Turbulent dynamics of pipe flow captured in a reduced model: puff relaminarisation and localised 'edge' states

ASHLEY P. WILLIS
AND
RICH R. KERSWELL
Department of Mathematics, University of Bristol, Bristol, U.K., BS8 1TW.

(Received ?? and in revised form ??)

Fully 3-dimensional computations of flow through a long pipe demand a huge number of degrees of freedom making it very expensive to explore parameter space and difficult to isolate the structure of the underlying dynamics. We therefore introduce a '$2\frac{1}{2}$-dimensional' model of pipe flow which is a minimal 3-dimensionalisation of the axisymmetric case: only sinusoidal variation in azimuth plus azimuthal shifts are retained, yet the same dynamics familiar from experiments are found. In particular the model retains the subcritical dynamics of fully resolved pipe flow, capturing realistic localised 'puff'-like structures which can decay abruptly after long times, as well as global 'slug' turbulence. Relaminarisation statistics of puffs reproduce the memoryless feature of pipe flow and reiterate the existence of a critical Reynolds number above which they become permanent, provided that the pipe is long enough to permit localisation. In a short pipe, exact travelling-wave solutions are found nearby to flow trajectories on the boundary between laminar and turbulent flow, with visitations occurring more frequently as the flow rate is increased. In a long pipe, the attracting state on the laminar-turbulent boundary is a localised structure which resembles a smoothened puff. This 'edge' state remains localised even for Reynolds numbers where the turbulent state is global suggesting that a slug is a destabilised puff rather than being a completely separate state.

1. Introduction

Laminar flow through a pipe is possible under controlled laboratory conditions up to flow rates well beyond those at which turbulence is typically observed. Pfenniger (1961) achieved laminar flow at Reynolds numbers, $Re := UD/\nu$, as high as 100 000, where U is the mean axial speed, D the diameter and ν the kinematic viscosity, indicating that, rather than the transition to turbulence being via a linear instability, some other mechanism must be responsible. Given an initial disturbance of sufficiently large amplitude, self-sustained turbulence is observed for Re of approximately 2000. This turbulent flow exhibits distinct spatial structures at different flow rates. For Re up to around 2250 the region of turbulence remains localised, with a length of approximately $20\,D$, and is referred to as a 'puff' (Wygnanski & Champagne 1973). At larger flow rates these puffs slowly delocalise by splitting into two or more puffs. At much larger Re, of around 2800, the disturbances develop into a rapidly expanding active region of turbulence, referred to as a 'slug'. No explanation has been offered that predicts such a progression in flow regimes and many issues remain unresolved.

The dynamics of perturbations at transitional Reynolds numbers is believed to be strongly influenced by a rapidly increasing number of branches of exact solutions that have been found to appear at these Re (Faisst & Eckhardt 2003; Wedin & Kerswell 2004; Kerswell 2005; Pringle & Kerswell 2007). At $Re < 1750$, puffs are observed to suddenly and unexpectedly decay in experiments (Peixinho & Mullin 2006), and it has been suggested that the turbulent state wanders between these unstable solutions before relaminarising (Faisst & Eckhardt 2004). The same data also suggest that the mean lifetime for a puff becomes infinite at $Re = 1750$ indicating that the puffs become permanent states at this point (Peixinho & Mullin 2006). This critical Re has been reproduced to within 7% using numerical computations which adopted the experimental protocol for initiating the puffs and worked within a long periodic pipe of $50\,D$, so as to realistically capture the puff structure (Willis & Kerswell 2007*b*). However, experiments using a different way of initiating the puffs and designed to capture longer puff transients claim that no such critical Re exists (Hof *et al.* 2006). Simulations also presented there in a short $\approx 6\,D$ pipe appear to support this conclusion. The obvious question is then whether numerically-simulated turbulence which fills a short pipe has the same relaminarisation characteristics as localised puff turbulence captured in longer numerical domains. A complete statistical study using fully-resolved 3-dimensional computations across a spectrum of periodic pipe domains remains prohibitively expensive, whereas a survey using a realistic model system could provide a clarifying demonstration of difference.

Evidence has also emerged recently that the many of the exact solutions known thus far sit on a separatrix between laminar and turbulent states, forming an 'edge' to the chaotic region of phase space (Schneider *et al.* 2007*b*; Kerswell & Tutty 2007; Willis & Kerswell 2007*a*; Duguet *et al.* 2007) (see Itano & Toh (2001); Wang *et al.* (2007); Viswanath (2007) for similar observations in channel and plane Couette flow). Schneider et al. (2007) have examined the dynamics of flow restricted to lie in this separatrix in a short $5\,D$ pipe, finding at long times a chaotic attractor apparently centred on a simple travelling wave solution (Pringle & Kerswell 2007). In such a short pipe, the turbulence naturally fills the pipe when triggered and the laminar-turbulent boundary endstate, or 'edge' state, is also a global state. In a longer ($\geqslant 25D$) pipe, however, localised puffs are the naturally triggered state at low Re which raises the issue of how the corresponding 'edge' state looks and how its variation with Re. For example, is it initially localised and does it lose localisation at the same Re as the turbulent puff? Again, a realistic model system can suggest probable answers to these questions quickly.

The use of model systems is well established in plane Couette flow, which exhibits the same abrupt subcritical transition behaviour as pipe flow. Several approaches have been designed to reduce the number of degrees of freedom of this problem in order develop more tractable models. The minimal flow unit introduced by Jiménez & Moin (1991) has been useful in identifying the key components that lead to self-sustaining turbulence in a very small domain (Hamilton *et al.* 1995). The model by Lagha & Manneville (2007) severely truncates the degrees of freedom in the cross-stream direction, but captures spanwise and streamwise spatial structures observed in plane Couette flow. Using this model, attempts have been made to measure the lifetime of localised turbulence and to determine the characteristic structures seen during the relaminarisation process itself. Currently, such calculations would be prohibitively expensive for fully 3-dimensional models. Severe truncation to only a few Fourier modes in the tilted cross-wise direction has also proven useful in determining the origin of oblique bands in plane Couette flow (Barkley & Tuckerman 2007).

The aim of this article is to establish a model system which preserves the rich dynamics of pipe flow but reduces the number of degrees of freedom of the system, so considerably

that the two current issues mentioned above can be probed: (*a*) Do the relaminarisation statistics for numerically-simulated turbulent puffs differ in character between short and long pipes? In particular, do long-pipe simulations indicate a critical Re for sustained puffs whereas short-pipe simulations not? (*b*) What does the attracting 'edge' state in the laminar-turbulent boundary look like in a long pipe? Is it localised like a puff and, if so, does it delocalise at the same Re as a puff?

Previous attempts to find such a reduced model have focused on axisymmetric pipe flow (Patera & Orszag 1981) and helical pipe flow (Landman 1990*a*,*b*), but in both cases the subcritical dynamics of pipe flow is not retained. We briefly revisit these calculations to search afresh for evidence of turbulent transients before introducing a new $2\frac{1}{2}$-dimensional model which retains the salient features of fully 3-dimensional pipe flow. The presentation starts by discussing the formulation used for the calculations performed throughout this and earlier work (Willis & Kerswell 2007*b*,*a*).

2. Formulation

For lengths scaled by the diameter D and velocities by the fixed mean axial speed U, the non-dimensional form of the Navier–Stokes equation is

$$(\partial_t + \boldsymbol{u}\cdot\boldsymbol{\nabla})\boldsymbol{u} = -\boldsymbol{\nabla}p + \frac{32}{Re}\,(1+\beta)\,\hat{\boldsymbol{z}} + \frac{1}{Re}\boldsymbol{\nabla}^2\boldsymbol{u}, \tag{2.1}$$

where the non-dimensional variable β is the fractional pressure, additional to the laminar flow, required to maintain a steady U. A Reynolds number Re_p, based on the applied pressure gradient, is given by $Re_p = Re\,(1+\beta)$. Our numerical formulation is based on the potential formulation of Marqués (1990), which is further re-expressed to ease numerical solution. An averaging operator is introduced in the axial direction, z, which is periodic over a length $2\pi/\alpha$,

$$P_z\,(\,\cdot\,) = \frac{\alpha}{2\pi}\int_0^{2\pi/\alpha} (\,\cdot\,)\;\mathrm{d}z. \tag{2.2}$$

The velocity, $\boldsymbol{u}$, is then expressed in terms of a potential $\psi = \psi(r,\theta,z)$, the axially independent flow $h = h(r,\theta)$ and a purely axially-dependent potential $\phi = \phi(r,\theta,z)$,

$$\boldsymbol{u} = h\,\hat{\boldsymbol{z}} + \boldsymbol{\nabla}\wedge(\hat{\boldsymbol{z}}\psi) + \boldsymbol{\nabla}\wedge\boldsymbol{\nabla}\wedge(\hat{\boldsymbol{z}}\phi), \tag{2.3}$$

such that $P_z\,\phi = 0$. Writing the nonlinear terms as $\boldsymbol{b} = (\boldsymbol{u}\cdot\boldsymbol{\nabla})\,\boldsymbol{u}$, the governing equations become

$$\begin{aligned}
(\partial_t - \tfrac{1}{Re}\nabla^2)\,h &= -P_z\,\hat{\boldsymbol{z}}\cdot\boldsymbol{b},\\
(\partial_t - \tfrac{1}{Re}\nabla^2)\,\nabla^2\nabla_h^2\phi &= -(1-P_z)\,\hat{\boldsymbol{z}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{\nabla}\wedge\boldsymbol{b},\\
(\partial_t - \tfrac{1}{Re}\nabla^2)\,\nabla_h^2\psi &= \hat{\boldsymbol{z}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{b},
\end{aligned} \tag{2.4}$$

where $\nabla_h^2 := \nabla^2 - \partial_{zz}$. For a boundary condition $\boldsymbol{u} = \boldsymbol{g}(\theta,z)$, conditions on the potentials are

$$\begin{gathered}
h = 0, \quad \phi = 0, \quad -\partial_r\psi = g_\theta, \quad -\nabla_h^2\phi = g_z, \\
\tfrac{1}{r}\,\partial_\theta\psi + \partial_{rz}\phi = g_r, \quad \partial_{rz}\nabla_h^2\psi - \tfrac{1}{r}\partial_\theta\nabla^2\nabla_h^2\phi = Re\,\hat{\boldsymbol{r}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{b} - \partial_{zz}\,\hat{\boldsymbol{r}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{g}
\end{gathered} \tag{2.5}$$

(see appendix A for details). Note that $\hat{\boldsymbol{r}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{b} = 0$ on the boundary unless an internal body force is added. Variables are expanded in Fourier modes,

$$A(r,\theta,z) = \sum_{k,m} A_{km}(r)\;\exp(\mathrm{i}\alpha kz + \mathrm{i}m\theta). \tag{2.6}$$

As the variables are real, their coefficients satisfy the property $A_{km} = A^*_{-k,-m}$, where * indicates the complex conjugate, and therefore only coefficients with $m \geqslant 0$ are kept. Numerical truncation in k and m is discussed in following sections. The operator P_z picks out $k = 0$ modes and $(1 - P_z)$ retrieves $k \neq 0$ modes. In addition to the boundary conditions (2.5), regularity at the axis imposes symmetries on the Fourier modes across the axis. For the potentials, each mode is even(odd) in r if m is even(odd).

The system for h is simple to solve as it is second order and has two boundary conditions on h, one at the boundary and a symmetry condition at the axis. The system for ψ and ϕ is more difficult to invert as it is coupled through the boundary condition. To enable numerical solution we reformulate the system for ψ and ϕ into a set of five equations, each second-order in r. Each Fourier mode for ψ,ϕ is expanded as the superposition

$$\begin{aligned} \psi(r) &= \bar{\psi}(r) + a\,\psi^H(r), \\ \phi(r) &= \bar{\phi}(r) + b\,\phi^H(r), \end{aligned} \tag{2.7}$$

where the coefficients a and b are scalars. Subscripts k and m have been dropped. The barred and superscripted functions solve two distinct systems. Firstly,

$$\begin{aligned} (\partial_t - \tfrac{1}{Re}\nabla^2)\,\bar{\phi}_2 &= -(1 - P_z)\,\hat{\boldsymbol{z}} \cdot \boldsymbol{\nabla} \wedge \boldsymbol{\nabla} \wedge \boldsymbol{b}, \qquad (2.8) \\ \nabla^2 \bar{\phi}_1 &= \bar{\phi}_2, \\ \nabla_h^2 \bar{\phi} &= \bar{\phi}_1, \\ (\partial_t - \tfrac{1}{Re}\nabla^2)\,\bar{\psi}_1 &= \hat{\boldsymbol{z}} \cdot \boldsymbol{\nabla} \wedge \boldsymbol{b}, \qquad (2.9) \\ \nabla_h^2 \bar{\psi} &= \bar{\psi}_1, \end{aligned}$$

with boundary conditions

$$\bar{\phi}_2 = \bar{\phi} = 0, \quad -\bar{\phi}_1 = g_z, \quad \partial_r \bar{\psi}_1 = 0, \quad \begin{cases} \bar{\psi} = 0 & \text{if } m = 0 \\ -\partial_r \bar{\psi} = g_\theta & \text{if } m \neq 0 \end{cases}, \tag{2.10}$$

where $\nabla^2 \equiv (1/r)\partial_r + \partial_{rr} - m^2/r^2 - \alpha^2 k^2$ and $\nabla_h^2 \equiv (1/r)\partial_r + \partial_{rr} - m^2/r^2$. This time-dependent system is written in matrix-vector form, according to the time and radial discretisation, then inverted sequentially for $\bar{\phi}_2 \to \bar{\phi}_1 \to \bar{\phi}$ and $\bar{\psi}_1 \to \bar{\psi}$. The second homogenised system is

$$\begin{aligned} (\partial_t - \tfrac{1}{Re}\nabla^2)\,\phi_2^H &= 0, \qquad (2.11) \\ \nabla^2 \phi_1^H &= \phi_2^H, \\ \nabla_h^2 \phi^H &= \phi_1^H, \\ (\partial_t - \tfrac{1}{Re}\nabla^2)\,\psi_1^H &= 0, \qquad (2.12) \\ \nabla_h^2 \psi^H &= \psi_1^H, \end{aligned}$$

with boundary conditions

$$\phi_2^H = 1, \quad \phi_1^H = \phi^H = 0, \quad \partial_r \psi_1^H = 1, \quad \begin{cases} \psi^H = 0 & \text{if } m = 0 \\ \partial_r \psi^H = 0 & \text{if } m \neq 0 \end{cases}. \tag{2.13}$$

As this system has no time-dependent $\boldsymbol{b}$, solutions with superscript H may be pre-computed. The original boundary conditions on ϕ and ψ are satisfied upon reconstruction from the barred and superscripted variables. Two boundary conditions are satisfied by construction for all cases: as $\bar{\phi}$ and ϕ^H satisfy trivial boundary conditions, then $\phi = \bar{\phi} + b\,\phi^H = 0$ on the boundary; similarly $-\phi_1 = g_z$ is satisfied automatically.

For axisymmetric modes, $m = 0$, the system is of lower order as ψ always appears as

$\partial_r\psi$, including in the boundary condition. The condition involving $\boldsymbol{b}$ is not required for this case (see appendix A). The simplest solution is to add the boundary condition $\psi = 0$ so that we may invert for ψ for all modes. When $m = 0$ the remaining two boundary conditions which couple the potentials are satisfied by selecting scalars a and b according to the following evaluated on the boundary:

$$a = -(\partial_r\bar{\psi} + g_\theta)/\partial_r\psi^H, \quad b = -(\partial_{rz}\bar{\phi} - g_r)/\partial_{rz}\phi^H \quad \text{if } m = 0. \tag{2.14}$$

For non-axisymmetric modes, $m \neq 0$, the condition $-\partial_r\psi = g_\theta$ is satisfied automatically. The last two conditions are satisfied by solving the system for a and b evaluated on the boundary,

$$\begin{bmatrix} \frac{1}{r}\partial_\theta\psi^H & \partial_{rz}\phi^H \\ -\partial_{rz}\psi_1^H & \frac{1}{r}\partial_\theta\phi_2^H \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = - \begin{bmatrix} \frac{1}{r}\,\partial_\theta\bar{\psi} + \partial_{rz}\bar{\phi} - g_r \\ Re\,\hat{\boldsymbol{r}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{b} - \partial_{zz}\,\hat{\boldsymbol{r}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{g} \end{bmatrix} \quad \text{if } m \neq 0. \tag{2.15}$$

As this only requires the inversion of a 2×2 matrix and the H-functions are pre-computed, this is an inexpensive way to ensure all boundary conditions are simultaneously satisfied to machine precision.

Both finite-difference and Chebyshev expansions have been used in radius. The latter is better at low radial resolution, but the former involves only banded matrices (a 9-point stencil is used), requiring less memory, and is faster for high radial resolutions. Time discretisation is second-order using Crank–Nicolson for the diffusion term and an Euler predictor step for the non-linear terms. Information from a Crank–Nicolson corrector step is used to control the timestep size. Nonlinear terms, $\boldsymbol{b}$, are evaluated using the pseudo-spectral method and are dealiased using the $\frac{3}{2}$-rule. The code was tested to reproduce eigenvalues about the laminar state, eigenvalues about nonlinear travelling wave solutions from Wedin & Kerswell (2004), by direct comparison with a primitive variable code (Kerswell & Tutty 2007) during the relaminarisation of a perturbed travelling wave, and to calculate the turbulent statistics of Eggels *et al.* (1994).

3. The absence of turbulence in previous models

The original calculations for axisymmetric pipe flow (Patera & Orszag 1981) and for helical flow (Landman 1990*a*) found no evidence for turbulence or even long transients at $Re \leqslant 4000$. Here we show that this conclusion extends to Re as large as 10^5 and for huge initial disturbances suggesting that there are no exact unstable solutions beyond Hagen-Poiseuille flow within these dynamical subspaces.

The axisymmetric model is straightforward to simulate using the numerical algorithm described above by time stepping only modes with $m = 0$. In helical flow, variations in θ and z are reduced to a dependency in the one variable $\xi = \theta + \alpha z$, where α is the pitch of the helix and periodicity over $L = 2\pi/\alpha$ is preserved. In this scenario the flow may be expanded $\boldsymbol{u}(r,\xi) = \sum_q \boldsymbol{u}_q(r)\exp(\mathrm{i}q\xi)$, corresponding to taking only modes $k = m \to q$ in our formulation. Letting $\boldsymbol{u}'$ be the deviation from the laminar flow, random initial disturbances of the form

$$\boldsymbol{u}' = \sum_{k^2+m^2\neq0} r^2(\tfrac{1}{2}-r)^2\,(\alpha^2k^2+m^2)^{-\frac{1}{2}}\,\boldsymbol{a}_{km}\exp(\mathrm{i}\alpha kz + \mathrm{i}m\theta), \tag{3.1}$$

were applied to the flow (after projection onto the space of solenoidal functions to enforce incompressibility), where the components of $\boldsymbol{a}_{km}$ were random numbers in $\mathbb{C}$ s.t. $|a_{km}| = 1$.

Figure 1 shows time evolution of random disturbances at $Re = 10\,000$ for a pipe of length 2π diameters. Initial disturbances were normalised for $E'_{k\neq0}/E_0$ up to 0.4, where

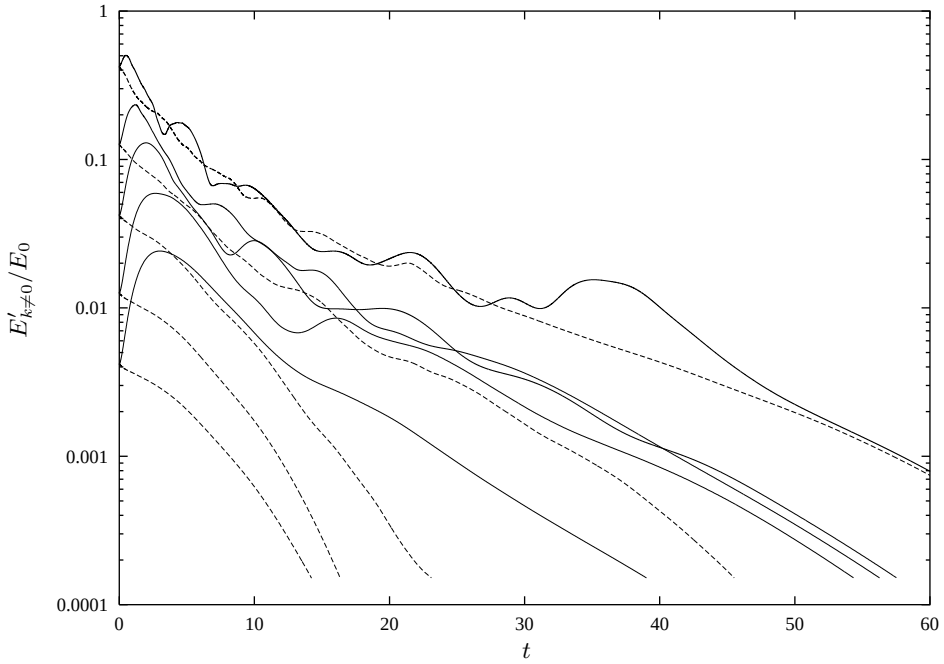

FIGURE 1. Decay of axisymmetric (dashed) and helical (solid) perturbations, initially of energy up to 40% of the laminar profile ($Re = 10,000$ and pipe $2\pi\, D$ long).

$E'_{k\neq0}$ is the energy of the axially-dependent modes and E_0 is the energy of the laminar profile; i.e. disturbances of up to 40% of the energy of the laminar flow were considered. By way of comparison, the turbulent test case of Eggels *et al.* (1994) has only $E'_{k\neq0}/E_0 \approx 0.014$ at $Re = 5300$. The number of finite difference points in radius and truncation of the Fourier modes was $(160, \pm 256)$ in (r, z) and (r, ξ); the large radial resolution was required to stably solve for such high initial energies. Axisymmetric disturbances show almost no sign of nonlinear interactions at this Re and decay almost monotonically. Five other sets of runs for other random disturbances showed similar behaviour. Helical flow is slightly more promising showing occasional moments of growth against a dominant decay.

Helical flow exhibits a linear instability at relatively small rotation rates about the axis, which occurs at longer L for larger Re (Mackrodt 1976). This supercritical bifurcation and the subsequent nonlinear waves (Toplosky & Akylas 1988; Landman 1990*b*) provided an excellent test of the helical code. Measuring the rotation rate by $R_\Omega = \frac{1}{4}\,\Omega\, Re$, where the angular velocity Ω is in units U/D, there is a supercritical bifurcation to helical waves at $R_\Omega = -52.43$, $Re = 4000$ and $L = 16\pi$. Figure 2 shows evolution of five random initial disturbances of $E'_{k\neq0}/E_0$ up to 0.3 in the presence of rotation; truncation at $(80, \pm 128)$ in (r, ξ). For $R_\Omega = -50$ the flow quickly returns the the parabolic profile. At $R_\Omega = -100$ and -300, well beyond the linear instability, the flow rapidly returns to a finite-amplitude helical wave flow. This supercritical behaviour persists for modest rotations so that no long term turbulent transients can be generated. Barnes & Kerswell (2000) have shown that these helical waves themselves undergo a supercritical Hopf bifurcation so that solutions cannot obviously be traced back to non-rotating flow. Figure 2 suggests that disconnected branches which could lead to subcritical turbulence in rotating helical pipe flow are unlikely to exist.

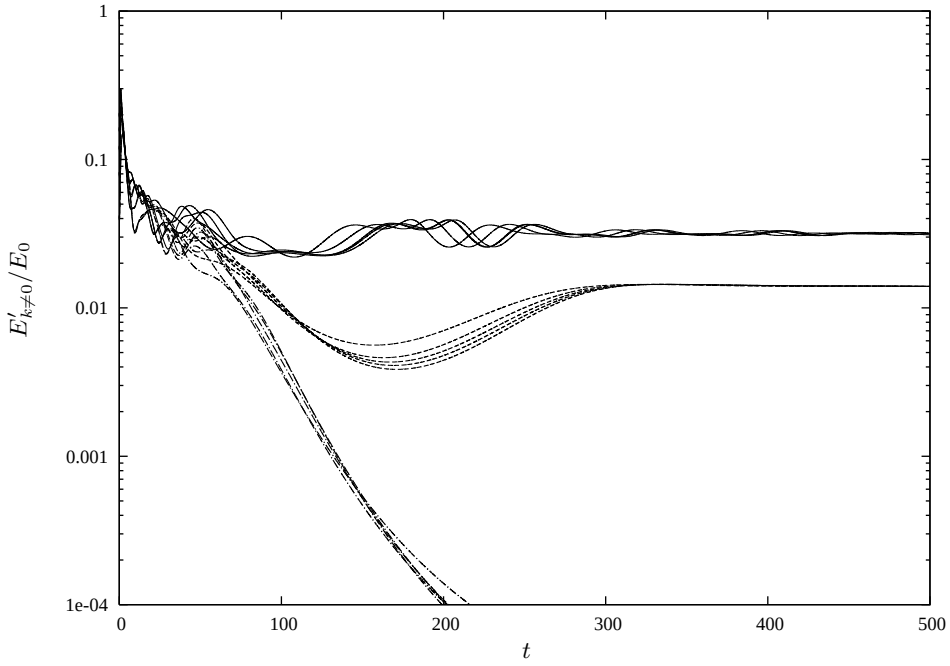

FIGURE 2. Large helical disturbances scaled from 4% to 20% of E_0 rapidly settle to parabolic flow for $R_\Omega = -50$ and to laminar helical waves for $R_\Omega = -100$, -300 ($Re = 4000$, $L = 16\pi$).

Our calculations suggest that rotating helical pipe flow follows the supercritical route to turbulence via a sequence of supercritical bifurcations, rather than the abrupt subcritical behaviour of 3-dimensional non-rotating pipe flow. Having seen strong decay at Re approximately five times that for which turbulence is observed in the laboratory, it appears that neither dynamics restricted to helical or axisymmetric subspaces are relevant for the observed transition.

4. A $2\frac{1}{2}$-dimensional model

We now introduce a third model which has high resolution in the cross-stream (radial) and streamwise (axial) directions, but only a few modes in the spanwise (azimuth) direction. The model was chosen to preserve a high radial resolution as streak features close to the wall appear to be important in the self-sustaining process as do detachments from the wall during the relaminarisation stages of low Reynolds number turbulence. High axial resolution was retained to allow the possibility of localised turbulent structures. This left only the azimuthal direction in which to reduce the number of degrees of freedom: only Fourier modes $m = 0, \pm m_0$ were considered, which corresponds to a sinusoidal variation in azimuth, an azimuthal shift of the sinusoid and a mean mode.

4.1. *Spatial characteristics*

Localised structures, surprisingly similar to puffs, were captured by the $2\frac{1}{2}$-dimensional model. Figure 3 compares a puff structure from a simulation fully resolved in azimuth (all m up to ± 24) with a 'puff' from the $2\frac{1}{2}$-dimensional model ($m = -3, 0, 3$). The plots are of the correct aspect ratio but only half of the computational domain is shown. Puffs from the model appear to be similar to resolved puffs in both length and structure, having

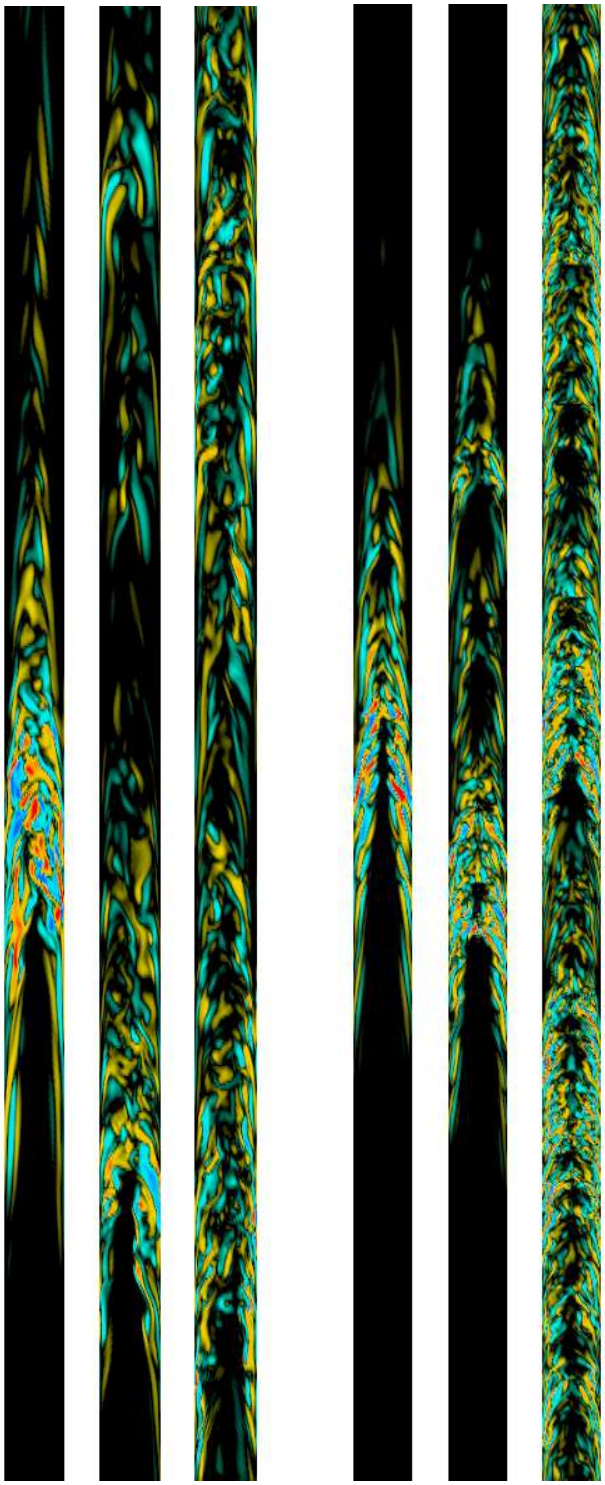

FIGURE 3. Axial vorticity in the (r, z)-plane, 1:1 aspect ratio, flow up the page; only $25\,D$ of a periodic $50\,D$ domain shown. *From left to right*: fully 3-dimensional simulations at $Re = 2000$, 2300, 2700 and $2\frac{1}{2}$-dimensional simulations at $Re = 2600$, 3200, 4000. Presented for both models are localised 'puffs', the early stages of delocalisation by the generation of a second puff downstream, and global slug-turbulence at larger Re.

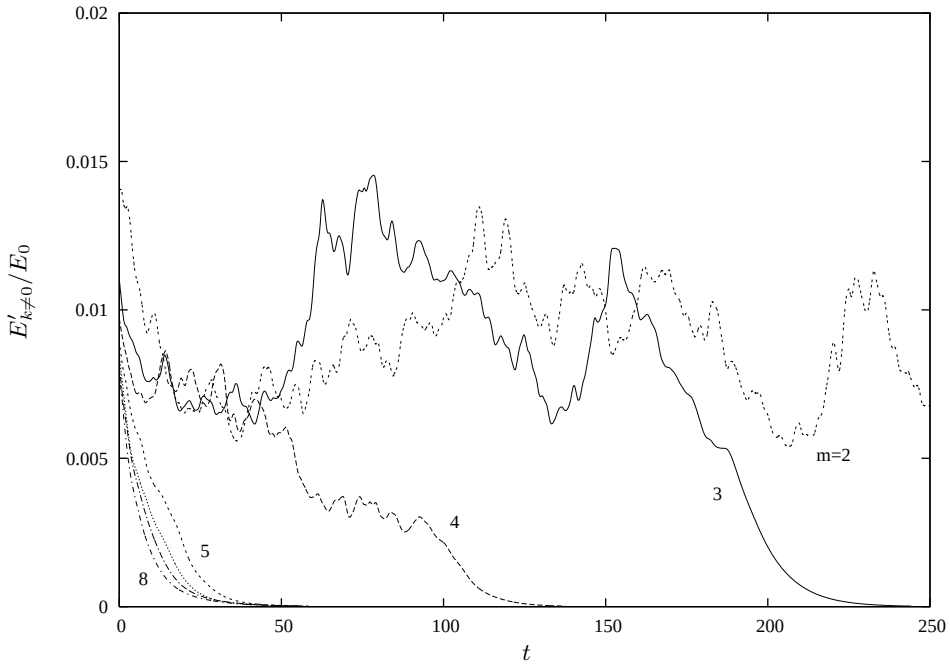

FIGURE 4. Initial trials of the $2\frac{1}{2}$-D model for several m-fold rotational symmetries.

a smooth upstream region close to the wall, an active turbulent region, and a dissipative region downstream. Only modest radial resolutions were required to observe such structures: a spectral resolution of 35 Chebyshev modes was used for the calculations of this section. Several calculations were performed with a lower resolution of 25 radial modes, but puff structures tended to elongate, requiring a longer pipe and thus offsetting the reduction in computation times. Axial resolution was chosen to approximately match the spectral drop-off in r (approximately 4 orders in the magnitude of the spectral coefficients, or 8 orders in the power spectrum) and was ± 384 for $L = 16\pi \approx 50$ diameters. Axial resolution was changed proportionally for other L considered in the rest of this section, hence keeping the smallest resolved scale fixed. Puffs were found to translate within 2% of U, slightly faster than in 3 dimensions, where they travel approximately 10% slower. Also shown in figure 3 is that the transition from localised to global turbulence is gradual, as observed experimentally. At larger Re the puff becomes delocalised, splitting into two or more localised turbulent regions with relatively laminar regions in between. At much larger Re the proportion of vigorous turbulence is seen to increase, as recorded by Gilbrech & Hale (1965).

4.2. *Temporal characteristics*

Another important feature captured by the model is that localised puffs may survive for long times before a sudden decay as observed experimentally (Peixinho & Mullin 2006). Typical transients for different m_0 are shown in figure 4 which indicates that higher rotational symmetries tend to decay more quickly. As structures of 3-fold rotational symmetry are the most frequently observed for transitional Re (Duggleby *et al.* 2007; Willis & Kerswell 2007*a*; Schneider *et al.* 2007*a*), $m_0 = 3$ was chosen for analysis of the lifetimes of disturbances. For this m_0, azimuthal length scales are also comparable to the radial length scale in the model. Figure 5 shows the probability distribution

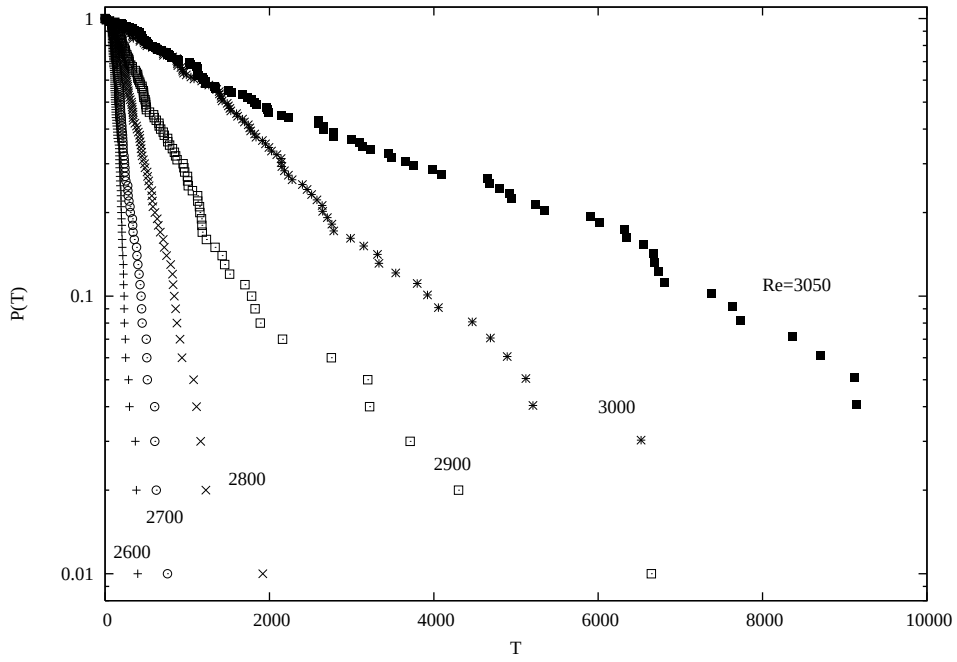

FIGURE 5. Probability of transient surviving to time $T\,(D/U)$ in the $m_0 = 3$ model for a pipe of approximately $100\,D$ length ($L = 32\pi$); 100 observations per Re.

function for puff lifetimes based upon 100 runs at each of several Re in a pipe ≈ 100 diameters long ($L = 32\pi$). Sets of initial puff conditions were generated from snapshots of a long run at a sufficiently large Re, similar to the annealing procedure adopted in Peixinho & Mullin (2006) and Willis & Kerswell (2007*b*). No dependence on the initial condition was observed, however, other than in the very early times of the transient. For the model, times of order $10^5\,D/U$ could be achieved, significantly longer than achieved in Willis & Kerswell (2007*b*) and in less computation time. The log-plot clearly shows an exponential distribution indicating a memoryless process.

From the observation times at each Re, a median time, τ, may be calculated. As mentioned above, whether τ diverges to infinity or not at a finite Re is a matter of ongoing debate. In laboratory experiments a pipe for which $D = 20$mm, $L = 785\,D$, Peixinho & Mullin (2006) found evidence that $\tau \sim 1/(Re_c - Re)$ with $Re_c = 1750$, independent of several initial conditions. In contrast, experiments by Hof *et al.* (2006) for $D = 4$mm, $L = 7500\,D$, and using a different method to initiate the puff, found that $\tau \sim \exp(c_1 Re)$, for some constant c_1. Numerical experiments by Willis & Kerswell (2007*b*,*c*) using well-resolved puffs in a $50\,D$ periodic pipe, however, show lifetimes to be significantly different from the exponential scaling and the simple power -1 was clearly seen with $Re_c = 1870$, overestimating the experimental value of 1750 Peixinho & Mullin (2006) by only 7%. Computational and experimental limitations have confined observations of τ to $O(10^3 - 10^4)$, but the effects of noise accumulated on longer time-scales is an open issue.

As the error in τ cannot be directly read from the log-plot, figure 5, median times were analysed using the bootstrapping method adopted in Willis & Kerswell (2007*c*). In this method samples of size N are generated by resampling from the original N observations, with equal probability of selecting each. This is repeated 100 000 times and

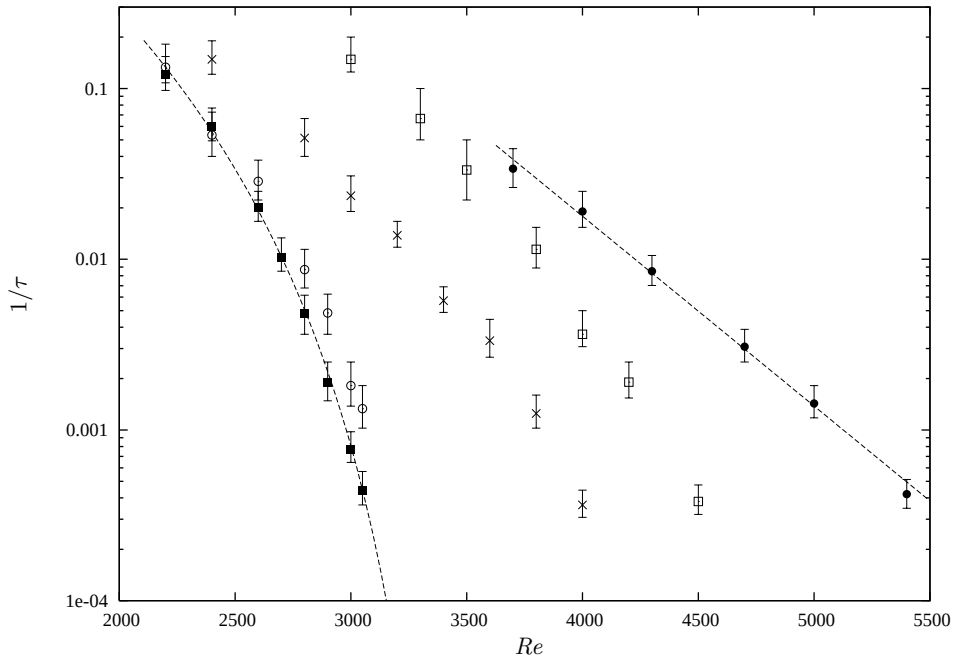

FIGURE 6. Sensitivity of the lifetime of transients to the pipe length. 100 observations per point for lengths in diameters, 32π, 16π, 8π, 4π, 2π (from left to right).

the distribution of the means of the samples is used to infer a confidence interval for the mean of the original data set. When all data fall within a maximum observation time, the confidence intervals generated by this method are close to those expected from the central limit theorem: a 95% confidence interval for the median time is approximately $\tau \pm 2\,\tau/\sqrt{N}$. When data is truncated a mean cannot be directly calculated. The bootstrapping method easily accommodates for such data by further resampling.

Figure 6 shows the maximum likelihood estimator for $1/\tau$ and 95% confidence intervals for several lengths of pipe. Lifetimes an order of magnitude larger than the calculations of Willis & Kerswell (2007*b*) were possible in a pipe twice as long, $L \approx 100$. The difference between results for the long pipes, $L \approx 50$ and $L \approx 100$ shows that lifetimes are sensitive to even small deviations from the laminar profile, caused here by the periodicity. This problem is unfortunately exacerbated in the model as puffs begin to slowly delocalise near the Re for which τ becomes large, preventing measurements to even smaller $1/\tau$. The gentle-curved fit for the long pipe is $\tau \sim 1/(Re_c - Re)^\beta$ with $\beta = 5.0$ and $Re_c = 3450$. Lagha & Manneville (2007) have also found that a $1/(Re_c - Re)^\beta$ scaling is much harder to distinguish from the $\exp(c_1 Re)$ scaling in their heavily truncated model for plane Couette flow. What is clear, however, is that as the pipe is shortened, the localised structure is more strongly affected by the periodicity, and this enhances the probability of relaminarisation. In the shorter pipes the scaling is more clearly exponential. Particularly noteworthy is the large jump in the probability of relaminarisation seen going from $L = 16\pi$ to $L = 8\pi$ when localised structures are no longer seen in the shorter pipe. For $L = 4\pi$ and $L = 2\pi$ the disturbance fills the pipe completely. Nevertheless, the probability of relaminarisation continues to increase and is clearly sensitive to the size of the domain. While the short periodic pipe could be considered as a simple model for slug flow, finite observed lifetimes appear to originate from the finiteness of the domain.

Thus the existence of finite lifetimes from such data cannot be extrapolated to the real case with confidence as it has been in Hof *et al.* (2006).

4.3. *Characteristics of the laminar–turbulent boundary: short pipe*

Given an initially laminar flow, small perturbations decay back to the laminar flow and larger perturbations develop into turbulence for sufficiently large Re. This naturally leads to the question of what characterises the dividing set of flows, for which a small perturbation may lead to either laminar or turbulent flow. Itano & Toh (2001) used a shooting method to find such a boundary in channel flow and discovered that the flow trajectory on the boundary settled upon what they thought was a travelling wave solution, but later identified as a near approach to a travelling wave on a periodic orbit (Toh & Itano 2003). This orbit is stable within the manifold of flows on this laminar-turbulent boundary or 'edge'. A similar situation is found in plane Couette flow (Lagha *et al.* 2007), where a single simple attractor is found. Pipe flow exhibits different characteristics, however, with Schneider *et al.* (2007*b*) finding a chaotic attractor in which trajectories pass nearby to exact travelling wave solutions (Mellibovsky & Meseguer 2007; Duguet *et al.* 2007). When the laminar-turbulent boundary dynamics is restricted within certain symmetry subspaces, however, simple attractors do emerge (Duguet *et al.* 2007).

Pipe flow is not so different from channel flow and yet they display different dynamics on the boundary. It is not difficult to imagine therefore, that the severe truncation of our model could also lead to a loss of chaotic behaviour on the boundary. In this section we show that the model preserves the chaotic end state for trajectories on the boundary and that exact solutions exist. This motivates extension of the results to long computational domains, where undirected calculations would be prohibitively expensive.

Boundary or edge trajectories for the model are shown in figure 7, for $L = \pi$ and over a range of Re. While a difference between the edge and developed turbulence may be seen by a rapid increase in energy, a clearer measure appears to be β related to the pressure required to maintain the fixed flux (see equation 2.1), or equivalently the friction. The pressure gradient is less than 10% greater when on the edge than the laminar value and is also smooth in time, whereas after a sudden increase to turbulence it is as rapidly varying as the energy.

At the lowest Re shown in figure 7 the energy of the edge is highly variable in time. At the next $Re = 6500$ a period of slow variation is observed. Duguet *et al.* (2007) have recently demonstrated that the edge can be used to find exact solutions by identifying phases where the flow has a relatively simple temporal behaviour. Such an episode is marked by a triangle where the instantaneous flow was used as an initial condition for a Newton–Raphson code. This converged to the exact solution pictured in figure 8(a), confirming that travelling wave solutions do exist for the model. The flow exhibits fast streaks towards the walls and slow streaks are shed towards the centre. The state found is also very similar to that marked by a square at $Re = 10\,000$ when the trajectory for a while appears to settle towards a steady (translating) state. As Re increases the variability on the edge surprisingly decreases, more so than can be explained by the increasing viscous time. The viscous time D^2/ν scales as Re in our time units, and longer trajectories at larger Re have been shown in figure 7 to compensate. A possible explanation for the decreasing variability is that the eigenvalues of the unstable directions from the travelling waves decrease with increasing Re, thus enabling closer and longer visitations.

The variability of the boundary in terms of energy or any other chosen amplitude measure also, of course, indicates the range of such attributes for initial conditions which will trigger turbulence. The fact that it is easier to trigger turbulence as Re increases (e.g.

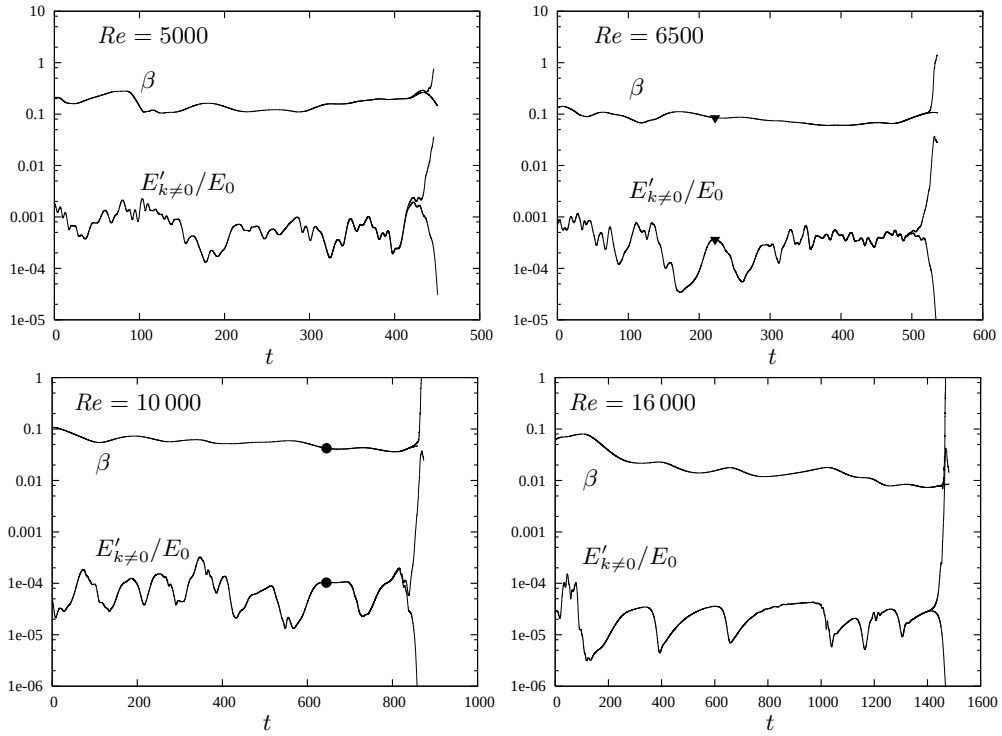

FIGURE 7. Edge trajectories for $L = \pi$. Exact travelling wave solution found nearby to the triangle at $Re = 6500$ (shown in figure 8) has very similar structure to that found at the circle for $Re = 10000$.

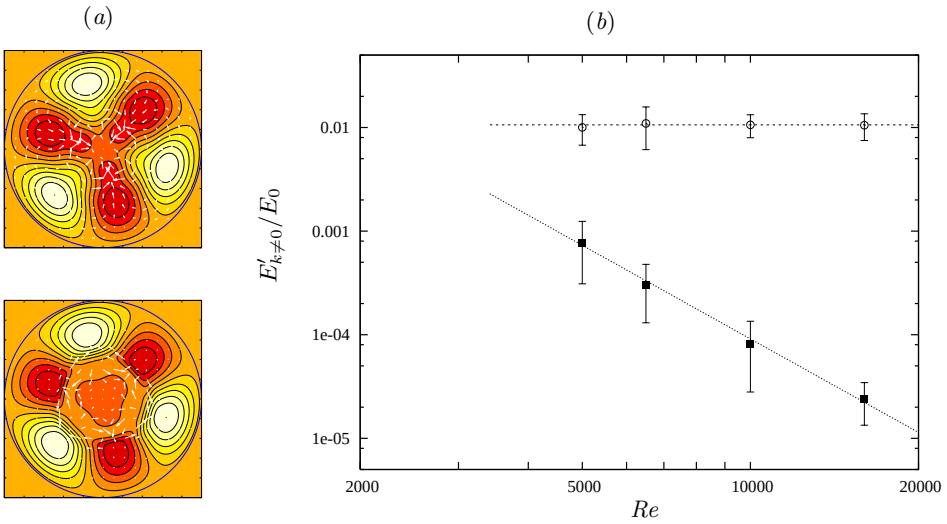

FIGURE 8. (a) Cross-sections of an exact travelling-wave solution $L = \pi\, D$, $Re = 6500$, wavespeed $c = 1.563\, U$ at $z = 0, \frac{1}{2}L$. Dark regions are slower than the laminar profile and white regions faster. (b) Energy on the boundary (squares) and of turbulence (circles). The slope for the boundary case is -3, implying that the amplitude, $A \sim Re^{-1.5}$. For turbulence $E'_{k\neq 0}/E_0 \approx 0.011$ which for the 3-dimensional test case was ≈ 0.014 at $Re = 5300$.

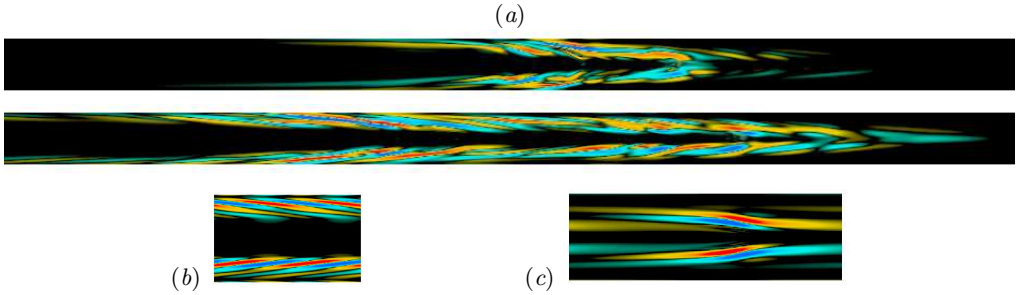

FIGURE 9. (*a*) Snapshots of axial vorticity at two different times from a laminar-turbulent boundary trajectory. $L = 16\pi\,D$, $Re = 4000$. $20\,D$ of $\approx 50\,D$ shown. The disturbance is localised but with an extended upstream region. (*b*) An exact solution shown over two axial periods, $L = 1.694\,D$, discovered using the method of Pringle & Kerswell (2007). (*c*) The exact solution of figure 8 found by the method of Duguet *et al.* (2007), $L = \pi$.

Hof *et al.* 2003) is reflected in the decrease in the mean boundary energy as Re increases seen in figure 8(*b*). Error bars in this figure represent the mean and one standard deviation above and below this mean but nevertheless, a clean scaling emerges for the amplitude of $A \sim Re^{-1.5}$. Mellibovsky & Meseguer (2006) found this scaling when considering streamwise perturbations, and Peixinho & Mullin (2007) in laboratory experiments with obliquely oriented jets.

4.4. *Characteristics of the laminar–turbulent boundary: long pipe*

The bisection procedure (Itano & Toh 2001; Schneider *et al.* 2007*b*; Duguet *et al.* 2007) for isolating the laminar-turbulent boundary can equally be applied to flow within a long pipe although the computational demands becomes increasingly intensive. The $2\frac{1}{2}$-dimensional model is ideal for a reconnaissance of likely behaviour, including identification of the form the edge state takes. With this motivated, trajectories on the laminar-turbulent boundary were found for a pipe ≈ 50 diameters long using the pressure gradient to distinguish between edge and turbulent flow states. Interestingly, the attracting edge state which emerges looks exactly like a turbulent puff except at the trailing edge (upstream region): compare figures 9(*a*) and 3. In a turbulent puff this is the most energetic part possessing a fluctuation energy level comparable to homogenised slug turbulence (see Willis & Kerswell (2007*a*)) at higher Re. However, the edge state is noticeably smoother even at $Re = 4000$ (figure 9) compared to a puff at $Re = 2600$ (figure 3) which has finer scales. The similarities between the two are the strong wall structures slanting into the axis at the upstream region, a region where the axial vorticity reaches the axis (the trailing edge region), followed by a gradual relaminarisation/decay downstream. Interestingly it is the upstream region that waxes and wanes (compare the two axial vorticity snap shots of figure 9) rather than the passive-looking downstream wake.

The edge state remains localised just like the turbulent puff up to $Re = 3000$, but surprisingly also remains localised for much higher Re when the puff has given way to (global) slug turbulence. To emphasize this, the localised edge state shown in figure 9 at $Re = 4000$ was generated starting from the global disturbance of figure 3. The fact that the edge state remains puff-like throughout the puff-to-slug transitional Re range suggests that slug turbulence is destabilised puff-turbulence rather than being a separate state occupying a different part of phase space. This is certainly consistent with simulations in which a puff state smoothly evolves into a slug by slowly expanding upstream as well as downstream. A corollary of this, of course, is that a puff and a slug cannot coexist at a

given Re: to our knowledge there are no reported experimental observations to contradict this claim.

Figure 9 also shows the axial vorticity for a travelling wave with shift-and-reflect symmetry, found using the method of Pringle & Kerswell (2007) and that for the travelling wave shown in figure 8 (which has no special symmetry). Both have axial vorticity slanted from the wall into the central axis reminiscent of the turbulent puff and edge state. A recent search for coherent fast-streak states within turbulent puffs has indicated that the flow transiently resembles travelling wave states upstream and downstream of the trailing edge region (Willis & Kerswell 2007*a*). Given that the energetic trailing edge region is absent in the edge state, there seems an even higher likelihood of seeing coherent states there.

Although the findings in this $2\frac{1}{2}$-dimensional model are only suggestive of what may occur in the fully 3-dimensional setting, they are sufficiently interesting to motivate a fully 3-dimensional long-pipe computation. This is currently underway.

5. Discussion

In this paper, we have described the numerical formulation used to simulate transitional pipe flow (Willis & Kerswell 2007*a*,*b*; Duguet *et al.* 2007). This is based upon the poloidal-toroidal potential decomposition of the velocity field discussed by Marqués (1990), where the difficulty of coupled boundary conditions has been by-passed by influence matrix methods here. Reducing the system to five simple second-order equations, the method is accurate, relatively simple to implement and computationally efficient. This has then been used here to explore dynamical subspaces in the hope of finding a reduced system to aid understanding.

No evidence for turbulent transients has been found in axisymmetric pipe flow, confirming an earlier investigation (Patera & Orszag 1981), or in helical pipe flow consistent with the work by Landman (1990*a*,*b*). Rotating pipe flow displays a classic supercritical bifurcation route to turbulence (Mackrodt 1976; Toplosky & Akylas 1988; Landman 1990*a*,*b*; Barnes & Kerswell 2000) which has no bearing on the non-rotating situation. A brief search in rotating helical pipe flow failed to find evidence for any disconnected subcritical branches of solution which may have reached back to the non-rotating limit. A $2\frac{1}{2}$-dimensional model, which represents a minimal 3-dimensionalisation of the axisymmetric limit, does, however, possess a subcritical transition scenario and all the important spatio-temporal characteristics of fully resolved pipe flow at Re of the same order of magnitude. Localised disturbances, structurally similar to turbulent puffs, are found in the model at low Re (≈ 2600), which slowly delocalise at intermediate Re (≈ 3200), and rapidly expand into slugs at high Re (≈ 4000). Exact unstable travelling wave solutions also exist within the model and appear to underpin the dynamics in phase space.

Within this $2\frac{1}{2}$-dimensional model, the relaminarisation statistics of the puffs have been examined in pipes of varying lengths. For a $100\,D$ pipe (long enough to capture a puff properly), the median lifetime τ is best fit by the relationship $\tau \sim (Re_c - Re)^{-5}$ indicating a puff becomes sustained at $Re_c = 3450$. On the other hand, in a short pipe of length $2\pi\,D$ where turbulence is either spontaneously global or absent, an exponential dependency $\tau \sim \exp(c_1 Re)$ is suggested implying that puffs are never permanent. These extreme cases clearly indicate the sensitivity of the relaminarisation statistics to the (periodic) computational domain. Counter intuitively, noise in the form of intrusive boundary conditions leans in the direction of *relaminarising* the flow rather than maintaining the turbulence. It is also worth remarking that longer transients in the $100\,D$ pipe (i.e. Re closer to 3450) could in principle be calculated given the computational savings

available in the model. However, we found that the puffs start to delocalise for $Re \approx 3200$ indicating that in fact, by $Re = 3450$, the puff has become unstable to a slug-like turbulence. This is, of course, exactly the same issue faced by puffs in the real system: they delocalise to become slugs for $Re = 2250$-2500 (Willis & Kerswell 2007*a*). No claim has been made anywhere in the literature that slugs are anything other than permanently sustained once generated. While issues still surround the differing experimental data sets (Peixinho & Mullin 2006; Hof *et al.* 2006), the need to resolve the spatial inhomogeneity of the puff state properly in numerical experiments is clear.

The $2\frac{1}{2}$-dimensional model has also presented an opportunity to probe the possible dynamics on the laminar-turbulent boundary in long pipe flow. Calculations indicate that the attracting state in this set is a localised puff-like structure which is smoother and less energetic in the trailing edge region than its turbulent puff counterpart. Also intriguingly, this end state remains localised way beyond in Re when the puff has delocalised. This tends to suggest that the turbulent puff still exists as a solution but has become unstable to a slug state. The variability of the flow on the laminar-turbulent boundary also highlights the variability in initial conditions which can trigger turbulence. Just focusing on the mean energy gives an amplitude scaling $A \sim Re^{-1.5}$ consistent with some numerical computations (Mellibovsky & Meseguer 2006) in a short pipe and laboratory experiments with a carefully specified jet configuration, designed to excite a coherent vortex (Peixinho & Mullin 2007). Clearly, exploring how far this realisation can be usefully developed is a promising area for future research.

In conclusion, we have introduced a model of pipe flow severely truncated in its azimuthal degrees of freedom but otherwise fully resolved in the others. This notwithstanding, the remaining system captures all of the rich dynamical behaviour observed in pipe flow but obtained at a fraction of the computational cost for the full 3-dimensional situation. It therefore presents a very accessible arena in which to test ideas quickly before deciding to invest a considerable effort in the full 3-dimensional system.

Acknowledgments: Many thanks to Yohann Duguet and Chris Pringle for finding exact solutions in the model. This research was funded by the EPSRC under grant GR/S76144/01.

Appendix A. Boundary conditions

The Navier–Stokes equation (2.1) plus the boundary condition $\boldsymbol{u} = \boldsymbol{g}(\theta, z)$ are equivalent to (2.4) provided that on the boundary

$$\hat{\boldsymbol{n}} \cdot \boldsymbol{\nabla} \wedge \left[(\partial_t - \frac{1}{Re}\boldsymbol{\nabla}^2)\boldsymbol{u} + \boldsymbol{b} \right] = 0, \tag{A 1}$$

where $\hat{\boldsymbol{n}}$ is its normal (see Marqués 1990). This condition ensures the term in the square brackets is equal to a gradient, such as the pressure. If not imposed, this term may be any $\chi\hat{\boldsymbol{z}}$ where $\nabla_h^2\chi = 0$. If this is not a gradient, then an unknown body force is introduced. For the axisymmetric case χ is constant, the curl of $\chi\hat{\boldsymbol{z}}$ is then zero, and therefore the condition is redundant. Otherwise, from the diffusion term in (A 1), using the properties $\boldsymbol{\nabla}^2\boldsymbol{u} = -\boldsymbol{\nabla} \wedge \boldsymbol{\nabla} \wedge \boldsymbol{u}$ and $\boldsymbol{\nabla} \wedge \boldsymbol{\nabla} \wedge (\hat{\boldsymbol{z}} f) = \boldsymbol{\nabla}(\partial_z f) - \hat{\boldsymbol{z}}\nabla^2 f$, one finds

$$\hat{\boldsymbol{r}} \cdot \boldsymbol{\nabla} \wedge \boldsymbol{\nabla}^2\boldsymbol{u} = \hat{\boldsymbol{r}} \cdot \boldsymbol{\nabla} \wedge \boldsymbol{\nabla} \wedge (\hat{\boldsymbol{z}}\nabla_h^2\psi) + \hat{\boldsymbol{r}} \cdot \boldsymbol{\nabla} \wedge \boldsymbol{\nabla} \wedge \boldsymbol{\nabla} \wedge (\hat{\boldsymbol{z}}\nabla_h^2\phi) + \partial_{zz}\hat{\boldsymbol{r}} \cdot \boldsymbol{\nabla} \wedge \boldsymbol{u}$$

$$= \partial_{rz}\psi_1 - \frac{1}{r}\partial_\theta\phi_2 + \partial_{zz}\hat{\boldsymbol{r}}\cdot\boldsymbol{\nabla}\wedge\boldsymbol{g},$$

which leads to the last condition of (2.5). The other conditions simply express $\boldsymbol{u} = \boldsymbol{g}(\theta, z)$ with the gauge freedom $\phi = 0$ on the boundary.

The simplified system for the axially averaged flow, $h(r, \theta)$, arises because there exists a closed circuit c that is not simply-connected running along the axial direction. This has an associated condition

$$\int_c \left[(\partial_t - \frac{1}{Re}\boldsymbol{\nabla}^2)\boldsymbol{u} + \boldsymbol{b}\right]\cdot \mathrm{d}l = 0,$$

which together with P_z on the second curl of the Navier–Stokes equations leads to the governing equation for h. The boundary condition $h = 0$ assumes $P_z(g_z) = 0$, which otherwise would correspond to a translating pipe or moving frame.

REFERENCES

Barkley, D. & Tuckerman, L. S. 2007 Mean flow of turbulent-laminar patterns in plane couette flow. *J. Fluid Mech.* **576**, 109–137.

Barnes, D. R. & Kerswell, R. R. 2000 New results in rotating Hagen–Poiseuille flow. *J. Fluid Mech.* **417**, 103–126.

Duggleby, A., Ball & Paul 2007 Dynamics of propogating turbulent pipe flow structures. Part II: Relaminarisation. *Phys. Fluids submitted (http://arxiv.org/abs/physics/0608259)* .

Duguet, Y., Willis, A. P. & Kerswell, R. R. 2007 Transition in pipe flow: the saddle structure on the boundary of turbulence. *J. Fluid Mech. submitted (http://arXiv.org/abs/0711.2175)* .

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 dns and experiment. *J. Fluid Mech.* **268**, 175–209.

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

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

Gilbrech, Donald A. & Hale, J. C. 1965 Further results on the transition from laminar to turbulent flow. In *Developments in mechanics* (ed. Simon Ostrach & Robert H. Scanlan), , vol. 2, pp. 3–15. Pergamon.

Hamilton, J. M., Kim, J. & Waleffe, F. 1995 Regeneration mechanisms of near-wall turbulence structures. *J. Fluid Mech.* **287**, 317–348.

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

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

Jiménez, J. & Moin, P. 1991 The mimimal flow unit in near-wall turbulence. *J. Fluid Mech.* **225**, 213–240.

Kerswell, R. R. 2005 Recent progress in understanding the transition to turbulence. *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.

Lagha, M. & Manneville, P. 2007 Modelling transitional plane Couette flow. *Eur. Phys. J.* **58**, 433–447.

Lagha, M., Schneider, T. M., Lillo, F. De & Eckhardt, B. 2007 Laminar-turbulent boundary in plane Couette flow. *preprint* .

Landman, M. J. 1990*a* On the generation of helical waves in circular pipe flow. *Phys. Fluids A* **2**, 738–747.

Landman, M. J. 1990*b* Time-dependent helical waves in pipe flow. *J. Fluid Mech.* **221**, 289–310.

Mackrodt, P. A. 1976 Stability of Hagen-Poiseuille flow with superimposed rigid rotation. *J. Fluid Mech.* **73**, 153–164.

Marqués, F. 1990 On boundary conditions for velocity potentials in confined flows: Application to Couette flow. *Phys. Fluids A* **2**, 729–737.

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 dynamics on the critical threshold. In *Proceedings of the 15th Int. Couette–Taylor workshop*. Le Havre, France.

Patera, A. T. & Orszag, S. A. 1981 Finite-amplitude stability of axisymmetric pipe flow. *J. Fluid Mech.* **112**, 467–474.

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

Peixinho, J. & Mullin, T. 2007 Finite amplitude thresholds for transition in pipe flow. *J. Fluid Mech.* **582**, 169–178.

Pfenniger, W. 1961 Transition in the inlet length of tubes at high reynolds numbers. In *Boundary layer and flow control* (ed. G. V. Lachman), pp. 970–980. Pergamon.

Pringle, C. C. T. & Kerswell, R. R. 2007 Asymmetric, helical, and mirror-symmetric traveling waves in pipe flow. *Phys. Rev. Lett.* **99**, 207450.

Schneider, T. M., Eckhardt, B. & Vollmer, J. 2007*a* Statistical analysis of coherent structures in transitional pipe flow. *Phys. Rev. E* **75**, 066313.

Schneider, T. M., Eckhardt, B. & Yorke, J. A. 2007*b* Turbulence transition and the edge of chaos in pipe flow. *Phys. Rev. Lett.* **99**, 034502.

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

Toplosky, N. & Akylas, T. R. 1988 Nonlinear spiral waves in rotating pipe flow. *J. Fluid Mech.* **190**, 39–54.

Viswanath, D. 2007 The dynamics of transition to turbulence in plane couette flow. *preprint (http://arxiv.org/abs/physics/0701337)* .

Wang, J., Waleffe, F. & Gibson, F. 2007 Lower branch coherent states in shear flows. *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. 2007*a* Coherent structures in local and global pipe turbulence. *submitted (http://arxiv.org/abs/0706.3330)* .

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

Willis, A. P. & Kerswell, R. R. 2007*c* Reply to comment on 'critical behaviour in the relaminarisation of localised turbulence in pipe flow'. *http://arxiv.org/abs/0707.2684* .

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**, 281–351.