dx.doi.org/10.1017/jfm.2015.268>.
- [31] Luchini P. Universality of the turbulent velocity profile. *Phys Rev Lett* 2017;118(22):224501.