J. Fluid Mech. (2007), *vol.* 584, *pp.* 69–102. © 2007 Cambridge University Press
doi:10.1017/S0022112007006301 Printed in the United Kingdom

Recurrence of travelling waves in transitional pipe flow

R. R. KERSWELL[1] AND O. R. TUTTY[2]

[1]Department of Mathematics, University of Bristol, Bristol BS8 1TW, UK

[2]School of Engineering Sciences, University of Southampton, Southampton, SO17 1BJ, UK

(Received 30 October 2006 and in revised form 26 February 2007)

The recent theoretical discovery of families of unstable travelling-wave solutions in pipe flow at Reynolds numbers lower than the transitional range, naturally raises the question of their relevance to the turbulent transition process. Here, a series of numerical experiments are conducted in which we look for the spatial signature of these travelling waves in transitionary flows. Working within a periodic pipe of $5D$ (diameters) length, we find that travelling waves with low wall shear stresses (lower branch solutions) are on a surface in phase space which separates initial conditions which uneventfully relaminarize and those which lead to a turbulent evolution. This dividing surface (a separatrix if turbulence is a sustained state) is then minimally the union of the stable manifolds of all these travelling waves. Evidence for recurrent travelling-wave visits is found in both $5D$ and $10D$ long periodic pipes, but only for those travelling waves with low-to-intermediate wall shear stress and for less than about 10 % of the time in turbulent flow at $Re = 2400$. Given this, it seems unlikely that the mean turbulent properties such as wall shear stress can be predicted as an expansion solely over the travelling waves in which their individual properties are appropriately weighted. Instead the onus is on isolating further dynamical structures such as periodic orbits and including them in any such expansion.

1. Introduction

Wall-bounded shear flows are of great practical importance, yet their transition to turbulence is still poorly understood. Typically, the laminar-flow solution is linearly stable (e.g. plane Couette flow, pipe flow) or if linearly unstable, only well beyond the regime where transition occurs (e.g. channel flow). As a result, transition is an abrupt process triggered by a perturbation of sufficient amplitude. Generically, this could be expected to lead to an intermediate state of reduced symmetry, but, in fact, the flow always immediately becomes temporally and spatially complicated. A new direction in rationalizing this phenomenon revolves around identifying alternative solutions (beyond the laminar state) to the governing Navier–Stokes equations. In the past few years, such solutions in the form of steady states or travelling waves have been found in plane Couette flow (Nagata 1990; Clever & Busse 1992, 1997; Waleffe 1998), channel flow (Itano & Toh 2001; Waleffe 2001, 2003), an autonomous wall flow (Jiménez & Simens 2001), and most recently, pipe flow (Faisst & Eckhardt 2003; Wedin & Kerswell 2004). Invariably, these solutions are saddle points in phase space. The idea is that through their stable and unstable manifolds and the interconnections between them together with the appearance of periodic orbits, phase space becomes sufficiently complicated to support 'turbulent' trajectories.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

Gathering supporting evidence for this scenario is in its infancy, especially for spatially extended systems, but progress is being made. In reduced numerical models and spatially confined systems, Eckhardt and coworkers have reported statistical evidence to suggest that these saddles collectively produce a chaotic repellor, at least at low Reynolds numbers (Schmiegel & Eckhardt 1997; Schmiegel 1999; Eckhardt *et al.* 2002; Faisst & Eckhardt 2004; Eckhardt *et al.* 2007). Other work has concentrated on establishing connections between specific flow behaviour and underlying nonlinear solutions present. In channel flow, for example, Itano & Toh (2001) interpret wall turbulent 'bursting' events with flow along the unstable manifold of a travelling-wave solution. They also managed to isolate a periodic-looking solution on the basin boundary of the turbulence by continually adjusting a trajectory such that it neither relaminarized nor became turbulent (Toh & Itano 2003) (see Skufca, Yorke & Eckhardt 2006 for an equivalent calculation in a model system). Jiménez *et al.* (2005) studied both channel flow and plane Couette flow in an effort to relate near-wall turbulent events to the large number of known nonlinear solutions. They concluded that the turbulence stayed close to the upper-branch travelling waves, as far as comparing simple statistics of the flow field such as maximum (over space) wall-normal and streamwise components were concerned. Another approach has focused upon identifying isolated periodic solutions directly from numerically integrated turbulent trajectories using a Newton–Raphson technique. In both the case of plane Couette flow (Kawahara & Kida 2001) and highly-symmetric forced box turbulence (van Veen, Kida & Kawahara 2006), the authors claim to find one periodic orbit which seems to share the same mean properties as the turbulent attractor.

In pipe flow, the only work so far aiming to establish the physical relevance of the recently discovered travelling waves, has been experimental (Hof *et al.* 2004, 2005). By analysing the flow structure across turbulent pipe flow (both of 'puff' and 'slug' type – see Wygnanski & Champagne 1973), good correspondence was found, at least occasionally, with the outer symmetrically arranged ring of fast 'streaks' (streamwise velocity anomalies) which is one of the dominant features of the travelling waves. The match is less clear, however, with regard to the complementary slow streaks centred around the pipe axis as well as with the smaller cross-stream velocities (e.g. figures 2E and 2F of Hof *et al.* 2004). The purpose of this paper is to build on this work by carrying out a detailed quantitative study which can explore how closely the travelling waves are reproduced or 'visited' in phase space and the frequency of such visits using direct numerical simulations. The decision as to whether the flow has 'visited' a travelling wave is made by individually comparing the spatial structure of the instantaneous flow with that of each of the full set of travelling waves currently known to exist in the system. A concurrent numerical study by Schneider, Eckhardt & Vollmer (2007) has chosen instead to focus on identifying coherent streak structures near the outer pipe wall by examining the azimuthal correlation of the downstream velocity there. Their results on how often the flow adopts a coherent streak structure near the wall provide a nice complementary upper estimate on how often we find that travelling waves are visited. If it emerges that turbulent pipe flow can be understood as an effectively random switching between the neighbourhoods of these travelling waves, then an appropriately weighted expansion across the 'active' travelling waves visited may provide a useful predictor of the turbulent flow properties. This presumes that some version of periodic orbit theory developed in low-dimensional dynamical systems (e.g. Cvitanovic 1988; Artuso, Aurell & Cvitanovic 1990*a*, *b*) may carry over to this very high (formally infinite) dimensional setting.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

The structure of the paper is as follows. Section 2 begins by briefly describing the numerical method used to solve the Navier–Stokes equations before discussing reasonable measures chosen to quantify if, and how well, the flow approaches a travelling wave (TW) solution. There is a certain amount of arbitrariness in this choice because the TWs are fully nonlinear solutions, not obviously orthogonal under any inner product. Hence, some experimentation has been necessary before a final choice on the exact 'correlation' functions to evaluate has been made. Given also that the TWs are parameterized continuously by their axial wavelength (albeit over a finite range), it has been convenient to impose a strict periodicity in the pipe to discretize the TWs which can exist in the system. As a result, a periodic pipe of length $5D$ has been used for the majority of the results. Even then, 37 TWs of two-, three- and four-fold rotational symmetry about the axis can be found at a Reynolds number of 2400. These are briefly described in § 3, together with their stability. In § 4, we show numerical evidence that some of the TWs are visited, but not all and for only part of the time. In § 5, the statistical frequency of these visits is quantified by examining the correlation data from across a number of runs. Finally, a discussion follows in § 6.

2. Formulation

2.1. *Numerics*

The Navier–Stokes equation and solenoidal condition for the flow of an incompressible Newtonian fluid along a circular straight pipe under the action of an imposed pressure gradient are

$$\partial_t \boldsymbol{u} + \boldsymbol{u}\cdot\nabla\boldsymbol{u} + \frac{1}{\rho}\nabla p = \nu\nabla^2\boldsymbol{u}, \qquad \nabla\cdot\boldsymbol{u} = 0, \tag{2.1}$$

where ν is the kinematic viscosity, p the pressure and ρ the constant density. Non-dimensionalizing the system using U the mean axial speed and the pipe radius $D/2$, where D is the diameter, gives rise to the Reynolds number $Re := UD/\nu$. A constant mass flow flow rate – or equivalently Re – is maintained along the pipe at all times. A numerical solution for the primitive variables (velocity and pressure) was developed in cylindrical coordinates (s, θ, z) using finite differences in the radial direction (s) and Fourier modes for z and θ. The time stepping was performed using the third-order Runge–Kutta scheme of Nikitin (2006). Hereinafter a quoted numerical resolution of (N, M, K) corresponds to $N+1$ equally-spaced radial points (i.e. a grid step of $1/N$ where $0 \leqslant s \leqslant 1$) and Fourier expansions in θ and z of wavenumbers $-M/2, \ldots, M/2$ and $(-K/2, \ldots, K/2)\pi/L$, respectively, where L diameters is the nominal length over which periodicity is imposed. As discussed below, the main choice of pipe length was $L = 5$: for this geometry, a coarse grid was $(25, 32, 30)$, an intermediate grid $(50, 48, 40)$ and a fine grid $(50, 60, 60)$: equivalent grids in longer pipes were used (e.g. a fine $10D$-grid was $(50, 60, 120)$). Commonly in studies of this type, the grid is stretched in the radial direction as the highest resolution is required near the wall. However, here there was little difference between the results obtained with uniform and non-uniform grids. This is not surprising given the nature of the flow at the Reynolds number considered in this study, as can be seen in figure 3.

In addition to checks of specific components of the code using analytic test solutions, a series of calculations was performed using as initial conditions a TW solution plus a perturbation in the form of the leading (unstable) eigenfunction. Good agreement was obtained for the growth of the disturbance and that predicted from the eigenvalue.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

Importantly, the code was also cross-validated with another time-stepping code based on velocity potentials (Willis & Kerswell 2007).

2.2. *Travelling waves*

The TWs so far identified (Faisst & Eckhardt 2003; Wedin & Kerswell 2004) are arranged into symmetry classes of m-fold rotational symmetry about the axis and then continuously parameterized by their axial wavenumber across a finite range. TWs of one-fold to six-fold symmetry have been found (Kerswell 2005), but only two-, three- and four-fold TWs are currently known to exist below $Re = 2485$. Within each symmetry class, the TWs appear through saddle node bifurcations so that close to the saddle node point there is a well-defined upper- and lower-branch solution for a given wavenumber. For higher Re, the solution surfaces typically kink and fold back on themselves so that multiple pairs of branches can coexist at the same wavenumber (e.g. figure 10 of Wedin & Kerswell 2004).

Imposing a pipe periodicity, immediately reduces the continuum of TWs present down to a discrete number which can fit into the pipe. This is a crucial simplification which means the matching procedure adopted below can monitor all the TWs available to the flow. The main choice of a $5D$ long pipe was a compromise between the need to keep the number of TWs to a manageable size (helped by a shorter pipe) and the need to have a dynamical system which could support turbulent behaviour at a value of Re where the TWs are fully resolvable (helped by a longer pipe). This pipe length has also been studied before (Eggels *et al.* 1994; Faisst & Eckhardt 2004) and sustained turbulence predicted for $Re > 2250$ (Faisst & Eckhardt 2004). This short length, of course, precludes capturing turbulent spatiotemporal features such as 'puffs' (Wygnanski & Champagne 1973) which typically extend over $20D$, but does allow an examination of 'temporal' turbulence which, when triggered, fills the whole pipe.

Figures 1 and 2 show the result of tracing out all the two-, three- and four-fold solution branches in the friction factor–axial wavelength plane at $Re = 2400$ where the friction factor (Schlichting 1968) is defined as

$$\Lambda := -\frac{1}{\rho}\frac{\mathrm{d}p}{\mathrm{d}z}\Big/\frac{U^2}{2D}, \tag{2.2}$$

with $\mathrm{d}p/\mathrm{d}z$ being the mean pressure gradient ($\Lambda_{lam} := 64/Re$ is the laminar value). The curves are similar but less contorted at $Re = 2000$ (Kerswell 2005) and $Re = 2200$ (not shown). The vertical dotted lines drawn at axial wavenumbers $\alpha = 0.625n$ where $n = 1, 2, \ldots, 5$ indicate the TWs which fit in the pipe which is actually taken to be $\pi/0.625 = 5.0265$ diameters long. The m-fold symmetry class, the letter label and wavenumber are used to identify the TWs in what follows. For example, the TW with wavenumber 1.25 and lowest friction factor in figure 1 is the $3b_1.25$ TW. For this pipe geometry and Re, there are 37 TWs (6 two-fold, 22 three-fold and 9 four-fold rotationally symmetric TWs) which can be numerically resolved and used to match against the flow.

2.3. *Matching*

As fully nonlinear solutions, the TWs do not possess any simple orthogonality with respect to an inner product. Therefore establishing when a directly numerically simulated flow, $\boldsymbol{u}_{DNS}$, approaches a chosen TW velocity field, $\boldsymbol{u}_{TW}$, is not a straightforward case of projection. Given this, a number of *ad hoc* 'correlation' functions were developed and tested to measure how close the flow comes momentarily

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

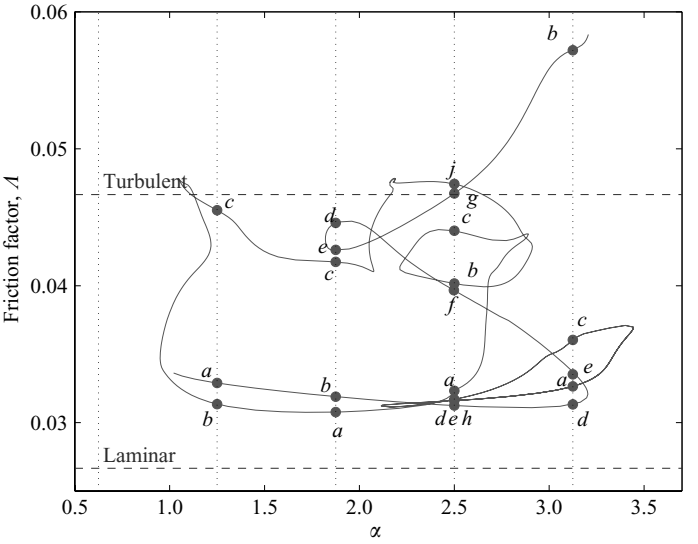


FIGURE 1. Solution branch for three-fold rotationally symmetric travelling waves plotted on a friction factor versus axial wavenumber plot at $Re = 2400$. The lower dashed line represents the lower bound given by the Hagen–Poiseuille solution ($\Lambda_{lam} = 64/Re$) and the upper dashed line corresponds to the $Re = 2400$ value of the log-law parameterization of experimental data $1/\sqrt{\Lambda} = 2.0\log(Re_m\sqrt{\Lambda}) - 0.8$ (see Schlichting 1968, equation (20.30)). The solution branch is shown only as far as it is assured to be resolved (hence the loose ends: the main mapping resolution was (8, 30, 6) in the truncation nomenclature of Wedin & Kerswell 2004). The dotted vertical lines indicate the wavenumbers ($\alpha = 0.625n$ in units of $2/D$, $n = 1, 2, 3, 4, 5$) which fit into a pipe of length $\pi/0.625\,D$ long. The letters are used to label each allowable TW together with the wavenumber.

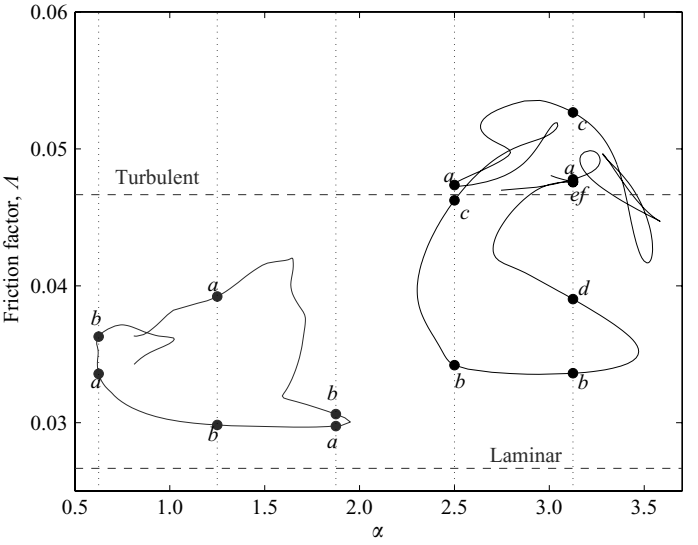


FIGURE 2. The equivalent of figure 1 but for the two-fold (left) and four-fold (right) rotationally symmetric travelling waves. Typical truncations used to resolve this solutions were (9, 25, 7) for the two-fold case and (6, 40, 5) for the four-fold case.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

to a TW. The approach was to construct an integral based on the velocity fields $\boldsymbol{u}_{DNS}$ and $\boldsymbol{u}_{TW}$ over one wavelength $2\pi/\alpha$ of the TW near the middle of the pipe. The correlation functions found most useful and thereby adopted were a normalized inner product based on the total velocity fields,

$$I_{tot}(t) = \max_{\theta_0, z_0} \left[\frac{\langle \boldsymbol{v}_{DNS}, \boldsymbol{v}_{TW} \rangle}{\sqrt{\langle \boldsymbol{v}_{DNS}, \boldsymbol{v}_{DNS} \rangle} \sqrt{\langle \boldsymbol{v}_{TW}, \boldsymbol{v}_{TW} \rangle}} \right], \tag{2.3}$$

once the mean profile of the TW

$$\overline{w}_{TW}(s) := \frac{\alpha}{4\pi^2} \int_0^{2\pi/\alpha} \int_0^{2\pi} w_{TW}(s, \theta, z)\, \mathrm{d}\theta\, \mathrm{d}z \tag{2.4}$$

had been subtracted from both velocity fields

$$\begin{aligned} \boldsymbol{v}_{DNS} &:= \boldsymbol{u}_{DNS}(s, \theta_0 + \theta, z_0 + z) - \overline{w}_{TW}(s)\widehat{z}, \\ \boldsymbol{v}_{TW} &:= \boldsymbol{u}_{TW}(s, \theta, z) - \overline{w}_{TW}(s)\widehat{z}, \end{aligned}$$

and an inner product using the cross-stream velocity components

$$I_{uv}(t) = \max_{\theta_0, z_0} \left[\frac{\langle \boldsymbol{u}^{\perp}_{DNS}, \boldsymbol{u}^{\perp}_{TW} \rangle}{\sqrt{\langle \boldsymbol{u}^{\perp}_{DNS}, \boldsymbol{u}^{\perp}_{DNS} \rangle} \sqrt{\langle \boldsymbol{u}^{\perp}_{TW}, \boldsymbol{u}^{\perp}_{TW} \rangle}} \right]. \tag{2.5}$$

Here, $\boldsymbol{u}^{\perp} = (u, v, 0)$ is the cross-stream velocity part of $\boldsymbol{u}$, and

$$\langle \boldsymbol{u}_1, \boldsymbol{u}_2 \rangle := \frac{\alpha}{2\pi^2} \int_0^{2\pi/\alpha} \int_0^{2\pi} \int_0^1 \boldsymbol{u}_1 \cdot \boldsymbol{u}_2\, s\, \mathrm{d}s\, \mathrm{d}\theta\, \mathrm{d}z \tag{2.6}$$

is the (usual) inner product. The phase optimization over θ_0 and z_0 is necessary because rotated and translated travelling waves are still solutions owing to the symmetries of the flow.

One obvious alternative measure would be to use the velocity fields minus the laminar flow, as is commonly done in investigations of this type (including results presented below). However, it has the major disadvantage that it can produce large values of I_{tot}, even when the flows do not match in any real sense. For example, figure 13 shows I_{tot} for a match to TW $2a_1.25$ for a run starting from TW $3b_3.125$. The three-fold symmetry of the starting condition persists for some time (until $t \approx 130D/U$), and during this initial phase the correlation for the match to TW $2a_1.25$ with two-fold symmetry is low (below 0.1), as might be expected. However, if the velocity minus the laminar value was used for the correlation, then I_{tot} would have a value of around 0.7 throughout this initial stage. Further, once the initial symmetry has broken, the match to $2a_1.25$ produces values of I_{tot} of around 0.6, as shown in figure 13. Using the laminar profile as a reference produces correlations of around 0.9. The relatively high value of the correlation coefficient during the initial phase, even though the flows have different symmetries, can be explained by the shape of the characteristic mean profile for the TWs. In common with turbulent flow in general, these show an increase in the streamwise velocity near the wall, and a decrease in the centre of the pipe (see Wedin & Kerswell 2004, figure 22) compared with the laminar profile. The high value of the correlation coefficient when using the laminar profile as the reference reflects this commonality in shape. Adopting the mean profile of the TW as reference avoids these misleadingly high correlations, but results in what may appear to be a relatively low value of I_{tot} even when there is a good match between the DNS flow and the TW.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

However, I_{tot} is not in itself a sufficient measure for a match. Even when the mean profile of the TW has been subtracted, the streamwise (streak) velocities are typically an order of magnitude larger than the cross-stream velocities and hence their matching contribution tends to dominate I_{tot}, to the extent that there is little difference between I_{tot} and I_w (the equivalent of I_{uv} but using the streamwise velocity component only). Hence, the introduction of the second measure I_{uv}.

By design, I_{tot} and I_{uv} can only take values in the interval $[-1, 1]$ with a value 1 indicating a perfect match. The phase optimization over θ_0 and z_0 (carried out by systematically evaluating all the options over the $[0, 2\pi) \times [-\pi/\alpha, \pi/\alpha)$ grid) in practice ensured that the correlations were never very negative, typically lying in the interval $[-0.2, 0.2]$. Experience indicated that there is evidence for a TW visit if I_{tot} and I_{uv} reach values of 0.5 and above (although more on this below).

A number of other measures were tried. An inner product using the streamwise vorticity, which provides a single measure of the cross-stream flow, was found to shadow I_{uv}, although at a lower level. The symmetry assumed for the TWs (see (2.7) below) ensures that the cross-stream velocity is zero along the axis of the pipe. This would not be expected to occur over an extended length of the pipe in fully turbulent flow. Hence, an inner product over a subset of the domain excluding the central portion could be suitable. This was investigated, and it was found that restricting the domain to $1/2 \leqslant s \leqslant 1$ produced somewhat higher correlations, particularly for I_{uv}, but, again, a similar pattern of behaviour. Hence, these other measures could be used to produce essentially the same results by adjusting the level of correlation that would be regarded as giving a good match.

Correlation functions such as those introduced here measure shape only, and not amplitude. In theory, it would be possible to have a high correlation between velocity fields, but a significant difference in the values of the velocity. The value of the perturbation kinetic energy and the mean wall shear stress can be used to ensure that $\boldsymbol{u}_{DNS}$ and $\boldsymbol{u}_{TW}$ are similar, or to choose the best match when there is more than one candidate. Some of the TWs are highly correlated (values greater then 0.95 have been observed), so that, at a specific time, using these measures, there can be a good match of the flow to more than one TW. One such case will be considered below.

The matching was performed by maximizing the value of I_{uv} over all possible values of θ_0 and z_0. I_{tot} was then calculated for the same orientation. I_{tot} is not suitable as the primary measure as it is so heavily dominated by the streamwise component that the cross-stream structure of the flow would, in effect, be discounted when choosing the 'best' match.

2.4. *Travelling wave stability*

An obvious way to start the DNS runs is to use a TW together with some small perturbation as an initial condition. If this perturbation is unstructured, for example, by relying on numerical discretization errors, the flow takes a long time to exit the neighbourhood of the TW and wastes CPU time. A better strategy is to find the unstable eigendirections of the TW and to use the most unstable eigenfunction with a small amplitude as the perturbation. Four 'lower' branch TWs – 2b_1.25, 3a_2.5, 3h_2.5 and 4b_3.125 – and four 'upper' branch TWs – 2a_1.25, 3b_3.125, 3j_2.5 and 4c_3.125 – were selected as starting TWs. The distinction between upper or lower branch solutions can be ambiguous when the solution surface is as convoluted as in figure 1. Here, we consider that TWs with high friction factors are upper-branch solutions and those with low friction factors are lower, branch solutions. The eight choices made, represent extreme and therefore unambiguous examples under this categorization.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

Branch	TW	Number of unstable eigenvalues	Largest growth rate (in units of U/D)	Resolution $(\mathscr{M}, \mathscr{N}, \mathscr{L})$
Lower	$2b_1.25$	$1r$	1.1×10^{-1}	$(9, 25, 7)$
	$3a_2.5$	$2r$	2.1×10^{-1}	$(8, 30, 6)$
	$3h_2.5$	$3r+2c$	1.7×10^{-1}	$(8, 30, 6)$
	$4b_3.125$	$1r+2c$	2.6×10^{-1}	$(6, 40, 5)$
Upper	$2a_1.25$	$4c$	6.6×10^{-2}	$(9, 25, 8)$
	$3b_3.125$	$2c$	3×10^{-3}	$(8, 30, 5)$
	$3j_2.5$	$6c$	1.7×10^{-1}	$(8, 30, 6)$
	$4c_3.125$	$6c$	3.7×10^{-1}	$(6, 40, 5)$

TABLE 1. The stability properties of typical upper and lower travelling waves at $Re = 2400$: r and c indicate the number of real and complex eigenvalues, respectively. The resolution $(\mathscr{M}, \mathscr{N}, \mathscr{L})$ is the same for the travelling waves and the stability calculation and indicates the azimuthal, radial and axial resolution, respectively (see Wedin & Kerswell 2004 for details). The unstable eigenvalues all correspond to disturbances possessing the same shift-and-reflect symmetry as the travelling wave.

The stability properties of the eight chosen TWs are given in table 1. The travelling waves all possess the shift-and-reflect symmetry $\mathscr{S}$

$$\mathscr{S} : (s, \theta, z) \to (s, -\theta, z + \pi/\alpha), \qquad \mathscr{S} : (u, v, w, p) \to (u, -v, w, p), \tag{2.7}$$

so permitted linear disturbances can be partitioned into those either symmetric or antisymmetric with respect to $\mathscr{S}$. When checked, the TWs were invariably stable to antisymmetric disturbances so table 1 concentrates exclusively on the situation in the $\mathscr{S}$-symmetric subspace. The number of unstable directions is strikingly small given the large degrees of freedom involved (e.g. $O(15\,000)$) and the size of the growth rates – $O(0.1\ U/D)$ indicate inertial instabilities. Figure 3 shows the structure of the most unstable eigenfunctions for the lower branch TWs $2b_1.25$ and $3a_2.5$. The $3a_2.5$ TW gives a particularly clear example of how an unstable eigenfunction is concentrated in the regions of maximum shear in the TW streak velocity.

3. Results

A series of runs were performed by taking as initial conditions each of the selected eight upper and lower TWs perturbed by a small amount of their most unstable eigenfunction. Since this perturbation can be added or subtracted, 16 runs were in fact done. This protocol highlighted a fundamental difference between the upper- and lower-branch TWs. For all the lower branch TWs tested, starting the run in one sense along the TWs most unstable manifold invariably led to an uneventful gradual relaminarisation, whereas starting in the other sense always produced a turbulent evolution (see figure 4). Both signs of perturbation, in contrast, produced a turbulent trajectory for the upper-branch TWs. This implies that the four lower-branch TWs (and by implication other lower-branch TWs) and their stable manifolds are part of a boundary dividing regions of phase space which lead to the two different types of behaviour. At least part of their unstable manifolds are normal to this surface, directed towards either the laminar or turbulent states. The fact that lower-branch solutions may be embedded in such a dividing surface has been suggested before

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

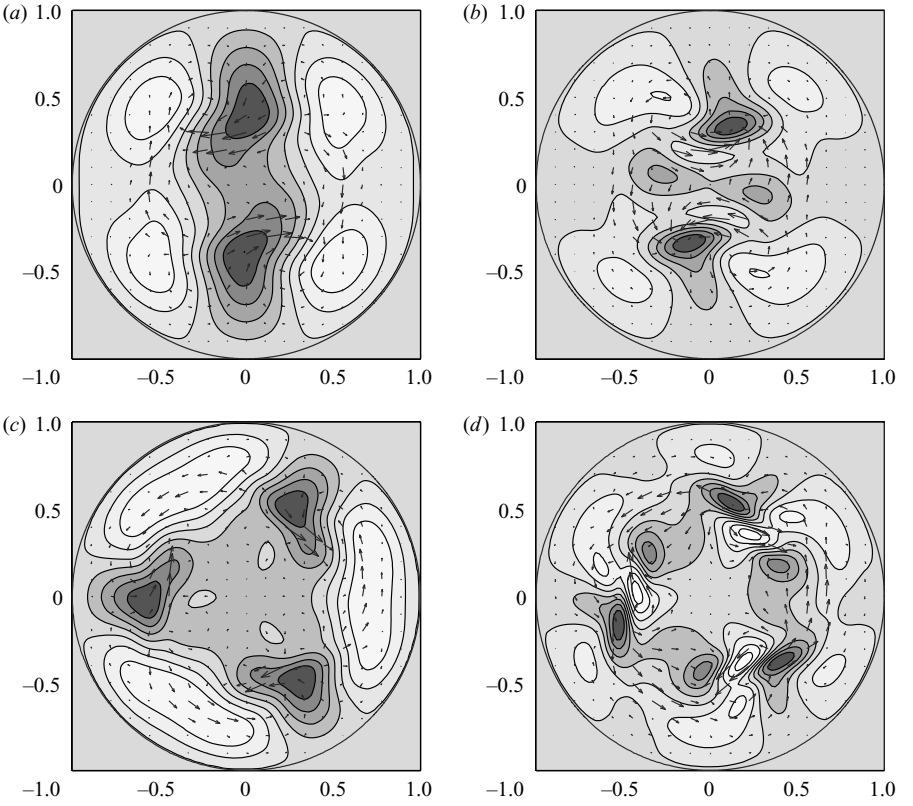


FIGURE 3. The travelling waves 2b_1.25 (*a*) and 3a_2.5 (*c*) and their most unstable eigenfunctions ((*b*) for 2b_1.25 and (*d*) for 3a_2.5). The arrows indicate the cross-stream velocities (larger arrows corresponding to larger speeds) for all four plots. The shading represents the axial velocity for the eigenfunctions whereas for the travelling waves, the axial velocity differential away from the laminar flow corresponding to the same mass flux is plotted (dark most negative – slow streaks; and light most positive – fast streaks). The same contour levels are used throughout with the eigenfunctions renormalized so that their largest axial velocity is set to the largest absolute contour level (the shading outside the pipe indicates 0: contours levels range from $-0.417U$ to $0.266U$ in 8 steps).

(e.g. Waleffe 2003) despite only one known example in channel flow (Itano & Toh 2001). The results here provide a more systematic verification of this idea. The surface – formally a separatrix if the turbulence is a sustained state – is undoubtedly more than just a union of lower-branch TWs and their stable manifolds. In plane Couette flow, for example, Kawahara (2005) has found a periodic orbit embedded in the separatrix.

In the current situation, this dividing surface has deliberately not been called a 'separatrix' as the second initial observation when doing these runs at $Re = 2400$ with a $5D$ pipe is that the turbulence is only transitory. In other words, the laminar state is still the global attractor at $Re = 2400$ in a $5D$ periodic pipe although the flow can experience a long but ultimately finite turbulent episode. Not surprisingly, the length of these turbulent transients seems to depend sensitively on the exact numerical resolution used. This is particularly true with regard to the azimuthal and axial resolution – reducing this too far can produce what looks to be sustained turbulence since the excitation of small scales and the ensuing enhanced dissipation

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

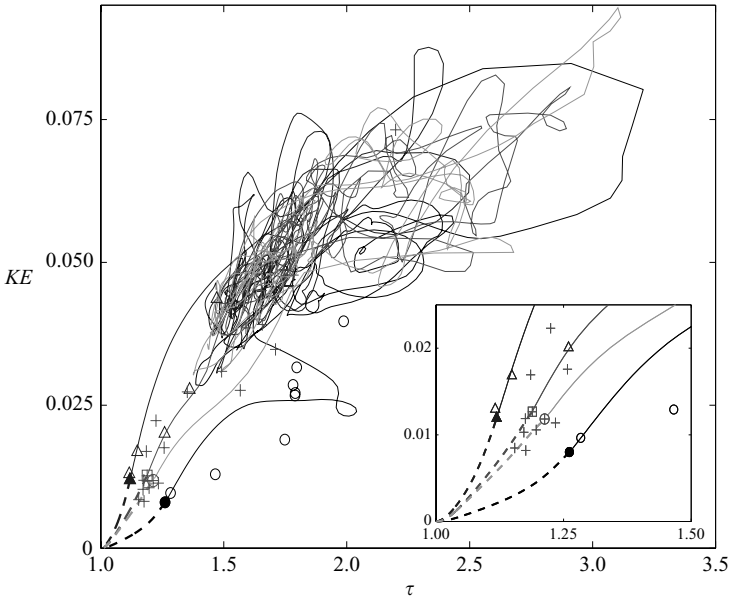


FIGURE 4. The disturbance kinetic energy per unit mass (see (3.1)), in units of U^2 versus wall shear stress τ (see (3.2)) in units of $-8\rho U^2/Re$ for $\boldsymbol{u}_{DNS}$ starting at the four lower branch TWs considered: $2b_1.25$ (▲), $3a_2.5$ (⊞), $3h_2.5$ (⊕) and $4b_3.125$ (●). The solid line indicates the turbulent evolution for one sign of the eigenvalue perturbation for each TW and the thick dashed line traces out the uneventful relaminarization for the other. The laminar state is represented by the point (1, 0). All the TWs present are also plotted: △, two-fold TWs; +, three-fold TWs; ○, four-fold TWs. The inset is a blow-up near the laminar point at (1, 0) to highlight the relaminarization.

which tends to relaminarize the flow are suppressed (Orszag & Kells 1980). We observed that turbulence remains transient when reducing the resolution from the working resolution (50, 60, 60) to (25, 30, 30) or (50, 24, 24), but looks to be sustained at (50, 16, 16) over a time greater than $3000D/U$: the resolution of Faisst & Eckhardt (2004) which predicts sustained turbulence at $Re = 2250$ is somewhere in between these last two choices. The issue of exactly when (or indeed if) pipe flow turbulence becomes sustained is an area of much current interest (Hof *et al.* 2006; Peixinho & Mullin 2006; Lagha & Manneville 2007; Willis & Kerswell 2007).

The two runs started around the lower branch TW $4b_3.125$ illustrate the general behaviour well. Figure 5 plots the disturbance kinetic energy (per unit mass)

$$KE := \frac{1}{2\pi L}\int_0^{2L}\int_0^{2\pi}\int_0^1 \tfrac{1}{2}(\boldsymbol{u}-\boldsymbol{u}_{lam})^2\, s\, \mathrm{d}s\, \mathrm{d}\theta\, \mathrm{d}z \tag{3.1}$$

of the flow versus the mean wall shear stress

$$\tau := \frac{1}{2\pi LD}\int_0^{2L}\int_0^{2\pi} \rho\nu \left.\frac{\partial w}{\partial s}\right|_{s=1} \mathrm{d}\theta\, \mathrm{d}z \quad \left[= -\tfrac{1}{8}\rho U^2\, \Lambda = \frac{\Lambda}{\Lambda_{lam}}\,(-8\rho U^2/Re)\right] \tag{3.2}$$

for the flow evolution together with all 37 of the TWs present. Strictly, when comparing with values for a TW, the kinetic energy and mean wall shear should be calculated for the section of the pipe containing the best match. However, for a pipe of this length, flow structures tend to persist over the full length of the pipe, particularly in the near-wall region containing the streaks. The wall shear stress

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

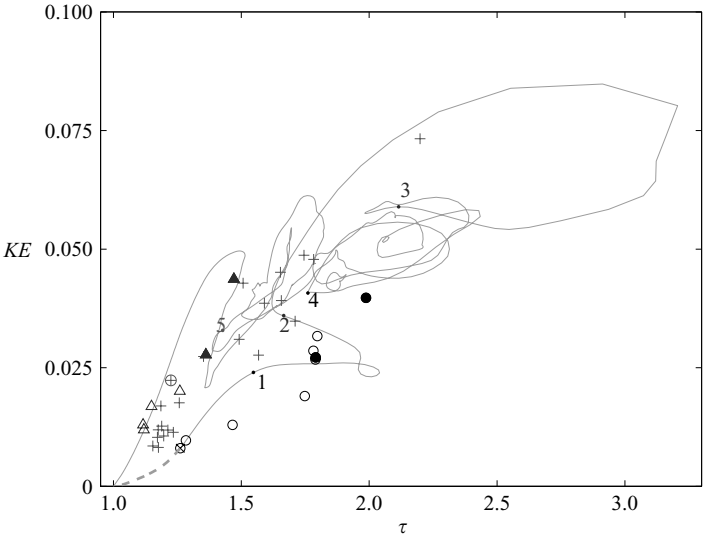


FIGURE 5. The disturbance kinetic energy per unit mass, in units of U^2 versus wall shear stress τ in units of $-8\rho U^2/Re$ for $\boldsymbol{u}_{DNS}$ starting at TW $4b_3.125$ ($\oplus$). The solid line indicates the turbulent evolution for one sign of the eigenvalue perturbation and the thick dashed line traces out the uneventful relaminarization for the other (respectively, $4b_3.125(+/-)$). The laminar state is represented by the point (1, 0). All the TWs present are also plotted: $\triangle$, two-fold TWs; $+$, three-fold TWs; $\circ$, four-fold TWs. Filled symbols indicate TWs which appear to be visited by $\boldsymbol{u}_{DNS}$ and the numbered dots correspond to the times of visits as determined by examining I_{tot} and I_{uv}. In chronological order: 1 – $4f_3.125$ (lower $\bullet$); 2 – $2b_0.625$ ($\blacktriangle$ *furthest* from 3); 3 – $2a_1.25$ (closer $\blacktriangle$); 4 – $4c_3.125$ (upper $\bullet$) and 5 – $3a_3.125$ ($\oplus$).

depends only on the streamwise velocity, which also provides the dominant component of the kinetic energy perturbation. As a result, there is usually little difference in these quantities for the full pipe and a section corresponding to one of the travelling waves. Larger differences were observed in longer pipes.

For one sign of the eigenfunction perturbation, the flow tamely relaminarizes whereas for the other it executes a long turbulent transient. During this latter evolution, there is evidence of close visits to at least five TWs. The first (labelled '1' in figure 5) is to $4f_3.125$ and occurs during the early stages as the flow trajectory moves away from $4b_3.125$ (see figure 6). The quality of this match is extremely high, suggesting that there may be a heteroclinic connection between the two TWs (see also figure 4 which shows the trajectory passing straight over the open circle representing $4f_3.125$). The fact that even at $t = 0$ there is already a considerable correlation with $4f_3.125$ indicates that $4b_3.125$ is structurally similar to $4f_3.125$, an observation which is also true for some other groupings of TWs within the same rotational symmetry class. Figure 6 also shows the correlation signal for $4c_3.125$ which has a very close visit after 120 D/U (labelled '4'). Taken together, the correlation functions for the other four-fold TWs indicate that the flow retains its four-fold symmetry until about 130 D/U (this was verified by examining the transient solution), whereupon it switches to a predominantly three-fold symmetry. Figure 7 shows this switch-over well via the correlation function I_{tot} for $3a_3.125$. This TW is the best candidate for a close visit (labelled '5' at $t \sim 250$ D/U) over all the three-fold symmetric TWs. Examining the instantaneous flow field at point 5 (figure 8) shows clear evidence of three equally spaced fast streaks around the outside of the pipe like $3a_3.125$ which are fairly

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

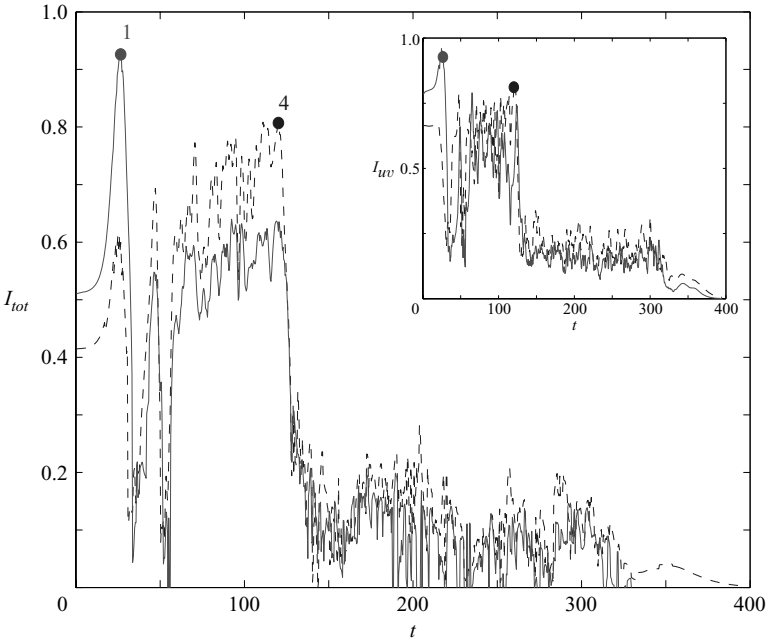


FIGURE 6. I_{tot} and I_{uv} (inset) as a function of time for TWs $4f_3.125$ (solid line) and $4c_3.125$ (dashed line) starting at TW $4b_3.125$. Dots label the times of probable closest visits to each TW (times '1' and '4' coincide with those in figure 5).

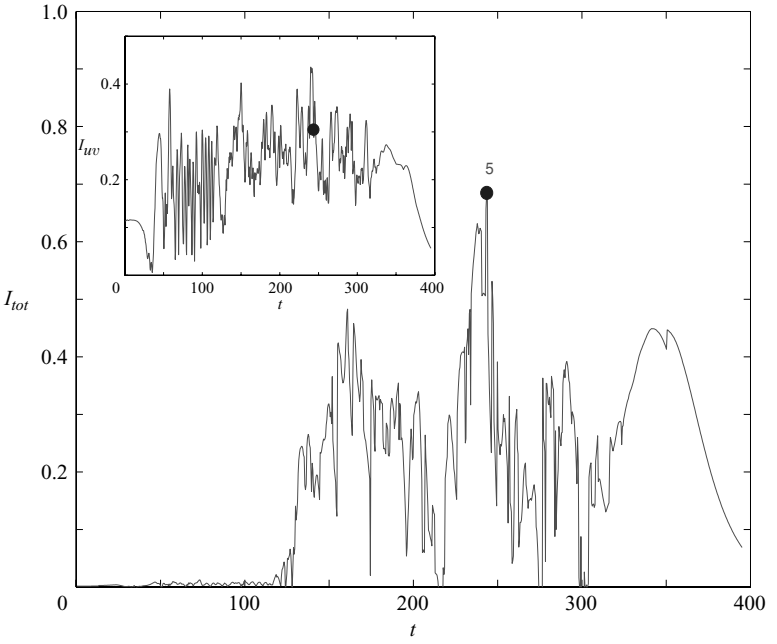


FIGURE 7. I_{tot} and I_{uv} (inset) as a function of time for TW $3a_3.125$ starting at TW $4b_3.125$. The dot labels the time of closest visit (corresponding to '5' in figure 5).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

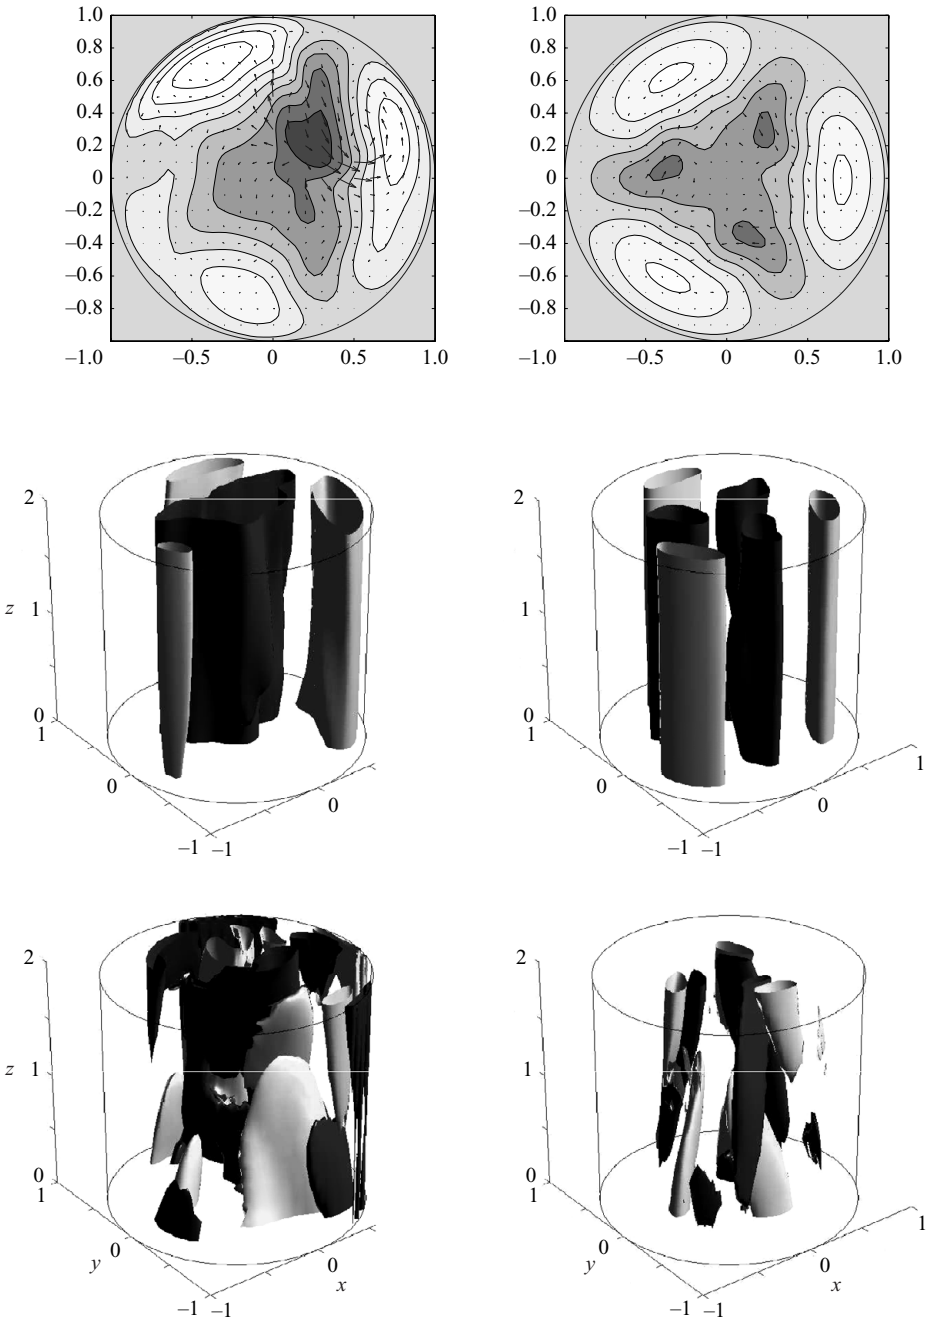


FIGURE 8. Comparison plots of the DNS flow (left-hand column) and the TW $3a_3.125$ (right-hand column) at point 5 in figure 7. The top row shows the velocity fields at the streamwise position of maximum $I_{tot}(z) + I_{uv}(z)$ for a cross-section of the pipe with the integration in the inner product performed over s and θ. The shading represents the axial velocity perturbation from laminar flow with contours from -0.55 (dark) to 0.5 (light) for the DNS, and -0.4 to 0.35 for the TW, with a step of 0.15. The arrows indicate the cross-stream velocity, scaled on magnitude (maximum $0.15\,U$). The middle row shows the streak structure over the wavelength of the TW, with contours of axial velocity at $\pm 0.3\,U$ (light/dark). The bottom row shows the axial vorticity, with contours at $\pm 0.6\,U/D$ (light/dark).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

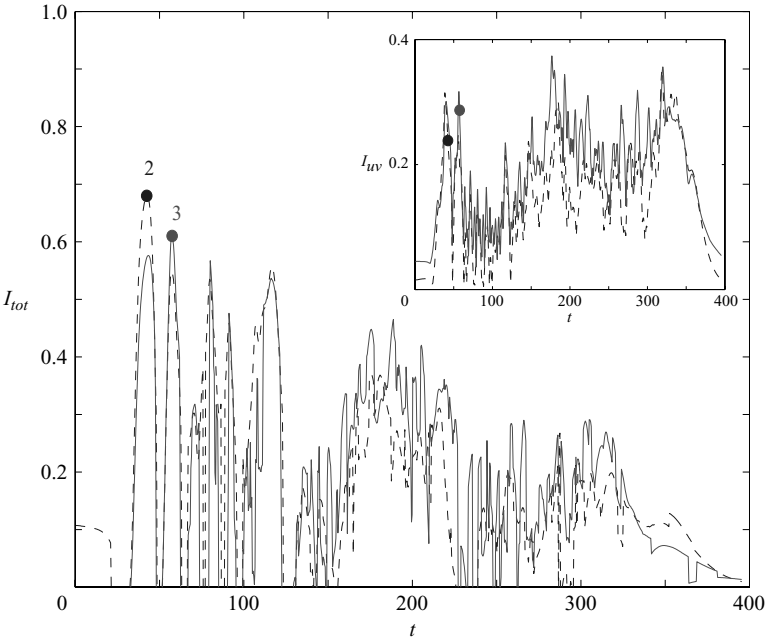


FIGURE 9. I_{tot} and I_{uv} (inset) as a function of time for TWs $2b_0.625$ (dashed line) and $2a_1.25$ (solid line) starting at TW $4b_3.125$. Dots label the times of probable closest visits to each TW (times '2' and '3' coincide with those in figure 5).

streamwise independent. Surprisingly, there is little similarity between the inner slow streak structures. The axial vorticity is a convenient, but particularly discriminating, way of probing the cross-stream velocity fields as it involves taking derivatives. Given the low value of I_{uv}, the poor comparison is expected; however, there is evidence that the DNS flow and the TW have roughly the same wavelength of variation.

During the dominantly four-fold symmetric phase, there is also evidence for visits to $2b_0.625$ and $2a_1.25$ (labelled '2' and '3') shown in figure 9. A further way to characterize the extent of all these visits is to compare the kinetic energy and wall shear stress of the flow across the length of the pipe used for matching with the values associated with the TW. These data are collected in figure 5. The relative closeness of these points to the corresponding values for the TW (e.g. '5' and $3a_3.125$) is further supporting evidence of a visit.

Runs started at upper-branch TWs with either sign of unstable eigenfunction lead to turbulent-looking trajectories. Figure 10 shows one of these runs starting with $3b_3.125$. Just as in the case of $4b_3.125$, the initial rotational symmetry of the flow lingers for a substantial time. During this phase, there is clear evidence of close visits to $3b_2.5$, $3c_2.5$ and $3j_2.5$ (see figure 11). The first visit is particularly significant as the structural overlap between the initial TW ($3b_3.125$) and the visited TW ($3b_2.5$) is low: by point '1', both I_{tot} and I_{uv} for $3b_2.5$ have risen from ≈ 0.4 to near 0.8 in just over $50\,D/U$. Figure 12, which compares the initial condition, the DNS flow at point '1' and the TW $3b_2.5$, confirms that the flow makes a significant adjustment to match with the new TW.

After a time of about $125\,D/U$ (see figure 13) the flow switches to being more two-fold symmetric. The flow then successively visits $2a_1.25$ and $2b_1.875$ before relaminarizing at a time $350\ D/U$. Point 4 gives a very good comparison between the

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

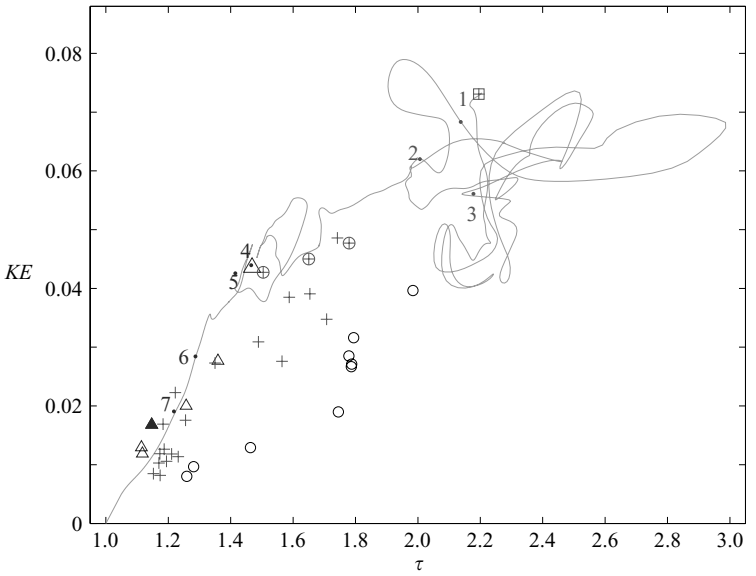


FIGURE 10. The disturbance kinetic energy per unit mass in units of U^2 versus wall shear stress τ in units of $-8\rho U^2/Re$ for $\boldsymbol{u}_{DNS}$ starting at the upper branch TW $3b_3.125(+)$ ($\boxplus$). The laminar state is represented by the point $(1, 0)$. All the TWs present are also plotted: $\triangle$, two-fold TWs; $+$, three-fold TWs; $\bigcirc$, four-fold TWs. The numbered dots correspond to the times of visits as determined by examining I_{tot} and I_{uv}. In chronological order: 1 – $3b_2.5$ (rightmost $\oplus$); 2 – $3j_2.5$ (leftmost $\oplus$); 3 – $3c_2.5$ (middle $\oplus$); 4 and 6 – $2a_1.25$ (large triangle); 5 and 7 – $2b_1.875$ (solid triangle).

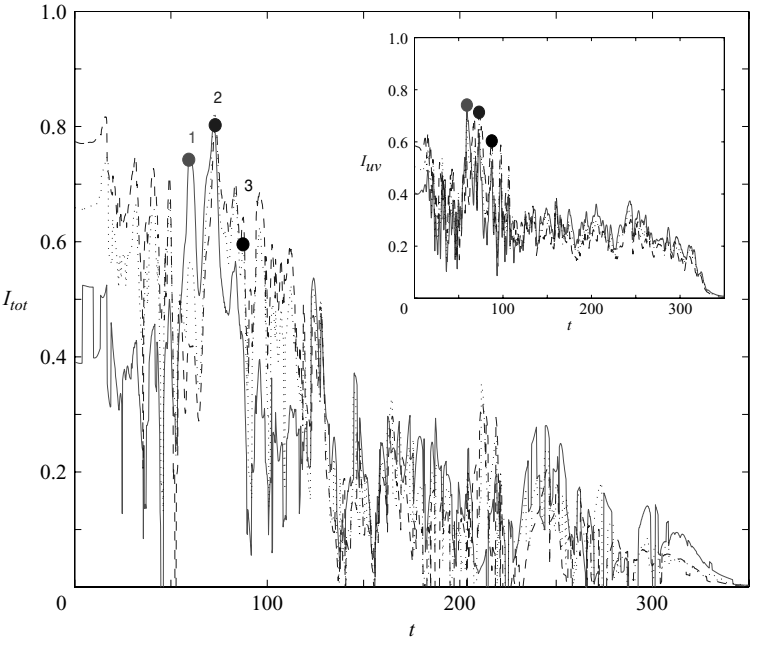


FIGURE 11. I_{tot} and I_{uv} (inset) as a function of time for TWs $3b_2.5$ (solid line), $3c_2.5$ (dotted line) and $3j_2.5$ (dashed line) starting at TW $3b_3.125$. The dots label the times of closest visits (numbers correspond to points in figure 10).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

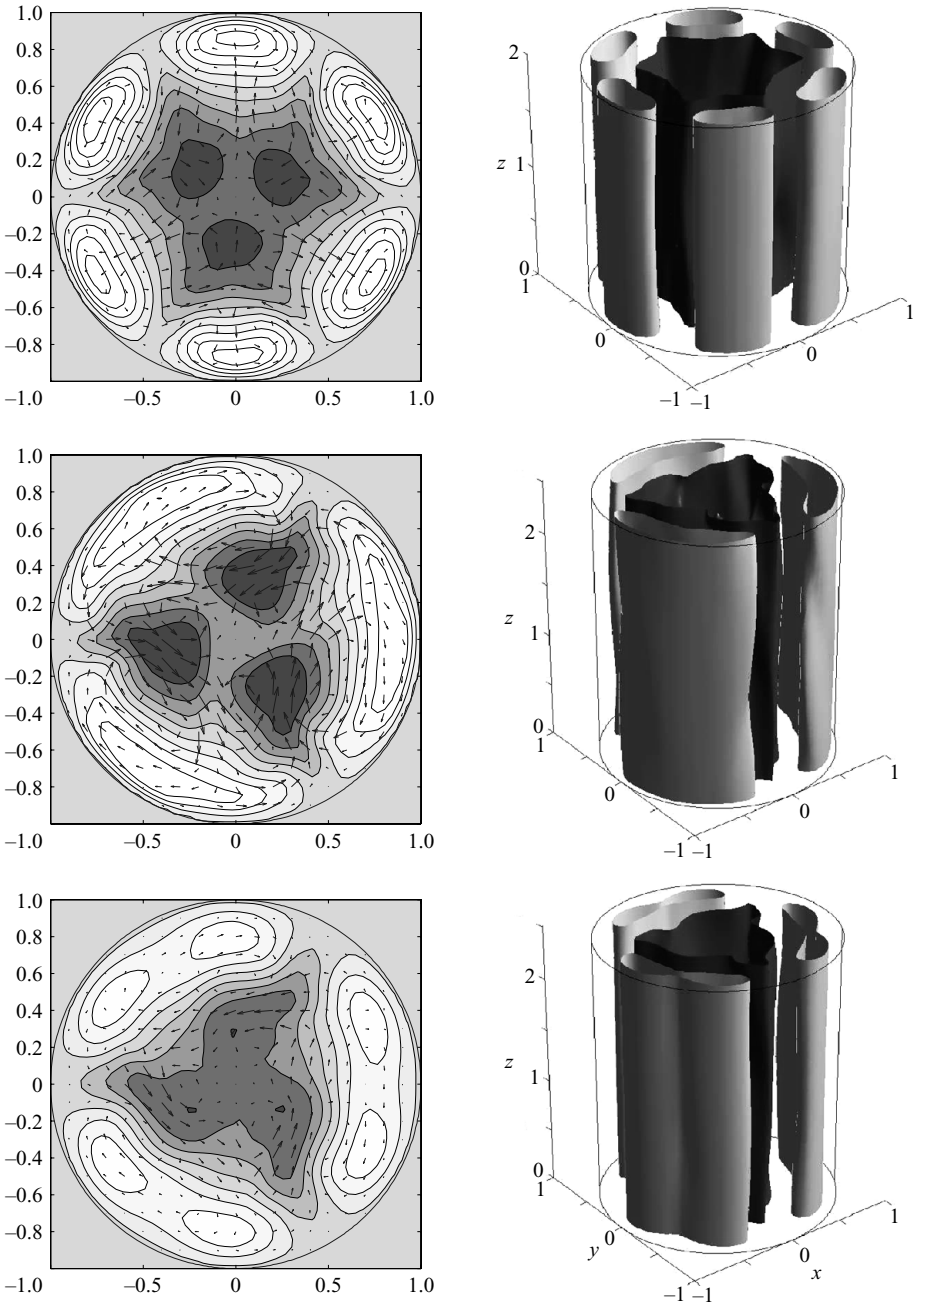


FIGURE 12. Comparison plots of the starting TW $3b_3.125$ (top), the DNS flow (middle) and the TW $3b_2.5$ (bottom) at point 1 in figure 11. The left-hand column shows the velocity field at a slice where the DNS flow most matches $3b_2.5$. The shading represents the axial velocity perturbation from laminar flow, with contours from -0.55 (dark) to 0.65 (light) (top), -0.55 to 0.5 (middle), and -0.55 to 0.35 (bottom), with a step of 0.15. The arrows indicate the cross-stream velocity, scaled on magnitude (maximum $0.135\,U$). The right-hand column shows the axial structure over a wavelength. Two contours of the axial velocity are shown at $\pm 0.3\,U$ (light/dark).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

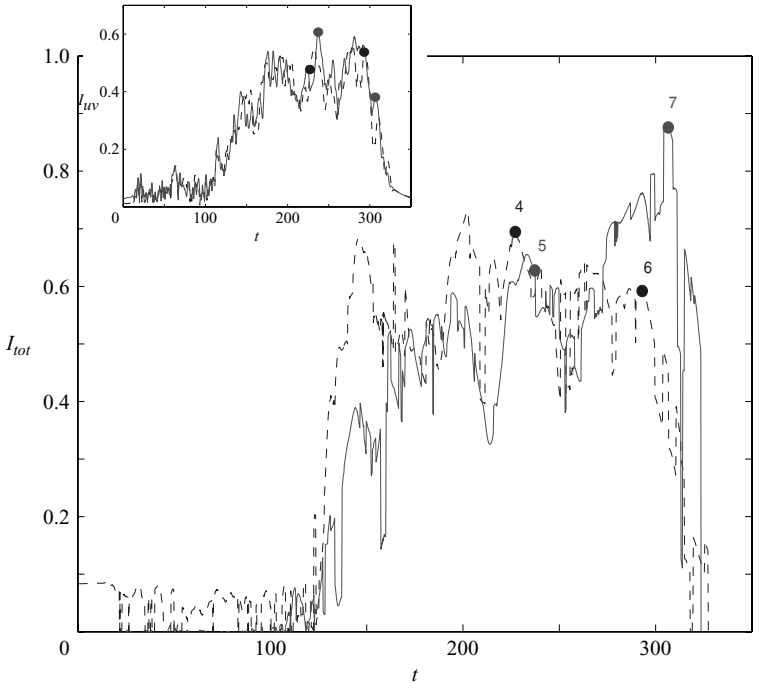


FIGURE 13. I_{tot} and I_{uv} (inset) as a function of time for TWs $2a_1.25$ (dashed line) and $2b_1.875$ (solid line) starting at TW $3b_3.125$. The dots label the times of closest visits (numbers correspond to points in figure 10).

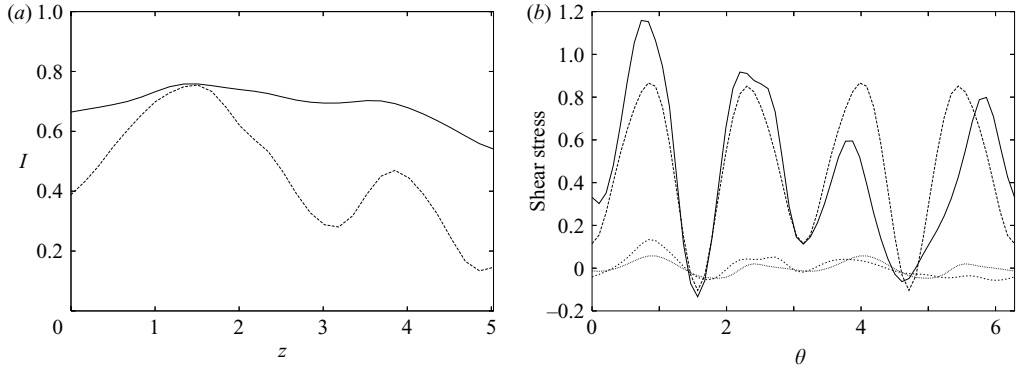


FIGURE 14. (*a*) The correlations $I_{tot}(z)$ (upper line) and $I_{uv}(z)$ (lower line) over the wavelength of TW $2a_1.25$ with the highest correlation of I_{uv} at point 4 in figure 13. $I_{tot}(z)$ and $I_{uv}(z)$ give the correlations for a cross-section with integration in the inner product performed over s and θ. (*b*) The azimuthal distribution of the wall shear stress in units of $-8\rho U^2/Re$ at the axial position of maximum $I_{tot}(z)$ and $I_{uv}(z)$ near $z = 1.5$. From the top at the left of the plot, the lines correspond to the axial stress from the DNS solution and that from the TW (minus the laminar value of one), and the azimuthal stress from the DNS solution and that from the TW.

the DNS flow and the TW $2a_1.25$ because the correlation functions I_{tot} and I_{uv} are simultaneously large within the comparison wavelength: see figure 14 at $z \approx 1.5$. The velocity plots (see figure 15) are very similar, even up to the same vortex in the top left-hand quadrant of the cross-sectional snapshots.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

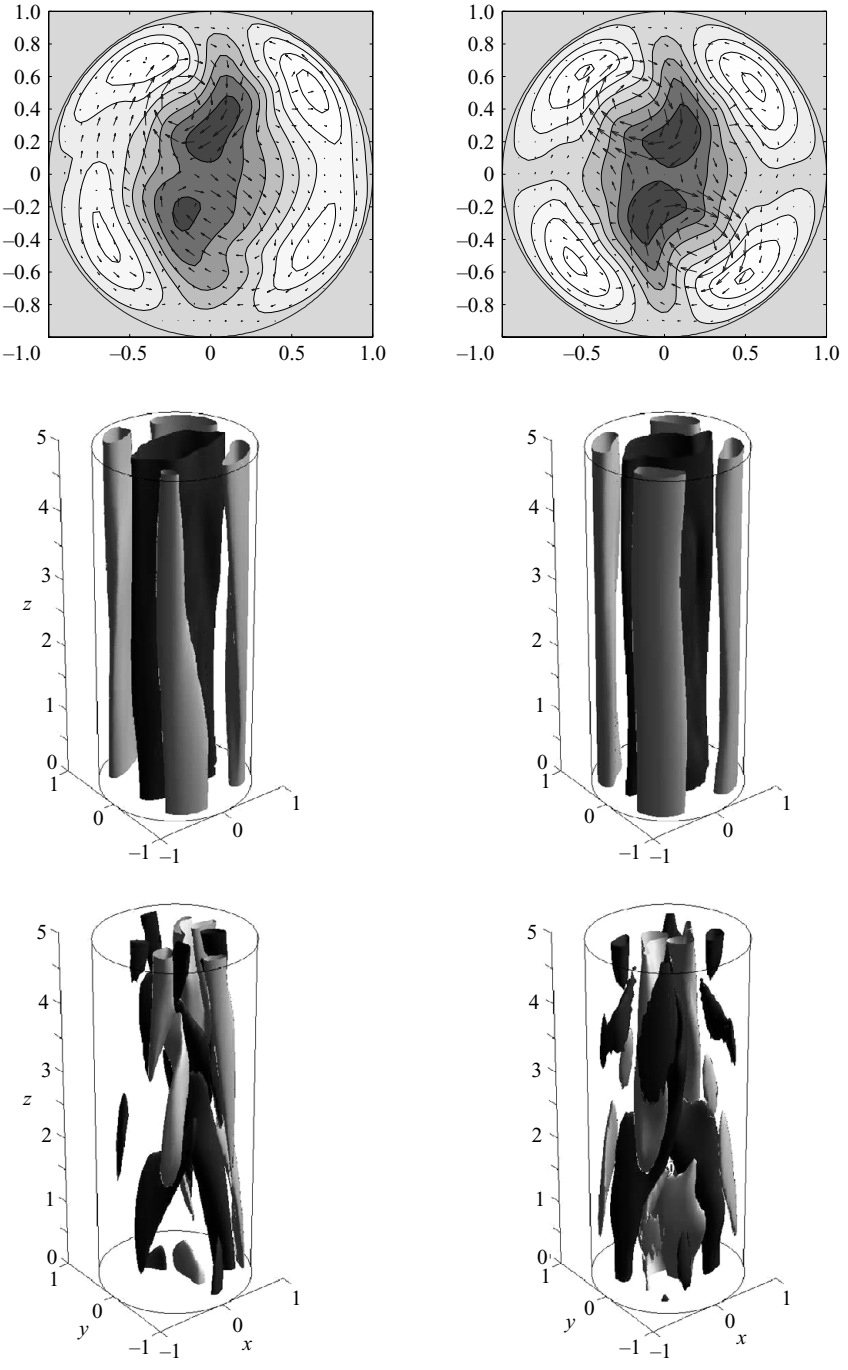


FIGURE 15. Comparison plots of the DNS flow (left-hand column) and the TW $2a_1.25$ (right-hand column) at point 4 in figure 13. The top row shows the velocity fields at the streamwise position of maximum I_{tot} shown in figure 14 ($z \approx 1.5$). The shading represents the axial velocity perturbation from laminar flow with contours from -0.55 (dark) to 0.5 (light), with a step of 0.15. The arrows indicate the cross-stream velocity, scaled on magnitude (maximum $0.098\,U$). The middle row shows the streak structure over the wavelength of the TW, with contours of axial velocity at $\pm 0.3\,U$ (light/dark). The bottom row shows the axial vorticity, with contours at $\pm 0.8\,U/D$ (light/dark).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

In figure 10, point 5 is flagged as a match to TW $2b_1.875$ as this TW has the largest correlation to the flow at this point, with $I_{tot} = 0.63$ and $I_{uv} = 0.61$. However, this is not the only TW that satisfies the criterion for a good match. Specifically, TW $2a_1.25$ has $I_{tot} = 0.58$ and $I_{uv} = 0.49$. Also, $2a_1.25$ is closer to point 5 than $2b_1.875$ in figure 10. That both of these TWs show a high correlation with the flow is not surprising as the TWs also show a high correlation (matching $2b_1.875$ to $2a_1.25$ gives $I_{tot} = 0.53$ and $I_{uv} = 0.62$). Figure 16 shows a comparison between the flow at point 5 and the two travelling waves. In general, figure 16 suggests that there is better agreement between TW $2a_1.25$ and the flow than $2b_1.875$, despite the fact that it has a lower correlation as measured by I_{tot} and I_{uv}. That this can occur is not unexpected since I_{tot} and I_{uv} are measures of shape and not amplitude. Figure 16 and the position of the TWs in figure 10 suggests there is a better agreement in amplitude between $2a_1.25$ and the flow than with $2b_1.875$.

No significant correlations were found with any four-fold TWs as the correlation functions $I_{tot} \leqslant 0.3$ and $I_{uv} \leqslant 0.45$ throughout the evolution. This was a persistent finding in all the other runs where a two-fold or three-fold symmetric TW was used to initiate the flow. Even when the flow was started by a four-fold symmetric TW, correlations with other four-fold symmetric TWs would be significant only in the initial phase of the flow evolution where the seeding symmetry is still present. In contrast, certain two-fold and three-fold TWs were consistently visited regardless of how the flow was initiated. The fact that the four-fold symmetric TWs have a higher wall-shear-stress-to-kinetic-energy ratio than the flow seems to adopt is indicated by figures 5 and 10 and equivalent plots for other starting TWs (not shown). The flow never seems to visit the part of phase space where the four-fold TWs are, unless specifically inserted, whereas some of the two-fold and three-fold symmetric TWs are in the active part of phase space populated by turbulent trajectories. This is probably a feature of the transitional Reynolds number used here as evidence for four-fold TWs has been found experimentally in fully turbulent flow at $Re = 3000$ (Hof *et al.* 2004). It may be that the part of phase space populated by the turbulent dynamics expands to encompass some four-fold symmetric waves by the time Re increases to this value.

Information about which TWs are visited in each of the 12 turbulent runs (8 starting from upper-branch TWs and 4 from lower-branch TWs) is summarized in tables 2 and 3. The criterion used to indicate a visit is $I_{tot} > \lambda$ and $I_{tot} + I_{uv} > 2\lambda$ with $\lambda = 0.5$ and $\lambda = 0.6$ for a closer visit. There is a strong correlation between the rotational symmetry of the starting TW and the waves subsequently visited, even when the initial transient – defined as the period when I_{tot} for the initial TW decreases (typically of $O(30D/U)$ in duration) – is not considered. For instance, runs 1, 2, 5 and 10 (see table 2) show no evidence of visits to TWs with different rotational symmetry. However, across this suite of runs, all the TWs are visited at least once except for the three TWs $3a_1.25$, $3b_1.25$ and $3c_1.25$. Table 3 shows the more discerning results of considering only times when the flow leaves the rotational symmetry class of the initial TW. Examples of when this occurs have already been discussed in run 4 (at about $130\,D/U$: see figure 7) and in run 7 (at about $125\,D/U$; see figure 13). This makes clear that the four-fold rotationally symmetric TWs are never visited except when the run is specifically started near one of these TWs (runs 4, 11 and 12) whereas two-fold and three-fold symmetric TWs are visited regardless of the starting symmetry (e.g. runs 3, 6, 7 and 9).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

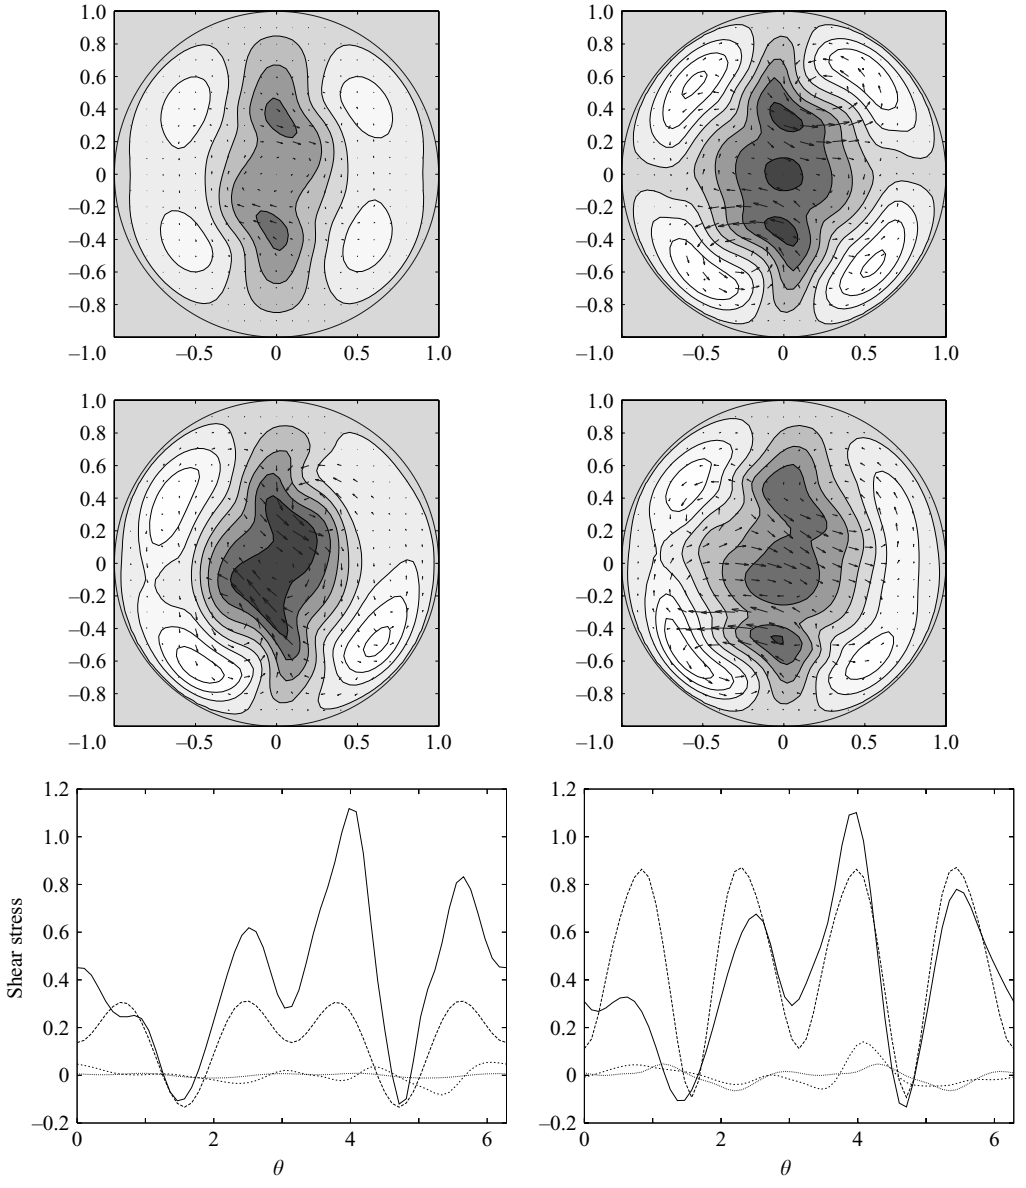


FIGURE 16. Comparison of the flow at point 5 in figure 10 with TWs $2b_1.875$ (left-hand side) and $2a_1.25$ (right-hand side). The top row shows the velocity fields for the TWs at the streamwise position of maximum correlation over a cross-section of the pipe, the middle row the DNS flow, and the bottom row the wall shear stress (in units of $-8\rho U^2/Re$) at the same position. The shading represents the axial velocity perturbation from laminar flow with contours from −0.55 (dark) to 0.5 (light) (right column), −0.4 to 0.2 (top left), and −0.7 to 0.65 (middle left), with a step of 0.15. The arrows indicate the cross-stream velocity, scaled on magnitude (maximum $0.162\,U$). For the wall shear stress, the top two lines show the axial stress (minus the laminar value of one) and the bottom two lines the azimuthal stress. For both the axial and azimuthal stress, the line with the largest peak is for the DNS solution, and the other (more regular) line from the TW.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

TW	Run	1	2	3	4	5	6	7	8	9	10	11	12
2*a*_0.625		○											○
2*b*_0.625												○	○
2*a*_1.25		●		○		★	★	●		○			
2*b*_1.25		★				●	●	○					
2*a*_1.875		●		○		●	●	●					
2*b*_1.875		●		○		●	●	●					
3*a*_1.25													
3*b*_1.25													
3*c*_1.25													
3*a*_1.875			●	○						○	○		
3*b*_1.875											○		
3*c*_1.875			●	●				○	○	●	●		
3*d*_1.875			○	○							○		
3*e*_1.875			○	○				○		○	○		
3*a*_2.5			★	●				○	●	○	○		
3*b*_2.5			●	●			●	●	●	●	●		
3*c*_2.5			●	●			○	●	●	●	●		
3*d*_2.5			○	●	○		○	○	●	○	●		
3*e*_2.5								○			○		
3*f*_2.5			●	●				●	●	○	●		
3*g*_2.5			●	●				●	●	○	●		
3*h*_2.5			●	★				○	○	○	○		
3*j*_2.5			●	●				●	●	★	★		
3*a*_3.125			○	●	○		○	●	●	○	●		
3*b*_3.125			●	●				★	★	●	●		
3*c*_3.125			●	●	○		○	●	●	●	●		
3*d*_3.125				○									
3*e*_3.125			○	○				○	○		○		
4*a*_2.5					●							●	●
4*b*_2.5					○							○	○
4*c*_2.5					○								○
4*a*_3.125					●							●	●
4*b*_3.125					★							○	
4*c*_3.125					●							★	★
4*d*_3.125					●							●	●
4*e*_3.125					●							●	●
4*f*_3.125					●							●	●

TABLE 2. Summary of visits for the 12 transitional runs (4 starting at lower-branch TWs + 8 starting at upper-branch TWs) which have turbulent episodes (data set A). TWs visited using the criterion $I_{tot} \geqslant 0.5$ and $I_{tot} + I_{uv} \geqslant 1$ at least once in a given run have the symbol ○ entered under that column against them (★ indicates the starting TW). ● indicates that the higher threshold of $I_{tot} \geqslant 0.6$ and $I_{tot} + I_{uv} \geqslant 1.2$ is satisfied. The initial transient as the flow moves away from the starting TW – typically of $O(30\,D/U)$ duration – is not considered. The code for the run numbers is: 1, 2*b*_1.25(−) (where (−) indicates the sign of the eigenfunction perturbation); 2, 3*a*_2.5(+); 3, 3*h*_2.5(+); 4, 4*b*_3.125(+); 5, 6, 2*a*_1.25(+/−); 7, 8, 3*b*_3.125(+/−); 9, 10, 3*j*_2.5(+/−); 11, 12, - 4*c*_3.125(+/−).

3.1. *Frequency of visits*

Given the evidence above that the TWs *are* visited, the next question is how frequently. To answer this, 'visit' statistics were compiled across a number of different runs which exhibited turbulent episodes. Ideally, these should be composed from one very long – e.g. $O(10\,000\,D/U)$ – turbulent run but, as discussed already, none could be generated

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

TW Run	1	2	3	4	5	6	7	8	9	10	11	12
2*a*_0.625	○	0.83			0.86		0.98	0.86				
2*b*_0.625	0.79	0.91			0.91		0.89	0.87				
2*a*_1.25	○		○		★	★	●	0.89	○			
2*b*_1.25	★	0.81	0.94			0.95	○					
2*a*_1.875	○	0.91	○		0.88	○	●	0.75	0.87			
2*b*_1.875	○		○		0.92	○	●	0.86	0.97	0.95		
3*a*_1.25												
3*b*_1.25												
3*c*_1.25												
3*a*_1.875		0.77										
3*b*_1.875												
3*c*_1.875												
3*d*_1.875												
3*e*_1.875												
3*a*_2.5		★		0.86		0.96				0.71		
3*b*_2.5	0.89	○	0.91	0.98		●		0.88		0.85		
3*c*_2.5		○		0.87	0.85	○						
3*d*_2.5		○	0.75	○		○				0.95		
3*e*_2.5												
3*f*_2.5												
3*g*_2.5						0.73						
3*h*_2.5		0.90	★	0.83		0.93						
3*j*_2.5				0.83					★	★		
3*a*_3.125		○	0.88	○	0.88	○				0.97		
3*b*_3.125							★	★				
3*c*_3.125	0.87	○	0.90	○	0.86	○				0.87		
3*d*_3.125												
3*e*_3.125												
4*a*_2.5												
4*b*_2.5												
4*c*_2.5												
4*a*_3.125												
4*b*_3.125				★								
4*c*_3.125											★	★
4*d*_3.125												
4*e*_3.125												
4*f*_3.125												

TABLE 3. As for table 2 except now only times when the flow has left the symmetry class of the starting TW are considered (data set B). Also, numerical values are recorded of $I_{tot} + I_{uv}$ in all near-miss instances where $I_{tot} > 0.5$, but ○ ($\lambda = 0.5$) or ● ($\lambda = 0.6$) is not warranted.

at this Re where the TWs are available. As a result, two strategies for quantifying the TW visit frequency were undertaken. The first involved piecing together all the turbulent episodes produced during the suite of $5D$ runs described above. The second was to generate longer turbulent data sets in a double-length $10D$ pipe; for relatively short pipes, the length of pipe appears to have a significant effect on the time for which turbulence is sustained, as can be seen from table 4. These latter runs were almost exclusively initiated by randomly selected velocity fields taken from a long turbulent coarse-grid run (see table 4). The coarse grid is too underresolved to reproduce the

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

Pipe lengt	Run	Turbulent duration in D/U	Initiation	Resolution (N, M, K)
$5D$	1	433 (673)	$2b_1.25(-)$	(50,60,60)
	2	327 (554)	$3a_2.5(+)$	(50,60,60)
	3	340 (517)	$3h_2.5(+)$	(50,60,60)
	4	147 (251)	$4b_3.125(+)$	(50,60,60)
	5	195 (385)	$2a_1.25(+)$	(50,60,60)
	6	397 (609)	$2a_1.25(-)$	(50,60,60)
	7	153 (265)	$3b_3.125(+)$	(50,60,60)
	8	263 (400)	$3b_3.125(-)$	(50,60,60)
	9	90 (257)	$3j_2.5(+)$	(50,60,60)
	10	18 (133)	$3j_2.5(-)$	(50,60,60)
$10D$	1	879	T	(50,48,80)
	2	1131	T	(50,48,80)
	3	390	T	(50,48,80)
	4	488	T	(50,48,80)
	5	349	T	(50,48,80)
	6	1446	T	(50,48,80)
	7	1182	T	(50,48,80)
	8	229	$2b_1.25(-)$	(50,60,120)
	9	335	T	(50,60,120)
	10	390	T	(50,60,120)
	11	586	T	(50,60,120)
	12	642	T	(50,60,120)
	13	558	T	(50,60,120)
	14	265	T	(50,60,120)
	15	640	T	(50,60,120)
	16	1135	T	(50,60,120)

TABLE 4. The various runs used to produce the 'visit' statistics. The $5D$ data comes from initiating the code with a disturbed TW and two 'turbulent' duration times are given. The first is measured from when the flow leaves the initial symmetry class of the TW (used for data set B) whereas the second, larger figure (in parentheses) is based on when the correlation function for the starting TW stops decreasing which is a much weaker condition that occurs earlier (used for data set A). The $10D$ runs are almost exclusively started using randomly selected velocity fields from a long turbulent run generated using a coarse $(25, 32, 60)$ grid (a strategy labelled as T). Note that some of the $10D$ runs were still turbulent when these data were gathered, whereas all the $5D$ runs had relaminarized.

TWs accurately and hence calculate correlation functions directly, but was adequate for initiating turbulence in finer grid runs.

For the $5D$ runs started at a TW, it was necessary to decide when the flow had left the neighbourhood of the TW and had started to behave in a turbulent fashion. Two different criteria were adopted for this. In the first (data set A), the start of the turbulent phase was taken as the time at which the correlation with the starting TW had reached a minimum. This turned out to be more stringent a criterion (i.e. produced a later time) than requiring that the flow trajectory merely enter the part of the (KE, τ)-plane occupied by the turbulent state. In the second (data set B), the start of the turbulent phase was taken as when the flow broke out of the initial rotational symmetry class of the starting TW (e.g. for run 4 this occurs at about $130\,D/U$: see figure 7). In both cases, the end of the turbulent phase was estimated well by when the flow trajectory left this 'turbulent' (KE, τ)-area. The subsequent relaminarization phase was easily identified with both the perturbation kinetic energy and wall stress

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

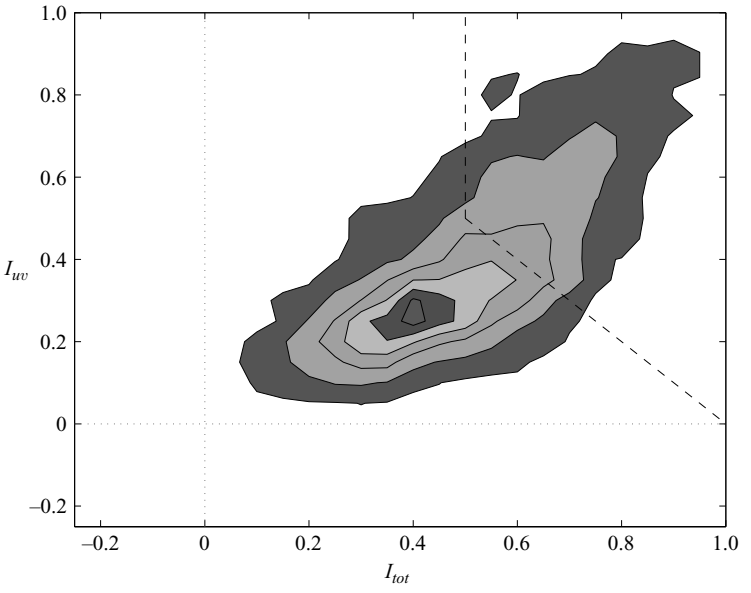


FIGURE 17. Joint p.d.f. of I_{tot} and I_{uv} calculated over the turbulent episodes of all $5D$ runs started at the eight upper- and lower-branch TWs (the values of I_{tot} and I_{uv} at a given time are selected by finding the TW with largest value of I_{tot}). Contours are drawn at $0.01, 0.1, 0.3, 0.5, 0.7$ and 0.9 of the maximum value. The dashed lines cordon off the visit region defined by $I_{tot} > 0.5$ and $I_{tot} + I_{uv} > 1$.

decaying monotonically down to their laminar values (e.g. see figures 5 and 10). These strategies produced correlation data sets lasting 4154 D/U (A) and 2363 D/U (B). For each, a joint probability density function was then computed for I_{tot} and I_{uv} corresponding to the TW with the largest I_{tot} at a given time (see figures 17 and 18). There is a clear positive correlation between I_{tot} and I_{uv} and significant evidence for recurrent TW visits. Comparing the two p.d.f.s, there are many close visits during the interval after the initial transient, but before the rotational symmetry class of the flow changes. This is presumably the result of the flow percolating out of the specific TW region in which it is initially inserted. As a result, the p.d.f. from data set A is likely to overestimate the visit frequency.

Using the criterion that a visit occurs if $I_{tot} > \lambda$ and $I_{tot} + I_{uv} > 2\lambda$, the percentage visit time is plotted against the quality of the visit λ in figure 19. To assess how close a visit must be (i.e. what λ is required) to be able to predict the instantaneous wall shear stress, figures 20 and 21 plot the instantaneous DNS wall stress (over the matching wavelength) against the wall shear stress associated with the visited TW for three different values of λ. The choice of seeking maximal I_{tot} to identify the closest TW was taken to facilitate the comparison, as this should provide the best match between the DNS streak structure at the pipe wall with that of the TW. Even with this, the correlation results for set A indicate that only $\lambda \geqslant 0.7$ level visits are really good enough to start making predictions of the wall shear stress from the visited TW. Since there are no $\lambda = 0.7$ visits in set B, $\lambda = 0.6$ is used to indicate closer visits, and better stress correlation is evident at this level compared to the $\lambda = 0.5$ results. Since the abscissa is discretized over the 37 τ_{TW} values, each vertical strip of data in these figures indicates that that TW has been visited. Figure 20 shows a visit bias to TWs with larger wall stress ($\leqslant -1.5 \times 8\rho U^2/Re$) whereas figure 21 is oppositely skewed

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

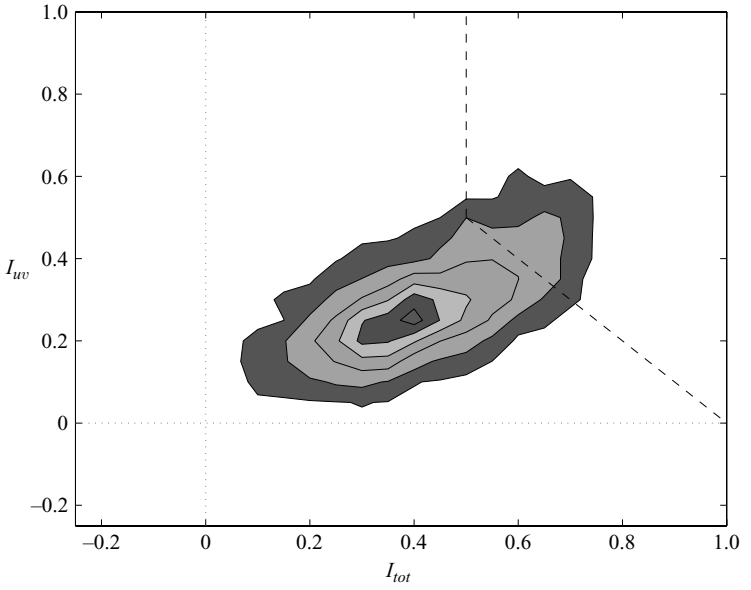


FIGURE 18. Joint p.d.f. of I_{tot} and I_{uv} calculated over the turbulent episodes of all $5D$ runs started at the eight upper- and lower-branch TWs using time episodes after the flow leaves the initial symmetry class (the values of I_{tot} and I_{uv} at a given time are selected by finding the TW with largest value of I_{tot}). Contours are drawn at 0.01, 0.1, 0.3, 0.5, 0.7 and 0.9 of the maximum value. The dashed lines cordon off the visit region defined by $I_{tot} > 0.5$ and $I_{tot} + I_{uv} > 1$.

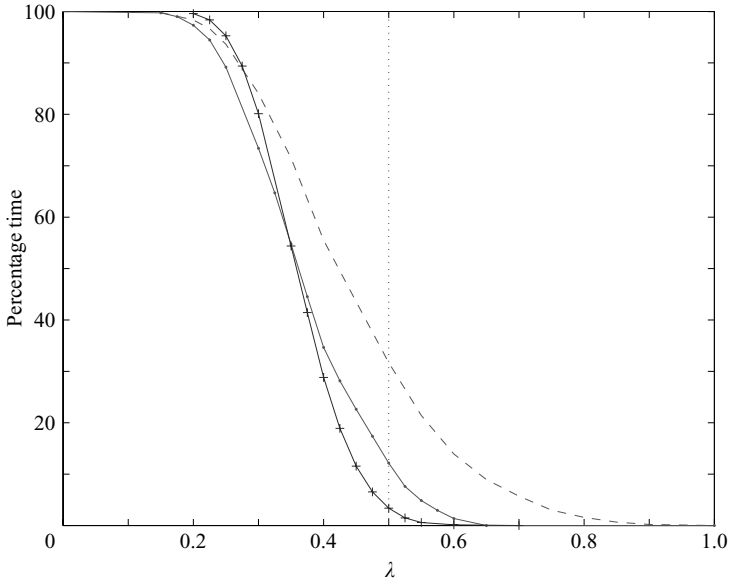


FIGURE 19. The percentage of time that a turbulent flow 'visits' a TW based on the criterion $I_{tot} > \lambda$ and $I_{tot} + I_{uv} > 2\lambda$ as a function of λ. Statistics gathered from the $5D$ runs started at a TW with only the initial transient subtracted (dashed line) and considering only times where the flow has left the initial symmetry class of the TW (solid line with dots), together with turbulent data compiled from various runs at $10D$ (solid line with crosses).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

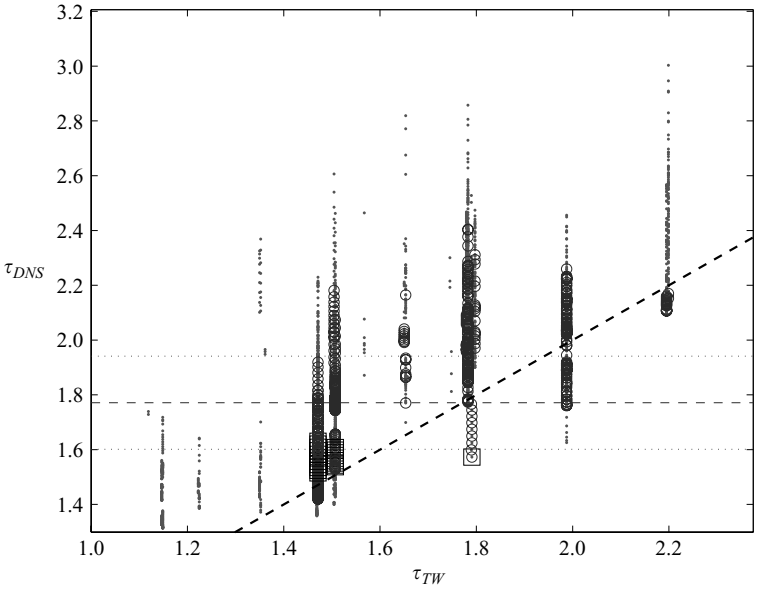


FIGURE 20. A plot of the wall shear stress τ_{DNS} versus the wall shear stress τ_{TW} of the TW being visited (τ_{DNS} is measured across the matching wavelength of the TW) for the $5D$ data set A runs. The visit criteria $I_{tot} > \lambda$ and $I_{tot} + I_{uv} > 2\lambda$, where $\lambda = 0.5$ (dots) indicate approximate visits, $\lambda = 0.7$ (circles) close visits and $\lambda = 0.9$ (squares) very close visits. The thick diagonal dashed line indicates a perfect match $\tau_{DNS} = \tau_{TW}$. The horizontal dashed line is the mean wall shear stress across the whole pipe, the dotted lines indicate one standard deviation either side of this mean and the limits of the vertical axis have been set to the maximum (3.2) and minimum (1.3) wall shear stress values (in units of the laminar stress $-8\rho U^2/Re$). Each vertical strip of points represents one TW – the close visits only occur for upper branch solutions where $\tau_{TW} > 1.3$.

($\geqslant -1.5 \times 8\rho U^2/Re$) to lower wall stress TWs. A possible explanation for this is set A is dominated in a 2-to-1 ratio with runs started by upper-branch TWs compared to lower-branch TWs. In the initial adjustment phase where the flow trajectory gradually leaves the vicinity of the initial TW, the flow visits other TWs and these are more likely to be upper- rather than lower-branch solutions in the neighbourhood of an initial upper-branch solution. The reason lower-branch TWs feature more in set B may be because the turbulent episodes considered were not long enough to desensitize the visit statistics from the final relaminarization process where the flow preferentially passes by lower-branch TWs. This could indicate that efforts to remove this phase from the data may not have been wholly successful.

Further runs were carried out in a $2\pi/0.625\ D$ ($\approx 10D$) pipe in a search for more sustained turbulent data. Extending the spatial domain by a factor of 2 allows a whole new set of TW wavelengths to fit into the pipe, roughly doubling the number of allowed TWs. The correlation calculations were not extended to encompass this enlarged set for two reasons. First, the new TWs are interspersed within the $5D$ set which already sample the available solutions well. Therefore, the new TWs should have structures very similar to the existing TW set. Secondly, the increase in the numerical overhead: currently calculating the correlation functions takes around 50 % of the CPU time. However, to compensate for this omission, the visit frequencies based on half the admissible TWs could justifiably be doubled. Initial conditions were randomly selected from an apparently sustained coarse turbulent run and used in

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

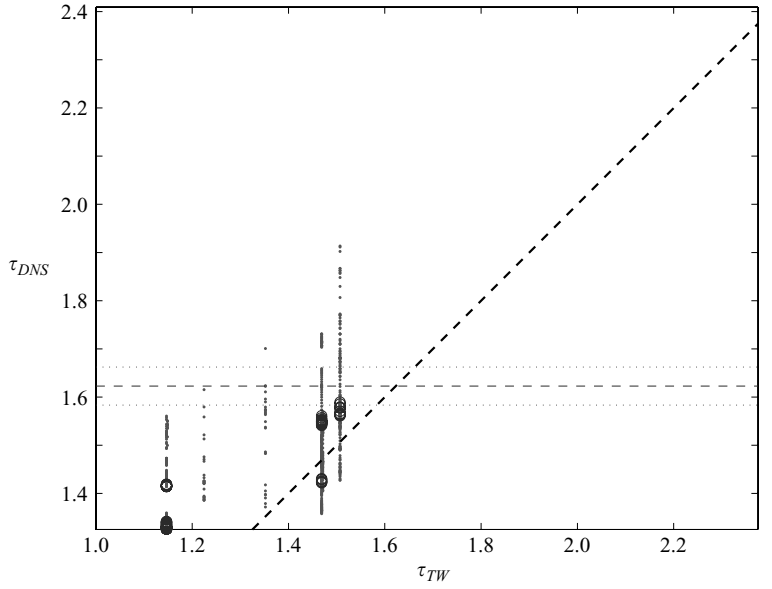


FIGURE 21. A plot of the wall shear stress τ_{DNS} versus the wall shear stress τ_{TW} of the TW being visited (τ_{DNS} is measured across the matching wavelength of the TW) for the $5D$ data set B runs. The visit critera are $I_{tot} > \lambda$ and $I_{tot} + I_{uv} > 2\lambda$, where $\lambda = 0.5$ (dots) indicate approximate visits, $\lambda = 0.6$ (circles) close visits. There are no very close ($\lambda = 0.9$) visits. The thick diagonal dashed line indicates a perfect match $\tau_{DNS} = \tau_{TW}$. The horizontal dashed line is the mean wall shear stress across the whole pipe, the dotted lines indicate one standard deviation either side of this mean and the limits of the vertical axis have been set to the maximum (2.4) and minimum (1.325) wall shear stress values (in units of the laminar wall stress $-8\rho U^2/Re$). Each vertical strip of points represents one TW. For $\lambda = 0.5$, the TWs visited are $2a_1.25$, $2b_1.875$, $3b_2.5$, $3a_3.125$ and $3c_3.125$. The TWs closely visited ($\lambda = 0.6$) are $2a_1.25$, $2b_1.875$ and $3b_2.5$.

intermediate- and fine-grid runs (see table 4). The statistics across the accumulated $5865\,D/U$ long intermediate grid turbulent data were similar to those across the accumulated $4780\,D/U$ long fine-grid turbulent data, so the sets were merged to give one large data set over $10\,000\,D/U$ in duration. The joint p.d.f. of (I_{tot}, I_{uv}) shown in figure 23 resembles that from the $5D$ data set B as does the wall stress comparison during visits shown in figure 24. Again, the most visited TWs look to be those with lower wall stresses and essentially the same subset of TWs are visited closely – $2a_1.25$, $2b_1.875$ and $3b_2.5$ in the $5D$ data compared to $2a_1.25$, $2b_1.875$, $3b_2.5$ and $3c_3.125$ in the $10D$ data. Figure 22 shows an example of a visit to $2a_1.25$ during a typical turbulent episode in this suite of runs. Significantly, no four-fold symmetric TWs were visited in the $10D$ runs, reinforcing the observation made in the $5D$ work that these TWs are not in the same part of phase space as that populated by a turbulent flow.

Figure 19 brings together the results of the frequency analysis by comparing the visit percentage as a function of visit quality (λ) for the three data sets. Doubling the $10D$ frequency data (as only half the possible TWs are considered) brings the observed visit frequency more into line with the results from the $5D$ data set B at $\lambda = 0.5$. Taken together, they suggest that TWs are visited for approximately 10 % of the time in turbulent pipe flow. Using $I_{tot} > \lambda$ as the criterion for a visit would increase the frequency of visits, by, in effect, ignoring the cross-stream component of the flow and considering the streamwise (streak) structure only.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

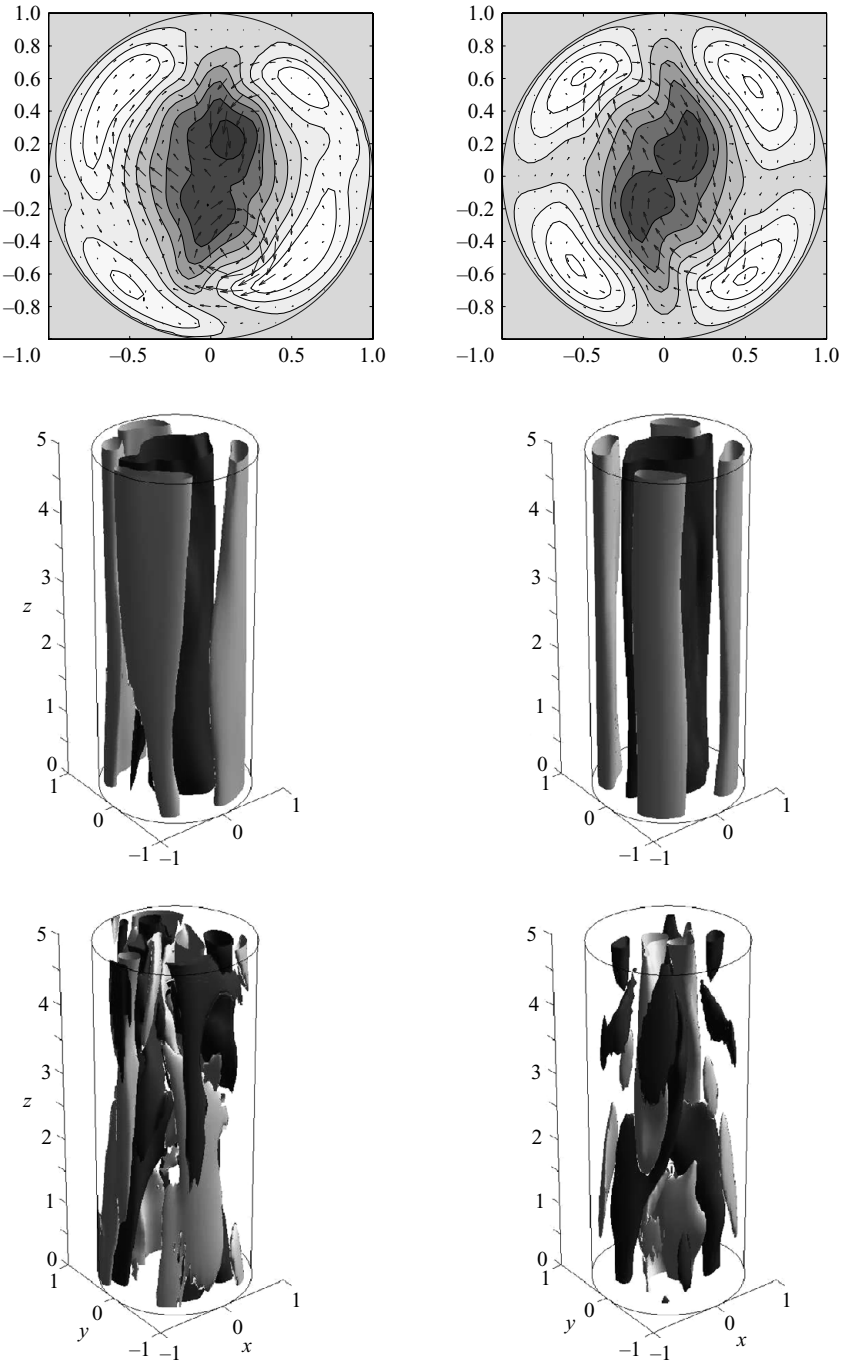


FIGURE 22. Comparison plots of the DNS flow (left-hand column) and the TW $2a_1.25$ (right-hand column) captured during a turbulent fine grid run in a $10D$ pipe. The top row shows the velocity fields at the streamwise position of maximum correlation over a cross-section of the pipe. The shading represents the axial velocity perturbation from laminar flow with contours from -0.7 (dark) to 0.35 (light) for the DNS, and -0.55 to 0.5 for the TW, with a step of 0.15. The arrows indicate the cross-stream velocity, scaled on magnitude (maximum $0.13\,U$). The middle row shows the streak structure over the wavelength of the TW, with contours of axial velocity at $\pm 0.3\,U$ (light/dark). The bottom row shows the axial vorticity, with contours at $\pm 0.8\,U/D$ (light/dark).

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

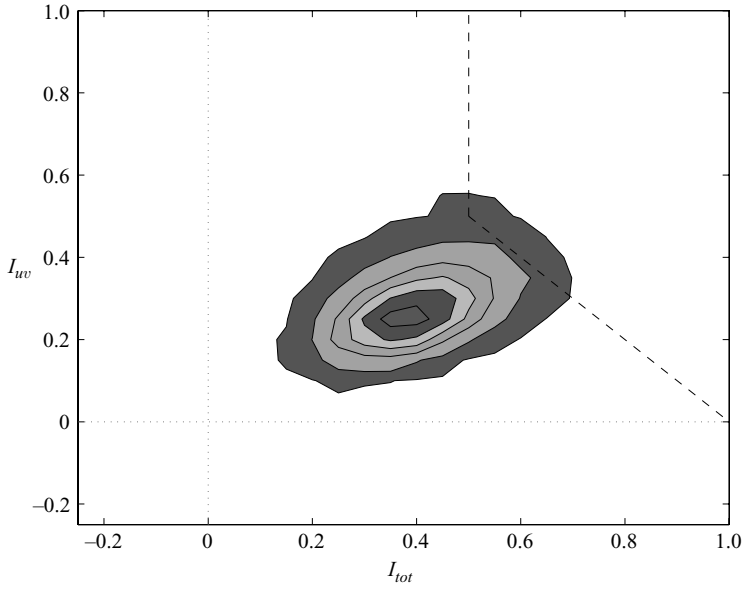


FIGURE 23. Joint p.d.f. of I_{tot} and I_{uv} calculated over the turbulent episodes of all $10D$ runs (the values of I_{tot} and I_{uv} at a given time are selected by finding the TW with largest value of I_{tot}). Contours are drawn at 0.01, 0.1, 0.3, 0.5, 0.7 and 0.9 of the maximum value. The dashed lines cordon off the visit region defined by $I_{tot} > 0.5$ and $I_{tot} + I_{uv} > 1$.

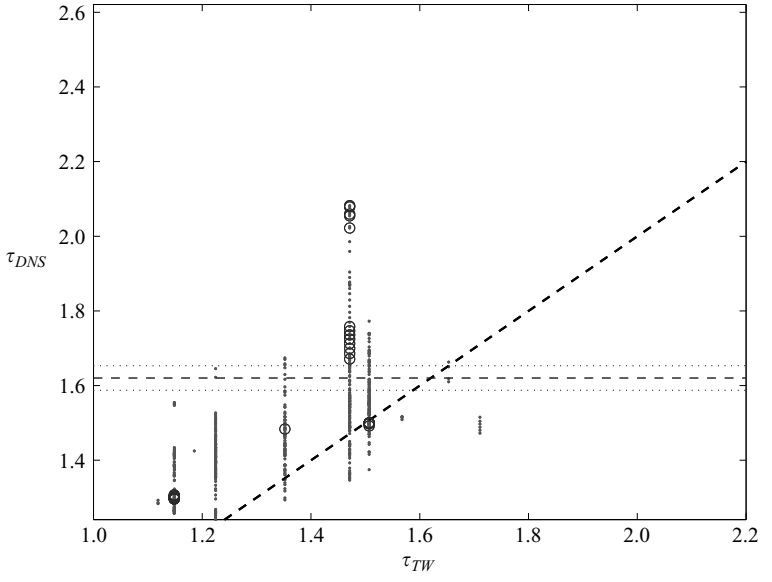


FIGURE 24. A plot of the wall shear stress τ_{DNS} versus the wall shear stress τ_{TW} of the TW being visited (τ_{DNS} is measured across the matching wavelength of the TW) for the $10D$ runs. The visit criteria are $I_{tot} > \lambda$ and $I_{tot} + I_{uv} > 2\lambda$, where $\lambda = 0.5$ (dots) indicates approximate visits and $\lambda = 0.6$ (circles) closer visits (there are no $\lambda = 0.7$ visits). The thick diagonal dashed line indicates a perfect match $\tau_{DNS} = \tau_{TW}$. The horizontal dashed line is the mean wall shear stress across the whole pipe, the dotted lines indicate one standard deviation either side of this mean and the limits of the vertical axis have been set to the maximum (2.62) and minimum (1.24) wall shear stress values (in units of the laminar wall stress $-8\rho U^2/Re$). Each vertical strip of points represents one TW. For $\lambda = 0.5$, the TWs visited are $2a_1.25$, $2b_1.25$, $3c_1.25$, $2a_1.875$, $2b_1.875$, $3c_1.875$, $3b_2.5$, $3c_2.5$, $3d_2.5$, $3a_3.125$ and $3c_3.125$. The TWs closely visited ($\lambda = 0.6$) are $2a_1.25$, $2b_1.875$, $3b_2.5$ and $3c_3.125$.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

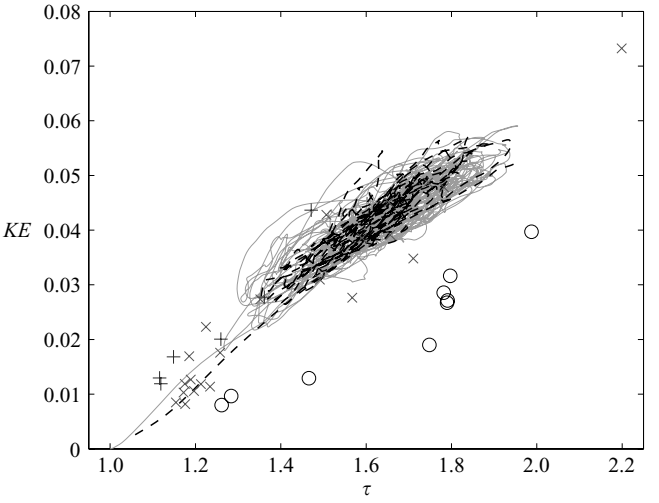


FIGURE 25. The disturbance kinetic energy per unit mass in units of U^2 versus wall shear stress τ in units of $-8\rho U^2/Re$ for $\boldsymbol{u}_{DNS}$ for two $10D$ runs started from a coarse grid turbulent run. The solid line is for an intermediate $(50, 48, 80)$ grid run and the dashed line a fine $(50, 60, 120)$ grid run. The laminar state is represented by the point $(1, 0)$. All the TWs are also plotted: +, two-fold TWs, ×, three-fold TWs, and ○, four-fold TWs.

Figure 25 shows the perturbation kinetic energy and the mean wall shear stress for a $10D$ pipe for one of the intermediate- and fine-grid runs started from coarse-grid turbulent data. This shows that the turbulent flow in this pipe is concentrated in a small patch of the (KE, τ)-space, with both grids occupying the same region of the space. The other intermediate- and fine-grid turbulent runs are also concentrated in this region. Coarse grid turbulent runs occupy a somewhat different region of the (KE, τ)-space with lower values of both variables. Both the intermediate and fine grid have the same radial resolution ($N = 50$). A test run was performed with the same grid in θ and z as the intermediate grid, but an increased resolution in s ($N = 75$). Again, the flow occupied the same region in the (KE, τ)-space.

Figure 26 shows power spectra of the surplus kinetic energy and the wall shear stress for turbulent flow in a $10D$ pipe. The spectra were generated by taking the data for the four longest fine-grid runs that were started from coarse-grid turbulent data, with a sample time of $782.1D/U$. A Hanning window was applied, with the windowed data scaled so that the variance matches that of the original data, and the spectra were generated by averaging over the four samples. The spectra show a peak, corresponding to a period of approximately $98D/U$ which suggests a periodic orbit embedded in the turbulent dynamics (the intermediate grid data produced spectra consistent with those for the fine grid, while the coarse grid also produced a peak, but at a higher frequency). This is currently under investigation.

4. Discussion

At this point it is worth cataloguing the achievements of this investigation.

(*a*) Using the travelling waves already known as starting points, all related travelling-wave solutions (TWs) of varying wavelengths and azimuthal symmetries which exist at $Re = 2400$ have been traced out (figures 1 and 2). They naturally partition into three distinct classes of particular rotational symmetry about the

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

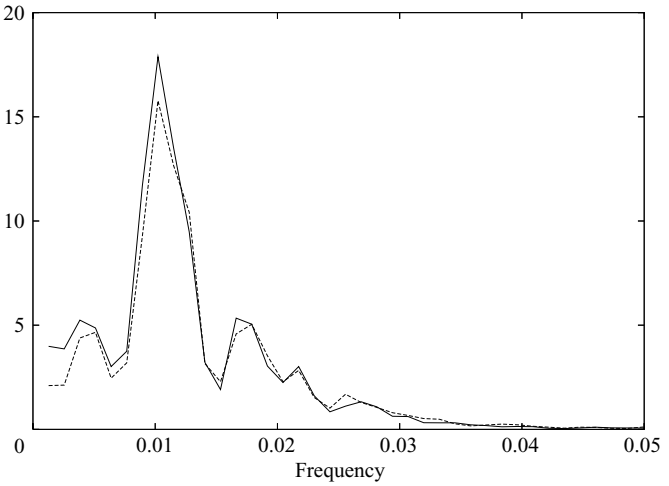


FIGURE 26. The power spectra of disturbance kinetic energy (multiplied by 5000) in units of U^3D (dashed) and wall shear stress in units of $4UD/Re^2$ (solid) versus frequency. The maximum in the spectra corresponds to a period of approximately 98 in units of D/U.

axis – two-fold, three-fold and four-fold symmetric TWs – and 37 in total fit into a $\pi/0.625\,(\approx 5)\,D$ periodic pipe.

(*b*) The linear stability of four 'lower branch' TWs (TWs with some of the lowest wall stresses of the 37) and four 'upper branch' solutions (TWs with the some of the highest wall stresses of the 37) has been carried out. All are inertially unstable and therefore saddles in phase space with growth rates typically of the size $O(0.1\,U/D)$ and all have very low dimensional unstable manifolds.

(*c*) All the lower-branch solutions considered appear to sit on a surface which separates initial conditions which uneventfully relaminarize and those which lead to a turbulent-looking evolution. The fact that lower-branch solutions may be part of such a dividing surface has been suggested before, but here we provide systematic evidence of multiple embeddings. The surface – formally a separatrix if the turbulence is a sustained state – then minimally appears as a union of lower-branch TWs and their stable manifolds. In contrast, initial conditions near all the upper-branch solutions tested invariably become turbulent.

(*d*) Turbulence in a $5D$ periodic pipe at $Re = 2400$ may be long-lived, but ultimately appears only transient. Turbulence also seems transient in a $10D$ pipe at $Re = 2400$, but the transients have, on average, a longer life. The mean lifetime of a turbulent episode is sensitive to the numerical resolution used, with turbulence appearing to last longer, if not sustained, in underresolved calculations.

(*e*) Clear numerical evidence is found that travelling waves are visited during transitional flows which corroborates the experimental observations of Hof *et al.* (2004). As in the experiments, the dominant recognizable feature of these visits is the appearance of an outer ring of equally spaced fast streaks. The more variable slow-streak structure towards the pipe axis is less clearly reproduced.

(*f*) A number of different correlation functions were experimented with to measure how 'close' a given perturbation velocity field is to that of a TW. Two, I_{tot} and I_{uv}, were chosen and a visit criterion – $I_{tot} > \lambda$ and $I_{tot} + I_{uv} > 2\lambda$ – developed based on a 'quality of visit' parameter λ. After examining the velocity matches at various different

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

levels of λ, a value of 0.5 was taken as indicating a 'visit', with values of 0.6 and higher indicating a 'close' visit. The correlations do not measure the amplitude of the perturbation velocity, and ancillary measures (the mean wall shear stress and the disturbance kinetic energy per unit mass) were used to verify the match in particular cases.

(*f*) Turbulent data in both $5D$ and $10D$ pipes indicate that some two-fold and three-fold symmetric TWs are recurrently visited whereas others are not. In particular, no evidence was found for visits to four-fold symmetric TWs at $Re = 2400$. The visited TWs correspond to low-to-intermediate wall stress solutions which are embedded in the same part of phase space as the turbulent dynamics. Other TWs are clearly in very different phase space locations and hence are never visited unless the flow is specifically inserted there initially (e.g. all the four-fold TWs). The fact that a four-fold symmetric wave has been observed experimentally (Hof *et al.* 2004), albeit at $Re = 3000$, suggests that by then the turbulent part of phase space has expanded to encompass some of the four-fold TWs.

(*g*) Based on the correlation functions, I_{tot} and I_{uv}, used and the visiting criterion adopted ($\lambda = 0.5$), numerical evidence suggests that travelling waves are visited only for about 10 % of the time in turbulent pipe flow at $Re = 2400$. These figures are broadly in line with the study by Schneider *et al.* (2007) which quantifies the frequency of coherent streak states near the pipe wall. They find coherent streak structures for 24 % of the time at $Re = 2200$ and 20 % of the time at $Re = 2500$, consistent with the thinking that travelling-wave visits are a strict subset of their coherent streak states. Above and beyond issues surrounding how a 'visiting' episode is determined, the exact visiting frequency found during this work must also be qualified by the fact that other currently unknown travelling waves may well exist, which would increase this statistic. In fact, while this work was being completed a whole new family of asymmetric travelling waves has been found at $Re = 2400$ (Pringle & Kerswell 2007) which would, of course, increase the visit frequency.

The last finding answers the original motivating question for this study. The fact that a turbulent flow is apparently visiting a TW only 10 % of the time implies that it is of limited use to view turbulence purely as the random switching between the neighbourhoods of TWs. For, say, predicting the average turbulent wall stress, an appropriately weighted sum of all the relevant TW wall stresses seems unlikely to work, given that so much time is spent away from the TWs in phase space. One possible way out of this conclusion is that other phase space objects such as periodic orbits (perhaps glimpsed in figure 26) should be included in any such expansion. Alternatively, if the time spent away from coherent flow features is truly significant, then some sort of weighted mixture of statistics describing the coherent and incoherent flow phases seems inescapable.

R. R. K. acknowledges the support of EPSRC under grant GR/S76144/01.

REFERENCES

Artuso, R., Aurell, E. & Cvitanovic, P. 1990*a* Recycling of strange sets: I cycle expansions. *Nonlinearity* **3**, 325–360.

Artuso, R., Aurell, E. & Cvitanovic, P. 1990*b* Recycling of strange sets: II applications. *Nonlinearity* **3**, 361–386.

Clever, R. M. & Busse, F. H. 1992 Three-dimensional convection in a horizontal fluid layer subjected to a constant shear. *J. Fluid Mech.* **234**, 511–527.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

Clever, R. M. & Busse, F. H. 1997 Tertiary and quaternary solutions for plane Couette flow. *J. Fluid Mech.* **344**, 137–153.

Cvitanovic, P. 1988 Invariant measurement of strange sets in terms of cycles. *Phys. Rev. Lett.* **61**, 2729–2732.

Eckhardt, B., Faisst, H., Schmiegel, A. & Schumacher, J. 2002 Turbulence transition in shear flows. *Advances in Turbulence IX: Proceedings of the Ninth European Turbulence Conference, Barcelona* (ed. I. P. Castro, P. E. Hancock & T. G. Thomas), p. 701.

Eckhardt, B., Schneider, T., Hof, B. & Westerweel, J. 2007 Turbulence transition in pipe flow. *Annu. Rev. Fluid Mech.* **39**, 447–468.

Eggels, J. G. M., Unger, F., Weiss, M. H., Westerweel, J., Adrian, R. J., Friedrich, R. & Nieuwstadt, F. T. M. 1994 Fully developed turbulent pipe flow: a comparison between direct numerical simulation and experiment. *J. Fluid Mech.* **268**, 175–209.

Faisst, H. & Eckhardt, B. 2003 Traveling waves in pipe flow. *Phys. Rev. Lett.* **91**, 224502.

Faisst, H. & Eckhardt, B. 2004 Sensitive dependence on initial conditions in transition to turbulence in pipe flow. *J. Fluid Mech.* **504**, 343–352.

Hof, B., van Doorne, C. W. H., Westerweel, J., Nieustadt, F. T. M., Faisst, H., Eckhardt, B., Wedin, H., Kerswell, R. R. & Waleffe, F. 2004 Experimental observation of nonlinear traveling waves in turbulent pipe flow. *Science* **305**, 1594–1597.

Hof, B., van Doorne, C. W. H., Westerweel, J. & Nieuwstadt, F. T. M. 2005 Turbulence regeneration in pipe flow at moderate Reynolds numbers. *Phys. Rev. Lett.* **95**, 214502.

Hof, B., Westerweel, J., Schneider, T. M. & Eckhardt, B. 2006 Finite lifetime of turbulence in shear flows. *Nature* **443**, 59–62.

Itano, T. & Toh, S. 2001 The dynamics of bursting process in wall turbulence. *J. Phys. Soc. Japan* **70**, 703–716.

Jiménez, J. & Simens, M. P. 2001 Low-dimensional dynamics in a turbulent wall flow. *J. Fluid Mech.* **435**, 81–91.

Jiménez, J., Kawahara, G., Simens, M. P., Nagata, M. & Shiba, M. 2005 Characterization of near-wall turbulence in terms of equilibrium and 'bursting' solutions. *Phys. Fluids* **17**, 015105.

Kawahara, G. 2005 Laminarization of minimal plane Couette Flow: Going beyond the basin of attraction of turbulence. *Phys. Fluids* **17**, 041702.

Kawahara, G. & Kida, S. 2001 Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. *J. Fluid Mech.* **449**, 291–300.

Kerswell, R. R. 2005 Recent progress in understanding the transition to turbulence in a pipe. *Nonlinearity* **18**, R17–R44.

Lagha, M. & Manneville, P. 2007 Modeling transitional plane Couette flow. *Eur. Phys. J.* (submitted).

Nagata, M. 1990 Three-dimensional finite-amplitude solutions to plane Couette flow: bifurcation from infinity. *J. Fluid Mech.* **217**, 519–527.

Nikitin, N. 2006 Third-order-accurate semi-implicit Runge–Kutta scheme for incompressible Navier–Stokes equations. *Intl J. Numer. Meth. Fluids* **51**, 221–233.

Orszag, S. A. & Kells, L. C. 1980 *J. Fluid Mech.* **96**, 159–205.

Peixinho, J. & Mullin, T. 2006 Decay of turbulence in pipe flow. *Phys. Rev. Lett.* **96**, 094501.

Pringle, C. & Kerswell, R. R. 2007 Asymmetric, helical and mirror-symmetric travelling waves in pipe flow. *Phys. Rev. Lett.* (submitted) (arXiv:physics/0703210).

Schlichting, H. 1968 Boundary Layer Theory. McGraw-Hill.

Schmiegel, A. 1999 Transition to turbulence in linearly stable shear flows. PhD thesis Philipps–Universitat Marburg.

Schmiegel, A. & Eckhardt, B. 1997 Fractal stability border in plane Couette flow. *Phys. Rev. Lett.* **277**, 197–225.

Schneider, T. M., Eckhardt, B. & Vollmer, J. 2007 Statistical analysis of coherent structures in transitional pipe flow. *Phys. Rev.* E (submitted).

Skufca, J., Yorke, J. A. & Eckhardt, B. 2006 The edge of chaos in a parallel shear flow. *Phys. Rev. Lett.* **96**, 174101.

Toh, S. & Itano, T. 2003 A periodic-like solution in channel flow. *J. Fluid Mech.* **481**, 67–76.

van Veen, L., Kida, S. & Kawahara, G. 2006 Periodic motion representing isotropic turbulence. *Fluid Dyn. Res.* **38**, 19–46.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press

Waleffe, F. 1998 Three dimensional coherent states in plane shear flows. *Phys. Rev. Lett.* **81**, 4140–4143.

Waleffe, F. 2001 Exact coherent structures in channel flow. *J. Fluid Mech.* **508**, 333–371.

Waleffe, F. 2003 Homotopy of exact coherent structures in plane shear flows. *Phys. Fluids* **15**, 1517–1534.

Wedin, H. & Kerswell, R. R. 2004 Exact coherent structures in pipe flow: travelling wave solutions. *J. Fluid Mech.* **508**, 333–371.

Willis, A. P. & Kerswell, R. R. 2007 Critical behaviour in the relaminarisation of localised turbulence in pipe flow. *Phys. Rev. Lett.* **98**, 014501.

Wygnanski, I. J. & Champagne, F. H. 1973 On transition in a pipe. Part 1. The origin of puffs and slugs and the flow in a turbulent slug. *J. Fluid Mech.* **59**, 281–351.

https://doi.org/10.1017/S0022112007006301 Published online by Cambridge University Press