

Stabilisation and drag reduction of pipe flows by flattening the base profile

Elena Marensi^{1,†}, Ashley P. Willis¹ and Rich R. Kerswell²

¹School of Mathematics and Statistics, University of Sheffield, Sheffield S3 7RH, UK

²Department of Applied Mathematics and Theoretical Physics, Centre for Mathematical Sciences, University of Cambridge, Cambridge CB3 0WA, UK

(Received 10 June 2018; revised 1 October 2018; accepted 11 December 2018;
first published online 28 January 2019)

Recent experimental observations (Kühnen *et al.*, *Nat. Phys.*, vol. 14, 2018*b*, pp. 386–390) have shown that flattening a turbulent streamwise velocity profile in pipe flow destabilises the turbulence so that the flow relaminarises. We show that a similar phenomenon exists for laminar pipe flow profiles in the sense that the nonlinear stability of the laminar state is enhanced as the profile becomes more flattened. The flattening of the laminar base profile is produced by an artificial localised body force designed to mimic an obstacle used in the experiments of Kühnen *et al.* (*Flow Turbul. Combust.*, vol. 100, 2018*a*, pp. 919–943) and the nonlinear stability measured by the size of the energy of the initial perturbations needed to trigger transition. Significant drag reduction is also observed for the turbulent flow when triggered by sufficiently large disturbances. In order to make the nonlinear stability computations more efficient, we examine how indicative the minimal seed – the disturbance of smallest energy for transition – is in measuring transition thresholds. We first show that the minimal seed is relatively robust to base profile changes and spectral filtering. We then compare the (unforced) transition behaviour of the minimal seed with several forms of randomised initial conditions in the range of Reynolds numbers $Re = 2400$ – $10\,000$ and find that the energy of the minimal seed after the Orr and oblique phases of its evolution is close to that of a critical localised random disturbance. In this sense, the minimal seed at the end of the oblique phase can be regarded as a good proxy for typical disturbances (here taken to be the localised random ones) and is thus used as initial condition in the simulations with the body force. The enhanced nonlinear stability and drag reduction predicted in the present study are an encouraging first step in modelling the experiments of Kühnen *et al.* and should motivate future developments to fully exploit the benefits of this promising direction for flow control.

Key words: transition to turbulence

1. Introduction

It is widely established that turbulent wall flows exert a much higher friction drag than laminar flows. Since the flow regime in oil and gas pipelines is generally turbulent, larger pumping forces are needed, compared to the laminar case, to maintain

[†] Email address for correspondence: e.marensi@sheffield.ac.uk

the desired flow rates with a consequent increase in energy consumption and carbon emissions. A great deal of research effort is thus directed towards the design of efficient control strategies to either reduce the turbulent drag or to delay the onset of turbulence. Transition to turbulence in pipe flows is a fully nonlinear problem because the laminar state is linearly stable to any infinitesimal disturbance. Therefore, if one wishes to control or delay transition, it is paramount to understand which kind of small (but finite-amplitude) disturbances are most effective in initiating the transition process. A useful tool that has recently been employed to tackle this challenge is the so-called minimal seed, i.e. the disturbance of lowest energy capable to trigger transition. However, the question of how representative the minimal seed is of typical ambient disturbances remains unanswered. To address this issue, we compare the transition behaviour of the minimal seed with that of different random initial disturbances in the range of Reynolds numbers $Re = 2400\text{--}10\,000$. We find that the energy of the minimal seed after the initial nonlinear unpacking phase is quite close to that of a localised random disturbance which just triggers transition. Suitable initial conditions are thus generated to investigate the stabilising effect of a simple model for the presence of a baffle in the core of the flow.

Before discussing the formulation (§ 2) and results (§ 3), we provide a short review of the problem of transition in pipe flows and the different control strategies used to avoid or suppress turbulence.

1.1. Transition in pipe flows and calculation of the minimal seed

The enigma of how laminar flow through a pipe undergoes the transition to turbulence has been intriguing and challenging scientists for over a century, since the pioneering experiments of Reynolds (Reynolds 1883). Despite many pieces of the puzzle being brought together in the past years (refer, for example, to Kerswell 2005; Eckhardt *et al.* 2007; Willis *et al.* 2008; Mullin 2011, for comprehensive reviews), a full understanding of the problem still eludes us.

All theoretical and numerical evidence indicates that the laminar state is linearly stable to any infinitesimal disturbance, although a rigorous proof is still lacking. In the absence of a linear instability of the laminar state from which a sequence of bifurcations may be initiated, transition can only be triggered by finite-amplitude background disturbances. For $Re > 3000$ the observed transition process is abrupt and catastrophic and it rapidly results in a complex and highly disordered state (Darbyshire & Mullin 1995). At transitional Reynolds numbers in the range $1800 < Re < 3000$, instead, turbulence first appears in localised patches of disordered motion, known as puffs, which coexist with the laminar flow (Wynanski & Champagne 1973; Avila *et al.* 2011). Depending on the level of background noise in the experiment, the flow rate at which transition occurs can be varied by more than an order of magnitude. This fact already puzzled Reynolds in 1883 who, in one set of experiments, found a transitional Reynolds number $Re_c \approx 2000$, while in another set of experiments with minimised level of background disturbance, found $Re_c \approx 13\,000$. This value was pushed to 10^5 by Pfenninger (1961) with a very tightly controlled environment of his experiments. Reynolds' lower critical value has been confirmed in other experiments (e.g. Wynanski & Champagne 1973; Darbyshire & Mullin 1995; Avila *et al.* 2011) with current estimates in the range 1760–2300.

At lower Reynolds numbers ($Re \approx 2000$), the critical Reynolds number is somewhat dependent on the definition of 'transition', but at larger Re ($Re > 3000$) where the transition is clear, it is widely recognised that the influence of background disturbances

becomes of great concern. The critical amplitude A_c for the onset of turbulence is expected to decrease with increasing Reynolds number and its behaviour can be characterised by $A_c \sim Re^{-\gamma}$, $\gamma > 0$. A key question is thus: what is the exponent γ and, more importantly, can this value be predicted theoretically? Trefethen *et al.* (2000) proposed a renormalisation of the amplitude by the average velocity in order to cast different experimental results in terms of a single definition of A_c , suggesting lower and upper bounds for $\gamma \in [1.2, 1.8]$. The experiments carried out by Peixinho & Mullin (2007) provided a critical exponent $\gamma \in [1.3, 1.5]$ when the flow was perturbed using push–pull disturbances and $\gamma = 1$ when the flow was perturbed by small impulsive jets. The latter scaling had previously been found in the experiments of Hof, Juel & Mullin (2003) and was later confirmed numerically by Mellibovsky & Meseguer (2009) in their ‘impulsive scenario’ with the flow being perturbed by a local impulsive forcing. In the ‘autonomous scenario’, instead, where the flow was perturbed by an initial array of streamwise vortices with random noise superimposed on it, Mellibovsky & Meseguer (2009) obtained critical exponents $\gamma \in [1, 1.5]$ much closer to those of Peixinho & Mullin (2007) for the push–pull disturbances.

The ultimate goal of these studies is to provide a characterisation of the basin of attraction of the laminar flow, i.e. the subset of initial conditions which asymptotically converge to the laminar state. However, these methods are impractical at finding the smallest possible solution capable of just kicking the system away from the laminar state, as they require a large number of simulations/experiments. Recent developments have been achieved using variational methods to construct fully nonlinear optimisation problems that seek the minimal seed (Pringle & Kerswell 2010; Cherubini *et al.* 2012; Pringle, Willis & Kerswell 2012; Duguet *et al.* 2013; Cherubini & Palma 2014); see Kerswell (2018) for a review. From a dynamical-systems point of view, the minimal seed represents the closest (in a chosen norm) point of approach of the laminar–turbulent boundary, or ‘edge’, to the basic state in phase space, as shown in figure 1. If transition is regarded as undesirable, such perturbation will be considered the ‘most dangerous’ disturbance. Previous studies in a pipe (Pringle & Kerswell 2010; Pringle *et al.* 2012, referred to as PK10 and PWK12, respectively, throughout the paper) have revealed important characteristics of the minimal seed, such as its fully localised nature and its three-phase evolution consisting of the Orr mechanism, the oblique phase and the lift-up, during which the flow gradually unwrap to give rise to a large, predominantly streamwise independent final state. However, a link between the critical initial energies of the minimal seed and those of disturbances that can typically be generated in a laboratory has not been provided yet. This will be the focus of the first part of the paper, with the outcomes being summarised in the key graph, figure 6, where the scaling $E_c = E_c(Re)$ for the minimal seed and several forms of randomised initial conditions are compared.

1.2. Control of pipe flows

Several different control strategies have been designed in the past fifty years to reduce the wall friction of fully turbulent flows (refer, for example, to Lumley & Blossey 1998; Kasagi, Suzuki & Fukagata 2009; Quadrio 2011, for reviews). In light of the central role of streamwise vortices in the drag and shear stress production (e.g. Kim, Moin & Moser 1987; Waleffe 1997), Choi, Moin & Kim (1994) proposed an ‘opposition control technique’ aimed at actively counteracting vortices or selected velocity components to reduce the skin-friction drag on the wall. Xu, Choi & Sung (2002) applied suboptimal opposition control (Lee, Kim & Choi 1998) to pipe flow at




FIGURE 1. (Colour online) Schematic of the problem in phase space. The laminar–turbulent edge Σ separates initial conditions that trigger turbulence (characterised by an initial energy E_0 greater than the critical value E_c) from those that decay back to the laminar state ($E_0 < E_c$). The minimal seed (MS) is the point infinitesimally close to the boundary just capable of triggering turbulence.

a wall Reynolds number $Re_\tau = 150$ and achieved drag reduction of approximately 13 to 23%. Both passive (e.g. riblets) and active (oscillations or generation of travelling waves) methods have been employed to inhibit the near-wall turbulence production. Drag reductions of 25 to 40% were obtained in fully turbulent pipe flows using spanwise wall oscillations (Choi & Graham 1998; Quadrio & Sibilla 2000; Choi, Xu & Sung 2002; Duggleby, Ball & Paul 2007; Zhou & Ball 2008). Zhou & Ball (2008) also considered streamwise oscillations but found this method to be less effective than spanwise oscillations. In their experimental and numerical study, Auteri *et al.* (2010) were able to achieve 33% drag reduction by imposing streamwise-modulated waves of spanwise velocity travelling forward in the streamwise direction. Willis, Hwang & Cossu (2010) also found it possible to reduce drag by forcing large-scale streaks in pipe flow and reported a power saving of up to 11%.

More recently, large-scale control methods that completely relaminarise fully turbulent flows by manipulating the mean profile have been successfully employed (Hof *et al.* 2010; Kühnen *et al.* 2018*b*). These methods target the mean shear in order to counteract its crucial role as energy source in near-wall turbulence (Schoppa & Hussain 2002). Based on the observation that the streamwise vorticity of a turbulent puff is mainly produced at the trailing edge by the fast incoming flow, Hof *et al.* (2010) developed both experimental and numerical methods to flatten the velocity profile at the upstream edge of the puff to intercept this mechanism and successfully relaminarise the puff. Their idea was further developed by Kühnen *et al.* (2018*b*), who were able to achieve a complete and final collapse of turbulence by appropriate distortions of the mean profile, with the friction losses reduced by as much as 90%. Compared to some of the other strategies presented so far (for example, the opposition control method, which requires a knowledge and detailed manipulation of the fully turbulent velocity field), this approach is much simpler to implement as it only requires a steady open-loop manipulation of the streamwise

velocity component. Experimentally, full relaminarisation could be obtained by: (i) increasing the turbulence level near the wall by vigorously stirring the flow with rotors or via wall-normal injection of additional fluid through several small holes in the pipe wall; or (ii) accelerating the flow close to the wall via streamwise injection of fluid through an annular gap at the wall; or (iii) by means of a movable pipe wall segment. The common and key ingredient to all these relaminarisation techniques is a flattened streamwise velocity profile achieved by a deceleration of the flow in the bulk region and/or an acceleration of the flow close to the wall. The important role of the mean-flow distortion was confirmed numerically by adding a global body force to the equations of motions such that the resulting velocity profile was more ‘plug shaped’. The efficiency of the control mechanism was directly related to the suppression of the lift-up mechanism (reviewed recently by Brandt 2014), measured by the linear transient growth. All disturbances schemes were shown to lead to a reduction of the linear transient growth, that is, the modified profile was shown to suppress the energy transfer from the mean flow to the streamwise vortices and to inhibit the streak–vortex interaction. An analogous relaminarisation problem, although in a different flow configuration, was recently studied by He, He & Seddighi (2016) using direct numerical simulations. They considered a vertical pipe heated from below and modelled the buoyancy effect with a streamwise body force of variable shape, while keeping the mass-flux constant. In this configuration, known as buoyancy-aided flow, the body force caused a flattening of the mean flow which led to suppression of turbulence and, as a consequence, heat transfer deterioration. A new perspective on the flow relaminarisation phenomenon was proposed, which relies on the quantification of the flow rate of the perturbation flow induced by the body force.

Most of the literature pertaining to the control of shear flows is devoted to suppressing fully turbulent flow. However, delaying (or preventing) transition to turbulence, thus avoiding the worst of turbulence *in toto*, is even more desirable. Nevertheless little literature is available on this subject. Suppressing the energy growth of initial perturbations to delay or prevent transition requires an understanding of how the basin of attraction of the laminar flow is modified in the presence of the control. So far, theoretical work has focused on investigating the sensitivity of the linearised Navier–Stokes equations around the laminar state in order to design suitable controls (Jovanović 2008). This approach has had some success in mitigating turbulence transition using both open-loop and feedback-based approaches (Kim & Bewley 2007). For example, Högberg, Bewley & Henningson (2003) used direct numerical simulation to demonstrate that linear feedback control strategies can significantly expand the laminar state’s basin of attraction of plane Couette flow for a range of Reynolds numbers. In channel flows, Moarref & Jovanović (2010) performed a perturbation analysis in the wave amplitude of the linearised Navier–Stokes equations to ‘design’ travelling waves which significantly reduce the sensitivity of the flow. However, as pointed out by Bewley (2001), due to the finite-amplitude nature of transition in shear flows, a fully nonlinear approach is required to probe the sensitivity of the laminar state to finite-amplitude disturbances. In a proof-of-concept study, Rabin, Caulfield & Kerswell (2014) showed how an optimisation approach could be used to design a more nonlinearly stable plane Couette flow through manipulation of the boundary conditions. By spanwise oscillating one boundary (with amplitude A and frequency ω), these authors showed that E_c could be increased by 40% through careful choice of A and ω .

Our study is motivated by a recent experimental observation (Kühnen *et al.* 2018a) that manipulation of the flow in a pipe with a baffle can lead to full relaminarisation

for flow rates up to three times ($Re = 6000$) that for which turbulence typically appears in the presence of ambient perturbations ($Re > 2000$). Our focus is on theoretically capturing the phenomenon observed in experiments so that the process can then be optimised.

2. Formulation

We consider the problem of constant mass-flux fluid flow through a straight cylindrical pipe of length L and diameter D . The flow is described using cylindrical coordinates $\{r, \theta, z\}$ aligned with the pipe axis. Length scales are non-dimensionalised by the radius of the pipe $D/2$ and velocity components by the laminar centreline velocity $2\bar{W}$, where \bar{W} is the constant bulk velocity. Unless otherwise specified, energies are given as ‘absolute energies’, i.e. not scaled by the energy of the laminar flow in the same domain. We consider a perturbation $\tilde{\mathbf{u}} = \{\tilde{u}, \tilde{v}, \tilde{w}\}$ superimposed on the laminar Hagen–Poiseuille flow (HPF) $\mathcal{W}(r)\hat{\mathbf{z}} = (1 - r^2)\hat{\mathbf{z}}$ so that the full velocity field is given by $\mathbf{u}_{tot} = \mathcal{W}(r)\hat{\mathbf{z}} + \tilde{\mathbf{u}}(r, \theta, z, t)$. The problem is governed by the Navier–Stokes and continuity equations

$$NS = \frac{\partial \tilde{\mathbf{u}}}{\partial t} + \mathcal{W} \frac{\partial \tilde{\mathbf{u}}}{\partial z} + \tilde{u} \mathcal{W}' \hat{\mathbf{z}} - \tilde{\mathbf{u}} \times \nabla \times \tilde{\mathbf{u}} + \nabla \tilde{p} - \frac{4\beta}{Re} \hat{\mathbf{z}} - \frac{1}{Re} \nabla^2 \tilde{\mathbf{u}} = 0, \quad \nabla \cdot \tilde{\mathbf{u}} = 0, \quad (2.1)$$

where the prime indicates total derivative, $Re = \bar{W}D/\nu$ is the Reynolds number and $\beta = \beta(\tilde{\mathbf{u}})$ is a correction to the pressure gradient such that the mass flux remains constant. The parameter $1 + \beta$ is an observed quantity in experiments and is defined as the ratio of the observed dissipation \mathcal{D} (Pope 2000) (or pressure gradient $\langle \partial p / \partial z \rangle$) and the corresponding laminar value \mathcal{D}_{lam} (or laminar pressure gradient $\langle \partial p / \partial z \rangle_{lam}$), namely

$$1 + \beta = \frac{\mathcal{D}}{\mathcal{D}_{lam}} = \frac{\langle \partial p / \partial z \rangle}{\langle \partial p / \partial z \rangle_{lam}}, \quad (2.2)$$

where the angle brackets indicate the volume integral

$$\langle \dots \rangle = \int_0^L \int_0^{2\pi} \int_0^1 \dots r \, dr \, d\theta \, dz. \quad (2.3)$$

Periodic boundary conditions are imposed in the streamwise direction and no-slip/no-penetration conditions on the pipe wall.

The formulation of the nonlinear variational problem closely follows PWK12 and the reader is referred to their §2 for a detailed explanation. In its simplest form the problem can be stated as follows: among all (incompressible) initial conditions of a given perturbation energy E_0 , we seek the disturbance that gives rise to the largest energy growth after a time T .

$$G(T, E_0) = \max_{\tilde{\mathbf{u}}_0 | E_0} \frac{\langle \tilde{\mathbf{u}}(\mathbf{x}, T)^2 \rangle}{\langle \tilde{\mathbf{u}}(\mathbf{x}, 0)^2 \rangle}, \quad (2.4)$$

where $\tilde{\mathbf{u}}_0 = \tilde{\mathbf{u}}(\mathbf{x}, 0)$. To accomplish this, a functional $\mathcal{L}(\mathbf{q}, \mathbf{q}^\dagger; E_0)$ is defined, where \mathbf{q} is the state vector and \mathbf{q}^\dagger is the adjoint state vector. The functional is maximised subject to the constraints imposed through the Lagrange multipliers, namely the constraints of the three-dimensional Navier–Stokes equations and of the initial disturbance energy. Numerically, the gradient of the functional with respect to every independent variable

is imposed to vanish using the iterative optimisation algorithm described in PWK12. Unless otherwise specified, the following criterion for convergence is implemented: the algorithm is said to have converged when the relative step, i.e. the size of the step taken in the ascent direction of \mathcal{L} relative to the initial energy E_0 , is smaller than 10^{-5} (this tolerance is chosen from experience gathered in previous calculations of nonlinear optimals, namely PK10 and PWK12).

To find the minimal seed, the initial energy E_0 is gradually increased and the variational problem solved until the critical initial energy E_c is reached where turbulence is just triggered. Ideally, for asymptotically long times $T = T_{opt}$, G is expected to approach a step function in E_0 , with the jump at the critical value E_c . In practice, two conjectures proposed by PWK12 are exploited. For asymptotically large T_{opt} , the initial energy E_{fail} at which the algorithm first fails to converge corresponds to E_c (conjecture 1) and the converged nonlinear optimal (NLOP) approaches the minimal seed as E_0 approaches E_c (conjecture 2).

The calculations are carried out using the open source code `Openpipeflow` (Willis 2017), with a flow variable q discretised in the domain $\{r, \theta, z\} = [0, 1] \times [0, 2\pi] \times [0, 2\pi/\alpha]$ using Fourier decomposition in the azimuthal and streamwise directions and finite differences in the radial direction, i.e.

$$q(r_n, \theta, z) = \sum_{k < |K|} \sum_{m < |M|} q_{n,k,m} e^{i(\alpha kz + m\theta)}, \quad (2.5)$$

where $n = 1, \dots, N$ and α is the streamwise wavenumber. The radial points are clustered close to the wall. For a pipe of length $5D$ at $Re = 2400$, we use $N = 64$, $K = 36$, $M = 32$ and time step $\Delta t = 0.01$, with the discretisation appropriately refined as the Reynolds number is increased to keep the resolution unaltered. Unless otherwise specified, throughout the paper we use $L = 5D$ as in PWK12.

3. Results and discussion

3.1. Robustness of the minimal seed

3.1.1. Changes to the base flow

Inspired by Kühnen *et al.*'s (2018*b*) experiments, we first study the minimal seed for transition with a modified base flow. By adding a suitable body force, the laminar base profile becomes more ‘plug-like’, i.e. flatter in the centre of the pipe, than the parabolic profile of HPF. Kühnen *et al.* (2018*b*) considered an initially fully turbulent state and showed that, by flattening the streamwise velocity profile, they were able to completely relaminarise the flow. Here, instead, our aim is to investigate how the basin of attraction of the laminar state is affected by flattening of the base flow and whether this modification affects the minimal seed, i.e. the closest point of approach of the edge to the laminar state shown schematically in figure 1.

Following Kühnen *et al.* (2018*b*) we use the following family of profiles for the base flow $\mathcal{W}(r; \delta, c)\hat{z}$ (see their equation (19)), where

$$\mathcal{W}(r; \delta, c) = (1 - \delta) \left[1 - \frac{\cosh(cr) - 1}{\cosh(c) - 1} \right]. \quad (3.1)$$

The parameter δ is the centreline difference between the laminar profile and the target profile and c is set by the constant mass-flux condition. The force $F = F(r)\hat{z}$ required to generate such a target velocity profile is obtained by substituting $\mathcal{W}(r; \delta, c)$ in the

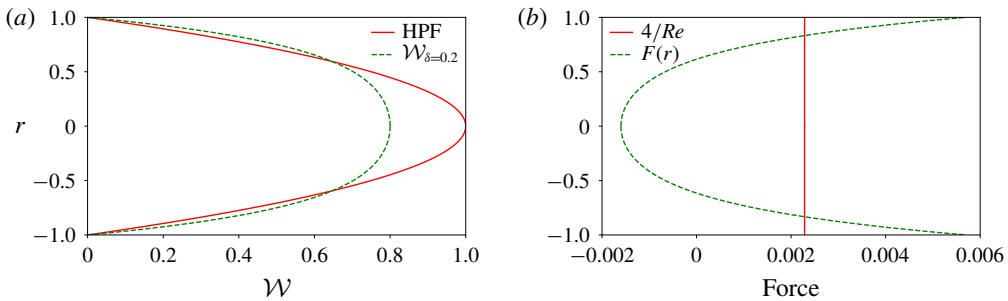


FIGURE 2. (Colour online) (a) Comparison of the parabolic profile (red solid line) and the forced profile (green dashed line) of the laminar flow. (b) The shape of the force chosen to decrease the centreline velocity of the laminar flow by 20% at $Re = 1750$ (green dashed line). The force is negative (pointing upstream) in the centre and positive (pointing downstream) near the wall. For comparison, the pressure gradient in a laminar flow $|\partial p/\partial z| = 4/Re$ is also shown (red solid line). Similar forced profiles and forcings are reported by Hof *et al.* (2010) (figure S2 of their supplementary material) and Kühnen *et al.* (2018b) (figure 7 of their extended data).

Case	Re	α	L/D	$T_{opt}(D/\bar{W})$	$N \times K \times M$
PK10	1750	2	0.5π	21.35	$60 \times 8 \times 16$
PWK12	2400	0.628	5	75	$64 \times 36 \times 32$

TABLE 1. Parameters and resolution for PK10 and PWK12 cases.

Navier Stokes equations, i.e. $F(r) := \hat{z} \cdot NS$. For example, the case with $\delta = 0.2$ and $c = 3.5935$ is shown in figure 2. The forcing decelerates the flow in the central part and accelerates it near the wall while the mass flux is kept fixed.

The effect of the forcing is studied for the parameters corresponding to the works by PK10 and PWK12 as summarised in table 1. As discussed in PWK12, the choice of parameters in PK10 was not ideal: the Reynolds number is close to the first appearance of turbulent state, the target time is short and as a result the algorithm struggles to discern between conditions that relaminarise and conditions that trigger turbulence, and in a tightly constrained geometry the basin boundary is highly fractal. Nevertheless it is useful here to show the effect of the global forcing.

Figure 3 shows the maximum growth (at the target time T_{opt}) as a function of E_0 for the cases with forced and parabolic base profile. The initial energy is gradually increased until E_c is reached.

Most of the data points have been verified by feeding the algorithm with at least two or three different initial conditions (for example, a snapshot from a turbulent run, another NLOP at a lower E_0 or a turbulence-inducing initial conditions at a higher E_0). The initial energies E_{fail} at which the optimisation algorithm first fails to converge are marked with a black cross in figure 3. According to the conjecture 1 of PWK12, these correspond to the critical initial energies E_c , where the edge touches the energy hypersurface at one velocity state. Due to the reasons mentioned above, in the PK10 case, especially for $\delta = 0.2$, convergence is sometimes not clear and deteriorates (becomes slower and slower) as E_0 is increased and approaches E_{fail} . The last data points before E_{fail} , for both values of δ , appear to show convergence, but there still

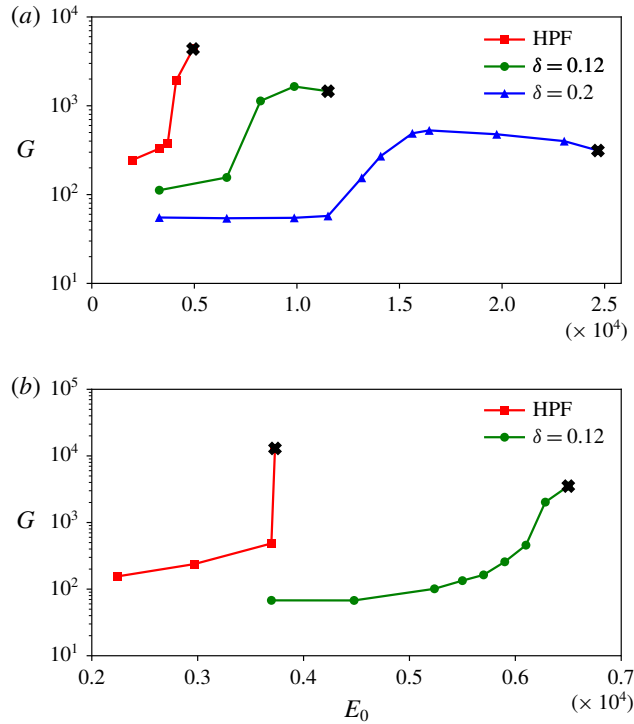


FIGURE 3. (Colour online) Comparison of G versus E_0 with parabolic and forced base profile for (a) PK10 and (b) PWK12 cases (refer to table 1). For each case, the thick black cross indicates the initial energy E_{fail} at which the optimisation algorithm first fails to converge, as we gradually increase E_0 . Up to the last data point before E_{fail} we were able to converge the algorithm. The critical energy of the minimal seed is thus bracketed between these two values of E_0 . Due to the convergence issues explained in the text and in PWK12, the search for E_c in the cases with a flattened base profile was not refined as much as in the unforced cases.

remains some doubt even after running the algorithm for more than 1000 iterations. Furthermore, the perturbations corresponding to $E_0 = E_{fail}$ decay immediately after the time T_{opt} is reached, because at this low Reynolds number turbulence is intermittent and appears only in the form of decaying puffs.

For the PWK12 case, convergence is clearer than in the PK10 case due to the larger domain, longer integration time and higher Reynolds number. However, in the forced case at $E_0 = 6.28 \times 10^{-5}$ (last data point before E_{fail}) neither a smooth convergence nor a clear increase of the residual was obtained, even after 1000 iterations. Note also that in the unforced case, G sharply increases when the critical initial energy is reached, while with a flattened base profile, the increase of G is much more gradual, as this case behaves similarly to cases where the system is close to the marginal Re (as in PK10 case discussed above).

Despite these convergence issues, figure 3 shows that by flattening the base profile, E_c is moved towards higher values of the initial energy and the maximum growth reached at time T_{opt} is decreased, i.e. the unforced curve $G = G(E_c)$ is shifted ‘down’ and ‘right’ as δ is increased. Therefore, the presence of the forcing expands the basin




FIGURE 4. (Colour online) (a) Time series of energy for the minimal seeds of PWK12 case (parameters provided in table 1) with forced (green dashed line) and unforced (red solid line) base profiles. In the forced case, the time evolution of the NLOP just before E_{fail} is also shown (green dash-dotted line). Although in this case the convergence of the optimisation algorithm was not clear, the decay of the energy straight after T_{opt} (indicated by the vertical dotted line) suggests that this disturbance does not lead to a turbulent episode. The cross-sections of the minimal seeds in (b) the forced (green dashed border) and (c) unforced (red solid border) cases are also shown. Contours indicate axial velocity perturbations (with the laminar flow subtracted off): white or light for positive perturbation, red or dark for negative, outside shade corresponds to zero. The arrows indicate cross-sectional velocities.

of attraction of the laminar base profile and reduces transient growth. For example, for the PWK12 case, the critical energy of the minimal seed moves from $E_0 = 3.73 \times 10^{-5}$ to $E_0 = 6.5 \times 10^{-5}$ and for $E_0 < E_c$ the NLOP of the forced case reaches less than half of the growth of the unforced case. The energy time series of the minimal seeds for the parabolic and forced cases are shown in figure 4(a). These initial conditions clearly lead to a turbulent episode which survives for at least double the optimisation time. The case $E_0 = 6.28 \times 10^{-5}$ for which convergence was critical is also shown. This initial state seems turbulent at the optimisation time but decays straight after.

Despite the critical initial energy being significantly increased with a flattened base profile, the fully localised structure of the minimal seed remains largely unchanged, as shown in the cross-sections of figures 4(b,c). Therefore, the structure of the minimal seed is found to be fairly robust to changes to the base flow. This is different from Rabin *et al.*'s (2014) study of oscillated plane Couette flow where, instead, qualitative changes in the structure of the minimal seed are found when compared to the non-oscillated case. In their study, however, the basic fluid response in the presence of spanwise oscillations becomes time dependent through the additional spanwise component (refer to their equation (2.1)), while in our case only the shape of the laminar flow profile is modified, its dimension (one) and dependencies (only radial) remain unchanged.

3.1.2. Filtering

In order to assess how robust the minimal seed is to smoothening, we perform the full optimisation procedure over perturbations in a lower-dimensional space for




FIGURE 5. (Colour online) Critical initial energy as a function of the filtering ratio \mathcal{F} for the unforced PWK12 case (parameters provided in table 1). Insets: cross-sections of the minimal seeds at different values of $\mathcal{F} = 0, 0.75$ and 0.93 . Contours indicate streamwise velocity beyond the laminar flow (white or light for positive, red or dark for negative) and arrows indicate cross-sectional velocities.

the set of parameters corresponding to the (unforced) PWK12 case. Specifically, at each iteration we project the initial condition onto a subspace where only the first $K_f \times M_f$ wavenumbers are retained. The resolution in the forward and backward steps is, however, unchanged (i.e. $K = 36$, $M = 32$). As a measure of how much we are truncating we introduce the filtering ratio \mathcal{F} defined as

$$\mathcal{F} = 1 - \sqrt{\frac{K_f M_f}{K M}}. \quad (3.2)$$

Figure 5 shows the critical energy of the minimal seed as a function of the filtering ratio, where $\mathcal{F} = 0$ means no filtering (i.e. fully resolved minimal seed) and $\mathcal{F} = 1$ would imply that no perturbation remains. Note that, for a fixed value of \mathcal{F} , values for K_f and M_f are chosen so that the corresponding ratios K_f/K and M_f/M are as close as possible, i.e. to avoid cases where the filtering in one direction is much higher than in the other direction. The critical initial energy of the minimal seed remains almost unchanged for $\mathcal{F} \leq 0.75$, that is, when only retaining 25% of the modes. Since we are restricting the initial condition to a subset of the energy hypersurface by adding the filtering constraint, one would expect the minimal seed to occur at a higher initial energy. Our results thus suggest that the edge is locally quite flat near the minimal seed. As \mathcal{F} is further increased, larger initial energies are needed in order to trigger turbulence. For example, for $\mathcal{F} = 0.93$ (i.e. only the first 3×2 modes retained) E_c is almost an order of magnitude larger than in the fully resolved case. For values of \mathcal{F} larger than this (for example we have tested $\mathcal{F} = 0.96$ where only 2×1 modes are retained) it seems that it is not possible to trigger transition. As shown by the cross-sections in the insets of figure 5, the structure of the minimal seed remains almost unchanged when the filter is applied, even for cases where the minimal seed occurs at larger E_c than in the unfiltered case (refer to the last inset to the right).

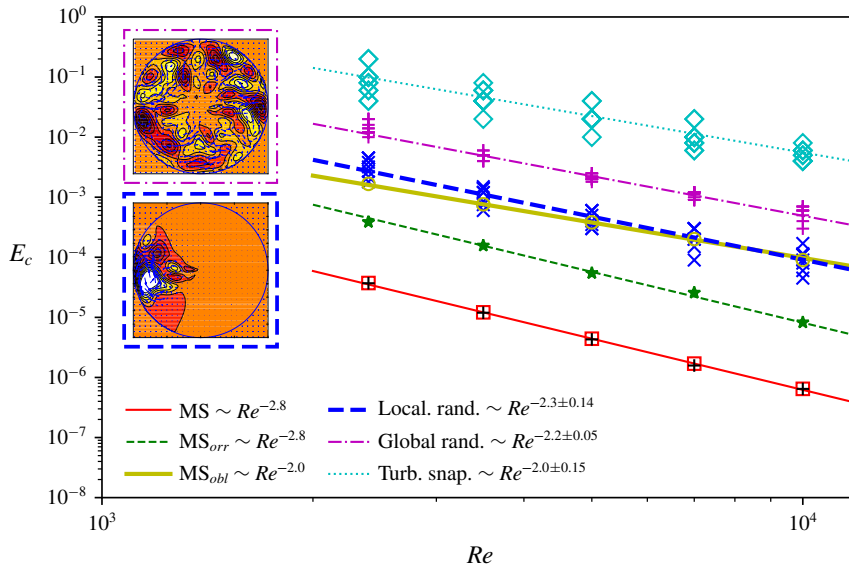


FIGURE 6. (Colour online) Energy thresholds E_c versus $Re = 2400, 3500, 5000, 7000$ and 10000 for different types of initial conditions: turbulent snapshots (cyan diamonds), artificially generated global (purple pluses) or localised (blue crosses) random fields and the minimal seed (black plus symbols correspond to the largest tested value of E_0 below which transition never occurs, while the red squares correspond to the smallest tested value of E_0 for which transition occurs at least once). The green stars and the yellow circles are the energies of the minimal seed after the Orr and oblique phases, respectively. The lines represent the power-law scalings obtained by fitting the last four data points. Insets: cross-sections of a typical global (purple dash-dotted border) and localised (dashed blue border) random initial condition. Contours indicate streamwise velocity beyond the laminar flow (white or light for positive, red or dark for negative) and arrows indicate cross-sectional velocities.

This study shows that the minimal seed is robust to quite severe spectral filtering, i.e. the small-scale structure of the minimal seed is not important. To confirm this, we also looked to see where the energy of the fully resolved minimal seed is distributed across the Fourier modes in θ and z . For $\mathcal{F} = 0.75$ and 0.93 we retain $\approx 96\%$ and 23% of the energy, respectively. It thus appears that, despite the fully localised structure of the minimal seed, most of the energy is carried by the large-scale modes.

It is possible that radial filtering may have a more significant effect on the energy thresholds, but the structure remains essentially similar, thus suggesting that truncation is not a simple route towards a new set of more ‘typical’ initial conditions. Conversely, it shows that the structure does not have to be perfectly formed to be the optimal. From this it seems reasonable that the minimal seed might be realised in a laboratory, but whether it would be ‘naturally’ realised among ‘ambient’ disturbances remains unclear at this stage.

3.2. Statistical study of transition to turbulence

To establish whether the minimal seed may be used to model typical ambient perturbations, we compare the critical initial energy of the minimal seed with that of

a random disturbance. The latter is generated by scattering energy randomly over the subset of wavenumbers $K \times M = 12 \times 9$. At $Re = 2400$ this truncation corresponds to $\mathcal{F} \approx 0.7$, before the rapid increase in E_c (figure 5), where the minimal seed could still be captured quite faithfully. An arbitrary complex amplitude A_n is generated for each of the spectral modes in the chosen subset and the radial dependence introduced in the form $\sum_{n=1}^{N_s} A_n r \sin(n\pi r)$ with $N_s = 5$ (the projection scheme integrated into the time stepping algorithm is used to ensure that the initial condition is solenoidal). These initial conditions are fed into direct numerical simulation (DNS) with time integration $T = 125(D/\bar{W})$. Five Reynolds numbers are considered: $Re = 2400, 3500, 5000, 7000$ and $10\,000$, with the numerical resolution appropriately enlarged for increasing Reynolds number. The same subset of wavenumbers used at $Re = 2400$ is employed for all the other Reynolds number considered. For each Re we consider 10 to 12 distinct random initial conditions and for each of them we gradually increase E_0 until turbulence is triggered and sustained for a time $T(D/\bar{W})$. The criterion to check for relaminarisation is $E_{3d} < 10^{-8}$, where E_{3d} is the energy associated with the streamwise-dependent modes only, as this quantity decays very rapidly when the flow relaminarises.

An analogous statistical study is carried out using 10 different snapshots from a turbulent run as initial condition (in a similar fashion to Schneider & Eckhardt (2008)). Both the random and snapshot initial conditions are global disturbances, while the minimal seed is fully localised. A third set of random localised initial conditions is obtained by scattering energy randomly over the wavenumbers as above, then by multiplying this global disturbance by a smooth spatial windowing function $\mathcal{B}(z)$ and $\mathcal{B}(\theta)$ so that the disturbance occupies only $1/5$ and $1/3$ of the streamwise and azimuthal domain, respectively. For the range of Reynolds numbers considered here we do not observe a strong localisation of the minimal seed in the radial direction and therefore we do not localise the random disturbance in this direction. The smoothing function is defined using equation 8 of Yudhistira & Skote (2011). For example, for the localisation in the streamwise direction we use:

$$\mathcal{B}(z) = g\left(\frac{z - z_{start}}{\Delta z_{rise}}\right) - g\left(\frac{z - z_{end}}{\Delta z_{fall}} + 1\right), \quad (3.3)$$

with

$$g(z^*) = \begin{cases} 0 & \text{if } z^* \leq 0 \\ \{1 + \exp[1/(z^* - 1) + 1/z^*]\}^{-1} & \text{if } 0 < z^* < 1 \\ 1 & \text{if } z^* \geq 1, \end{cases} \quad (3.4)$$

where $z_{start} = 2L/5$ and $z_{end} = 3L/5$ indicate the spatial extent over which the disturbance is non-zero, and $\Delta z_{rise} = \Delta z_{fall} = L/10$ are the rise and fall distances of the disturbance. A similar expression to (3.3) is used for $\mathcal{B}(\theta)$ with $\theta_{start} = 2\pi/3$, $\theta_{end} = 4\pi/3$ and $\Delta\theta_{rise} = \Delta\theta_{fall} = \pi/3$.

The critical initial energies of the minimal seeds are found with a two-digit accuracy for all the Reynolds numbers considered, except for $Re = 10\,000$ for which only one-digit accuracy is reached due to large computational cost of the simulations. The curves $E_c = Re^{-\gamma}$ obtained by fitting critical energies found for the turbulent snapshot, global and localised random sets of data points and for the minimal seed are shown in figure 6. The data points at $Re = 2400$ are not used to obtain the fittings as turbulence can still be transient at this relatively low Reynolds number. A power-law exponent $\gamma \approx 2.8$ for the minimal threshold energy is obtained. The same exponent

Source	Geometry	Disturbance	Energy exponent γ
Present study	HPF	Turb. snapshot	2.0
Present study	HPF	Global random	2.2
Present study	HPF	Localised random	2.3
Present study	HPF	Minimal seed	2.8
Hof <i>et al.</i> (2003)	HPF	Impulsive jets	2
Mellibovsky & Meseguer (2009)	HPF	Local impulsive forcing	2
Peixinho & Mullin (2007)	HPF	'Push-pull' with constant mass flux	2.8
Mellibovsky & Meseguer (2009)	PCF	Minimal seed	2.7
Duguet <i>et al.</i> (2013)	ASBL	Minimal seed	2
Cherubini <i>et al.</i> (2015)	PPF	Minimal seed	3.0
Huang <i>et al.</i>	—	Localised	2.5
Wall unit scaling (Re_τ)	—	Localised	2.5

TABLE 2. Power-law scalings $E_c = Re^{-\gamma}$ obtained in the present study and in the literature for different types of disturbances and in different geometries (HPF: Hagen–Poiseuille flow, PCF: plane Couette flow, ASBL: asymptotic suction boundary layer, PPF: plane Poiseuille flow). For asymptotically large Reynolds number we expect to see scaling with wall units, which leads to an exponent $\gamma = 2.5$.

was obtained experimentally by Peixinho & Mullin (2007) and later confirmed numerically by Mellibovsky & Meseguer (2009) using 'push-pull' perturbations with constant flow rate. Minimal energies for transition to turbulence were also calculated by Duguet *et al.* (2013) for plane Couette flow (refer to their figure 3), by Cherubini, De Palma & Robinet (2015) for the asymptotic suction boundary layer (refer to their figure 9) and, very recently, by Z. Huang, C. P. Caulfield and G. Xi (personal communication) for plane channel flow. Our exponent is close to the $\gamma \approx 2.7$ obtained by Duguet *et al.* (2013) for plane Couette flow and $\gamma \approx 3$ found by Huang *et al.* for plane Poiseuille flow. Furthermore, our study is reminiscent of figure 19 in Reddy *et al.* (1998), where the streamwise-vortex and oblique-wave transition scenarios for plane channel flows are compared to two-dimensional linear optimals and noise, although the power-law exponents are not reported in their study. The minimal seed curve $E_c = Re^{-2.8}$ is steeper than those pertaining to all the other forms of disturbances considered. The power-law exponent $\gamma \approx 2$ for the turbulent snapshot recalls the scaling found by Hof *et al.* (2003) and Peixinho & Mullin (2007) in their experiments using small impulsive jets to perturb the flow. For the random global and localised disturbances, larger exponents than for the turbulent snapshot are obtained, i.e. $\gamma \approx 2.2$ and 2.3, respectively (refer to table 2 for a summary of the different exponents γ discussed above). Most noticeably, the critical energy of the minimal seed is almost three orders of magnitude lower than that of the global disturbances considered here. With localisation of the random disturbances, the critical energy drops by almost an order of magnitude with respect to the global initial conditions, but it is still significantly larger than that of the minimal seed.

Although the minimal seed has a very peculiar structure, which might be unlikely to be generated in a laboratory, it evolves to a structure that looks much more familiar to a 'natural' disturbance over a relatively short time scale corresponding to the Orr and oblique-wave phases. The early stages of the minimal seed evolution are shown in the three-dimensional visualisations of figure 7 at $Re = 2400$.

The end of the Orr phase is identified by analysing the flow topology of the NLOP close to $t = 0$: the streaks that are initially tightly layered and inclined back into




FIGURE 7. (Colour online) Snapshots showing isocontours of streamwise perturbation velocity during the early stages of the time evolution of the minimal seed for $Re = 2400$. (a) $t = 0$, (b) $t = 1.1D/\bar{W}$ (end of the Orr phase) and (c) $t = 2.5D/\bar{W}$ (end of the oblique phase). The isocontours in each snapshot correspond to 50% of the maximum (light/yellow) and 50% of the minimum (dark/red) of the streamwise perturbation velocity in the pipe at that time.

the oncoming flow, are tilted away from the wall by the Orr mechanism and slightly separated. In a short pipe, as in the PK10 case, the evolution of the flow structure during the Orr phase is relatively clearer to visualise than for longer pipes (refer to figure 1 of PK10 and our figure 7). Therefore, there is some discretion in our choice of the energy size at the end of the Orr phase, but this has a negligible effect on our ensuing discussion. As in the short-pipe case of PK10, our data suggest that the Orr phase occurs in a very short time scale ($0 < t(D/\bar{W}) \lesssim 0.5-1$) and gives rise to an initial spurt of energy growth. By the end of the Orr phase, the helical modes starts to become dominant, thus signalling that the oblique-wave mechanism has come into play. In PK10 the end of the oblique-wave phase was clearly signalled by the ‘shoulder’ in the time evolution of the total perturbation energy E and the corresponding ‘bump’ in the time evolution of the three-dimensional energy E_{3d} , as reproduced in figure 8(a). For the present choice of $L = 5D$, which is approximately three times longer than that in PK10, the ‘shoulder’ in the energy time series is not clearly distinguishable and E_{3d} does not decay straight after the peak because of the longer length-scale modes (refer to red solid and green dashed lines in figure 8b). However, by defining a three-dimensional energy $\tilde{E}_{3d} = E_{3d}(k > 3)$ that includes only energy from streamwise wavenumbers $k > 3$, we are able to observe the ‘bump’ in \tilde{E}_{3d} (refer to the blue dash-dotted line in figure 8b), where we identify the end of the oblique phase. Both the time scales of the Orr ($0 < t(D/\bar{W}) \lesssim 0.5-1$) and oblique-wave ($0.5-1 \lesssim t(D/\bar{W}) \lesssim 2.5-3.5$) phases are found to be similar to the PK10 case and to remain almost unchanged with Reynolds number.

If we compare the energies of the minimal seed after the Orr and oblique phases with the random localised initial conditions (refer to figure 6), then the gap is

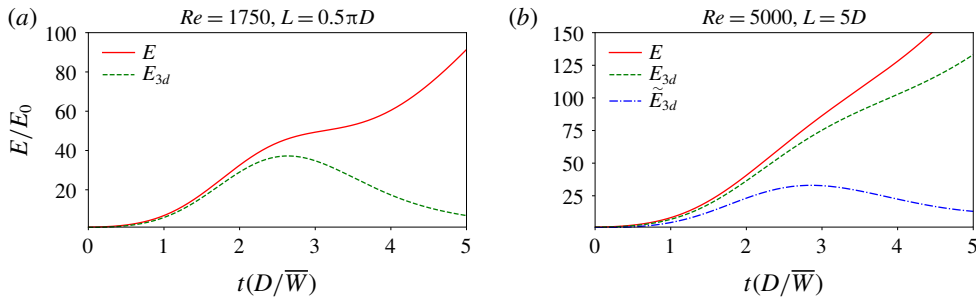


FIGURE 8. (Colour online) Time evolution of the total perturbation energy E (red solid line) and the energy associated with the streamwise-dependent modes only E_{3d} (green dashed line) in (a) the case of a short pipe $L = 0.5\pi D$ at $Re = 1750$ (from PK10) and in (b) the case of a longer pipe ($L = 5D$) at $Re = 5000$ (present study). In the latter case the time evolution of $\tilde{E}_{3d} = E_{3d}(k > 3)$ is also displayed (blue dash-dotted line) to show the characteristic ‘bump’ towards the end of the oblique phase.

reduced to less than an order of magnitude in the former case and, most noticeably, to practically zero in the latter case.

Comparing the scaling of the critical initial energy of the minimal seed with the scaling of its energy at the end of the Orr phase suggests that the growth produced via the Orr mechanism is independent of the Reynolds number, as expected due to the inviscid nature of the Orr process. The growth produced via the oblique-wave mechanism is almost of $O(Re)$. While the growth factor of $O(Re^2)$ for the maximum linear transient growth due to the lift-up mechanism is well documented (e.g. Schmid & Henningson 2012), to the best of our knowledge, this is the first time that scaling laws are obtained numerically for the Orr and oblique-wave mechanisms.

As shown by PWK12, the NLOP tracks the laminar–turbulent boundary Σ before either relaminarising or triggering turbulence. The two bracketing cases shown in figure 6 as black plusses (relaminarising disturbances) and red squares (turbulence-inducing disturbances) are further refined with bisections until the difference in the initial energies is less than 0.005%. These refined bracketing trajectories are shown in figure 9(a) for the range of Reynolds numbers considered and provide evidence of an edge tracked by the minimal seed. Our data (figure 9b) suggest that the energy of an edge state E_Σ decreases with increasing Re , approximately as Re^{-1} .

As suggested by Kerswell (2018) (refer also to appendix B of Kerswell, Pringle & Willis (2014)), the minimal seed couples together (via the nonlinear effects) the Orr, oblique-wave and lift-up mechanisms, which occur on different time scales and are uncoupled in the linearised dynamics. By ensuring that the energy of the preceding phase feeds into the following, the minimal seed is thus able to produce a much larger overall growth than any possible in the linearised problem. Our data support this picture and suggest that the oblique-wave process produces a growth of almost $O(Re)$, which is then further magnified by the lift-up mechanism up to an edge state whose energy scales approximately as $\sim Re^{-1}$. From this, it follows that the lift-up mechanism only produces an energy growth of approximately $O(Re)$, rather than the usually quoted growth factor of $O(Re^2)$. A possible explanation follows from the length scales of the minimal seed becoming finer (and thus the rolls experiencing more dissipation) as the Reynolds number increases. This is evidenced

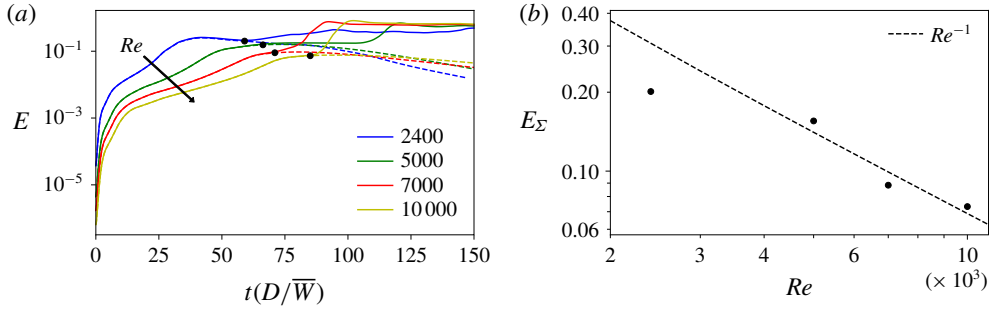


FIGURE 9. (Colour online) (a) Trajectories close to the edge (solid lines: ‘just above’, dashed line: ‘just below’) at different Reynolds numbers in the range 2400–10000. (b) Energy of the edge as a function of Re . At each Reynolds number, E_Σ is calculated as the point on the laminar trajectory where the normalised difference from the turbulent trajectory becomes greater than 5%. The values of E_Σ are also marked with black circles in (a). The dashed curve Re^{-1} is added to show that our scaling assumption is reasonable. This power-law exponent is not calculated by fitting the data but the curve is shifted so that it passes closer to the last three data points (consistent with figure 6).

by the cross-sections shown in figure 10 of the time evolution of the minimal seed up to the beginning of the lift-up phase for $Re = 2400$ and 10000 . Rolls advect the mean shear to drive high and low-speed streaks. The diffusion term for a roll of spanwise wavelength ℓ suggests that such a roll survives a time $\sim Re \ell^2$. For a shear of $O(1)$, the growth in amplitude of a streak is then $\sim Re \ell^2$. The usual argument with $\ell = O(1)$ then implies an energy growth $\sim Re^2$. Here, an energy growth $\sim Re$ suggests a length scale $\ell \sim Re^{-1/4}$.

It appears that we do not yet see scaling with wall units, $\ell^+ = Re_\tau \ell$ and $u^+ = Re/Re_\tau u$, where $Re_\tau \equiv u_\tau R/\nu$ is the wall Reynolds number (R is the radius of the pipe) and $u_\tau = \sqrt{\nu \partial \tilde{w} / \partial r}|_{wall}$ is the friction velocity. For a localised perturbation (l^+ constant) we have $Re_\tau \approx \sqrt{2Re}$. This would suggest a scaling $\ell \sim Re^{-1/2}$. In the early stages of the minimal seed dynamics, however, assuming that the energy in wall units E^+ of a localised perturbation is constant leads to a energy scaling $E \sim Re^{-2.5}$. At finite Reynolds number we observe a scaling of $E \sim Re^{-2.8}$.

The energy of the minimal seed at the end of the oblique phase and the critical initial energy of the localised random disturbances are similar over the range studied. Furthermore, their flow topology is also broadly similar, as can be seen by comparing the cross-section shown in the bottom inset of figure 6 for a typical random localised initial condition with those shown in the rightmost panel of figure 10 for the minimal seed at the end of the oblique phase. In this sense, we will regard the minimal seed at the end of the oblique phase as a reasonable proxy for the transition threshold for random disturbances.

3.3. Control via a body forcing

Motivated by the recent experiments performed by Hof *et al.* (2010), Kühnen *et al.* (2018a,b), as discussed in § 1.2, we study the effect of adding a localised body force that mimics the presence of a baffle in the core of the flow. To study the influence of the baffle on the transition threshold, following the results of the previous section,




FIGURE 10. (Colour online) (a) Time evolution of the minimal seed for $Re = 10000$. From left to right: $t=0$, $t=1.25D/\bar{W}$ (end of the Orr phase) and $t=3.75D/\bar{W}$ (end of the oblique phase). (b) Time evolution of the minimal seed (rotated by π) for $Re=2400$. From left to right: $t=0$, $t=1.1D/\bar{W}$ (end of the Orr phase) and $t=2.5D/\bar{W}$ (end of the oblique phase). The cross-sections (at $t=0$) and right half-sections (end of oblique phase), ten contours are used between the extremes of the corresponding streamwise velocity perturbations. The central half-sections (end of Orr phase) are scaled with the extremes of the corresponding streamwise velocity perturbations at the end of the oblique phase (right half-sections). The cross-sectional velocities (indicated with arrows) are scaled differently for visualisation reasons.

we use the minimal seed at the end of the oblique phase as an initial condition, and verify that it measures transition similarly to random localised disturbances.

Consider a mesh of stationary point objects in the flow and assume that each point experiences a drag proportional to the total velocity. The baffle is then approximated by the following forcing

$$\mathbf{F}(r, \theta, z, t) = -A \mathcal{B}(z) \mathbf{u}_{tot}(r, \theta, z, t), \quad (3.5)$$

where A is the (scalar constant) amplitude of the forcing and $\mathcal{B}(z)$ is a (scalar) smoothed step-like function (refer to equation (3.3)) that introduces a streamwise localisation of the force. The product $A \mathcal{B}(z)$ is a measure of the blockage by the fine mesh. The form of the forcing in (3.5) represents a primitive implementation of an immersed boundary method (refer, for example, to equation (2) in Fadlun *et al.* 2000). As a first approximation, we assume that the baffle is uniform in the radial direction. The streamwise modulation $\mathcal{B}(z)$ is fixed so that the baffle occupies a fifth of the pipe, and the smoothing effect is felt for a tenth of the pipe upstream and downstream of it. The only control parameter is thus A . Simulations are fed with the minimal seed at the end of the oblique phase and the random localised disturbance obtained in the unforced cases (refer to figure 6), with the energy gradually rescaled until transition is triggered. The time horizon is $T = 125(D/\bar{W})$, as in the statistical study presented in § 3.2. For the random localised disturbance we apply a random z -shift to the unforced field before feeding it into the simulations with forcing. For comparison, we have also calculated the effect of the baffle on the transition threshold using the (unforced) minimal seed as initial condition. In this instance, we consider the cases where the initial disturbance is centred in the baffle and half-length away from it.

Re	$(1 + \beta)_{turb}$	$(1 + \beta)_{lam}^{A=0.005}$	\mathcal{DR}_{lam}
2400	1.695	1.143	32.5 %
3500	2.250	1.202	46.5 %
5000	2.940	1.275	56.6 %
7000	3.783	1.354	64.2 %
10 000	4.944	1.48	70 %

TABLE 3. Effect of a forcing of amplitude $A = 0.005$ at different Reynolds numbers. The second and fourth columns are calculated using (3.6) and (3.7), respectively.

Our purpose is to investigate whether the flow can be kept laminar in the presence of the baffle and how much net energy can be saved. The presence of the baffle causes a pressure drop downstream, which is measured by $(1 + \beta)_{lam/turb}^A = \mathcal{D}_{lam/turb}^A / \mathcal{D}_{lam}$, where $\mathcal{D}_{lam/turb}^A$ is the observed value of the dissipation in the presence of the forcing in either the laminar or turbulent case and \mathcal{D}_{lam} is the corresponding laminar value in the unforced case. Hereinafter, the superscript ‘A’ indicates the forced case ($A > 0$) and the subscripts ‘lam’ or ‘turb’ refer to the flow being laminar ($E_0 < E_c$) or turbulent ($E_0 > E_c$) at the current value of $A \geq 0$. We use the turbulent dissipation \mathcal{D}_{turb} in the unforced case as a reference value to quantify the effect of the forcing. In the unforced case, $1 + \beta \equiv Re_p / Re$ where $Re_p = W_{cl}R/\nu$ is the Reynolds number for fixed pressure (W_{cl} is the centreline velocity of the laminar flow). From the Blasius formula (Blasius 1913) for the turbulent friction coefficient $C_f \equiv 16Re_p / Re^2 = 0.0791Re^{-0.25}$, it follows that

$$(1 + \beta)_{turb} \equiv \frac{\mathcal{D}_{turb}}{\mathcal{D}_{lam}} = \frac{Re_p}{Re} = \frac{0.0791}{16} Re^{0.75}. \quad (3.6)$$

For the forcing to be beneficial, the dissipation in the presence of the forcing in either the laminar or turbulent case must be lower than the turbulent dissipation in the unforced case, i.e. $\mathcal{D}_{lam/turb}^A < \mathcal{D}_{turb}$, or, equivalently $(1 + \beta)_{lam/turb}^A < (1 + \beta)_{turb}$. As a measure of how beneficial the forcing is, we define a ‘laminar’ and a ‘turbulent’ drag reduction as:

$$\mathcal{DR}_{lam/turb} = \frac{\mathcal{D}_{turb} - \mathcal{D}_{lam/turb}^A}{\mathcal{D}_{turb}} = \frac{(1 + \beta)_{turb} - (1 + \beta)_{lam/turb}^A}{(1 + \beta)_{turb}}. \quad (3.7)$$

As A is increased, the critical initial energy E_c can be pushed further from the corresponding value in the unforced case, but the pressure drop also increases. For example, in the case of a random localised initial condition, a forcing with $A = 0.005$ avoids turbulence being triggered for values of the initial energy where turbulence was first hit in the unforced case (i.e. the initial energies corresponding to the blue crosses in figure 6), and a considerable laminar drag reduction is obtained, as shown in table 3. However, a slight increase of E_0 would result in the flow to become turbulent again, with almost no drag reduction being achieved. Hence, with this very low choice of A , an almost null raise of E_c is achieved.

At $Re = 5000$ we perform a parametric study on A to find the optimum value which provides the largest E_c at the minimum cost, i.e. with the maximum drag reduction. The value of $Re = 5000$ is chosen to ensure that in the case without a control the turbulence is sustained, i.e. the chances of turbulence decaying randomly are practically nil. The results are summarised in figure 11. In (a) the critical

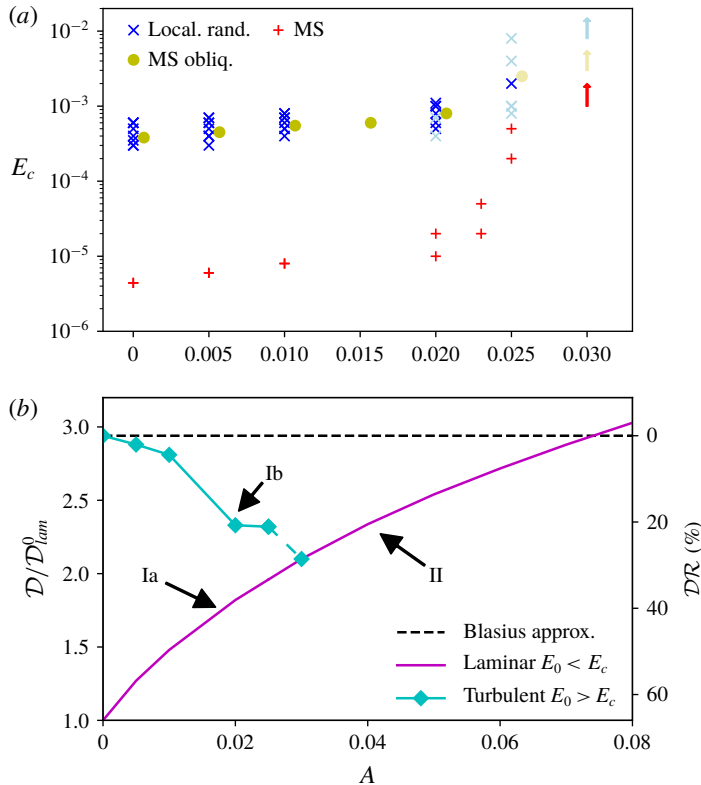


FIGURE 11. (Colour online) Effect of the forcing for different A at $Re = 5000$. (a) Critical initial energies versus forcing amplitude. The critical initial energies for $A = 0$ coincide with the data at $Re = 5000$ of figure 6 for the same types of disturbance. The blue crosses and yellow circles pertain to simulations fed with a random localised initial condition and the minimal seed at the end of the oblique phase, respectively. Note that the yellow circles are slightly shifted for visualisation reasons. In addition, calculations were performed with the minimal seed (with two different shifts applied) as initial disturbance (red plusses). The dark-coloured symbols indicate cases where the flow remains turbulent as E_0 is increased further from the first appearance of turbulence, while the light-coloured symbols denote cases where turbulence is intermittent and characterised by short lifetime. The arrows pointing upwards indicate that $E_c \rightarrow \infty$, i.e. a full collapse of turbulence is obtained. (b) Dissipation and drag reduction versus forcing amplitude. For $A < 0.03$ either laminar (Ia: $E_0 < E_c$) or turbulent (Ib: $E_0 > E_c$) drag reductions are possible, for $A \geq 0.03$, only laminar drag reduction (II) is achieved as turbulence is suppressed.

energies are shown as a function of A for the initial conditions considered here. The corresponding curves for the turbulent and laminar drag reductions are shown in (b). Figure 11(a) shows that, as A is gradually increased up to $A \approx 0.02$, the critical initial energy increases only slightly, but a considerable turbulent drag reduction is obtained. For example, at $A = 0.02$, we obtain $DR_{turb} = 20.7\%$ and $DR_{lam} = 37.8\%$.

For $A > 0.02$, transition starts to become intermittent, i.e. as E_0 is increased further from the first appearance of a turbulent episode, the flow either relaminarise or is characterised by short-lifetime turbulence. In these cases, the ‘critical initial energy’ is indicated by light-coloured symbols in figure 11(a) to distinguish them from

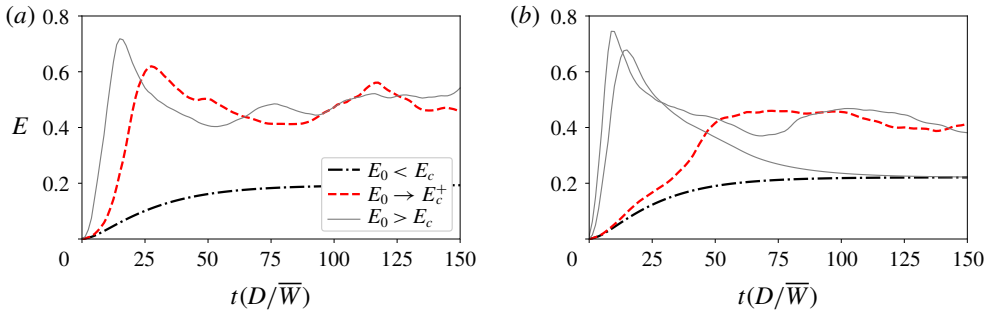


FIGURE 12. (Colour online) Time series of energy for simulations fed with a random localised disturbance in the presence of a forcing of amplitude (a) $A = 0.02$ and (b) $A = 0.025$. In (a) the flow remains turbulent as E_0 is increased further from the first appearance of turbulence (indicated with red dashed line) and turbulence is sustained, while in (b) turbulence is intermittent and characterised by short lifetime. In the latter case, for $E_0 > E_c$, the flow is found to either relaminarise or remain turbulent but with a rapid decay towards the end of the observation time window $T = 125(D/\bar{W})$. These scenarios are marked with dark and light symbols in figure 11, respectively.

the cases, indicated with dark-coloured symbols, where turbulence is sustained. An example of these two situations is shown in figure 12, where in (a) the flow remains turbulent for $E_0 > E_c$ and turbulence is sustained, while in (b) initial disturbances relaminarise and some trigger transition, with the disturbance decaying towards the end of the time window in the latter cases. The scenario where turbulence is short lived and intermittent is analogous to cases where the Reynolds number is close to the first appearance of turbulence ($Re = 1800\text{--}2000$), that is, the effect of increasing A is analogous to that of decreasing the Reynolds number. As A is further increased (e.g. at $A = 0.025$) this effect becomes more and more pronounced (for example, in the case of the random localised initial condition the ‘light crosses’ become dominant with respect to the ‘dark crosses’) until for $A \geq 0.03$ a full collapse of turbulence is obtained (no turbulence episodes are observed). Therefore, the forcing does significantly modify the basin of attraction of the laminar state by expanding it while making its fractal nature more evident until, first, the chaotic attractor transitions back to a chaotic saddle and finally the laminar state remains the only global attractor. The fact that collapse of the turbulent attractor is possible for a forcing of this form can be confirmed using an energy stability type of analysis (refer to appendix A): for large enough A , the laminar state becomes the global attractor.

At $A = 0.03$, where full relaminarisation is first obtained, the laminar drag reduction is still significant (approximately 30%), as shown in figure 11(b). Therefore we can conclude that this is the optimum choice of forcing amplitude A_{opt} . For $A \geq 0.03$ it is not possible to determine \mathcal{DR}_{turb} as turbulence is not observed and thus the curve $\mathcal{DR}_{turb}(A)$ is connected onto $\mathcal{DR}_{lam}(A)$. For $A < 0.03$ we can have either laminar or turbulent drag reduction (i.e. the dynamics either sits on curve Ia or Ib in figure 11), while for $A \geq 0.03$, we only have laminar drag reduction (the dynamics sits on curve II) due to the relaminarisation (destabilisation of turbulent state). The forcing is found to be beneficial up to $A_c = 0.073$, where \mathcal{DR}_{lam} becomes negative, i.e. for $A > A_c$ the cost of the control due to the pressure drop downstream of the baffle becomes greater than the gain due to the relaminarisation.

Figure 11 provides further confirmation that the energy of the minimal seed at the end of the oblique phase is a reasonable proxy to measure transition threshold for random localised disturbances. In this sense, the minimal seed at the end of the oblique phase is a useful tool, potentially enabling us to characterise the critical energy using a single simulation in place of a more expensive statistical study.

4. Conclusions

Nonlinear variational methods have been used to seek the minimal seed, i.e. the initial perturbation of lowest energy that triggers transition to turbulence. The minimal seed represents the most dangerous disturbance to the basic state and, as a result, is of fundamental interest either from the viewpoint of triggering transition efficiently or, oppositely, in designing flow control strategies. We have shown that the structure of the minimal seed is fairly robust to changes in the base flow and to spectral filtering. In the first case, the minimal seed was calculated with a prescribed base flow characterised by a flatter profile in the centre of the pipe as compared to the unforced parabolic profile. In the second case, we projected the initial condition onto a subspace where only a fraction of the streamwise and azimuthal modes were retained. The critical initial energy of the minimal seed was shown to increase with the flatter base profile and with severe spectral filtering (less than 10% of the modes retained), but the structure of the minimal seed was found to remain largely unchanged in both cases.

In order to generate initial conditions that may be considered to model ambient perturbations, we compared the transition behaviour of the minimal seed with that of scaled turbulent snapshots and artificially generated global and localised random disturbances. The random disturbances were obtained by scattering energy randomly over a subset of wavenumbers (the smallest subset for which the critical initial energy of the minimal seed was found to remain unchanged in the previous analysis) and the transition thresholds calculated by gradually increasing the initial energy until turbulence was triggered. Power-law scalings of the critical initial energy $E_c = Re^{-\gamma}$ were obtained with γ in the range 2–3 for different forms of disturbances and $\gamma \approx 2.8$ for the minimal seed. The critical initial energy of the minimal seed was found to be approximately two orders of magnitude lower than that of a localised random disturbance, thus suggesting that, despite being robust, the minimal seed is also quite special. However, when we considered the energy of the minimal seed after the initial nonlinear unpacking phase (composed of the Orr and oblique-wave phases), which occurs in a relatively short time scale, the energy gap with the random localised disturbance became negligible. In this sense, the minimal seed at the end of the oblique phase can be regarded as a reasonable proxy for measuring transition thresholds. Energy growth factors of approximately $O(1)$ and nearly $O(Re)$ for the Orr and oblique-wave mechanisms, respectively, were numerically obtained in the present study for the first time, and the overall picture of the nonlinear growth undergone by the minimal seed to trigger transition discussed. The Orr phase is inviscid and thus the growth produced via this mechanism is independent of the Reynolds number. The oblique-wave process produces a growth in energy of nearly $O(Re)$, which is then seeded to the lift-up mechanism. The disturbance grows by a factor of $O(Re)$ via the lift-up mechanism, rather than the usually quoted growth factor $O(Re^2)$ (a reasonable explanation for this is suggested by providing evidence that the structures become smaller and smaller as the Reynolds number increases), up to an edge state whose energy is shown to scale as $O(Re^{-1})$.

This analysis prepared us with initial conditions for the study of stabilised pipe flows, where a body force was added to the governing equations to mimic the presence of a baffle in the core of the flow, as in the recent experiments by Hof *et al.* (2010) and Kühnen *et al.* (2018a,b). A parametric study on the effect of the amplitude A of the forcing (corresponding, roughly speaking, to a level of blockage in the pipe due to the baffle) at $Re = 5000$ was performed by feeding the simulations with a random localised disturbance and with the minimal seed at the end of the oblique phase found in the unforced case. This confirmed that the minimal seed evolved until the end of the oblique phase is a good proxy for a localised random perturbation, i.e. the critical energy for transition is similar under a variety of forcing situations. An optimum value of the forcing amplitude, $A_{opt} = 0.03$, which provides a full collapse of turbulence ($E_c \rightarrow \infty$) with a drag reduction of approximately 30%, was obtained. Significant drag reductions were found to be possible even in cases where a full collapse of turbulence was not achieved. The forcing was found to be beneficial up to $A_c = 0.073$, for values greater than which the cost of the control due to the pressure drop downstream of the baffle exceeded the energy gain. Although it is not possible at this stage to obtain meaningful estimates of A in laboratory experiments (e.g. Kühnen *et al.* 2018a), due to the artificial forcing used here, this study has shown that modifying the core of the flow by inserting an obstacle is an efficient way of delaying or suppressing turbulence. This method is potentially very attractive as it is passive (no energy input) and very easy to implement. Therefore, we hope that the encouraging results presented here will motivate future developments, such as more realistic modelling of the experimental baffle, to fully exploit the benefits of this control method.

Acknowledgements

This work was funded by EPSRC grant EP/P000959/1. Fruitful discussions with B. Hof are kindly acknowledged. We would like to thank Z. Ding, J. Kühnen, C. Pringle, P. Ricco and D. Wise for insightful comments on a first draft of the manuscript. E.M. would like to acknowledge the assistance of D. Wise in producing some of the figures and dealing with technical issues with the computer clusters. This work would not have been possible without the use of the computing facilities of N8 HPC Centre of Excellence, funded by the N8 consortium and EPSRC (grant EP/K000225/1). The Centre is coordinated by the Universities of Leeds and Manchester.

Appendix A. Energy stability analysis of the forced flow

Consider the Navier–Stokes equations with the forcing given by (3.5):

$$\frac{\partial \mathbf{u}_{tot}}{\partial t} + \mathbf{u}_{tot} \cdot \nabla \mathbf{u}_{tot} + \nabla p = \frac{1}{Re} \nabla^2 \mathbf{u}_{tot} - A \mathcal{B}(z) \mathbf{u}_{tot}, \quad (\text{A } 1)$$

where $\mathbf{u}_{tot}(\mathbf{x}, t) = \mathbf{U}(r, z) + \tilde{\mathbf{u}}(\mathbf{x}, t)$, $\mathbf{U}(r, z)$ is some steady basic state (the laminar response to an axisymmetric baffle) and $\tilde{\mathbf{u}}(\mathbf{x}, t)$ is a possibly large perturbation. The latter is governed by:

$$\frac{\partial \tilde{\mathbf{u}}}{\partial t} + \tilde{\mathbf{u}} \cdot \nabla \mathbf{U} + \mathbf{U} \cdot \nabla \tilde{\mathbf{u}} + \tilde{\mathbf{u}} \cdot \nabla \tilde{\mathbf{u}} + \nabla \tilde{p} = \frac{1}{Re} \nabla^2 \tilde{\mathbf{u}} - A \mathcal{B}(z) \tilde{\mathbf{u}}. \quad (\text{A } 2)$$

Dotting with $\tilde{\mathbf{u}}$ and integrating over the volume gives

$$\frac{\partial}{\partial t} \langle \tilde{\mathbf{u}}^2 \rangle = \langle |\nabla \tilde{\mathbf{u}}|^2 \rangle \left\{ \frac{\langle \tilde{\mathbf{u}} \cdot [-A \mathcal{B}(z) - \nabla \mathbf{U}] \cdot \tilde{\mathbf{u}} \rangle}{\langle |\nabla \tilde{\mathbf{u}}|^2 \rangle} - \frac{1}{Re} \right\}. \quad (\text{A } 3)$$

For the disturbance to decay (i.e. $\partial\langle\tilde{\mathbf{u}}^2\rangle/\partial t < 0$), the amplitude of the forcing needs to be sufficiently large so that, for $Re \rightarrow \infty$,

$$\langle\tilde{\mathbf{u}} \cdot [-A\mathcal{B}(z) - \nabla\mathbf{U}] \cdot \tilde{\mathbf{u}}\rangle < 0 \quad \forall \tilde{\mathbf{u}}. \quad (\text{A } 4)$$

Provided $\|\nabla\mathbf{U}\|$ is bounded as $A \rightarrow \infty$, we find a critical value of the amplitude

$$A_{crit} = \max_{\tilde{\mathbf{u}}} \frac{-\langle\tilde{\mathbf{u}} \cdot \nabla\mathbf{U} \cdot \tilde{\mathbf{u}}\rangle}{\langle\tilde{\mathbf{u}} \cdot \mathcal{B}(z) \cdot \tilde{\mathbf{u}}\rangle}, \quad (\text{A } 5)$$

where $\langle\tilde{\mathbf{u}} \cdot \mathcal{B}(z) \cdot \tilde{\mathbf{u}}\rangle$ is positive definite. Therefore, a forcing with amplitude $A > A_{crit}$ can stabilise any perturbation or, in other words, the steady basic state becomes a global attractor.

REFERENCES

- AUTERI, F., BARON, A., BELAN, M., CAMPANARDI, G. & QUADRIO, M. 2010 Experimental assessment of drag reduction by traveling waves in a turbulent pipe flow. *Phys. Fluids* **22** (11), 115103.
- AVILA, K., MOXEY, D., DE LOZAR, A., AVILA, M., BARKLEY, D. & HOF, B. 2011 The onset of turbulence in pipe flow. *Science* **333**, 192–196.
- BEWLEY, T. R. 2001 Flow control: new challenges for a new renaissance. *Prog. Aerosp. Sci.* **37** (1), 21–58.
- BLASIUS, H. 1913 Das ähnlichkeitsgesetz bei reibungsvorgängen in flüssigkeiten. In *Mitteilungen über Forschungsarbeiten auf dem Gebiete des Ingenieurwesens*, pp. 1–41. Springer.
- BRANDT, L. 2014 The lift-up effect: the linear mechanism behind transition and turbulence in shear flows. *Eur. J. Mech. (B/Fluids)* **47**, 80–96.
- CHERUBINI, S., DE PALMA, P. & ROBINET, J.-CH. 2015 Nonlinear optimals in the asymptotic suction boundary layer: transition thresholds and symmetry breaking. *Phys. Fluids* **27** (3), 034108.
- CHERUBINI, S., DE PALMA, P., ROBINET, J.-CH. & BOTTARO, A. 2012 A purely nonlinear route to transition approaching the edge of chaos in a boundary layer. *Fluid Dyn. Res.* **44** (3), 031404.
- CHERUBINI, S. & PALMA, P. D. 2014 Minimal perturbations approaching the edge of chaos in a couette flow. *Fluid Dyn. Res.* **46** (4), 041403.
- CHOI, H., MOIN, P. & KIM, J. 1994 Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.* **262**, 75–110.
- CHOI, J.-I., XU, C.-X. & SUNG, H. J. 2002 Drag reduction by spanwise wall oscillation in wall-bounded turbulent flows. *AIAA* **40** (5), 842–850.
- CHOI, K.-S. & GRAHAM, M. 1998 Drag reduction of turbulent pipe flows by circular-wall oscillation. *Phys. Fluids* **10** (1), 7–9.
- DARBYSHIRE, A. G. & MULLIN, T. 1995 Transition to turbulence in constant-mass-flux pipe flow. *J. Fluid Mech.* **289**, 83–114.
- DUGGLEBY, A., BALL, K. S. & PAUL, M. R. 2007 The effect of spanwise wall oscillation on turbulent pipe flow structures resulting in drag reduction. *Phys. Fluids* **19** (12), 125107.
- DUGUET, Y., MONOKROUSOS, A., BRANDT, L. & HENNINGSON, D. S. 2013 Minimal transition thresholds in plane Couette flow. *Phys. Fluids* **25** (8), 084103.
- ECKHARDT, B., SCHNEIDER, T. M., HOF, B. & WESTERWEEL, J. 2007 Turbulence transition in pipe flow. *Annu. Rev. Fluid Mech.* **29**, 447–468.
- FADLUN, E. A., VERZICCO, R., ORLANDI, P. & MOHD-YUSOF, J. 2000 Combined immersed-boundary finite-difference methods for three-dimensional complex flow simulations. *J. Comput. Phys.* **161** (1), 35–60.

- HE, S., HE, K. & SEDDIGHI, M. 2016 Laminarisation of flow at low Reynolds number due to streamwise body force. *J. Fluid Mech.* **809**, 31–71.
- HOF, B., DE LOZAR, A., AVILA, M., TU, X. & SCHNEIDER, T. M. 2010 Eliminating turbulence in spatially intermittent flows. *Science* **327** (5972), 1491–1494.
- HOF, B., JUEL, A. & MULLIN, T. 2003 Scaling of the turbulence transition threshold in a pipe. *Phys. Rev. Lett.* **91**, 244502.
- HÖGBERG, M., BEWLEY, T. R. & HENNINGSON, D. S. 2003 Linear feedback control and estimation of transition in plane channel flow. *J. Fluid Mech.* **481**, 149–175.
- JOVANOVIĆ, M. R. 2008 Turbulence suppression in channel flows by small amplitude transverse wall oscillations. *Phys. Fluids* **20** (1), 014101.
- KASAGI, N., SUZUKI, Y. & FUKAGATA, K. 2009 Microelectromechanical systems-based feedback control of turbulence for skin friction reduction. *Annu. Rev. Fluid Mech.* **41**, 231–251.
- KERSWELL, R. R. 2005 Recent progress in understanding the transition to turbulence in a pipe. *Nonlinearity* **18**, R17–R44.
- KERSWELL, R. R. 2018 Nonlinear nonmodal stability theory. *Annu. Rev. Fluid Mech.* **50** (1), 319–345.
- KERSWELL, R. R., PRINGLE, C. C. T. & WILLIS, A. P. 2014 An optimization approach for analysing nonlinear stability with transition to turbulence in fluids as an exemplar. *Rep. Prog. Phys.* **77** (8), 085901.
- KIM, J. & BEWLEY, T. R. 2007 A linear systems approach to flow control. *Annu. Rev. Fluid Mech.* **39**, 383–417.
- KIM, J., MOIN, P. & MOSER, R. 1987 Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.* **177**, 133–166.
- KÜHNEN, J., SCARSELLI, D., SCHANER, M. & HOF, B. 2018a Relaminarization by steady modification of the streamwise velocity profile in a pipe. *Flow Turbul. Combust.* **100** (4), 919–943.
- KÜHNEN, J., SONG, B., SCARSELLI, D., BUDANUR, N. B., RIEDL, M., WILLIS, A., AVILA, M. & HOF, B. 2018b Destabilizing turbulence in pipe flow. *Nat. Phys.* **14** (4), 386–390.
- LEE, C., KIM, J. & CHOI, H. 1998 Suboptimal control of turbulent channel flow for drag reduction. *J. Fluid Mech.* **358**, 245–258.
- LUMLEY, J. & BLOSSEY, P. 1998 Control of turbulence. *Annu. Rev. Fluid Mech.* **30** (1), 311–327.
- MELLIBOVSKY, F. & MESEGUER, A. 2009 Critical threshold in pipe flow transition. *Phil. Trans. R. Soc. Lond. A* **367** (1888), 545–560.
- MOARREF, R. & JOVANOVIĆ, M. R. 2010 Controlling the onset of turbulence by streamwise travelling waves. Part 1. Receptivity analysis. *J. Fluid Mech.* **663**, 70–99.
- MULLIN, T. 2011 Experimental studies of transition to turbulence in a pipe. *Annu. Rev. Fluid Mech.* **43**, 1–24.
- PEIXINHO, J. & MULLIN, T. 2007 Finite-amplitude thresholds for transition in pipe flow. *J. Fluid Mech.* **582**, 169–178.
- PFENNINGER, W. 1961 Transition in the inlet length of tubes at high Reynolds numbers. In *Boundary Layer and Flow Control* (ed. G. V. Lachmann), pp. 970–980. Pergamon.
- POPE, S. B. 2000 *Turbulent Flows*. Cambridge University Press.
- PRINGLE, C. C. T. & KERSWELL, R. R. 2010 Using nonlinear transient growth to construct the minimal seed for shear flow turbulence. *Phys. Rev. Lett.* **105**, 154502.
- PRINGLE, C. C. T., WILLIS, A. P. & KERSWELL, R. R. 2012 Minimal seeds for shear flow turbulence: using nonlinear transient growth to touch the edge of chaos. *J. Fluid Mech.* **702**, 415–443.
- QUADRIO, M. 2011 Drag reduction in turbulent boundary layers by in-plane wall motion. *Phil. Trans. R. Soc. Lond. A* **369** (1940), 1428–1442.
- QUADRIO, M. & SIBILLA, S. 2000 Numerical simulation of turbulent flow in a pipe oscillating around its axis. *J. Fluid Mech.* **424**, 217–241.
- RABIN, S. M. E., CAULFIELD, C. P. & KERSWELL, R. R. 2014 Designing a more nonlinearly stable laminar flow via boundary manipulation. *J. Fluid Mech.* **738**, 1–12.

- REDDY, S. C., SCHMID, P. J., BAGGETT, J. S. & HENNINGSON, D. S. 1998 On stability of streamwise streaks and transition thresholds in plane channel flows. *J. Fluid Mech.* **365**, 269–303.
- REYNOLDS, O. 1883 An experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and the law of resistance in parallel channels. *Proc. R. Soc. Lond. A* **174**, 935–982.
- SCHMID, P. J. & HENNINGSON, D. S. 2012 *Stability and Transition in Shear Flows*, vol. 142. Springer Science & Business Media.
- SCHNEIDER, T. M. & ECKHARDT, B. 2008 Lifetime statistics in transitional pipe flow. *Phys. Rev. E* **78** (4), 046310.
- SCHOPPA, W. & HUSSAIN, F. 2002 Coherent structure generation in near-wall turbulence. *J. Fluid Mech.* **453**, 57–108.
- TREFETHEN, L. N., CHAPMAN, S. J., HENNINGSON, D. S., MESEGUER, A., MULLIN, T. & NIEUWSTADT, F. T. 2000 Threshold amplitudes for transition to turbulence in a pipe. *Numer. Anal. Rep. 00/17*. Oxford University Computer Laboratory.
- WALEFFE, F. 1997 On a Self-Sustaining Process in shear flows. *Phys. Fluids* **9**, 883–900.
- WILLIS, A. P. 2017 The Openpipeflow Navier–Stokes solver. *SoftwareX* **6**, 124–127.
- WILLIS, A. P., HWANG, Y. & COSSU, C. 2010 Optimally amplified large-scale streaks and drag reduction in turbulent pipe flow. *Phys. Rev. E* **82**, 036321.
- WILLIS, A. P., PEIXINHO, J., KERSWELL, R. R. & MULLIN, T. 2008 Experimental and theoretical progress in pipe flow transition. *Phil. Trans. R. Soc. Lond. 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.
- XU, C.-X., CHOI, J.-I. & SUNG, H. J. 2002 Suboptimal control for drag reduction in turbulent pipe flow. *Fluid Dyn. Res.* **30** (4), 217–231.
- YUDHISTIRA, I. & SKOTE, M. 2011 Direct numerical simulation of a turbulent boundary layer over an oscillating wall. *J. Turbul.* **12**, N9.
- ZHOU, D. & BALL, K. S. 2008 Turbulent drag reduction by spanwise wall oscillations. *Int. J. Eng. Trans.* **21** (1), 85.