

Stable manifolds and the transition to turbulence in pipe flow

D. VISWANATH¹ † AND P. CVITANOVIĆ²

¹Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA

²School of Physics, Georgia Institute of Technology, Atlanta, GA 30332, USA

(Received 12 January 2008 and in revised form 24 December 2008)

Lower branch travelling waves and equilibria computed in pipe flow and other shear flows appear intermediate between turbulent and laminar motions. We take a step towards connecting these lower branch solutions to transition by deriving a numerical method for finding certain special disturbances of the laminar flow in a short pipe. These special disturbances cause the disturbed velocity field to approach the lower branch solution by evolving along its stable manifold. If the disturbance were slightly smaller, the flow would relaminarize, and if slightly larger, it would transition to a turbulent state.

1. Introduction

The connection between the law of resistance to water flowing in a tube and the sinuous or direct nature of the internal motion of the fluid was established by Reynolds (1883). In the transitional regime, he observed ‘flashes’ and recorded them in figure 16 of his paper. Later experiments using hot-wire measurements have revealed the structure of puffs and slugs (Wynanski & Champagne 1973). Particularly intriguing are equilibrium puffs that maintain their spatial extent as they travel downstream with a characteristic speed (Wynanski, Sokolov & Friedman 1975). Such puffs are approximately 20 pipe diameters long and are observed for Re (Reynolds number) somewhat greater than 2000. Further, the structure of the puff is independent of the disturbance used to create it.

For a range of Re , the flow injected into the pipe assumes the familiar Hagen–Poiseuille laminar profile downstream. The flow does not have the laminar profile at the inlet. Therefore it is important to distinguish between disturbances at the inlet and disturbances to fully developed laminar flow (Willis *et al.* 2008). The early experiments used inlet disturbances, but in theoretical investigations such as this one, it has been common practice to consider disturbances to the laminar flow.

In their experiments to determine the dependence of the threshold for transition on Re , Darbyshire & Mullin (1995) used a constant mass-flux pipe and introduced disturbances to the laminar flow at a point sufficiently downstream from the inlet. They were able to determine thresholds over a range of Re , but figure 11 and other figures in their paper showed that certain disturbances above the threshold do not transition, while certain disturbances below the threshold do transition. Hof, Juel & Mullin (2003) determined that the threshold scaled as Re^α with $\alpha = -1$, when the laminar flow was disturbed by a single boxcar pulse of fluid injected at six different

† Email address for correspondence: divakar@umich.edu

points. Mellibovsky & Meseguer (2007) have reproduced $\alpha = -1$ in a numerical study that added a body force term to the Navier–Stokes equation to model the effect of the boxcar pulse of fluid. For different disturbances of the laminar flow, Peixinho & Mullin (2007) found $\alpha < -1$. In that experiment, the transition is sequential, with flow visualizations showing the disturbance changing its form as it travels downstream before leading to bigger structures. Following O’Sullivan & Breuer (1994), Peixinho & Mullin (2007) have pointed out that disturbances that lead to $\alpha < -1$ probably do not significantly distort the mean flow.

With regard to theory, Reynolds’s assertion that ‘there was small chance of discovering anything new or faulty’ in the Navier–Stokes equation has stood the test of time. Thanks to numerical computations, we now know that the incompressible Navier–Stokes equation adequately explains a remarkable wealth of phenomena related to transitional turbulence and fully developed turbulence. Although the Navier–Stokes equation can be solved numerically in certain regimes, the nature of the solutions of that equation has proved difficult to understand.

It is clear, however, that the nature of the solutions is quite different in the turbulent and transitional regimes. Fully developed turbulence is characterized by rapid decay of correlations and fine scales. Statistical theories that separate turbulent velocity fields into means and fluctuations have had significant successes (Narasimha 1989), even though coherent motions are present in certain regions of fully developed turbulence (Robinson 1991). In contrast, the transition problem seems to be fundamentally dynamical in nature.

Following the experiments of Darbyshire & Mullin (1995), Schmiegel & Eckhardt (1997) and Faisst & Eckhardt (2004) argued that there is no sharp boundary between initial conditions that trigger turbulence and those that do not and demonstrated computationally that the stability border for plane Couette flow is a fractal. One of their suggestions, namely that a chaotic saddle could be present for transitional Re , illustrates the dynamical nature of the transition problem.

More recently, Faisst & Eckhardt (2003) and Wedin & Kerswell (2004) computed a number of travelling-wave solutions of pipe flow. Their work was preceded by computations of somewhat similar solutions of channel flows by Nagata (1990) and Waleffe (1998). Hof *et al.* (2004) found streak patterns in puffs and slugs that appeared close to those of some pipe-flow travelling waves. The travelling waves were computed in short pipes, typically only a few pipe diameters long, while puffs are as long as 20 pipe diameters. Therefore the following question may be asked: do the experimentally observed structures correspond to the computed travelling waves?

The correlation functions, such as those of Schneider, Eckhardt & Vollmer (2007a), used to detect streak patterns in experimental or numerical flow fields look for m -fold rotational symmetry with respect to the pipe axis. Using such a correlation function, figure 5 of Schneider *et al.* (2007a) illustrates a transition from a four-streak state to a six-streak state within a spatial range of a single pipe radius. Willis & Kerswell (2008) found structures with $m = 3$ and $m = 4$ within streamwise distances of about 2 pipe diameter preceding the trailing edge of the puff and 5 pipe diameters following the trailing edge. Figure 5 of their paper gives some evidence that parts of the puff on either side of its trailing edge (but not at the trailing edge itself) visit travelling-wave solutions with $m = 3$ and $m = 4$. It could be significant that figures 7 and 23 of Wedin & Kerswell (2004) (which use different units) imply that some of the $m = 3$ and $m = 4$ travelling waves have wave speeds relatively close to that of the puff. Willis & Kerswell (2008) also found that the qualitative comparisons of Hof *et al.* (2004) which used slug cross-sections had significant problems. Instead of the trailing

edge of the puff, Eckhardt & Schneider (2008) used the centre of turbulence energy to fix a position within a moving puff. Their centre of turbulence energy is more precisely defined, and it moves with the puff in a quite regular manner. They found the axial correlation lengths around that centre to be quite short. While the question raised in the previous paragraph cannot be answered conclusively at the moment, these arguments suggest that short-pipe computations are of some relevance.

Another point to be mentioned is that puffs, in which streak patterns resembling those of some travelling waves have been detected, are observed experimentally only for $Re < 2800$. The transition experiments that measure thresholds (Hof *et al.* 2003; Peixinho & Mullin 2007) reach Re as high as 20 000. Thus the relevance of puffs to transition may seem limited. However, there is a possibility that puffs exist as solutions of the Navier–Stokes equation beyond the Re at which they are observed in experiments (Willis & Kerswell 2009).

Most of the lower branch solutions of pipe flow and the channel flows seem to be on the laminar–turbulent boundary (Itano & Toh 2001; Kawahara 2005; Kerswell 2005; Kerswell & Tutty 2007; Wang, Gibson & Waleffe 2007; Duguet, Willis & Kerswell 2008; Gibson, Halcrow & Cvitanović 2008; Schneider *et al.* 2008; Viswanath 2008*b*), which means that for some tiny disturbances of the lower branch solution, the disturbed state evolves and becomes laminar uneventfully. For other disturbances, the disturbed state evolves and becomes turbulent or undergoes a turbulent episode before it becomes laminar.

In this paper, we investigate if there are small disturbances of the laminar solution, for which the disturbed state evolves and hits a given lower branch solution. By small, we mean firstly that the magnitude of the disturbance should decrease algebraically with Re and secondly that the disturbance should not change the mean flow significantly. The existence of such a disturbance would establish that the flow can transition from laminar to turbulence by passing through the vicinity of the given lower branch travelling wave.

In fact, Kreiss, Lundbladh & Henningson (1994) found such disturbances when computing thresholds without fully realizing that they were hitting a lower branch equilibrium solution of plane Couette flow. But the situation they tackled is an especially simple one because the lower branch solution has a single unstable direction (Toh & Itano 2003; Wang *et al.* 2007; Schneider *et al.* 2008; Viswanath 2008*b*). We consider the asymmetric travelling wave computed by Pringle & Kerswell (2007). That travelling wave has two unstable directions both of which lie in a symmetric subspace, and thus it serves to illustrate that the method used for computing thresholds cannot be used to hit travelling waves that have more than one unstable direction.

The asymmetric travelling wave of Pringle & Kerswell (2007), which has two fast streaks located near one side of the pipe, is shown in figure 1. The preferential location of the streaks towards one side is also found in edge states that occur in transition computations (Schneider, Eckhardt & Yorke 2007*b*). Table 1 gives basic data for that travelling wave at four different Re . That data will be useful for judging the closeness of approaches to the travelling wave. The choice of units, the significance of I and D in table 1 and the meaning of the streamwise modes with $n = 0, \pm 1, \pm 2, \pm 3$ are explained in §2. More extensive data for the asymmetric travelling wave can be found elsewhere (Viswanath 2008*a*).

To find a small disturbance of the laminar solution that evolves into a given lower branch state, it is necessary to consider a linear superposition of disturbances whose dimension equals that of the unstable manifold of the lower branch state. That requirement follows from a consideration of the co-dimension of the stable manifold




FIGURE 1. Contour plots of z -averaged streamwise velocity with the laminar flow subtracted. Rolls are superposed. The contour levels are equispaced in the intervals $[-0.195, 0.170]$ and $[-0.191, 0.148]$, with the lighter regions being faster. The maximum magnitude of the vectors in superposed quiver plots are 0.0075 and 0.0038, respectively.

Re	$I = D$	ke	ke_d	ke_0	ke_1	ke_2	ke_3
2000	1.0881	0.9783	0.013	0.9778	4.4×10^{-4}	2.3×10^{-5}	2.2×10^{-7}
2500	1.0802	0.9790	0.012	0.9788	2.8×10^{-4}	1.2×10^{-5}	1.0×10^{-7}
3000	1.0755	0.9794	0.012	0.9792	1.9×10^{-4}	7.6×10^{-6}	5.9×10^{-8}
4000	1.0705	0.9796	0.012	0.9795	1.1×10^{-4}	3.8×10^{-6}	2.8×10^{-8}

TABLE 1. The kinetic energy of the travelling wave and the kinetic energy with laminar flow subtracted are denoted by ke and ke_d , respectively. The last four columns give the kinetic energy in modes with $n = 0, \pm 1, \pm 2, \pm 3$.

of the lower branch state. In addition, the disturbances that are linearly combined must be chosen carefully.

To hit the asymmetric travelling wave, we consider three different disturbances of the laminar solution. The first disturbance is obtained by extracting the rolls, which are formed by averaging the travelling wave in the streamwise direction and retaining only the radial and azimuthal components of the velocity field. The choice of rolls is related to the so-called lift-up mechanism (Landahl 1980). The two other disturbances are the two unstable eigenvectors of the travelling wave. If one of the eigenvectors is added to the travelling wave, it has the effect of either reinforcing or weakening the fast streaks. The other eigenvector seems to alter the location of the fast streaks. It must be remembered, however, that the disturbances are added to the laminar solution and not to the travelling wave.

In § 4, we show that disturbances of the laminar flow obtained by varying any two of these three disturbances evolve and hit the travelling wave. The choice of the unstable eigenvectors might seem puzzling, as our intention is to hit the travelling wave and not to move away from it. The reason that choice works is partially explained in §§ 3 and 4. Section 4 also shows that the magnitudes of the disturbances needed to hit the asymmetric travelling wave diminish algebraically with Re .

All our computations use a pipe that is π pipe diameters long, and the travelling waves are computed with 85 715 active degrees of freedom. The computations of relative periodic solutions (or modulated travelling waves) in plane Couette flow use triple the number of degrees of freedom (Viswanath 2007), although those computations are roughly 5 to 10 times as expensive with the same number of degrees of freedom. The computation of travelling waves is an insignificant part of the total computational expense, however, as will become clear in §4. We need to use a short pipe to keep the total computational expense manageable.

The pipe we use is too short to capture transitional structures such as puffs. To add to the earlier discussion of the relevance of short-pipe computations of travelling waves, we mention the work of Mellibovsky & Meseguer (2006) which seems to suggest that transition scenarios can be independent of pipe length. The logic which is used to find disturbances of the laminar solution that hit the asymmetric travelling wave makes fairly intricate use of the dynamical properties of the travelling wave. More work is needed to determine if the same logic is applicable to transition in pipes of more realistic length.

2. Preliminaries

The code for direct numerical simulation of pipe flow uses cylindrical coordinates with u , v and w being the components of the velocity in the radial (r), polar (θ) and axial (z) directions, respectively. The boundary conditions are no-slip at the walls and periodic in the z direction with constant mass flux. The length of the periodic domain in the z direction is denoted by $2\pi\Lambda$. We use $\Lambda = 1$ throughout. The units for distance and velocity are chosen so that the pipe radius is 1, and the Hagen–Poiseuille profile is given by $w = 1 - r^2$. The Reynolds number, Re , is based on the pipe radius, centreline velocity of the Hagen–Poiseuille flow and kinematic viscosity ν . The unit of mass is chosen so that the density of the fluid is 1. The units and boundary conditions follow those of Faisst & Eckhardt (2004).

Let $\bar{w}(r)$ denote the mean velocity in the axial direction and $\bar{v}(r)$ the mean velocity in the polar direction. The mass flux per unit area is given by $2 \int_0^1 r \bar{w}(r) dr$ and is equal to 1/2 for all velocity fields that obey the boundary condition. The pressure gradient necessary to maintain constant mass flux varies from instant to instant. For the Hagen–Poiseuille flow, it is $-4/Re$.

The spatial discretization is spectral. The radial component of the velocity u is represented as

$$u(r, \theta, z) = \sum_{n=-N}^{n=N} \sum_{m=-M}^{m=M} \hat{u}_{n,m}(r) \exp(im\theta) \exp(inz/\Lambda). \quad (2.1)$$

For the velocity field to be regular at $r = 0$, the coefficients $\hat{u}_{n,m}(r)$ must be even functions of r for odd m and odd functions of r for even m . Thus the functions can be reconstructed by storing their values at $r = \cos(i\pi/L)$, $i = 0, 1, \dots, (L-1)/2$. We assume L odd so that there is no point at $r = 0$ (Trefethen 2000). The radial component of vorticity is denoted by ξ . It is represented in the same way u is represented. The other quantities used to represent the velocity field are $\bar{v}(r)$, which is an odd function of r , and $\bar{w}(r)$, which is an even function of r . The velocity field is constructed using u , ξ , \bar{v} , \bar{w} and the divergence-free condition. The advection term is dealiased using the Orszag 3/2 rule. All the computations use $(N, M, L) = (16, 18, 81)$.

The rate of energy dissipation per unit mass is given by $2D/Re$, where D is the integral of

$$\frac{1}{4\pi^2\Lambda} \left(\frac{1}{r^2} \left(u^2 + v^2 - 2\frac{\partial u}{\partial\theta}v + 2u\frac{\partial v}{\partial\theta} \right) + \sum_{U=u,v,w} \left(\frac{\partial U}{\partial r} \right)^2 + \left(\frac{\partial U}{\partial z} \right)^2 + \frac{1}{r^2} \left(\frac{\partial U}{\partial\theta} \right)^2 \right) \quad (2.2)$$

over the volume of the pipe. In its more familiar form, D is the integral of the sum of the norms of gradients of the three components of the velocity field (Wedin & Kerswell 2004). The term under the summation in (2.2) gives $|\nabla U|^2$ for a scalar field $U(r, \theta, z)$. The terms outside the summation in (2.2) arise as cross-terms when that operator is applied to $u \cos \theta - v \sin \theta$ and $u \sin \theta + v \cos \theta$. The explicit form of (2.2) displays the $1/r^2$ singularities that are hidden in vector notation. Because those singularities cancel at $r = 0$, the numerical evaluation of D in a spectral code is a delicate matter. The rate of energy input per unit mass is given by $2I/Re$, where

$$I = -\frac{Re}{4\pi^2\Lambda} \int \nabla \cdot (p\mathbf{u}), \quad (2.3)$$

with p being pressure and with the integral being over the volume of the pipe. For the Hagen–Poiseuille laminar flow, both D and I evaluate to 1.

Figure 1 shows the asymmetric travelling-wave solution first computed by Pringle & Kerswell (2007). To compute that travelling wave, we added rolls which approximate the pattern in figure 1 to the laminar solution and evolved the velocity field to allow the streaks to develop. The resulting velocity field was used as the initial guess for the generalized minimal residual–hookstep (GMRES-hookstep) method, which converged without a hitch. The number of active degrees of freedom in the representation of a velocity field is $(L-2) + ((2N-1)(2M-1)-1)(L-3)/2$. The method uses translation operators to handle the invariance of the pipe-flow equation with respect to shifts along z and rotations along θ . These operators are given by

$$\begin{aligned} \mathcal{T}_1 u(r, \theta, z) &= \sum_{m,n} im \hat{u}_{n,m}(r) \exp(im\theta) \exp(inz/\Lambda) \\ \mathcal{T}_2 u(r, \theta, z) &= \sum_{m,n} (in/\Lambda) \hat{u}_{n,m}(r) \exp(im\theta) \exp(inz/\Lambda), \end{aligned} \quad (2.4)$$

where the indices m, n correspond to the representation (2.1). A detailed description of the GMRES-hookstep method can be found elsewhere (Viswanath 2007, 2008a).

The equations of pipe flow are unchanged by the shift–reflect symmetry:

$$\begin{aligned} u(r, \theta, z) &\rightarrow u(r, -\theta, z + \pi\Lambda), \\ v(r, \theta, z) &\rightarrow -v(r, -\theta, z + \pi\Lambda), \\ w(r, \theta, z) &\rightarrow w(r, -\theta, z + \pi\Lambda). \end{aligned} \quad (2.5)$$

The velocity field of the travelling wave of figure 1 is also unchanged by this discrete symmetry.

The magnitudes of disturbances and the norms of velocity fields are given in §4 and other places, using the square root of kinetic-energy norm. The kinetic energy, which is reported in tables such as table 1, is normalized to be 1 for laminar flow.

To conclude this section, we mention a technical point about pipe-flow simulation using spectral codes that appears not to have been discussed in the literature. Once the advection term is computed, the equations for evolving the modes decouple for pairs (m, n) such that the resulting equations depend only upon r for a fixed (m, n) . The

Re	λ_1	λ_2	λ_3
2000	0.03247	0.00897	-0.00594
2500	0.03049	0.00725	-0.02282 + i0.02041
3000	0.02861	0.00631	-0.01978 + i0.01664
4000	0.02529	0.00531	-0.01536 + i0.01190

TABLE 2. The only unstable eigenvalues are λ_1 and λ_2 ; λ_3 has the greatest real part among the stable eigenvalues whose eigenvectors lie in the shift-reflect invariant subspace.

decoupled equations will have terms with the factor $m^2/r^2 + n^2/\Lambda^2$ in the denominator, and because of that factor the terms will have a singularity at the point $r = -im\Lambda/n$ in the complex plane. If the number $2N$ of grid points in the streamwise direction is increased while keeping the pipe length $2\pi\Lambda$ fixed, that singularity moves closer to the real line with greater values of n now being allowed. When the singularity moves closer to the real line, one has to use more grid points in the r direction to solve the decoupled equations with the same level of accuracy (Trefethen 2000).

3. Unstable manifold of the travelling wave

To find disturbances of the laminar solution that evolve and hit the asymmetric travelling wave, it is essential to understand the unstable directions and the unstable manifold of that travelling wave. Suppose we disturb the laminar solution using rolls of the appropriate form and some ‘noise’, the magnitude of which is a fixed fraction of that of the rolls, to introduce streamwise dependence. The disturbed state will evolve and develop streaks. At the point of closest approach to the travelling wave, we can think of the evolving velocity field as the travelling wave plus two components, one of which is a combination of the stable eigenvectors of the travelling wave with the other being a combination of the unstable eigenvectors. The stable eigenvectors will decay under evolution. However, the component along the unstable eigenvectors will be amplified and will take the evolving velocity field away from the travelling wave. To ensure that the disturbed state hits the travelling wave, the disturbance has to be arranged in such a way that the evolving velocity field is free of the unstable directions as it approaches the travelling wave.

Such a disturbance is easiest to arrange if the travelling wave has only one unstable direction. The component along that direction at the point of closest approach can be eliminated by simply varying the magnitude of the initial disturbance. However, the asymmetric travelling wave has two unstable directions as shown in table 2. The two unstable eigenvalues λ_1 and λ_2 decrease with increasing Re at rates given by $Re^{-0.41}$ and $Re^{-0.87}$, respectively (Viswanath 2008a). Table 2 also shows the leading stable eigenvalue.

Figure 2 gives a more complete idea of the spectrum of the linearization around the asymmetric travelling wave. The spectra at different Re were computed using the Arnoldi iteration. Some of the interior eigenvalues near the centres of the circles in figure 2 are omitted. But we are certain that no unstable eigenvalues are omitted. In addition, we have verified that none of the eigenvalues in the figure is spurious.

Data for the unstable eigenvectors is given in table 3. Most of the kinetic energy of the travelling waves themselves is in the $n = 0$ (or mean) mode, as shown in table 1. Much of the kinetic energy remains in $n = 0$ for the λ_1 eigenvector, although $n = 1$




FIGURE 2. For eigenvalues λ of the travelling wave, the plots show $\exp(10\lambda)$ as circles for easier visualization. If the corresponding eigenvector lies in the shift-reflect invariant subspace, the circle is solid.

Re	λ_1				λ_2			
	ke_0	ke_1	ke_2	ke_3	ke_0	ke_1	ke_2	ke_3
2000	6.9×10^{-1}	2.6×10^{-1}	4.9×10^{-2}	1.1×10^{-3}	1.2×10^{-1}	7.4×10^{-1}	1.4×10^{-1}	3.1×10^{-3}
2500	7.0×10^{-1}	2.6×10^{-1}	4.2×10^{-2}	7.6×10^{-4}	1.1×10^{-1}	7.7×10^{-1}	1.2×10^{-1}	2.3×10^{-3}
3000	7.0×10^{-1}	2.6×10^{-1}	3.8×10^{-2}	6.2×10^{-4}	1.1×10^{-1}	7.8×10^{-1}	1.1×10^{-1}	1.9×10^{-3}
4000	7.0×10^{-1}	2.6×10^{-1}	3.2×10^{-2}	4.9×10^{-4}	1.1×10^{-1}	7.9×10^{-1}	9.4×10^{-2}	1.5×10^{-3}

TABLE 3. The kinetic energies in the $n = 0, \pm 1, \pm 2, \pm 3$ modes of the λ_1 and λ_2 eigenvectors. The eigenvectors are normalized to have kinetic energy equal to 1.




FIGURE 3. Contour plots of the z -averaged streamwise velocity for the two unstable eigenvectors at $Re = 2500$. If the eigenvectors are normalized to have unit kinetic energy, the level curves are equispaced in the intervals $[-1.38, 1.66]$ and $[-1.01, 0.34]$. The lighter regions correspond to higher values. At other values of Re , the signs of the eigenvectors are chosen to yield plots similar to the ones above.

now has more than a quarter of the kinetic energy. For the λ_2 eigenvector, the $n = 1$ mode dominates.

Figure 3 shows that the λ_1 eigenvector weakens the high speed streaks of the travelling wave. The effect of adding the λ_2 eigenvector to the travelling wave would




FIGURE 4. Plots of the mean streamwise flow $\bar{w}(r)$ at various values of Re . The eigenvectors are normalized to have kinetic energy equal to 1.

be to displace the high-speed streaks to a more symmetrical position. It must be noted, however, that the plot of the streaks of the λ_2 eigenvector is not as meaningful because the $n = 1$ mode is dominant. The two plots in figure 3 are used to assign positive and negative signs to the eigenvectors at different Re in a consistent manner. When comparing cross-sections of velocity fields to travelling-wave solutions (Hof *et al.* 2004; Eckhardt & Schneider 2008; Willis & Kerswell 2008), it may be worthwhile to look at the unstable eigenvectors of the travelling waves. The inevitable deviations from the streak patterns of the travelling waves may correlate with the streak patterns of the unstable eigenvectors.

Figure 4 shows the mean streamwise flow that corresponds to the λ_1 and λ_2 eigenvectors. To form an idea of the distortion to the mean flow of the laminar solution when those eigenvectors are added as disturbances, the plots in figure 4 must be scaled by a factor of $1/50$ or less.

Figure 5(a) shows that if the asymmetric travelling wave is disturbed with a small and positive multiple of the λ_1 eigenvector, the disturbed state evolves and becomes laminar uneventfully. That is unsurprising because the disturbance has the effect of weakening the high-speed streaks. In contrast, adding a negative multiple leads to what appears to be sustained turbulence at $Re = 2500$. It is easily noticeable that energy dissipation D is greater than energy input I when the plots in figure 5(a) spike up but is lesser when the plots dip down. Thus the kinetic energy of the velocity field as a whole decreases during the spikes but increases during the dips. The decrease of kinetic energy during a spike is well correlated with flattening of the mean velocity profile.

Figure 6(a) shows a schematic sketch of the unstable directions of the asymmetric travelling wave while distinguishing between directions that turn turbulent and ones that do not. Figure 5(b) corresponds to two trajectories close to the border between turbulent and laminar directions. In that figure, the trajectories near the border separate after $t = 300$. By refining the border, it appears that the point of separation can be deferred indefinitely with a view to locating edge states.

To better visualize the unstable manifold, we adopt a technique introduced by Gibson *et al.* (2008) with the aim of getting good phase-space visualizations of turbulent trajectories. The velocity field is a point in phase space, and the evolution of an initial velocity field with respect to the incompressible Navier–Stokes equation




FIGURE 5. (a) Plots of D (solid) and I (dashed) against time at $Re = 2000$ and $Re = 2500$. For both values of Re , perturbation of the travelling wave by a positive multiple of the λ_1 eigenvector shown in figure 4 leads to rapid laminarization (thin lines at lower left corner). Perturbation by a negative multiple leads to a long transient at $Re = 2000$ and what appears to be sustained turbulence at $Re = 2500$. (b) Plots close to the edge.

is a trajectory in that phase space. Because the phase space is infinite-dimensional and does not lend itself to plots directly, one has to use projections. An obvious projection would be to pick some Fourier–Chebyshev modes from the discretization of the velocity field. Although such projections have been employed, they have a number of shortcomings. The choices of the component of the velocity vector and of the mode of that component are both arbitrary. The component and the mode that are chosen capture only a partial aspect of the velocity field. As a result of these shortcomings, such projections look messy, and one cannot form a good idea of the dynamical structures in phase space from such projections.

Following Gibson *et al.* (2008), the projection we use picks a set of velocity fields that appears well suited to visualize trajectories on the unstable manifold. Let \mathbf{u}_{TW} be the velocity field of the travelling wave, and let \mathbf{u}' and \mathbf{u}'' be an orthonormal basis for its unstable space. The notion of orthogonality between velocity fields corresponds to the kinetic-energy norm. We will choose \mathbf{u}' to be the same direction as the leading eigenvector. With that choice the second eigenvector at $Re = 2500$ is approximately $-0.13\mathbf{u}' + 0.99\mathbf{u}''$. For each velocity field that satisfies the shift–reflect symmetry (2.5), we obtain a projection in terms of \mathbf{u}_{TW} , \mathbf{u}' and \mathbf{u}'' . The velocity fields \mathbf{u}' and \mathbf{u}'' satisfy the no-slip boundary condition and have zero mass flux.




FIGURE 6. (a) A schematic sketch of directions on the unstable manifold, where the two eigenvectors are labelled. The trajectories initiated in the dashed directions undergo turbulent episodes. (b) Projections of trajectories whose initial points are obtained by perturbing the travelling wave at $Re = 2500$ by a combination of the λ_1 and λ_2 modes. The initial points are all near $(0, 0)$. The trajectories that laminarize turn right, while those that transition to turbulence (dashed) turn left.

Given a velocity field \mathbf{u} that satisfies the shift–reflect symmetry, one can decompose it as $\mathbf{u} - \mathbf{u}_{TW} = c_0 \mathbf{u}' + c_1 \mathbf{u}'' + \mathbf{r}$, where the remainder \mathbf{r} is orthogonal to the plane of the eigenvectors. One could use c_0 and c_1 to represent \mathbf{u} in a plot, but that would be unsatisfactory. The problem is that one can translate \mathbf{u}_{TW} in the streamwise direction and obtain different velocity fields that stand for the same wave. In order to eliminate dependence on translations in the z direction, we shift the velocity field \mathbf{u} by s_z in the z direction and consider

$$\mathbf{u}(r, \theta, z + s_z) - \mathbf{u}_{TW} = c_0 \mathbf{u}' + c_1 \mathbf{u}'' + \mathbf{r}_{s_z}. \quad (3.1)$$

The shift s_z is chosen to minimize $|\mathbf{r}_{s_z}|$, and the axes of the projection, c_0 and c_1 , are the coefficients for that shift. The need to pick a shift s_z arises because the equations of pipe flow are unchanged by translations along z . The shift–reflect symmetry is broken by rotations in the θ direction. Since we have restricted ourselves to vector fields with the shift–reflect symmetry, shifts in θ are not considered in (3.1). The need to factor out continuous symmetries arises in ordinary differential equations (ODEs) (Gilmore & Letellier 2007) and partial differential equations (PDEs) such as the Kuramoto–Sivashinsky equation as well (Cvitanović, Davidchack & Siminos 2007).

Figure 6(b) shows trajectories on the unstable manifold using such a projection. The initial velocity fields were of the form $\mathbf{u}_{TW} + \epsilon a \mathbf{u}' + \epsilon b \mathbf{u}''$, with $\epsilon = \times 10^{-4}$ and a, b being scalars. In all, we considered 10 velocity fields corresponding to $(a, b) = (\pm 1, \pm 1)$, the four coordinate directions, and two directions along the second eigenvector. Since ϵ was small, all these velocity fields were very nearly on the unstable manifold. The distinction between trajectories that laminarize uneventfully and those that undergo a turbulent episode is clear in Figure 6(b).

In the next section, we return to such projections of the unstable manifold to partially justify arguments used to find disturbances of the laminar flow which evolve and hit the asymmetric travelling wave.

4. Hitting the travelling wave at $Re = 2000$ and $Re = 2500$

The f_r column of table 4 gives the norm of the rolls. The velocity field of the rolls, denoted by \mathbf{u}_r , is obtained by averaging the travelling wave in the streamwise direction and discarding the streamwise component of the velocity. The norms of the unstable eigenvectors are denoted by f_1 and f_2 . The velocity fields of the eigenvectors are denoted by \mathbf{u}_1 and \mathbf{u}_2 . The laminar solution is \mathbf{u}_L , and the travelling wave is \mathbf{u}_{TW} .

The distance δ of closest approach listed in table 4 is obtained as follows: The Navier–Stokes equation is integrated from the initial velocity field $\mathbf{u}_L + f_r \mathbf{u}_r + f_1 \mathbf{u}_1 + f_2 \mathbf{u}_2$. The initial velocity field has the shift–reflect symmetry, and so does \mathbf{u}_t , where \mathbf{u}_t is the velocity field at time t . Define

$$\delta(f_r, f_1, f_2) = \min_{t \geq 0} \min_{0 \leq s_z < 2\pi\Lambda} \|\mathbf{u}_t(r, \theta, z + s_z) - \mathbf{u}_{TW}\|. \quad (4.1)$$

To compare \mathbf{u}_t and \mathbf{u}_{TW} , one has to minimize over shifts s_z for the same reason as in (3.1). To find the minimizing shift s_z , we first try $s_z = \pi\Lambda k/N$, $0 \leq k < 2N$. Using that data, an interval that contains the minimum is found, and that interval is refined recursively to a depth equal to 30. We refer to the result of the inner minimum in (4.1) as the distance between \mathbf{u}_{TW} and \mathbf{u}_t . This method of finding that distance is expensive, with the cost of finding the distance being more than 20 times the cost of a single time step. However, it finds the distance with an accuracy of four or five digits.

Given the expense of finding the distance between \mathbf{u}_t and \mathbf{u}_{TW} , the distance being the inner minimum in (4.1), care has to be exercised in finding the outer minimum over t . If the distance is computed after every time step, the cost of the computation becomes prohibitive. The wall time for integrating a velocity field for a time interval of 100 is about an hour on an Opteron processor but becomes more than 20 hours if the distance to \mathbf{u}_{TW} is computed after every time step. For an initial waiting time when the streaks are still forming, we do not compute the distance at all. This waiting time is longer for larger Re . Thereafter the distance is computed every 100 time steps only, a time step being 0.01. As the distances vary smoothly as a function of time, we use polynomial extrapolation to predict if the distance function has a minimum within the next 100 time steps or not. If it is predicted to have a minimum within the next 100 time steps, we measure the distance every 10 time steps. If the distance function is predicted to have a minimum within the next 10 time steps, we measure the distance after every time step. The value of δ is the first local minimum found in this manner, and it is very probably also the global minimum over t . The time step can be successively decreased to get finer estimates of δ , but that was not implemented.

The times T at which the minima were attained are given in table 4; T is measured with a precision of 0.01 in only two lines of that table. The measurements in the other lines have a precision of 0.1. If the last four columns of table 4 are compared with the last four columns of table 1, the comparison confirms that the approach to the travelling wave is closer when δ is smaller. Figure 7 leaves no room for doubt that the first disturbances listed for $Re = 2000$ and $Re = 2500$ in table 4 evolve and hit the corresponding travelling waves. The smallest distance δ from the travelling wave is realized after time T , which is listed in table 4. Figure 7 shows that the plots of the distances from the travelling wave and the laminar solution both become flat around $t = T$. The disturbance moves away from the laminar solution rapidly at $t = 0$. In contrast, for heteroclinic connections there are two flat regions in t , which correspond to time spent in the neighbourhoods of the invariant solutions joined by the heteroclinic connection (Halcrow *et al.* 2009).

f_r	f_1	f_2	T	δ	ke_0	ke_1	ke_2	ke_3
7.084220×10^{-3}	-2.114944×10^{-2}	0	371.49	8.1×10^{-4}	9.8×10^{-1}	4.3×10^{-4}	2.2×10^{-5}	2.1×10^{-7}
1.242439×10^{-2}	1.890000×10^{-2}	0	219.70	8.9×10^{-3}	9.8×10^{-1}	3.3×10^{-4}	1.4×10^{-5}	1.0×10^{-7}
9.119378×10^{-3}	0	1.720663×10^{-2}	257.80	4.2×10^{-3}	9.8×10^{-1}	4.1×10^{-4}	2.0×10^{-5}	1.9×10^{-7}
1.017473×10^{-2}	0	-1.400000×10^{-2}	240.90	5.5×10^{-3}	9.8×10^{-1}	3.7×10^{-4}	1.7×10^{-5}	1.4×10^{-7}
9.048182×10^{-3}	3.923814×10^{-4}	1.757170×10^{-2}	281.40	2.9×10^{-3}	9.8×10^{-1}	4.2×10^{-4}	2.2×10^{-5}	2.0×10^{-7}
5.687034×10^{-3}	-1.643322×10^{-2}	0	323.84	2.9×10^{-3}	9.8×10^{-1}	2.5×10^{-4}	1.0×10^{-5}	7.7×10^{-8}
9.591329×10^{-3}	1.562075×10^{-2}	0	318.60	7.7×10^{-3}	9.8×10^{-1}	3.1×10^{-4}	1.5×10^{-5}	1.4×10^{-7}
7.049209×10^{-3}	0	1.420195×10^{-2}	268.00	4.8×10^{-3}	9.8×10^{-1}	2.4×10^{-4}	9.4×10^{-6}	6.9×10^{-8}
7.986395×10^{-3}	0	-1.115761×10^{-2}	309.80	6.0×10^{-3}	9.8×10^{-1}	2.2×10^{-4}	8.1×10^{-6}	5.5×10^{-8}
5.676511×10^{-3}	-1.646355×10^{-2}	8.768074×10^{-5}	292.70	3.7×10^{-3}	9.8×10^{-1}	2.9×10^{-4}	1.3×10^{-5}	1.2×10^{-7}

TABLE 4. Data at $Re = 2000$ (above the double line) and at $Re = 2500$ (below the double line). The various f are magnitudes of disturbances of the laminar solution; T is the time of closest approach to the travelling wave; δ is the distance from the travelling wave at that time. The last four columns give the kinetic energies in the $n = 0, \pm 1, \pm 2, \pm 3$ modes at the point of closest approach.




FIGURE 7. Distances from the laminar solution (dashed) and the travelling wave (solid). The three pairs of curves correspond to the first line with $Re = 2000$ in table 4 and the first two lines with $Re = 2500$. Data in that table may be used to deduce the assignment of curves above the rows in that table.

We are yet to explain the method used to find the numbers f_r , f_1 and f_2 in table 4. With each of those disturbances to the laminar solution, the disturbed state lands close to the stable manifold of the travelling wave and evolves to make a close approach of small δ to the travelling wave. Each line in table 4 was obtained by minimizing $\delta(f_r, f_1, f_2)$ in different ways. The manner of minimization will now be described.

Although δ has three arguments corresponding to three disturbances, each minimization was two-dimensional. The unstable manifold of the travelling wave at the Re under consideration is two-dimensional as shown by table 2. If the directions that correspond to translating and rotating the travelling wave are ignored, the co-dimension of the travelling wave is two. The inner minimization in (4.1) accounts for streamwise translations, while rotations around the pipe axis would break the shift–reflect symmetry. Thus the stable manifold is in effect a co-dimension two object.

Suppose there is some system of coordinates for the infinite-dimensional phase space in which the travelling wave is $(0, 0, 0, \dots)$. Suppose further that its unstable manifold is given by fixing all except the first two coordinates at zero and that its stable manifold is obtained by fixing the first two coordinates at zero. To disturb the laminar solution on to the stable manifold, the first two components must be zeroed out. Generically, it is impossible to zero out two components by varying the amplitude of a single disturbance. That is why we varied two disturbances.

All the rows of table 4 have $f_r > 0$. Adding the rolls to the laminar solution causes the flow to develop streaks of approximately the right form. But the key to hitting the travelling wave is to disturb the laminar state in such a way that the evolving velocity field is free from unstable directions as it approaches the travelling wave. If we were allowed to make disturbances near the travelling wave, we could simply eliminate the unstable directions. Since the disturbances are made to the laminar solution, we have to somehow guess the directions at $t = 0$ which evolve into unstable directions at $t = T$, the time of closest approach. Analysis can possibly suggest a better choice, but we simply used the unstable directions \mathbf{u}_1 and \mathbf{u}_2 to disturb the laminar flow. In




FIGURE 8. Similar to figure 6(a), but the thick lines show the projections of trajectories at $Re = 2500$ which are initialized using the disturbances in table 4.




FIGURE 9. The axes correspond to the T and δ columns in table 4. The plot is for $Re = 2000$ and shows the progress of a single minimization (first line of table 4). Each point corresponds to a certain stage in the sequence of optimizations used to find a disturbance of the laminar flow such that the disturbed flow evolves and hits the travelling wave.

addition the magnitude of the rolls themselves can be varied to get a direction that has a non-zero component along the unstable manifold at $t = T$.

Each row of table 4 that has $f_1 = 0$ and $f_2 = 0$ was obtained by fixing that disturbance at 0 and minimizing over the other two disturbances. In addition, the sign of the other disturbance that adds an eigenvector was prescribed. Thus there are four rows of that type for $Re = 2000$ and $Re = 2500$. For the last row with $Re = 2000$ or $Re = 2500$, f_r was fixed, while f_1 and f_2 were varied. Figure 8 shows that for disturbances at $Re = 2500$, the flow evolves to a state in which its dynamics is governed mainly by the unstable manifold of the travelling wave, thus partially supporting the reasoning used to find those disturbances.

When $\delta(f_r, f_1, f_2)$ is minimized numerically, the disturbances found at successive stages of the minimization give smaller δ but with larger values of T , the time of closest approach to the travelling wave, as shown in figure 9. For the theoretical ideal

$\delta = 0$, T would be infinite. Thus the numerical optimization becomes progressively more expensive.

A severer impediment to numerical minimization is the non-smooth dependence of δ on the disturbances when $\delta = 0$. Because the time to hit the travelling wave diverges, even a small change in the disturbances causes a big change in the value of δ .

The numerical optimization was implemented using MATLAB's *fmincon()*, which allows constraints to be placed on the values of f_r or f_1 or f_2 . The C++ code for computing the function $\delta(f_r, f_1, f_2)$ was invoked from MATLAB. The unconstrained version *fminunc()* was not used because it tends to take such large steps while varying f_r or f_1 or f_2 that the numerical integration of the Navier–Stokes equation becomes unstable. Because of the non-smoothness, a nonlinear least squares solver, such as MATLAB's *lsqnonlin()*, might be a better option than *fmincon()*; *lsqnonlin()* minimizes $\sqrt{|x+2|}$ from $x = 3$ with just 6 function evaluations, while *fmincon()* takes 63 function evaluations to find a slightly worse approximation to the minimum. It was not used, however, because it does not provide the facility to constrain the arguments.

The choice of initial guesses for the disturbances is not much of an issue because the numerical optimization is relatively efficient at the early stages. However, as $\delta = 0$ is approached, the optimization routine tries unrealistically large steps, necessitating a lot of wasteful backtracking. It is difficult to assess the quality of the search directions. Each row in table 4 required at least 200 hours of computing and often significantly more. The first rows with $Re = 2000$ and $Re = 2500$ required much more than 1000 hours to attain smaller values of δ . The computations required repeated manual intervention to reset the parameters to *fmincon()*, which is the reason we were able to run the numerical optimization longer for only two rows of table 4.

Table 5 shows that the magnitude of the disturbances of the laminar flow required to hit the travelling wave diminishes with Re . The quality of the approach to the travelling wave degrades with increasing Re . The quality of the approach can be assessed using the δ column of table 5 and by comparing the last four columns of that table to the last four columns of table 1. We are not certain why the numerical minimization has worse performance for increasing Re , although it could be because the non-smoothness issue gets worse as Re increases. The tendency of the eigenvalues to approach the imaginary axis at varying rates as Re is increased may have something to do with it.

5. Conclusion

The aim of this paper was to find a method to disturb the laminar flow such that the disturbed state evolves and hits a given travelling wave. For the asymmetric travelling wave of Pringle & Kerswell (2007), we showed that certain linear combinations of rolls and the two unstable directions generate disturbances of the laminar solution that evolve and hit the travelling wave. Numerical minimization was used to find linear combinations that achieve that effect. As the numerical minimization comes closer and closer to finding disturbances that evolve and hit the travelling wave, the minimization problem becomes non-smooth.

It is reasonable to conjecture that this method of disturbing the laminar flow so that it evolves and hits a given travelling wave is applicable to other lower branch solutions of pipe flow and the channel flows. However, the computational effort will increase with the number of unstable directions. As indicated in the text, it may

Re	f_r	f_1	f_2	T	δ	ke_0	ke_1	ke_2	ke_3
2000	9.119378×10^{-3}	0	1.720663×10^{-2}	257.80	4.2×10^{-3}	9.8×10^{-1}	4.1×10^{-4}	2.0×10^{-5}	1.9×10^{-7}
2500	7.049209×10^{-3}	0	1.420195×10^{-2}	268.00	4.8×10^{-3}	9.8×10^{-1}	2.4×10^{-4}	9.4×10^{-6}	6.9×10^{-8}
3000	6.466047×10^{-3}	0	-9.890996×10^{-3}	356.61	7.0×10^{-3}	9.8×10^{-1}	2.2×10^{-4}	9.4×10^{-6}	8.4×10^{-8}
4000	4.600000×10^{-3}	0	8.829559×10^{-3}	386.95	9.5×10^{-3}	9.8×10^{-1}	6.9×10^{-5}	1.5×10^{-6}	7.1×10^{-9}

TABLE 5. Data at various Re . The columns are as in table 4.

be possible to use analysis to find disturbances that work better than the unstable eigenvectors of the travelling wave.

It is yet unknown if spatially localized structures such as puffs in transitional pipe flow and turbulent spots in plane Couette flow correspond in their entirety to invariant solutions of the Navier–Stokes equation. The existence of such invariant solutions is a topic worthy of investigation. If such solutions are indeed found, the logic used to find disturbances that evolve and hit the asymmetric travelling wave will become applicable to transition in pipes of realistic lengths with consequences for experimental investigations.

The authors thank the referees and J. F. Gibson for helpful discussions. They are particularly grateful to a referee for catching a mad misstatement in the caption of figure 9. The Center for Advanced Computing at the University of Michigan provided computing facilities. DV was partly supported by NSF grants DMS-0407110 and DMS-0715510. PC was partly supported by NSF grant DMS-0807574.

REFERENCES

- CVITANOVIĆ, P., DAVIDCHACK, R. L. & SIMINOS, E. 2007 State space geometry of a spatio-temporally chaotic Kuramoto–Sivashinsky flow. *arXiv:0709.2944*.
- DARBYSHIRE, A. G. & MULLIN, T. 1995 Transition to turbulence in constant-mass-flux pipe flow. *J. Fluid Mech.* **289**, 83–114.
- DUGUET, Y., WILLIS, A. P. & KERSWELL, R. R. 2008 Transition in pipe flow: the saddle structure on the boundary of turbulence. *J. Fluid Mech.* **613**, 255–274. *arXiv:0711.2175*.
- ECKHARDT, B. & SCHNEIDER, T. 2008 How does flow in a pipe become turbulent? *Euro. Phys. J. B*, **64**, 457–462.
- 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.
- GIBSON, J. F., HALCROW, J. & CVITANOVIĆ, P. 2008 Visualizing the geometry of state space in plane Couette flow. *J. Fluid Mech.* **611**, 107–130. *arXiv:0705.3957*.
- GILMORE, R. & LETELLIER, C. 2007 *The Symmetry of Chaos*. Oxford University Press.
- HALCROW, J., GIBSON, J. F., CVITANOVIĆ, P. & VISWANATH, D. 2009 Heteroclinic connections in plane Couette flow. *J. Fluid Mech.* **611**, 365–376.
- HOF, B., VAN DOORNE, C. W. H. *et al.* 2004 Experimental observation of nonlinear traveling waves in turbulent pipe flows. *Science* **305**, 1594–1598.
- HOF, B., JUEL, A. & MULLIN, T. 2003 Scaling of the turbulence transition threshold in a pipe. *Phys. Rev. Lett.* **91**, 244502.
- ITANO, T., & TOH, S. 2001 The dynamics of bursting process in wall turbulence. *J. Phys. Soc Jpn* **70**, 701–714.
- KAWAHARA, G. 2005 Laminarization of minimal plane Couette flow: going beyond the basin of attraction of turbulence. *Phys. Fluids* **17**, 041702.
- KERSWELL, R. R. 2005 Recent progress in understanding the transition to turbulence in a pipe. *Nonlinearity* **18**, R17–R44.
- KERSWELL, R. R. & TUTTY, O. R. 2007 Recurrence of travelling waves in transitional pipe flow. *J. Fluid Mech.* **584**, 69–102.
- KREISS, G., LUNDBLADH, A. & HENNINGSON, D. S. 1994 Bounds for threshold amplitudes in subcritical shear flows. *J. Fluid Mech.* **270**, 175–198.
- LANDAHL, M. T. 1980 A note on an algebraic instability of inviscid parallel shear flows. *J. Fluid Mech.* **98**, 243–251.
- MELLIBOVSKY, F. & MESEGUER, A. 2006 The role of streamwise perturbations in pipe flow transition. *Phys. Fluids* **18**, 074104.
- MELLIBOVSKY, F. & MESEGUER, A. 2007 Pipe flow transition threshold following localized impulsive perturbations. *Phys. Fluids* **19**, 044102.

- NAGATA, M. 1990 Three dimensional finite amplitude solutions in plane Couette flow: bifurcation from infinity. *J. Fluid Mech.* **217**, 519–527.
- NARASIMHA, R. 1989 The utility and drawbacks of traditional approaches. In *Whither Turbulence? Turbulence at the Cross-Road* (ed. J. Lumley), pp. 13–48. Springer.
- O’SULLIVAN, P. L. & BREUER, K. S. 1994 Transient growth in circular pipe flow. Part 2. Nonlinear development. *Phys. Fluids* **6**, 3652–3664.
- PEIXINHO, J. & MULLIN, T. 2007 Finite-amplitude thresholds for transition in pipe flow. *J. Fluid Mech.* **582**, 169–178.
- PRINGLE, C. C. T. & KERSWELL, R. R. 2007 Asymmetric, helical, and mirror-symmetric traveling waves in pipe flow. *Phys. Rev. Lett.* **99**, 074502.
- REYNOLDS, O. 1883 An experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channels. *Phil. Trans. R. Soc. Lond.* **174**, 935–982.
- ROBINSON, S. K. 1991 Coherent motions in the turbulent boundary layer. *Annu. Rev. Fluid Mech.* **23**, 601–639.
- SCHMIEGEL, A. & ECKHARDT, B. 1997 Fractal stability border in plane Couette flow. *Phys. Rev. Lett.* **79**, 5250.
- SCHNEIDER, T. M., ECKHARDT, B. & VOLLMER, J. 2007a Statistical analysis of coherent structures in transitional pipe flow. *Phys. Rev. E* **75**, 066313.
- SCHNEIDER, T. M., ECKHARDT, B. & YORKE, J. A. 2007b Turbulence transition and edge of chaos in pipe flow. *Phys. Rev. Lett.* **99**, 034502.
- SCHNEIDER, T. M., GIBSON, J. F., LAGHA, M., DE LILLO, F. & ECKHARDT, B. 2008 Laminar-turbulent boundary in plane Couette flow. *Phys. Rev. E* **78**, 037301. arXiv:0805.1015.
- TOH, S. & ITANO, T. 2003 A periodic-like solution in channel flow. *J. Fluid Mech.* **481**, 67–76.
- TREFETHEN, L. N. 2000 *Spectral Methods in Matlab*. SIAM.
- VISWANATH, D. 2007 Recurrent motions within plane Couette turbulence. *J. Fluid Mech.* **580**, 339–358.
- VISWANATH, D. 2008a The critical layer in pipe flow at high Re . *Phil. Trans. R. Soc. A.* **367**, 561–576.
- VISWANATH, D. 2008b The dynamics of transition to turbulence in plane Couette flow. In *Mathematics and Computation, a Contemporary View: The Abel Symposium 2006* (H. Munthe-Kaas & B. Owren), vol. 3, pp. 109–127. Springer. arXiv:0701337.
- WALEFFE, F. 1998 Three-dimensional coherent states in plane shear flows. *Phys. Rev. Lett.* **81**, 4140–4143.
- WANG, J., GIBSON, J. F. & WALEFFE, F. 2007 Lower branch coherent states in shear flows: transition and control. *Phys. Rev. Lett.* **98**, 204501.
- 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. 2008 Coherent structures in localised and global pipe turbulence. *Phys. Rev. Lett.* **100**, 124501.
- WILLIS, A. P. & KERSWELL, R. R. 2009 Turbulent dynamics of pipe flow captured in a reduced model: puff relaminarisation and localised ‘edge’ states. *J. Fluid Mech.* **619**, 213–233.
- WILLIS, A. P., PEIXINHO, J., KERSWELL, R. R. & MULLIN, T. 2008 Experimental and theoretical progress in pipe flow transition. *Phil. Trans. R. Soc. A.* **366**, 2671–2684.
- 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**, 281–335.
- WYGNANSKI, I., SOKOLOV, M. & FRIEDMAN, D. 1975 On transition in a pipe. Part 2. The equilibrium puff. *J. Fluid Mech.* **69**, 283–304.