

Identifying eigenmodes of averaged small-amplitude perturbations to turbulent channel flow

A. S. Iyer¹, F. D. Witherden¹, S. I. Chernyshenko¹ and P. E. Vincent^{1,†}

¹Department of Aeronautics, Imperial College London, London SW7 2AZ, UK

(Received 10 August 2018; revised 20 June 2019; accepted 23 June 2019;
first published online 25 July 2019)

Eigenmodes of averaged small-amplitude perturbations to a turbulent channel flow – which is one of the most fundamental canonical flows – are identified for the first time via an extensive set of high-fidelity graphics processing unit accelerated direct numerical simulations. While the system governing averaged small-amplitude perturbations to turbulent channel flow remains unknown, the fact such eigenmodes can be identified constitutes direct evidence that it is linear. Moreover, while the eigenvalue associated with the slowest-decaying anti-symmetric eigenmode mode is found to be real, the eigenvalue associated with the slowest-decaying symmetric eigenmode mode is found to be complex. This indicates that the unknown linear system governing the evolution of averaged small-amplitude perturbations cannot be self-adjoint, even for the case of a uni-directional flow. In addition to elucidating aspects of the flow physics, the findings provide guidance for development of new unsteady Reynolds-averaged Navier–Stokes turbulence models, and constitute a new and accessible benchmark problem for assessing the performance of existing models, which are used widely throughout industry.

Key words: turbulence simulation, turbulence theory, general fluid mechanics

1. Introduction

1.1. Overview

Canonical flows play a special role in fluid dynamics. Their relative simplicity allows deep insight into mechanisms underlying real-world flow phenomena, and they also provide invaluable teaching material. In the laminar regime, canonical flows can often be described by explicit analytical formulae, which are exact solutions of the Navier–Stokes equations. In the turbulent regime some features of canonical flows can also be described by explicit analytical formulae, obtained semi-empirically by combining theoretical considerations with direct observations. The logarithmic velocity profile in a boundary layer is perhaps the most widely known example. The goal of the present study is to identify the slowest-decaying eigenmodes of averaged

† Email address for correspondence: p.vincent@imperial.ac.uk

small-amplitude perturbations to a turbulent channel flow – which is one of the most fundamental canonical flows.

The paper is structured as follows: § 1.2 describes the analogous laminar channel flow case, for which the slowest-decaying eigenmodes can be obtained analytically, § 1.3 describes the turbulent channel flow case and provides a problem specification, § 1.4 reviews previous related work, § 2 details our chosen methodology, § 3 identifies the slowest-decaying eigenmodes and finally § 4 draws conclusions.

1.2. Laminar channel flow case

Consider a plane channel that is infinite in the x - and z -directions, with parallel no-slip walls at $y = \pm 1$. Flow in such a channel is governed by the Navier–Stokes equation and the continuity equation

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u}, \tag{1.1}$$

$$\nabla \cdot \mathbf{u} = 0, \tag{1.2}$$

subject to the boundary condition

$$\mathbf{u}(x, \pm 1, z, t) = 0, \tag{1.3}$$

where $\mathbf{u} = (u, v, w)$ is the velocity, p is the pressure and Re is the Reynolds number.

If Re is low enough such that flow is laminar, a simple unsteady solution of (1.1)–(1.2) can be obtained by imposing an initial condition

$$\mathbf{u}(x, y, z, 0) = (u_{IN}(y), 0, 0), \tag{1.4}$$

that is independent of the wall-parallel coordinates x and z . In this case it is easy to verify that $\mathbf{u} = (u(y, t), 0, 0)$ at all times, the pressure gradient is independent of space and (1.1)–(1.2) reduces to the heat equation

$$\frac{\partial u}{\partial t} = -d + \frac{1}{Re} \frac{\partial^2 u}{\partial y^2}, \tag{1.5}$$

where

$$d = \frac{dp}{dx}. \tag{1.6}$$

To uniquely determine the solution, either the pressure gradient d or the mass flow rate must be prescribed. Assuming the latter, separation of variables leads to

$$u(y, t) = U(y) + u'(y, t), \quad d(t) = D + d'(t), \tag{1.7a,b}$$

where D is the constant pressure gradient associated with the steady state solution $U(y)$, and

$$u'(y, t) = \sum_{n=1}^{\infty} C_n e^{\lambda_n t} u_n(y), \quad d'(t) = \sum_{n=1}^{\infty} C_n e^{\lambda_n t} d_n, \tag{1.8a,b}$$

where C_n are determined by the form of u_{IN} , and $u_n(y)$, d_n and λ_n are solutions of the eigenvalue problem

$$\lambda_n u_n = -d_n + \frac{1}{Re} \frac{d^2 u_n}{dy^2}, \quad u_n(-1) = u_n(1) = 0, \quad \int_{-1}^{+1} u_n dy = 0. \quad (1.9a-c)$$

All eigenvalues of this problem are negative, and hence $u(y, t) \rightarrow U(y)$ as $t \rightarrow \infty$. Moreover the eigenvalues can be ordered such that $\lambda_{n+1} > \lambda_n$ for all n . The slowest-decaying eigenmode is given by

$$u_1(y) = \sin(\pi y), \quad d_1 = 0, \quad \lambda_1 = -\pi^2/Re, \quad (1.10a-c)$$

which is anti-symmetric. The second slowest-decaying eigenmode is given by

$$u_2(y) = \cos(\alpha y) - \cos(\alpha), \quad d_2 = \lambda_2 \cos \alpha, \quad \lambda_2 = -\alpha^2/Re, \quad (1.11a-c)$$

which is symmetric, where $\alpha = 4.4934$ is the smallest solution of

$$\sin(\alpha) = \alpha \cos(\alpha). \quad (1.12)$$

Consequently, the long-time behaviour of the solution for a generic initial condition is typically $u(y, t) \approx U(y) + C_1 e^{\lambda_1 t} u_1(y)$ (assuming $u_1(y)$ is excited by the initial condition), whereas the long-time behaviour for symmetric initial conditions, which would not excite $u_1(y)$, is typically $u(y, t) \approx U(y) + C_2 e^{\lambda_2 t} u_2(y)$ (assuming $u_2(y)$ is excited by the initial condition).

1.3. Turbulent channel flow case

If Re is high enough such that flow is turbulent, then the widely accepted triple decomposition (Hussain & Reynolds 1970) can be applied, and the velocity field can be written as

$$\mathbf{u} = \bar{\mathbf{U}} + \bar{\mathbf{u}}' + \mathbf{u}'', \quad (1.13)$$

where $\bar{\mathbf{U}}$ is a steady component, $\bar{\mathbf{u}}'$ is a coherent component and \mathbf{u}'' is a fluctuating component. Averaging the Navier–Stokes equation (1.1) and the continuity equation (1.2) leads to the unsteady Reynolds-averaged Navier–Stokes equation and continuity equation

$$\frac{\partial \bar{\mathbf{u}}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} = -\nabla \bar{p} + \frac{1}{Re} \nabla^2 \bar{\mathbf{u}} - \nabla \bar{\mathbf{u}}'' \mathbf{u}'', \quad (1.14)$$

$$\nabla \cdot \bar{\mathbf{u}} = 0, \quad (1.15)$$

subject to the boundary condition

$$\bar{\mathbf{u}}(x, \pm 1, z, t) = 0, \quad (1.16)$$

where $\bar{\mathbf{u}} = \bar{\mathbf{U}} + \bar{\mathbf{u}}'$ is the averaged velocity, \bar{p} is the averaged pressure and $\bar{\mathbf{u}}'' \mathbf{u}''$ is called the Reynolds stress tensor. In the case of a plane channel flow the averaging can be done in either of the wall-parallel directions, or over an ensemble of channels. According to the widely accepted ergodic hypothesis these averaging methods and

their combinations are equivalent, provided that the flow is statistically homogeneous in wall-parallel directions.

Analogous to the laminar case, one can attempt to obtain a simple solution by imposing an initial condition

$$\bar{\mathbf{u}}(x, y, z, 0) = (\bar{u}_N(y), 0, 0), \quad (1.17)$$

that is independent of the wall-parallel coordinates x and z . In this case it is easy to verify that $\bar{\mathbf{u}} = (\bar{u}(y, t), 0, 0)$ at all times, the averaged pressure gradient is independent of space and (1.14)–(1.15) reduces to

$$\frac{\partial \bar{u}}{\partial t} = -\bar{d} + \frac{1}{Re} \frac{\partial^2 \bar{u}}{\partial y^2} - \frac{\partial \overline{u''v''}}{\partial y}, \quad (1.18)$$

where

$$\bar{d} = \frac{d\bar{p}}{dx}. \quad (1.19)$$

Since (1.18) contains an unknown Reynolds stress term, it is not closed. We anticipate that the correct, yet unknown, closure exists. The objective of the current study is to ascertain whether in the limit of small-amplitude perturbations (1.18), with the correct closure, has solutions of a form analogous to (1.8), and in particular to identify the slowest-decaying anti-symmetric and symmetric eigenmodes that are analogous to (1.10) and (1.11). However, since the closure for (1.18) is unknown, the system cannot be solved directly. Therefore, the problem will be tackled via an extensive set of high-fidelity graphics processing unit (GPU) accelerated direct numerical simulations (DNS) of channel flow, from which averaged solutions will be obtained and analysed directly.

1.4. Related work

The notion of a small-amplitude perturbation implies two flows are considered, one of which is designated a base flow, and the other as a sum of the base flow and the small-amplitude perturbation. Within this context, the base flow is often considered to be either an evolving time-dependent turbulent flow field, or an averaged flow field.

Analysis of small-amplitude perturbations to time-dependent turbulent base flows usually involves calculation of Lyapunov exponents, characterising the growth rate of the difference between the base flow and a small variation thereof. Recent studies in this area include e.g. Nikitin (2018). When flow is turbulent, small-amplitude perturbations defined in this way usually grow very quickly and then saturate at a finite amplitude, becoming in effect the difference between two unrelated realisations of the same turbulent state. Such small-amplitude perturbations are not considered in the present study.

Analysis of small-amplitude perturbations to averaged base flows began with the seminal work of Malkus (1956), who attempted to build a theory of turbulence based on an assumption that average velocity profiles are marginally stable with respect to infinitesimal perturbations as described by the linearised Navier–Stokes equations, whilst at the same time maximising energy dissipation rate. It was later shown that this assumption is not true (Reynolds & Tiederman 1967). However, the approach inspired further lines of research, including those motivated by the

idea that patterns, also called organised structures, observed in developed turbulent flows, could be explained by studying instabilities of the averaged base flow. Indeed, patterns in a finite-amplitude solution that evolved from a small-amplitude random perturbation are sometimes similar to patterns of the fastest growing infinitesimal disturbance. This is often the case for free-shear turbulent flows, which are strongly affected by the inviscid instability of the inflectional average flow profile (Gaster, Kit & Wygnanski 1985; Roshko 1993), but not for near-wall turbulent flows where viscosity becomes important, since it is often impossible to find the fastest growing infinitesimal disturbance, as all disturbances eventually decay if their evolution is calculated using the molecular viscosity. It was observed, however, that transiently growing solutions of the linearised Navier–Stokes equations exhibit patterns similar to those seen in developed turbulent flows. To investigate this further, the full Navier–Stokes equations can be rewritten as linearised equations forced by nonlinear terms, and then these nonlinear terms can be modelled as statistically stationary stochastic forcing (Farrell & Ioannou 1993), so that the amplitude of the fluctuations remains time independent. The patterns of the fluctuating motion are then a result of the filtering properties of the linearised Navier–Stokes operator i.e. its selectivity in amplifying or suppressing particular patterns. These filtering, or selectivity, properties can be studied directly using techniques known as input–output analysis, as well as indirectly by considering the pseudospectra of the operator (Trefethen *et al.* 1993) or optimal perturbations with maximum transient growth (Butler & Farrell 1993). The development of this philosophy proved fruitful, with results including elucidation of how streak spacing depends on wall distance (Chernyshenko & Baig 2005), how turbulent drag reduction due to spanwise wall oscillations depends on the frequency of the oscillations (Moarref & Jovanović 2012), how the angle of near-wall streaks depends on time in flows with drag reduction caused by wall oscillations (Blesbois *et al.* 2013) and how drag reduction by in-plane wavy motion of the wall depends qualitatively (although not quantitatively) on the frequency and streamwise wavenumber of this motion (Duque-Daza *et al.* 2012). Further progress, during which the above-described paradigm became known as a resolvent framework, was made when nonlinear effects in the form of triadic interactions were considered (see the review by McKeon (2017)). Another way of including nonlinear effects was proposed via a statistical state dynamics approach, in which they are restricted to the interaction of the streamwise average flow and perturbations (see the review by Farrell, Gayme & Ioannou (2017)). The above references serve as an entry point into this thriving area of research. However, we note that it is only indirectly related to the problem considered in the present study.

An alternative line of research stemming from the seminal work of Malkus (1956) has focused on the behaviour of averaged small-amplitude perturbations to averaged base flows. This is directly related to the problem considered in the present study. Following the observation that average turbulent flow profiles are stable with a laminar viscosity, Hussain & Reynolds (1970) introduced the now widely accepted triple decomposition, leading to the unclosed system (1.18). Importantly, it remains unclosed even after linearisation under the assumption of a small coherent component. Davis (1974) investigated deriving turbulent closures for averaged small-amplitude perturbations directly from the Navier–Stokes equations, and Gritsun & Branstator (2007) utilised the general fluctuation-dissipation theorem to gain additional insight. However, a full theoretical solution for the evolution of averaged small-amplitude perturbations is yet to be found. This leaves only three routes for further progress: the use of semi-empirical closures, experiments and DNS. Reynolds & Hussain (1972)

themselves used a simple eddy-viscosity closure, and achieved good agreement with associated experimental data. However, more recent studies by Sen & Veeravalli (1998, 2000), Sen *et al.* (2007), which employed more complicated closures with anisotropic Reynolds stresses, produced very different results. In particular they observed growth as opposed to decay of averaged small-amplitude perturbations, although these results, too, were found to agree with their own experimental data. Such a contradiction highlights the challenges associated with making progress using semi-empirical closures and experiments. Finally, in terms of relevant DNS, Luchini, Quadrio & Zuccher (2006) studied the behaviour of a time-dependent channel flow subject to an impulse force, and more recently Russo & Luchini (2016) investigated a time-independent channel flow, where small perturbations were maintained by a constant body force. In the later study DNS results were found to disagree with those obtained using a turbulent closure, highlighting the fundamental difficulty of identifying the correct closure.

2. Methodology

2.1. Overview

Our approach was based on performing ensembles of turbulent channel flow DNS in a finite channel. Specifically, ensembles of statistically independent snapshots of a fully developed turbulent channel flow were modified with either anti-symmetric or symmetric net-zero mass flux perturbations to their streamwise velocity, that depended solely on the wall-normal direction. The evolution of the ensemble- and wall-parallel averages of these perturbations were then analysed, and as the averaged perturbations became small their slowest-decaying anti-symmetric and symmetric eigenmodes were identified and extracted.

It was assumed that the averaged perturbations exhibit an initial nonlinear transient, followed by a linear decay phase as their amplitude becomes small. In this linear phase, it was further assumed that averaged perturbations can be represented as a sum of eigenmodes, each decaying at different rates, such that eventually after some time t_s the slowest-decaying eigenmode dominates. It is clear that in order to identify this slowest-decaying eigenmode one must study the long-time behaviour of the averaged perturbations when $t > t_s$. However, this poses a challenge. Specifically, the averaged perturbations will inherently include averaging errors, which fluctuate irregularly as a function of time, but remain finite. Hence after some time t_e they are of the same order as the averaged perturbation. If the fluctuations are large enough such that $t_e < t_s$, i.e. the averaging errors dominate the solution before only a single slowest-decaying eigenmode dominates, then it will be impossible to identify such a slowest-decaying eigenmode, even if it exists. The averaging errors can be reduced by averaging over more data. Hence, to overcome this challenge the ensemble of turbulent channel flows must be made large; such that $t_e \gg t_s$. As will be shown, this requires a significant number of DNS to be undertaken, and indeed the approach has only recently become feasible with the emergence of high-fidelity GPU accelerated solvers.

The solver used to undertake the DNS here is described in § 2.2, the set-up for a single channel is described in § 2.3, the ensemble and averaging approach is described in § 2.4 and finally the procedure for identifying the slowest-decaying eigenmode is described in § 2.5.

2.2. Solver

All DNS were undertaken using PyFR (Witherden, Farrington & Vincent 2014; Witherden, Vermeire & Vincent 2015), which is based on the high-order flux

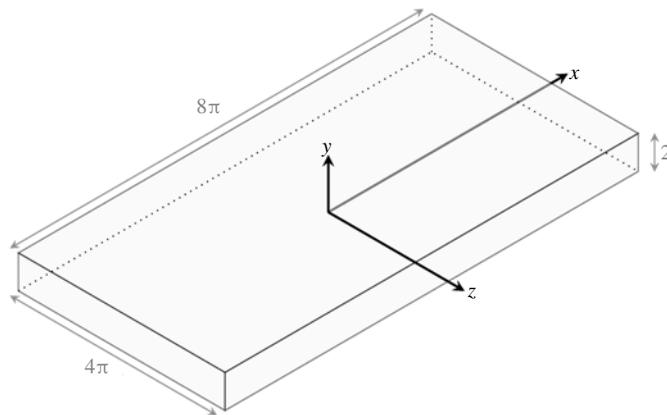


FIGURE 1. Schematic of channel domain.

reconstruction method of Huynh (2007), and can efficiently leverage the capabilities of modern massively parallel GPU platforms (Vincent *et al.* 2016). The compressible Navier–Stokes equations were solved at a Mach number of 0.1 to approximate incompressible conditions. A Rusanov Riemann solver was employed to calculate the inter-element fluxes, an explicit RK45[2R+] scheme (Kennedy, Carpenter & Lewis 2000) was used to advance the solution in time and all runs were performed using double precision arithmetic. Specifically, a patched version of commit ebf386bf4f30e9caaf1cbe8546169f6f200875a1 to the PyFR Git repository (e.g. <https://github.com/vincentlab/PyFR>) was employed. The patch, and the full version of the code, are provided as electronic supplementary material, available at <https://doi.org/10.1017/jfm.2019.520>.

2.3. Set-up

A schematic of the computational domain is shown in figure 1. It has a length of 8π in the x -direction, a height of 2 in the y -direction and a width of 4π in the z -direction. The origin is located at the centre of the domain. Periodic boundary conditions were applied in the x - and z -directions, and an adiabatic no-slip condition was applied at $y = \pm 1$. Flow was driven through the domain via a volumetric momentum source, coupled with a feedback controller to ensure a constant bulk velocity of 1 in the x -direction. The viscosity was set to achieve a target friction Reynolds number of $Re_\tau = 180$, which corresponds to a bulk Reynolds number based on the bulk velocity and channel half-height of $Re = 2767$. Time units were based on the bulk velocity and channel half-height.

The domain was meshed with a structured grid of $62 \times 19 \times 60$ hexahedral elements along the x -, y - and z -directions respectively. The mesh was uniform in the x - and z -directions, but non-uniform in the y direction, with resolution increased towards the channel walls. The flow solution within each element was represented using a fourth-order polynomial. Mesh independence, and comparisons with published direct numerical simulation data, are shown in appendix A and appendix B, respectively. The mesh is provided as electronic supplementary material.

2.4. Ensembles and averages

Consider an ensemble Π of statistically independent snapshots of a fully developed channel flow solution, obtained in the domain defined in §2.3. The x -velocities $u(x, y, z)_e$ from each instance e of Π can be averaged as follows to obtain

$$\bar{U}(y) = \frac{1}{N \times 8\pi \times 4\pi} \sum_{e=0}^N \int_{-4\pi}^{4\pi} \int_{-2\pi}^{2\pi} u(x, y, z)_e \, dx \, dz, \tag{2.1}$$

which in the limit $N \rightarrow \infty$ recovers the exact averaged x -velocity of the fully developed base flow. In the present study $\bar{U}(y)$ was calculated using $N = 768$ snapshots. These were obtained by running a series of independent statistically stationary channel flow simulations, from which solutions were extracted and stored every 4 time units.

Now consider that each snapshot in Π is perturbed with either an anti-symmetric function

$$\bar{u}'_{aN}(y) = 0.3 \sin(\pi y), \tag{2.2}$$

to create an ensemble Π_a of initial conditions with a purely anti-symmetric initial perturbation, or a symmetric function

$$\bar{u}'_{sN}(y) = (0.3 - 0.3297y^2) \cos(3\pi y/2), \tag{2.3}$$

to create an ensemble Π_s of initial conditions with a purely symmetric initial perturbation, where we note that both $\bar{u}'_{aN}(y)$ and $\bar{u}'_{sN}(y)$ have net-zero mass flux, depend solely on the wall-normal direction, are zero at the walls, but are otherwise arbitrary.

Finally consider restarting new simulations at $t = 0$ from the initial conditions in each ensemble $\Pi_{a/s}$. The resulting x -velocities $u_{a/s}(x, y, z, t)_e$ obtained from each instance e of $\Pi_{a/s}$ can be recorded, and their perturbations from the fully developed base flow $\bar{U}(y)$ can be averaged as follows to obtain

$$\bar{u}'_{a/s}(y, t) = \frac{1}{N \times 8\pi \times 4\pi} \sum_{e=0}^N \int_{-4\pi}^{4\pi} \int_{-2\pi}^{2\pi} u_{a/s}(x, y, z, t)_e - \bar{U}(y) \, dx \, dz, \tag{2.4}$$

which in the limit $N \rightarrow \infty$ recovers the exact averaged x -velocity perturbation after a purely anti-symmetric/symmetric initial perturbation. In the present study $\bar{u}'_{a/s}(y, t)$ were each calculated using simulation results from $N = 768$ initial conditions.

It will be demonstrated *a posteriori* that use of $N = 768$ snapshots to calculate $\bar{U}(y)$, and use of simulation results from $N = 768$ initial conditions to calculate $\bar{u}'_{a/s}(y, t)$, is sufficient to suppress averaging errors to the extent that single and hence slowest-decaying anti-symmetric/symmetric eigenmodes can be identified.

2.5. Identification of the slowest-decaying eigenmodes

We sought to identify the eigenvalue λ_a and eigenfunction $u_a(y)$ associated with the slowest-decaying anti-symmetric eigenmode by studying $\bar{u}'_a(y, t)$, since it was generated from a purely anti-symmetric initial perturbation. Similarly, we sought to identify the eigenvalue λ_s and eigenfunction $u_s(y)$ associated with the slowest-decaying symmetric eigenmode by studying $\bar{u}'_s(y, t)$, since it was generated from a purely

symmetric initial perturbation. To begin, the anti-symmetric component of $\bar{u}'_a(y, t)$, denoted $\bar{u}'_{aa}(y, t)$, was extracted. This filtered out any symmetric components possibly generated during the nonlinear transient, and thus facilitated identification of the slowest-decaying anti-symmetric eigenmode. Similarly the symmetric component of $\bar{u}'_s(y, t)$, denoted $\bar{u}'_{ss}(y, t)$, was extracted. This filtered out any anti-symmetric components possibly generated during the nonlinear transient, and thus facilitated identification of the slowest-decaying symmetric eigenmode. Subsequently, a single anti-symmetric/symmetric eigenmode was fitted to a series of time-windowed sections of $\bar{u}'_{aa/ss}(y, t)$, with a view to identifying if sections of $\bar{u}'_{aa/ss}(y, t)$ existed after the nonlinear transient phase, and after all but one of the eigenmodes had decayed, but before averaging errors became significant, where it could be well approximated by a single, and hence slowest-decaying anti-symmetric/symmetric eigenmode. Specifically, fitting was attempted in a series of windows $t \in [t^*, t^* + \Delta t^*]$, for an equispaced distribution of window start times t^* in the range 0 to 120, with candidate eigenvalues $\lambda^*_{a/s}(t^*)$ and ‘shifted’ candidate eigenfunctions $\hat{u}^*_{a/s}(y, t^*)$ associated with a single slowest-decaying anti-symmetric/symmetric eigenmode identified as solutions to the optimisation problem

$$\lambda^*_{a/s}(t^*), \hat{u}^*_{a/s}(y, t^*) = \underset{\lambda, \hat{u}(y)}{\operatorname{argmin}} \epsilon_{a/s}(t^*), \quad (2.5)$$

$$\epsilon_{a/s}(t^*) = \sqrt{\int_{-1}^1 \int_{t^*}^{t^* + \Delta t^*} (\operatorname{Re}\{\hat{u}(y)e^{\lambda(t-t^*)}\} - \bar{u}'_{aa/ss}(y, t))^2 dt dy}, \quad (2.6)$$

and associated relative errors $R^*_{a/s}(t^*)$ obtained as

$$R^*_{a/s}(t^*) = \frac{\min_{\lambda, \hat{u}(y)} \epsilon_{a/s}(t^*)}{\sqrt{\int_{-1}^1 \int_{t^*}^{t^* + \Delta t^*} \bar{u}'_{aa/ss}(y, t)^2 dt dy}}, \quad (2.7)$$

where the ‘shifted’ candidate eigenfunctions were represented by 95 degrees of freedom, using the same piecewise polynomial discretisation as for the DNS simulations, and minimisation was performed using the Levenberg–Marquardt algorithm via the LMFIT package (Newville *et al.* 2014). Note that for technical reasons the time origin of the eigenmode in (2.6) was shifted by t^* prior to fitting, hence why $\hat{u}^*_{a/s}(y, t^*)$ are referred to as ‘shifted’ candidate eigenfunctions. Once obtained, these can be ‘unshifted’ to obtain candidate eigenfunctions

$$u^*_{a/s}(y, t^*) = e^{-\lambda_{a/s} t^*} \hat{u}^*_{a/s}(y, t^*), \quad (2.8)$$

that can be directly compared.

If a single slowest-decaying anti-symmetric/symmetric eigenmode is identifiable then there should exist a range of t^* where:

- (i) The candidate eigenvalues $\lambda^*_{a/s}(t^*)$ are approximately independent of t^* , with a value that can be ascribed to the true eigenvalue $\lambda_{a/s}$.
- (ii) The relative errors $R^*_{a/s}(t^*)$ are small.
- (iii) The candidate eigenfunctions $u^*_{a/s}(y, t^*)$ are approximately independent of t^* , with a form that can be ascribed to the true eigenfunction $u_{a/s}(y)$.

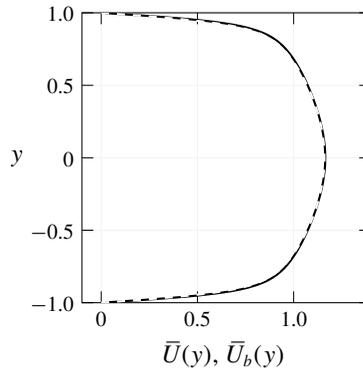


FIGURE 2. Plots of $\bar{U}(y)$, the averaged x -velocity of the unperturbed base flow (black line), and $\bar{U}_b(y)$, obtained by nonlinear least squares fitting of a rational polynomial to $\bar{U}(y)$ (dashed white line).

We note that, strictly, the above approach will identify slowest-decaying anti-symmetric/symmetric eigenmodes excited by the initial perturbations (2.2) and (2.3), respectively. It is highly likely that these are also indeed the slowest-decaying anti-symmetric/symmetric eigenmodes *per se*. However, to confirm, it would be necessary to repeat the procedure using data from additional ensembles with various initial perturbations, and check that the eigenmodes all agree. Given the cost associated with DNS, such a check was not feasible here.

Finally, we comment that by applying a control volume approach in the normal way to the full Navier–Stokes (1.1) and continuity equations (1.2), averaging over the x - and z -directions, and using the fact that the bulk velocity is constant, one can obtain an expression for the averaged pressure gradient perturbation associated with each eigenmode

$$\frac{d\bar{p}_{a/s}}{dx}(t) = \text{Re} \left\{ \frac{e^{\lambda_{a/s}t}}{2\text{Re}} \left. \frac{du_{a/s}(y)}{dy} \right|_{-1}^{+1} \right\}. \tag{2.9}$$

3. Results

3.1. Base flow

Figure 2 plots the averaged x -velocity of the unperturbed base flow $\bar{U}(y)$, along with an analytical fit

$$\bar{U}_b(y) = (1 - y)(1 + y) \left(\frac{0.36y^4 - 1.37y^2 + 1.16}{0.90y^4 - 1.87y^2 + 1} \right), \tag{3.1}$$

which was obtained by nonlinear least squares fitting of a rational polynomial to $\bar{U}(y)$.

3.2. Slowest-decaying anti-symmetric eigenmode

Figure 3 plots $\bar{u}'_{aa}(0.76, t)$, the anti-symmetric component of the averaged x -velocity perturbation, after a purely anti-symmetric initial perturbation, at $y = 0.76$. The evolution is non-oscillatory, which suggests that the slowest-decaying anti-symmetric

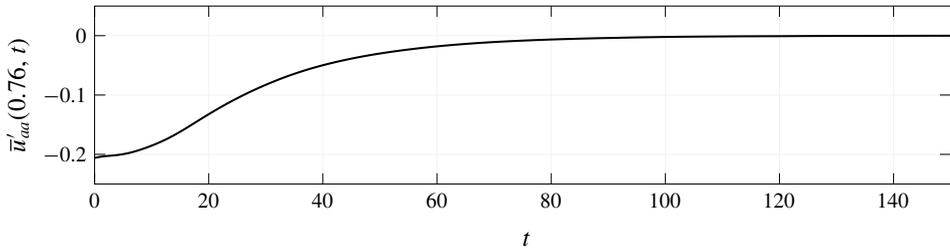


FIGURE 3. Plot of $\bar{u}'_{aa}(0.76, t)$, the anti-symmetric component of the averaged x -velocity perturbation, after a purely anti-symmetric initial perturbation, at $y = 0.76$.

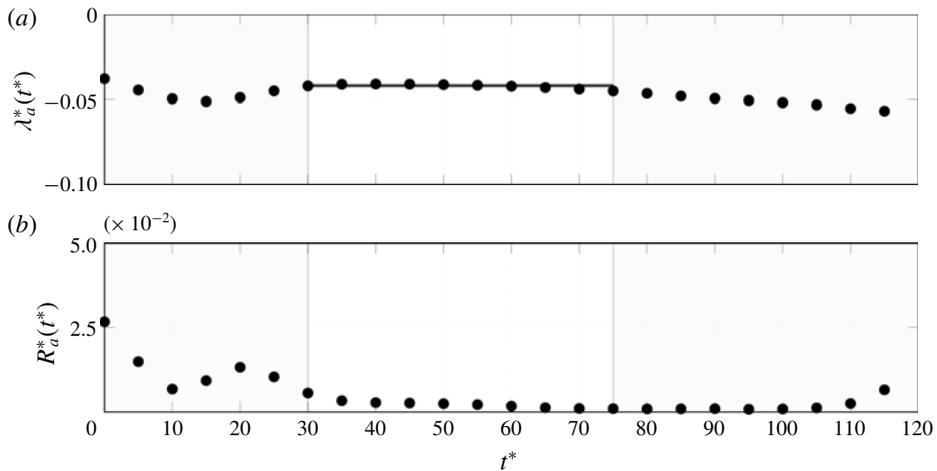


FIGURE 4. Plots of candidate eigenvalues $\lambda_a^*(t^*)$ and relative fitting errors $R_a^*(t^*)$ as a function of window start time t^* , obtained using the procedure outlined in § 2.5 with a window size of $\Delta t^* = 30$.

eigenmode is real valued. Therefore, we assume that the corresponding eigenvalue λ_a and eigenfunction $u_a(y)$ are real valued before proceeding with the analysis.

Figure 4 plots candidate eigenvalues $\lambda_a^*(t^*)$ and relative fitting errors $R_a^*(t^*)$ as a function of window start time t^* , obtained using the procedure outlined in § 2.5 with a window size of $\Delta t^* = 30$. Within $30 \leq t^* \leq 75$ the candidate eigenvalues $\lambda_a^*(t^*)$ are approximately constant and the relative fitting errors $R_a^*(t^*)$ are small. Figure 5 plots

$$\tilde{u}_a^*(y) = \frac{1}{4} \sum_{n=0}^3 \frac{u_a^*(y, 30 + 15n)}{\max_y \{u_a^*(y, 30 + 15n)\}}, \quad (3.2)$$

the average of four normalised candidate eigenfunctions within $30 \leq t^* \leq 75$, obtained using the procedure outlined in § 2.5, and a window size of $\Delta t^* = 30$. Figure 5 also plots their associated standard deviation, which is seen to be small.

We therefore conclude that within the interval $30 \leq t^* \leq 75$ the anti-symmetric component of the averaged x -velocity perturbation after a purely anti-symmetric initial perturbation can be described by a single, and hence slowest-decaying, anti-symmetric eigenmode. The associated eigenvalue can be approximated as

$$\lambda_a = -0.042, \quad (3.3)$$

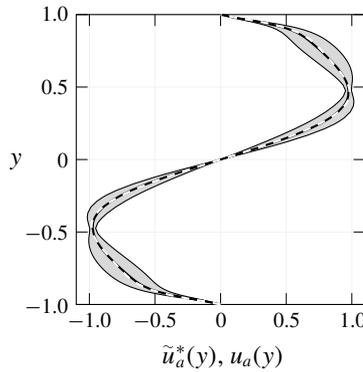


FIGURE 5. Plots of $\tilde{u}_a^*(y)$, the average of four normalised candidate eigenfunctions within $30 \leq t^* \leq 75$, obtained using the procedure outlined in § 2.5, and a window size of $\Delta t^* = 30$ (black line), along with their standard deviation (grey shaded region), and $u_a(y)$, obtained by least squares fitting of a rational polynomial to $\tilde{u}_a^*(y)$ (dashed white line). Note that $\tilde{u}_a^*(y)$ (black line) and $u_a(y)$ (dashed white line) are almost exactly superposed.

obtained from the average of $\lambda_a^*(t^*)$ within the interval $30 \leq t^* \leq 75$, which has a standard deviation of 0.001. The associated eigenfunction $u_a(y)$ can be approximated as

$$u_a(y) = (1 - y)(1 + y) \left(\frac{-5.52y^7 + 11.81y^5 - 9.77y^3 + 3.52y}{0.50y^4 - 1.47y^2 + 1} \right), \quad (3.4)$$

which is also plotted in figure 5, and was obtained by nonlinear least squares fitting of a rational polynomial to $\tilde{u}_a^*(y)$.

As a check, figure 6 compares $\tilde{u}'_{aa}(y, t)$, the anti-symmetric component of the averaged x -velocity perturbation after a purely anti-symmetric initial perturbation with the eigenmode $C_a u_a(y) e^{\lambda_a t}$ within $30 \leq t \leq 75$, where C_a is real and depends on the initial condition, and is chosen here via a nonlinear least squares fit. The contour plots are similar, providing further confidence that a single, and hence slowest-decaying, anti-symmetric eigenmode has been identified.

We note that the shape of the eigenfunction is similar to that of a sine curve, which is the corresponding anti-symmetric laminar solution defined by (1.10). However, the eigenvalue λ_a is found to be 11.8 times bigger than the analogous laminar value obtained from (1.10) with $Re = 2767$. This is in line with expectations since eddy viscosity is typically greater than molecular viscosity.

Figure 7 plots $\tilde{u}'_{as}(0, t)$, the symmetric component of the averaged x -velocity perturbation, after a purely anti-symmetric initial perturbation, at $y = 0$. We note that it is non-zero, with a maximum value that is of the same order as the square of the amplitude of $\tilde{u}'_{aa}(0.76, t)$ shown in figure 3, i.e. the initial purely anti-symmetric perturbation generates a symmetric component. This is in contrast to the laminar case, where an anti-symmetric perturbation will always remain purely anti-symmetric, since the governing equation for a unidirectional laminar flow is linear. It can be shown (see appendix C) that if the evolution of the averaged perturbation is governed by a symmetric linear operator, and the initial perturbation is anti-symmetric, then the averaged perturbation must remain purely anti-symmetric for all time. Our observation of the counter indicates that the full governing equations for the averaged perturbation

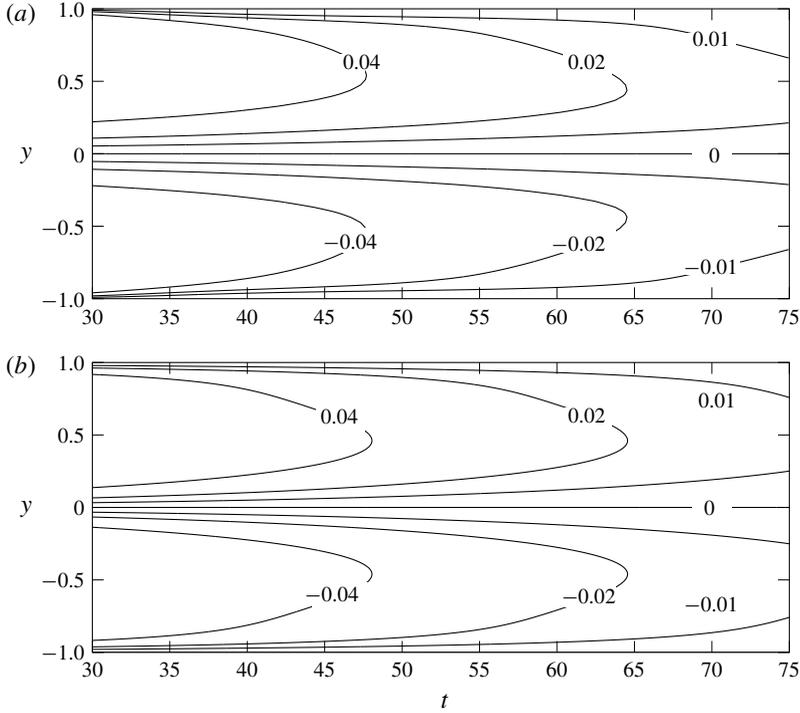


FIGURE 6. Contours of $\bar{u}'_{aa}(y, t)$, the anti-symmetric component of the averaged x -velocity perturbation, after a purely anti-symmetric initial perturbation (a) and the eigenmode $C_a u_a(y) e^{\lambda_a t}$ (b).

must be nonlinear, even for a uni-directional channel flow case. Consequently, any turbulence models should remain nonlinear even for a uni-directional flow, a corollary of which is the requirement that for models employing eddy viscosity, such viscosity cannot be independent of the flow. The well-known Prandtl model, in which the eddy viscosity

$$v_t = l(y) \frac{d(\bar{U} + \bar{u}')}{dy}, \quad (3.5)$$

where $l(y)$ is the mixing length, satisfies this requirement. However, the assumption that $v_t = v_t(y)$, often used in studies of small averaged perturbations, does not. This finding is in line with results from previous studies into the response of turbulent flow to constant forcing (Russo & Luchini 2016; Luchini 2018).

Finally, we note that in accordance with (2.9) the averaged pressure gradient perturbation associated with the anti-symmetric eigenmode is zero, since the eigenfunction $u_a(y)$ is anti-symmetric.

3.3. Slowest-decaying symmetric eigenmode

Figure 8 plots $\bar{u}'_{ss}(0, t)$, the symmetric component of the averaged x -velocity perturbation, after a purely symmetric initial perturbation, at $y = 0$. The evolution is oscillatory, which suggests that the slowest-decaying symmetric eigenmode is

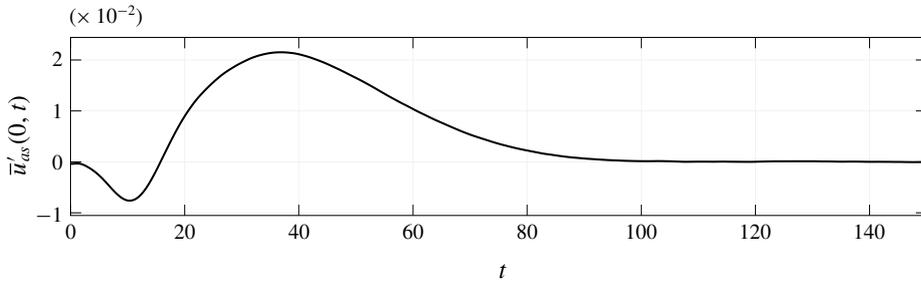


FIGURE 7. Plot of $\bar{u}'_{as}(0, t)$, the symmetric component of the averaged x -velocity perturbation, after a purely anti-symmetric initial perturbation, at $y=0$.

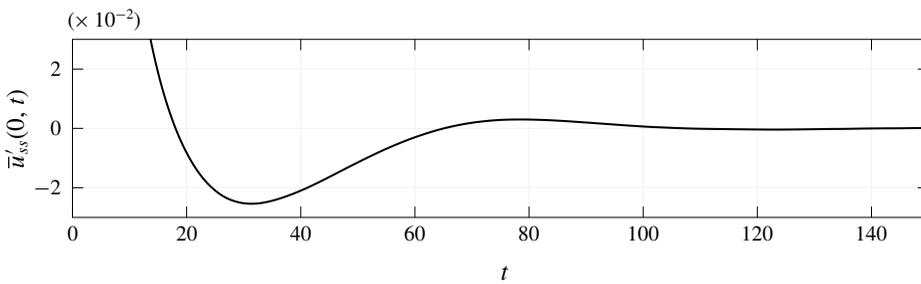


FIGURE 8. Plot of $\bar{u}'_{ss}(0, t)$, the symmetric component of the averaged x -velocity perturbation, after a purely symmetric initial perturbation, at $y=0$.

complex valued. Therefore, we assume that the corresponding eigenvalue λ_s and eigenfunction $u_s(y)$ are complex valued before proceeding with the analysis.

Figure 9 plots the real and imaginary parts of candidate eigenvalues $\lambda_s^*(t^*)$ and relative fitting errors $R_s^*(t^*)$ as a function of window start time t^* , obtained using the procedure outlined in § 2.5 with a window size of $\Delta t^* = 30$. Within $30 \leq t^* \leq 75$ the real and imaginary parts of candidate eigenvalues $\lambda_s^*(t^*)$ are approximately constant and the relative fitting errors $R_s^*(t^*)$ are small. Figure 10 plots the real and imaginary parts of

$$\tilde{u}_s^*(y) = \frac{1}{4} \sum_{n=0}^3 \frac{u_s^*(y, 30 + 15n)}{|u_s^*(0, 30 + 15n)|}, \tag{3.6}$$

the average of four normalised candidate eigenfunctions within $30 \leq t^* \leq 75$, obtained using the procedure outlined in § 2.5, and a window size of $\Delta t^* = 30$. Figure 10 also plots their associated standard deviations, which is seen to be small.

We therefore conclude that within the interval $30 \leq t^* \leq 75$ the symmetric component of the averaged x -velocity perturbation after a purely symmetric initial perturbation can be described by a single, and hence slowest-decaying, symmetric eigenmode. The associated eigenvalue can be approximated as

$$\lambda_s = -0.049 - i0.069, \tag{3.7}$$

obtained from the average of $\lambda_s^*(t^*)$ within the interval $30 \leq t^* \leq 75$, which has a standard deviation of $0.003 + i0.004$. The real and imaginary parts of the associated

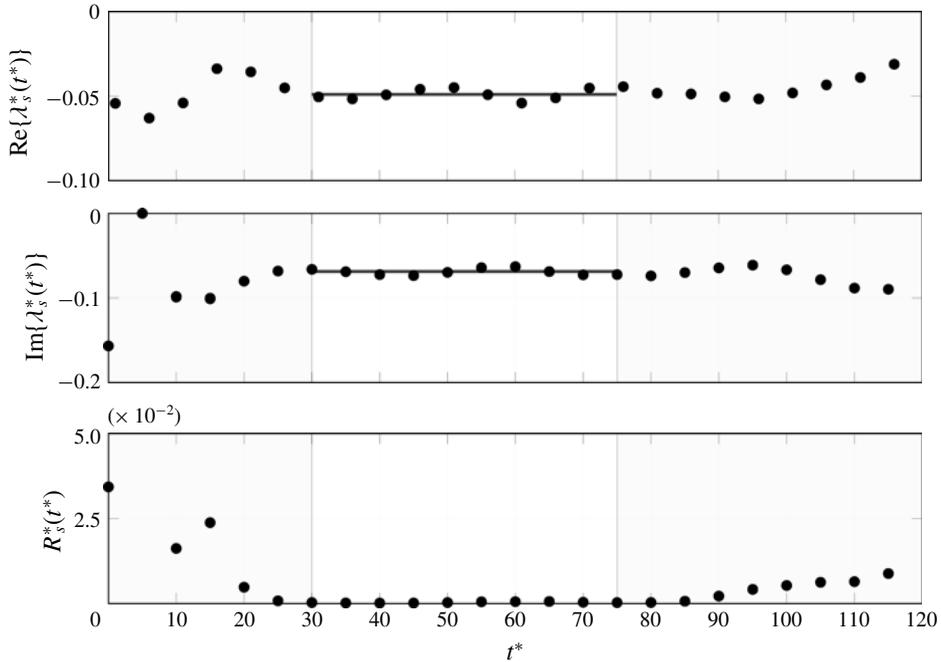


FIGURE 9. Plots the real and imaginary parts of candidate eigenvalues $\lambda_s^*(t^*)$ and relative fitting errors $R_s^*(t^*)$ as a function of window start time t^* , obtained using the procedure outlined in § 2.5 with a window size of $\Delta t^* = 30$.

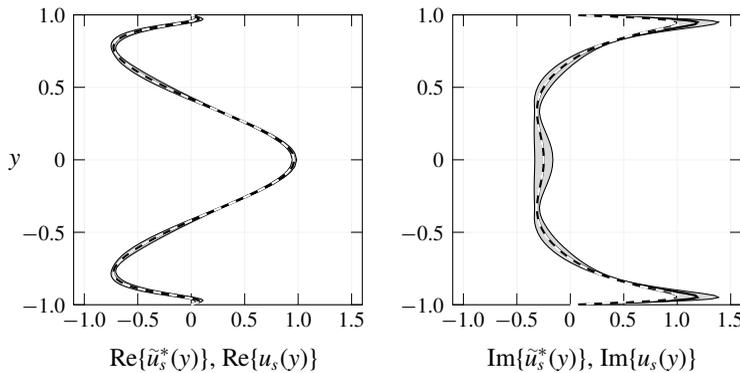


FIGURE 10. Plots of the real and imaginary parts of $\tilde{u}_s^*(y)$, the average of four normalised candidate eigenfunctions within $30 \leq t^* \leq 75$, obtained using the procedure outlined in § 2.5, and a window size of $\Delta t^* = 30$ (black lines), along with their standard deviations (grey shaded regions), and the real and imaginary parts of $u_s(y)$, obtained by least squares fitting of rational polynomials to the real and imaginary parts of $\tilde{u}_s^*(y)$, respectively (dashed white lines). Note that the real and imaginary parts of $\tilde{u}_s^*(y)$ (black lines) and the real and imaginary parts of $u_s(y)$ (dashed white lines) are almost exactly superposed.

eigenfunction $u_s(y)$ can be approximated as

$$\text{Re}\{u_s(y)\} = (1-y)(1+y) \left(\frac{-0.19y^6 + 6.20y^4 - 6.56y^2 + 0.96}{2.75y^6 - 4.23y^4 + 0.60y^2 + 1} \right), \quad (3.8)$$

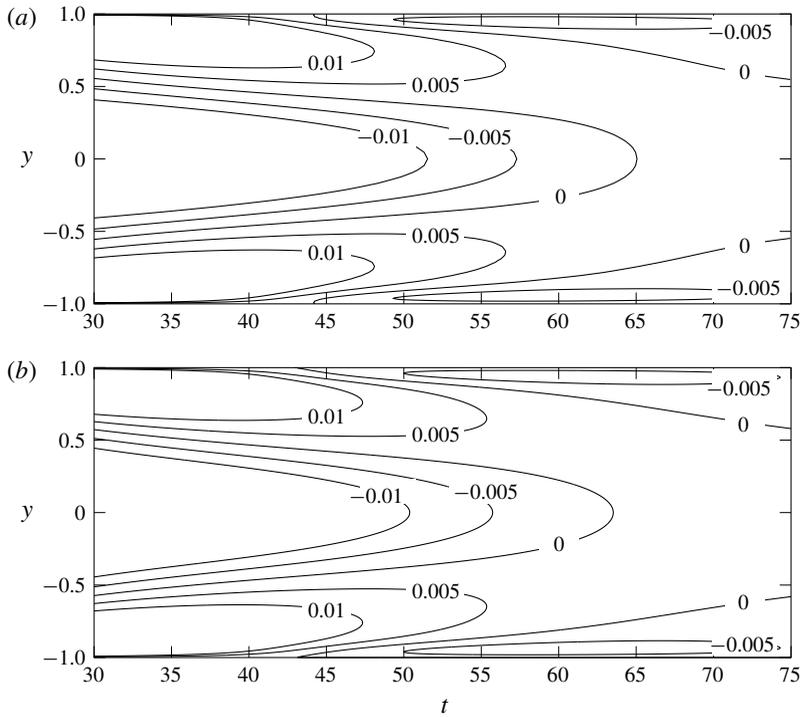


FIGURE 11. Contours of $\bar{u}'_{ss}(y, t)$, the symmetric component of the averaged x -velocity perturbation after a purely symmetric initial perturbation (a) and the real part of the eigenmode $\text{Re}\{C_s u_s(y)e^{\lambda_s t}\}$ (b).

$$\text{Im}\{u_s(y)\} = (1 - y)(1 + y) \left(\frac{3.66y^8 - 8.72y^6 + 7.44y^4 - 1.41y^2 - 0.25}{1.64y^6 - 2.11y^4 - 0.48y^2 + 1} \right), \quad (3.9)$$

which are also plotted in figure 10, and were obtained by nonlinear least squares fitting of rational polynomials to $\text{Re}\{\bar{u}'_s(y)\}$ and $\text{Im}\{\bar{u}'_s(y)\}$, respectively. We note that since the problem is real valued, the complex conjugate of λ_s and $u_s(y)$ will also form a valid eigenmode.

As a check, figure 11 compares $\bar{u}'_{ss}(y, t)$, the symmetric component of the averaged x -velocity perturbation after a purely symmetric initial perturbation with the real part of the eigenmode $\text{Re}\{C_s u_s(y)e^{\lambda_s t}\}$ within $30 \leq t \leq 75$, where C_s is complex and depends on the initial condition, and is chosen here via a nonlinear least squares fit. The contour plots are similar, providing further confidence that a single, and hence slowest-decaying, symmetric eigenmode has been identified.

We note that in contrast to the corresponding symmetric laminar solution defined by (1.11), the eigenvalue and eigenfunction are complex. This implies that the corresponding linear operator cannot be self-adjoint, even in the case of a uni-directional flow. Moreover, unlike for the anti-symmetric perturbation, the shape of the eigenfunctions differ substantially between the laminar and turbulent cases. In particular, the symmetric eigenfunction has a large magnitude approximately 12 wall units from the wall, in a region where turbulent fluctuations are known to attain their maximum value. We can hypothesise that the complex eigenvalue is a result of an interplay between two physical processes, one occurring in the bulk of the flow, and

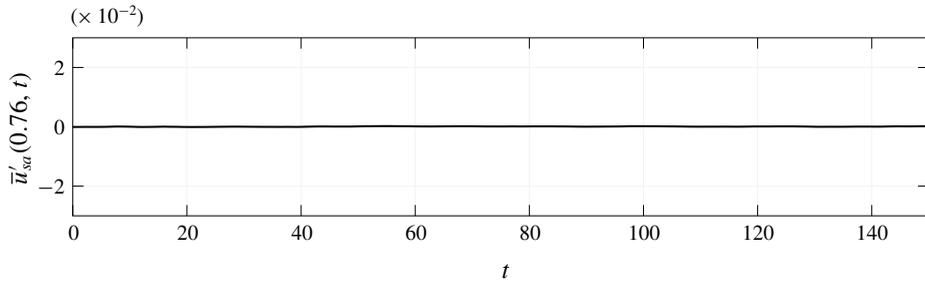


FIGURE 12. Plot of $\bar{u}'_{sa}(0.76, t)$, the anti-symmetric component of the averaged x -velocity perturbation, after a purely symmetric initial perturbation, at $y=0.76$.

another associated with turbulent fluctuations concentrated near the walls. Also, the real part of the eigenvalue $\text{Re}\{\lambda_s\}$ is found to be 6.72 times bigger than the analogous laminar value obtained from (1.11) with $Re = 2767$. This is in line with expectations since eddy viscosity is typically greater than molecular viscosity.

Figure 12 plots $\bar{u}'_{sa}(0.76, t)$, the anti-symmetric component of the averaged x -velocity perturbation, after a purely symmetric initial perturbation, at $y = 0.76$. We note that it is zero i.e. the initial purely symmetric perturbation does not generate an anti-symmetric component. This is expected since the mean flow, the initial perturbation, and the boundary conditions are all symmetric.

Finally, we note that in accordance with (2.9) the averaged pressure gradient perturbation associated with the symmetric eigenmode is

$$\begin{aligned} \frac{d\bar{p}_s}{dx}(t) &= \text{Re}\{e^{\lambda_s t}(0.00247 + i0.0104)\} \\ &= e^{-0.049t}(0.0104 \sin(0.069t) + 0.00247 \cos(0.069t)). \end{aligned} \quad (3.10)$$

4. Conclusion

Eigenmodes of averaged small-amplitude perturbations to a turbulent channel flow – which is one of the most fundamental canonical flows – have been identified for the first time via an extensive set of high-fidelity GPU accelerated DNS. Specifically, for a turbulent channel flow with $Re_\tau = 180$ and a constant bulk velocity, the slowest-decaying anti-symmetric eigenmode was found to have an eigenvalue $\lambda_a = -0.042$ and an eigenfunction $u_a(y)$ approximated by the analytical expression (3.4), and the slowest-decaying symmetric eigenmode was found to have an eigenvalue $\lambda_s = -0.049 - i0.069$ and an eigenfunction $u_s(y)$ approximated by the analytical expressions (3.8) and (3.9). While the system governing averaged small-amplitude perturbations to turbulent channel flow remains unknown, the fact such eigenmodes could be extracted constitutes direct evidence that it is linear. Moreover, given λ_s and $u_s(y)$ were found to be complex this unknown linear system cannot be self-adjoint, even for the case of a uni-directional flow. In addition to elucidating aspects of the flow physics, the findings provide guidance for development of new unsteady Reynolds-averaged Navier–Stokes turbulence models, and constitute a new and accessible benchmark problem for assessing the performance of existing models, which are used widely throughout industry.

Future studies can employ the approach adopted here to identify further eigenmodes of averaged small-amplitude perturbations to turbulent channel flow. These could

Configuration	$n_x \times n_y \times n_z$	p
42-13-40-p4	$42 \times 13 \times 40$	4
52-16-50-p4	$52 \times 16 \times 50$	4
62-19-60-p4	$62 \times 19 \times 60$	4
82-25-80-p4	$82 \times 25 \times 80$	4
62-19-60-p3	$62 \times 19 \times 60$	3
62-19-60-p5	$62 \times 19 \times 60$	5

TABLE 1. Configurations.

include e.g. those associated with higher Reynolds numbers, and those associated with perturbations that vary in wall-parallel directions; with the objective of obtaining a more extensive set of eigenmodes analogous to those given by the Orr–Sommerfeld equation for laminar channel flow.

Acknowledgements

The authors would like to thank the Engineering and Physical Sciences Research Council for their support via an Early Career Fellowship (EP/K027379/1 and EP/R030340/1), the Swiss National Supercomputing Centre for access to Piz Daint and Cambridge University for access to Wilkes.

Supplementary material

Supplementary material is available at <https://doi.org/10.1017/jfm.2019.520>.

Appendix A. Mesh independence

A.1. Overview

Simulations were performed on various meshes, and using various polynomial orders to represent the solution in each element. Full details of each configuration are provided in table 1, where n_x , n_y and n_z are the number of mesh elements in the x , y and z directions respectively, and p is the polynomial order used to represent the solution in each element. All meshes were uniform in the x - and z -directions, but non-uniform in the y -direction, with resolution skewed towards the channel walls.

A.2. Velocity

Figures 13 and 14 show plots of normalised time-averaged x -velocity

$$u^+ = u \sqrt{\frac{\rho}{\tau_w}}, \quad (\text{A } 1)$$

where ρ is fluid density and τ_w is mean wall shear stress, as a function of normalised wall-normal distance

$$y^+ = \frac{\Delta y}{\nu} \sqrt{\frac{\tau_w}{\rho}}, \quad (\text{A } 2)$$

where ν is fluid kinematic viscosity and Δy is wall-normal distance. Figure 13 shows variation with n_x , n_y and n_z for fixed p . The three finest meshes (52-16-50-p4, 62-19-60-p4, 82-25-80-p4) produce visually identical results, indicating mesh independence

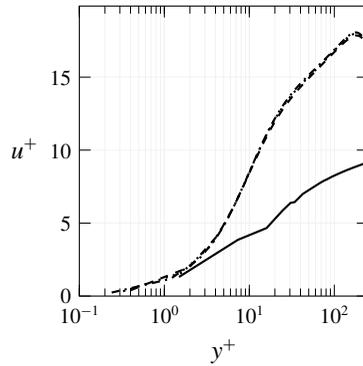


FIGURE 13. Plots of u^+ for configuration 42-13-40-p4 (solid line), 52-16-50-p4 (dashed line), 62-19-60-p4 (dotted line) and 82-25-80-p4 (dash-dotted line).

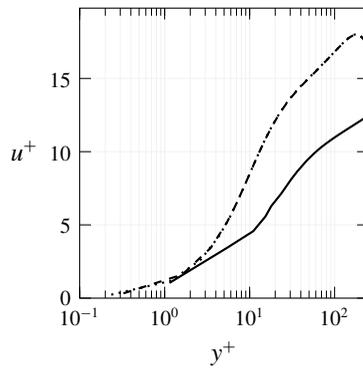


FIGURE 14. Plots of u^+ for configuration 62-19-60-p3 (solid line), 62-19-60-p4 (dashed line) and 62-19-60-p5 (dotted line).

has been achieved. Figure 14 shows variation with p for fixed n_x , n_y and n_z . The two highest polynomial orders (62-19-60-p4, 62-19-60-p5) also produce visually identical results, also indicating that mesh independence has been achieved. One can conclude that the 62-19-60-p4 configuration is sufficient to achieve mesh independence of the velocity.

A.3. Velocity fluctuations

Figures 15 and 16 show plots of x -velocity fluctuations u_{rms} , y -velocity fluctuations v_{rms} and z -velocity fluctuations w_{rms} , as a function of y . Figure 15 shows variation with n_x , n_y and n_z for fixed p . The three finest meshes (52-16-50-p4, 62-19-60-p4, 82-25-80-p4) produce visually identical results, indicating mesh independence has been achieved. Figure 16 shows variation with p for fixed n_x , n_y and n_z . The two highest polynomial orders (62-19-60-p4, 62-19-60-p5) also produce visually identical results, also indicating that mesh independence has been achieved. One can conclude that the 62-19-60-p4 configuration is sufficient to achieve mesh independence of the velocity fluctuations.

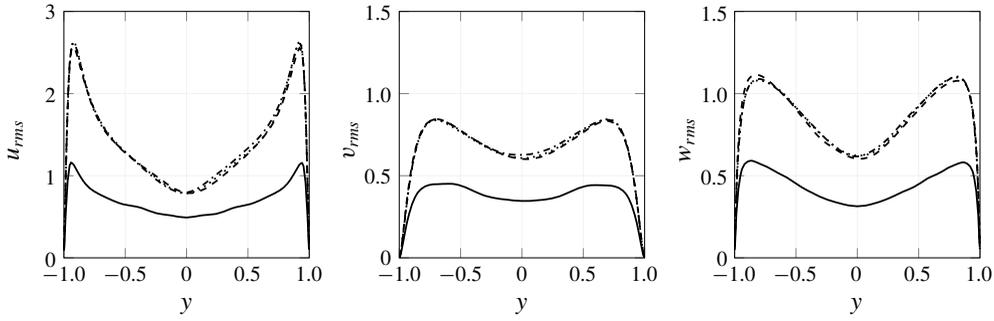


FIGURE 15. Plots of u_{rms} , v_{rms} and w_{rms} for configuration 42-13-40-p4 (solid lines), 52-16-50-p4 (dashed lines), 62-19-60-p4 (dotted lines) and 82-25-80-p4 (dash-dotted lines).

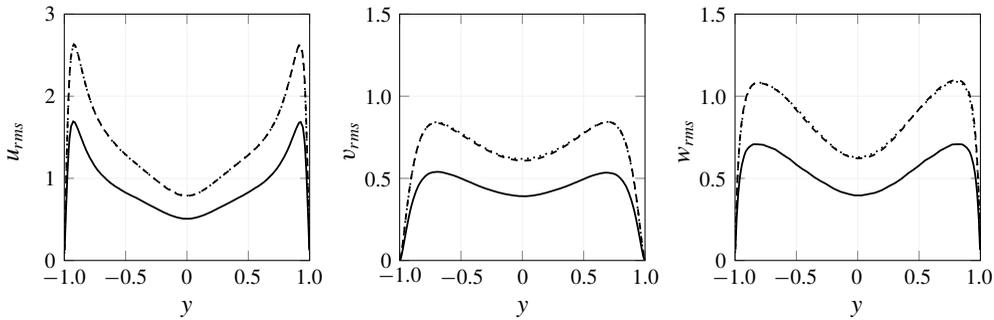


FIGURE 16. Plots of u_{rms} , v_{rms} and w_{rms} for configuration 62-19-60-p3 (solid lines), 62-19-60-p4 (dashed lines) and 62-19-60-p5 (dotted lines).

Appendix B. Comparison with published DNS data

Results obtained using the 62-19-60-p4 configuration were compared with DNS results at the same Re_τ obtained by Kim, Moin & Moser (1987). Figure 17 compares plots of u^+ as a function of y^+ , and u_{rms} , v_{rms} and w_{rms} as a function of y . There is good agreement in the inner, transition and outer layers. This provides further evidence that the 62-19-60-p4 configuration effectively resolves the flow physics.

Appendix C. Persistence of anti-symmetry for linear evolution operators

Consider defining an evolution operator E_t for an ensemble- and x - z -averaged perturbation such that

$$\bar{u}'(t, y) = E_t[\bar{u}'_0(y)], \tag{C 1}$$

where $\bar{u}'_0(y)$ is the initial perturbation, and $\bar{u}'(t, y)$ is the ensemble- x - z averaged perturbation at time t . Given the symmetry of the problem one can reason that

$$\bar{u}'(t, -y) = E_t[\bar{u}'_0(-y)], \tag{C 2}$$

and thus

$$\bar{u}'(t, y) + \bar{u}'(t, -y) = E_t[\bar{u}'_0(y)] + E_t[\bar{u}'_0(-y)]. \tag{C 3}$$

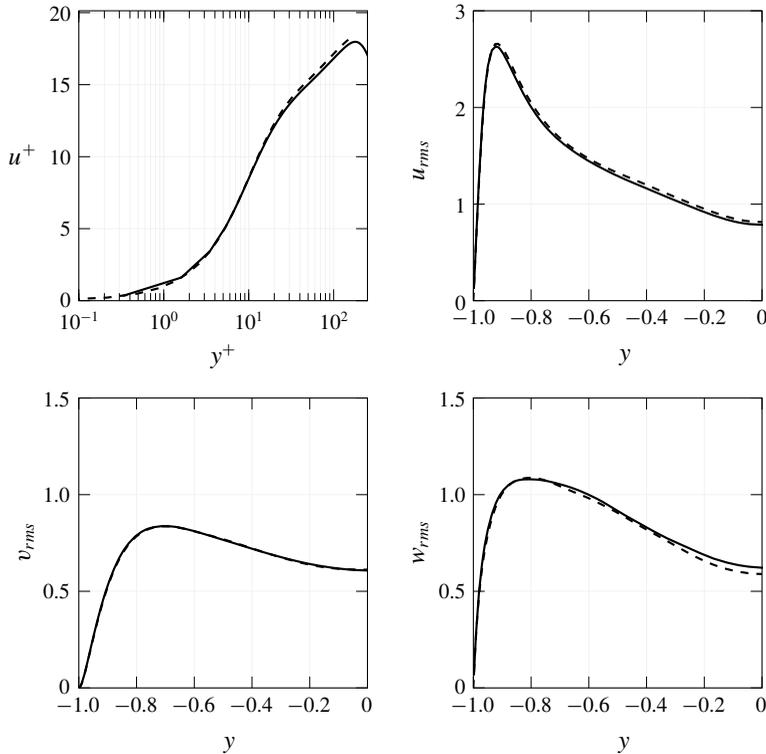


FIGURE 17. Plots of u^+ , u_{rms} , v_{rms} and w_{rms} obtained using the 62-19-60-p4 configuration (solid lines) and from previous DNS by Kim *et al.* (1987) (dashed lines).

If the evolution operator E_t is linear then

$$\bar{u}'(t, y) + \bar{u}'(t, -y) = E_t[\bar{u}'_0(y) + \bar{u}'_0(-y)], \quad (\text{C4})$$

and further, if the initial perturbation is anti-symmetric, $\bar{u}'_0(y) = -\bar{u}'_0(-y)$, then

$$\bar{u}'(t, y) + \bar{u}'(t, -y) = E_t[\bar{u}'_0(y) - \bar{u}'_0(y)] = E_t[0], \quad (\text{C5})$$

and thus since $E_t[0] = 0$ one obtains

$$\bar{u}'(t, y) = -\bar{u}'(t, -y). \quad (\text{C6})$$

Therefore, if E_t is a symmetric linear operator, and the initial perturbation is anti-symmetric, then the ensemble- and x - z -averaged perturbation must remain purely anti-symmetric for all time.

REFERENCES

- BLESBOIS, O., CHERNYSHENKO, S. I., TOUBER, E. & LESCHZINER, M. A. 2013 Pattern prediction by linear analysis of turbulent flow with drag reduction by wall oscillation. *J. Fluid Mech.* **724**, 607–641.

- BUTLER, K. M. & FARRELL, B. F. 1993 Optimal perturbations and streak spacing in wall-bounded turbulent shear flow. *Phys. Fluids A* **5** (3), 774–777.
- CHERNYSHENKO, S. I. & BAIG, M. F. 2005 The mechanism of streak formation in near-wall turbulence. *J. Fluid Mech.* **544**, 99–131.
- DAVIS, R. E. 1974 Perturbed turbulent flow, eddy viscosity and the generation of turbulent stresses. *J. Fluid Mech.* **63** (4), 673–693.
- DUQUE-DAZA, C. A., BAIG, M. F., LOCKERBY, D. A., CHERNYSHENKO, S. I. & DAVIES, C. 2012 Modelling turbulent skin-friction control using linearised Navier–Stokes equations. *J. Fluid Mech.* **702**, 403–414.
- FARRELL, B. F., GAYME, D. F. & IOANNOU, P. J. 2017 A statistical state dynamics approach to wall turbulence. *Phil. Trans. R. Soc. Lond. A* **375**, 20160081.
- FARRELL, B. F. & IOANNOU, P. J. 1993 Stochastic forcing of the linearized Navier–Stokes equations. *Phys. Fluids A* **5** (11), 2600–2609.
- GASTER, M., KIT, E. & WYGNANSKI, I. 1985 Large-scale structures in a forced turbulent mixing layer. *J. Fluid Mech.* **150**, 23–39.
- GRITSUN, A. & BRANSTATOR, G. 2007 Climate response using a three-dimensional operator based on the fluctuation dissipation theorem. *J. Atmos. Sci.* **64**, 2558–2575.
- HUSSAIN, A. K. M. F. & REYNOLDS, W. C. 1970 The mechanics of an organized wave in turbulent shear flow. *J. Fluid Mech.* **41** (2), 241–258.
- HUYNH, H. T. 2007 A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods. *AIAA Paper* 2007-4079, 1–42.
- KENNEDY, C. A., CARPENTER, M. H. & LEWIS, R. M. 2000 Low-storage, explicit Runge–Kutta schemes for the compressible Navier–Stokes equations. *Appl. Numer. Maths* **35** (3), 177–219.
- KIM, J., MOIN, P. & MOSER, R. 1987 Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.* **177**, 133–166.
- LUCHINI, P. 2018 An elementary example of contrasting laminar and turbulent flow physics. [arXiv:1811.11877](https://arxiv.org/abs/1811.11877).
- LUCHINI, P., QUADRIO, M. & ZUCCHER, S. 2006 The phase-locked mean impulse response of a turbulent channel flow. *Phys. Fluids* **18** (12), 121702.
- MALKUS, W. V. R. 1956 Outline of a theory of turbulent shear flow. *J. Fluid Mech.* **1** (5), 521–539.
- MCKEON, B. 2017 The engine behind (wall) turbulence: perspectives on scale interactions. *J. Fluid Mech.* **817**, P1.
- MOARREF, R. & JOVANOVIĆ, M. R. 2012 Model-based design of transverse wall oscillations for turbulent drag reduction. *J. Fluid Mech.* **707**, 205–240.
- NEWVILLE, M., STENSITZKI, T., ALLEN, D. B. & INGARGIOLA, A. 2014 *LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python*. Zenodo. doi:[10.5281/zenodo.11813](https://doi.org/10.5281/zenodo.11813).
- NIKITIN, N. 2018 Characteristics of the leading Lyapunov vector in a turbulent channel flow. *J. Fluid Mech.* **849**, 942–967.
- REYNOLDS, W. C. & HUSSAIN, A. K. M. F. 1972 The mechanics of an organized wave in turbulent shear flow. Part 3. Theoretical models and comparisons with experiments. *J. Fluid Mech.* **54** (2), 263–288.
- REYNOLDS, W. C. & TIEDERMAN, W. G. 1967 Stability of turbulent channel flow, with application to Malkus's theory. *J. Fluid Mech.* **27** (2), 253–272.
- ROSHKO, A. 1993 Instability and turbulence in shear flows. In *Theoretical and Applied Mechanics 1992*, pp. 1–24. Elsevier.
- RUSSO, S. & LUCHINI, P. 2016 The linear response of turbulent flow to a volume force: comparison between eddy-viscosity model and DNS. *J. Fluid Mech.* **790**, 104–127.
- SEN, P. K. & VEERAVALLI, S. V. 1998 On the behaviour of organised disturbances in a turbulent boundary-layer. *Sadhana* **23** (2), 167–193.
- SEN, P. K. & VEERAVALLI, S. V. 2000 Behaviour of organised disturbances in fully developed turbulent channel flow. *Sadhana* **25** (5), 423–437.
- SEN, P. K., VEERAVALLI, S. V., CARPENTER, P. W. & JOSHI, G. 2007 Organised structures in wall turbulence as deduced from stability theory-based methods. *Sadhana* **32** (1–2), 51–64.

- TREFETHEN, L. N., TREFETHEN, A. E., REDDY, S. C. & DRISCOLL, T. A. 1993 Hydrodynamic stability without eigenvalues. *Science* **261** (5121), 578–584.
- VINCENT, P., WITHERDEN, F., VERMEIRE, B., PARK, J. S. & IYER, A. 2016 Towards green aviation with python at petascale. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, p. 1. IEEE Press.
- WITHERDEN, F. D., FARRINGTON, A. M. & VINCENT, P. E. 2014 PyFR: an open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach. *Comput. Phys. Commun.* **185** (11), 3028–3040.
- WITHERDEN, F. D., VERMEIRE, B. C. & VINCENT, P. E. 2015 Heterogeneous computing on mixed unstructured grids with PyFR. *Comput. Fluids* **120**, 173–186.