

# Minimal seeds for shear flow turbulence: using nonlinear transient growth to touch the edge of chaos

Chris C. T. Pringle<sup>1†</sup>, Ashley P. Willis<sup>2</sup> and Rich R. Kerswell<sup>3</sup>

<sup>1</sup> Department of Earth Sciences, University of Bristol, Bristol BS8 1RJ, UK

<sup>2</sup> School of Mathematics and Statistics, University of Sheffield, Sheffield S3 7RH, UK

<sup>3</sup> Department of Mathematics, University of Bristol, Bristol BS8 1TW, UK

(Received 12 September 2011; revised 12 April 2012; accepted 16 April 2012;  
first published online 29 May 2012)

We propose a general strategy for determining the minimal finite amplitude disturbance that triggers transition to turbulence in shear flows. This involves constructing a variational problem that searches over all disturbances of fixed initial amplitude which respect the boundary conditions, incompressibility and the Navier–Stokes equations, to maximize a chosen functional over an asymptotically long time period. The functional must be selected such that it identifies turbulent velocity fields by taking significantly enhanced values compared to those for laminar fields. We illustrate this approach using the ratio of the final to initial perturbation kinetic energies (energy growth) as the functional and the energy norm to measure amplitudes in the context of pipe flow. Our results indicate that the variational problem yields a smooth converged solution provided that the initial amplitude is below the threshold for transition. This optimal is the nonlinear analogue of the well-studied (linear) transient growth optimal. At the critical threshold, the optimization seeks out a disturbance that is on the ‘edge’ of turbulence during the period. Above this threshold, when disturbances trigger turbulence by the end of the period, convergence is then practically impossible. The first disturbance found to trigger turbulence as the amplitude is increased identifies the ‘minimal seed’ for the given geometry and forcing (Reynolds number). We conjecture that it may be possible to select a functional such that the converged optimal below threshold smoothly converges to the minimal seed at threshold. Our choice of the energy growth functional is shown to come close to this for the pipe flow geometry investigated here.

**Key words:** nonlinear dynamical systems, nonlinear instability, turbulent transition

---

## 1. Introduction

Shear flows are ubiquitous in our everyday lives yet predicting their behaviour remains an outstanding and important issue both scientifically and economically. Typically such flows become turbulent even though there may be an alternative linearly stable ‘basic state’, which is the simplest solution consistent with the driving forces and boundary conditions. This bistability means that the problem of transition comes

<sup>†</sup> Email address for correspondence: [chris.pringle@bris.ac.uk](mailto:chris.pringle@bris.ac.uk)

down to understanding the laminar–turbulent boundary in phase space that divides initial conditions that lead to the turbulent state from those that relax back to the basic state. This boundary has more generally been labelled the ‘edge of chaos’, allowing for transient turbulence (Skufca, Yorke & Eckhardt 2006). There have been notable recent successes in tracking parts of this boundary which, because it is a hypersurface in phase space, can be approached by a simple bisection technique (Itano & Toh 2001; Skufca *et al.* 2006; Schneider, Eckhardt & Yorke 2007; Duguet, Willis & Kerswell 2008). As this technique is based upon integrating the governing Navier–Stokes equations forward in time, only the parts of the edge near (relative) attractors embedded in the edge are revealed by this tracking approach.

Interestingly, these attracting regions invariably seem to be on perturbation energy levels well above those known to be sufficient to trigger transition. Viswanath & Cvitanovic (2009) provide a good illustration of this in a short pipe of length  $\pi$  diameters, which they numerically simulate with 85715 degrees of freedom. By only mixing three fixed flow fields, they identify initial conditions that experience an  $O(10^4)$  magnification of the (perturbation) energy in approaching a travelling wave thought to be embedded in the edge relative attractor (Pringle & Kerswell 2007; Schneider *et al.* 2007): see their table 4 for Reynolds number  $Re = 2000$ . Duguet, Brandt & Larsson (2010a) tackle the same question in plane Couette flow, finding evidence for energy growth on the edge of over  $10^2$  at  $Re = 400$  for a pair of oblique waves (see their figure 9 where the plateau edge energy is  $O(5 \times 10^{-3})$ ). Initial conditions with low energies on the edge represent energy-efficient targets to trigger transition, as an infinitesimal perturbation of these states will lead to transition. The most efficient of all perturbations will be the flow field having the lowest energy  $E_c$  on the edge, hereafter called the ‘minimal seed’ for transition, which represents the closest (in perturbation energy norm) point of approach of the edge to the basic state in phase space. This represents the best way to disturb the basic state and as a result is of fundamental interest, either from the viewpoint of triggering transition efficiently or, conversely, in designing flow control strategies.

Currently, there are no accepted strategies for identifying minimal seeds beyond the impractical ‘brute force’ approach of surveying all initial conditions. The purpose of this paper is to continue to develop a new strategy initiated in Pringle & Kerswell (2010), hereafter referred to as PK10, based upon identifying finite-amplitude disturbance fields which, as they evolve via the full Navier–Stokes equations, maximize a key functional over a period of time. This key functional is taken here to be the energy growth of the disturbance over a time period, as suggested in PK10 and Cherubini *et al.* (2010, 2011) in the boundary layer context. The rationale behind this is the observation that the minimal seed must experience considerable energy growth as it evolves in time up to the attracting plateau on the edge; note that Cherubini *et al.* (2010, 2011) use the term ‘minimal seed’ differently, in a structural sense rather than the energetic sense adopted here. In the special case of a unique steady relative attractor on the edge – for instance, in a pipe  $\approx 2.5$  diameters long at  $Re = 2400$  and within the symmetric subspace  $\mathcal{R}_2 \cap \Omega_2 \cap \mathcal{S}$  (see § 3.6 Duguet *et al.* 2008), the unique edge relative attractor is C3.1.25 (later renamed N2 in Pringle, Duguet & Kerswell 2009) – the minimal seed will indeed be the optimal solution  $\mathbf{u}_0^*(\mathbf{x})$  to the following variational problem: which initial condition on the edge (label this set  $\Sigma$ ) will experience, for asymptotically long times  $T$ , the largest energy growth,

defined by

$$\mathcal{G}(T) := \max_{\mathbf{u}_0 \in \Sigma} \frac{\int \mathbf{u}(\mathbf{x}, T)^2 dV}{\int \mathbf{u}(\mathbf{x}, 0)^2 dV}, \quad (1.1)$$

where  $\mathbf{u}(\mathbf{x}, t)$  is the flow at time  $t$  evolved via the Navier–Stokes equations from the initial condition  $\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x})$  (incompressibility is tacitly assumed throughout) and  $\mathbf{u} = (u, v, w)$  is the perturbation velocity field obtained by subtracting the laminar state from the total velocity field (invariably used hereafter). The problem with pursuing this criterion is not the fact that the relative edge attractor may have a fluctuating energy, as these fluctuations are typically small compared to the total growth, but confining competitor fields to the hard-to-define edge set  $\Sigma$ . A more practical approach can be manufactured by turning the problem around to consider the largest energy growth  $G$  over all initial (incompressible) conditions of a given perturbation energy  $E_0$ , that is,

$$G(T; E_0) := \max_{\mathbf{u}_0: \int \mathbf{u}_0^2 dV = 2E_0} \frac{\int \mathbf{u}(\mathbf{x}, T)^2 dV}{\int \mathbf{u}(\mathbf{x}, 0)^2 dV}. \quad (1.2)$$

At precisely  $E_0 = E_c$ , where the edge touches the energy hypersurface at one velocity state, this optimization problem considers the growth of this state (the minimal seed) against the energy growth of all the other initial conditions below the edge. Given that these latter initial conditions lead to flows that grow initially but ultimately relax back to the basic state, it is reasonable to hypothesize that the minimal seed remains the optimal initial condition for the revised variational problem (1.2) for large enough  $T$ . *A priori*, the minimal seed energy  $E_c(Re; \text{geometry})$  is unknown but a very interesting quantity in its own right, as its behaviour indicates how the basin of attraction of the basic state shrinks with increasing  $Re$ . Hence, the variational problem (1.2) must be solved as an increasing function of  $E_0$  until  $E_c$  is reached. Knowing when this has occurred motivates the following hypothesis.

**Hypothesis.** For asymptotically large  $T$ , the minimal seed will be given by the flow field of initial energy  $E_0$  that experiences the largest energy growth such that for any initial energies exceeding this  $E_0 = E_c$ , the energy growth problem (1.2) fails to have a smooth solution.

This is in fact a very strong statement containing two separate but related conjectures, the first being a necessary condition for the second.

**Conjecture 1.** For  $T$  sufficiently large, the initial energy value  $E_{\text{fail}}$  at which the energy growth problem (1.2) first fails (as  $E_0$  is increased) to have a smooth optimal solution will correspond exactly to  $E_c$ .

**Conjecture 2.** For  $T$  sufficiently large, the optimal initial condition for maximal energy growth at  $E_0 = E_c - \epsilon^2$  converges (in the energy norm) to the minimal seed at  $E_c$  as  $\epsilon \rightarrow 0$ .

The idea behind the first conjecture is that the optimization algorithm (1.2), in exploring the  $E_0$ -hypersurface for the optimal solution, will detect any state on the energy hypersurface that leads to turbulence, given that this leads to the highest values

of  $G$ . Once the algorithm is dealing with a turbulent end state at time  $T$ , the extreme sensitivity of the final-state energy at  $T$  to changes in the initial condition, due to exponential divergence of adjacent states, will effectively mean non-smoothness and prevent convergence. Crucially, if true, this means that the failure of the algorithm solving (1.2) should identify  $E_c$  regardless of whether the minimal seed is the optimal solution of (1.2) for  $E_0 = E_c$  or not. The key feature for Conjecture 1 to hold is not the precise form of the functional being maximized, but the fact that the functional attains higher values for initial conditions that go turbulent than those in the basin of attraction of the basic state. Other plausible choices include the final dissipation rate, the total dissipation, which has recently been explored with success by Monokrousos *et al.* (2011) in the context of plane Couette flow, or more general Sobolev norms that emphasize strain rates. The second, stronger conjecture, however, proposes that the optimal initial condition converged at  $E_0 < E_c$  values approaches the minimal seed as  $E \rightarrow E_c$ . This implies that the energy growth functional is then a special choice which picks out the minimal seed. We present evidence in this paper to support both these conjectures.

The variational problem (1.2) in the limit of infinitesimally small energy  $E_0$  reduces to the well-known (linear) transient growth problem (Gustavsson 1991; Butler & Farrell 1992; Reddy & Henningson 1993; Trefethen, Trefethen & Reddy 1993; Schmid & Henningson 1994). In this case, the evolution of the initial condition  $\mathbf{u}_0$  is determined by the linearized Navier–Stokes equations and there are various ways to proceed: for a matrix-based approach see, for example, Reddy & Henningson (1993) and for a time-stepping approach see Luchini (2000). The optimal that emerges typically shows energy growth factors  $G$  that scale with  $Re^2$  (when optimization is also carried out over  $T$ ) due to the non-normality of the linearized evolution operator. It has been common practice to assume that this linear optimal (LOP) for vanishing  $E_0$  (together with some noise to break two-dimensionality) is a good approximation to the minimal seed, as the energy of the minimal seed is ‘small’ compared to that of the basic flow or target turbulent flow. It was shown in PK10, however, that the presence of nonlinearity in the variational problem is crucial in revealing new nonlinear optimals (NLOPs) that emerge ‘in between’, i.e. for  $0 < E_0 < E_c$  (the term ‘NLOP’ is used to describe the three-dimensional optimal, which, not being smoothly connected to the LOP as  $E_0$  increases, has distinctly different structure). The calculations of PK10 were directed more at demonstrating the feasibility of including nonlinearity in the transient growth calculation and showing the dramatic manner in which this alters the established linear result, than in identifying the minimal seed. In particular,  $T = T_{lin}$  was taken, where  $T_{lin}$  is the optimal growth time for the linear optimal (LOP) and is not asymptotically large. Numerical limitations of the simulation code used in PK10 (written from scratch as part of the first author’s thesis) also meant that getting close to the edge proved difficult. In this paper, we revisit those calculations using a highly parallel code (described in Willis & Kerswell (2009) and tested with the results of PK10) to probe the ‘gap’ between  $E_{fail}$  and  $E_c$  noticed there.

The plan of the paper is as follows. Section 2 formulates the variational problem (1.2) and describes briefly the iterative scheme used to solve it. Section 3 analyses for the first time the mechanism by which the nonlinear optimal (NLOP) that emerged in PK10 attains more growth than the linear optimal (LOP). The results of PK10 are then extended to higher energies in § 4 to re-examine the reported gap between  $E_{fail}$  and  $E_c$ . In § 5, a larger domain is studied using a longer optimization time to provide a first test of the conjectures discussed above. Our results are summarized and discussed in § 6 with a glossary of terms following at the end.

## 2. Formulation

The context for our exploration is the problem of constant mass-flux fluid flow through a cylindrical pipe. With length scales non-dimensionalized by half the pipe diameter  $\pi D/2$  and velocities by the mean axial velocity  $U$ , the laminar flow is given by

$$\mathbf{u}_{lam} = \mathcal{U}(s)\hat{\mathbf{z}} = 2(1 - s^2)\hat{\mathbf{z}}, \quad (2.1)$$

using cylindrical coordinates  $(s, \phi, z)$  aligned with the pipe axis. In keeping with the majority of published work on pipe flow, results are reported in time units of  $D/U$ . Energies are non-dimensionalized by the energy of the laminar flow in the same domain. We then consider a perturbation to this laminar profile such that the full velocity field is given by

$$\mathcal{U}(s)\hat{\mathbf{z}} + \mathbf{u}(s, \phi, z, t), \quad (2.2)$$

where  $\mathbf{u} = (u, v, w)$ , and for convenience define the volume integral

$$\langle \dots \rangle := \int \dots dV = \int_0^L \int_0^{2\pi} \int_0^1 \dots s ds d\phi dz. \quad (2.3)$$

In order to calculate the initial condition that produces the most energy growth, we use a variational approach pioneered for the linearized problem (Luchini & Bottaro 1998; Andersson, Berggren & Henningson 1999; Corbett & Bottaro 2000; Luchini 2000) but now recently extended to incorporate the full Navier–Stokes equations (PK10; Cherubini *et al.* 2010, 2011; Monokrousos *et al.* 2011; see also Zuccher, Luchini & Bottaro 2004 for earlier work using the boundary layer equations). The functional we choose is defined as

$$\begin{aligned} \mathcal{L} := & \left\langle \frac{1}{2} \mathbf{u}(\mathbf{x}, T)^2 \right\rangle - \lambda \left[ \left\langle \frac{1}{2} \mathbf{u}(\mathbf{x}, 0)^2 \right\rangle - E_0 \right] \\ & - \int_0^T \left\langle \mathbf{v} \cdot \left[ \frac{\partial \mathbf{u}}{\partial t} + \mathcal{U} \frac{\partial \mathbf{u}}{\partial z} + \mathcal{U}' u \hat{\mathbf{z}} - \mathbf{u} \times \nabla \times \mathbf{u} + \nabla p - \frac{1}{Re} \nabla^2 \mathbf{u} \right] \right\rangle dt \\ & - \int_0^T \langle \Pi \nabla \cdot \mathbf{u} \rangle dt - \int_0^T \Gamma(t) \langle \mathbf{u} \cdot \hat{\mathbf{z}} \rangle dt. \end{aligned} \quad (2.4)$$

This functional will be maximized by the same flow field as problem (1.2). It is equivalent to finding the flow field with greatest energy at time  $t = T$ , subject to four conditions applied through Lagrange multipliers – namely that the initial condition,  $\mathbf{u}(\mathbf{x}, 0)$  has kinetic energy  $E_0$  and that it evolves subject to the incompressible Navier–Stokes equations with fixed mass flux along the pipe. The last constraint introduces a slight subtlety in that the pressure field must be subdivided into a time-dependent constant pressure gradient part  $\Lambda(t)z$  that adjusts to maintain constant mass flux and a strictly (spatially) periodic part  $\hat{p}$  so that

$$p := \Lambda(t)z + \hat{p}(s, \phi, z, t). \quad (2.5)$$

The Lagrange multipliers  $\mathbf{v} = (v_s, v_\phi, v_z)$ ,  $\Pi$  and  $\Gamma$  are known as the adjoint variables. The function  $\mathcal{L}$  will be maximized when all of its variational derivatives are equal to zero. Taking variational derivatives (the boundary conditions on the flow field – non-slip at  $s = 1$  and periodicity across the pipe – are used to eliminate surface

terms) leads us to

$$\begin{aligned}
 \delta \mathcal{L} = & \langle \delta \mathbf{u}(\mathbf{x}, T) \cdot [\mathbf{u}(\mathbf{x}, T) - \mathbf{v}(\mathbf{x}, T)] \rangle + \langle \delta \mathbf{u}(\mathbf{x}, 0) \cdot [-\lambda \mathbf{u}(\mathbf{x}, 0) + \mathbf{v}(\mathbf{x}, 0)] \rangle \\
 & - \int_0^T \left\langle \delta \mathbf{v} \cdot \left[ \frac{\partial \mathbf{u}}{\partial t} + \mathcal{U} \frac{\partial \mathbf{u}}{\partial z} + \mathcal{U}' u \hat{\mathbf{z}} - \mathbf{u} \times \nabla \times \mathbf{u} + \nabla p - \frac{1}{Re} \nabla^2 \mathbf{u} \right] \right\rangle dt \\
 & + \int_0^T \left\langle \delta \mathbf{u} \cdot \left[ \frac{\partial \mathbf{v}}{\partial t} + \mathcal{U} \frac{\partial \mathbf{v}}{\partial z} - \mathcal{U}' v \hat{\mathbf{z}} + \nabla \times (\mathbf{v} \times \mathbf{u}) - \mathbf{v} \times \nabla \times \mathbf{u} + \nabla \Pi \right. \right. \\
 & \left. \left. + \frac{1}{Re} \nabla^2 \mathbf{v} - \Gamma(t) \hat{\mathbf{z}} \right] \right\rangle dt - \int_0^T \langle \delta \Pi \nabla \cdot \mathbf{u} \rangle dt - \int_0^T \delta \Gamma \langle \mathbf{u} \cdot \hat{\mathbf{z}} \rangle dt \\
 & + \int_0^T \langle \delta \hat{p} \nabla \cdot \mathbf{v} \rangle dt - \int_0^T \delta \Lambda(t) \langle \mathbf{v} \cdot \hat{\mathbf{z}} \rangle dt - \delta \lambda \left[ \left\langle \frac{1}{2} \mathbf{u}(\mathbf{x}, 0)^2 \right\rangle - E_0 \right]. \quad (2.6)
 \end{aligned}$$

The nine terms making up the variational derivative can physically be interpreted as meaning that to maximize  $\mathcal{L}$ : (i)  $\mathbf{u}(\mathbf{x}, T)$  and  $\mathbf{v}(\mathbf{x}, T)$  must satisfy a compatibility condition; (ii)  $\mathbf{u}(\mathbf{x}, 0)$  and  $\mathbf{v}(\mathbf{x}, 0)$  must satisfy an optimality condition; (iii)  $\mathbf{u}$  must evolve according to the Navier–Stokes equations (with boundary conditions  $\mathbf{u}(1, \phi, z, t) = \mathbf{0}$  and  $\mathbf{u}(s, \phi, 0, t) = \mathbf{u}(s, \phi, L, t)$ ); (iv)  $\mathbf{v}$  must evolve according to the adjoint Navier–Stokes equations (with natural boundary conditions  $\mathbf{v}(1, \phi, z, t) = \mathbf{0}$  and  $\mathbf{v}(s, \phi, 0, t) = \mathbf{v}(s, \phi, L, t)$ ); (v)  $\mathbf{u}$  is incompressible; (vi)  $\mathbf{u}$  has constant mass flux; (vii)  $\mathbf{v}$  is incompressible; (viii)  $\mathbf{v}$  has constant mass flux; and (ix) the initial kinetic energy is  $E_0$  (respectively as the terms appear in (2.6)).

In order to find a maximum to this problem, an iterative algorithm is employed, seeded by an initial flow field  $\mathbf{u}_0 := \mathbf{u}(\mathbf{x}, 0)$  of appropriate kinetic energy (similar shorthand is used henceforth, e.g.  $\mathbf{v}_0 := \mathbf{v}(\mathbf{x}, 0)$ ). By integrating this field forward in time in accordance with the Navier–Stokes equations we can ensure that conditions (iii) and (v) are met. The compatibility condition (ii) is satisfied by fixing  $\mathbf{v}_T = \mathbf{u}_T$ , which supplies a final condition for the adjoint-Navier–Stokes equations to be integrated backwards in time. This procedure generates  $\mathbf{v}_0$  and ensures conditions (iv) and (vi) are fulfilled. After this ‘forth-and-back’ time integration, the only outstanding Euler–Lagrange condition is that the variational (Fréchet) derivative

$$\frac{\delta \mathcal{L}}{\delta \mathbf{u}_0} = -\lambda \mathbf{u}_0 + \mathbf{v}_0 \quad (2.7)$$

should vanish. As this does not happen automatically,  $\mathbf{u}_0$  is moved in the ascent direction to increase  $\mathcal{L}$  and hopefully approach a maximum where it will vanish. An initial condition for the next iteration is given by

$$\mathbf{u}_0^{n+1} = \mathbf{u}_0^n + \frac{\epsilon}{\lambda} \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0^n}, \quad (2.8)$$

where  $\langle (\delta \mathcal{L} / \delta \mathbf{u}_0)^2 \rangle / \lambda^2 \sim \langle \mathbf{u}_0^2 \rangle$  is a convenient rescaling. The one remaining Lagrange multiplier  $\lambda$  is determined by arranging for the new initial condition to satisfy the initial energy constraint  $\langle (\mathbf{u}_0^{n+1})^2 / 2 \rangle = E_0$ . It is found that  $\lambda \sim \langle \mathbf{u}_T^2 \rangle / \langle \mathbf{u}_0^2 \rangle$ , and for the choice  $\epsilon = 1$  this strategy is equivalent to the power method in the linear case (e.g. Luchini 2000). In practice for the nonlinear case, the choice  $\epsilon = 1$  is usually too large. The following simple strategy for the adaptive selection of  $\epsilon$  was found to be effective.

- Select an initial value for  $\epsilon$ , e.g. 0.5.
- Let

$$d := \left\langle \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0^n} \cdot \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0^{n+1}} \right\rangle / \sqrt{\left\langle \left( \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0^n} \right)^2 \right\rangle \left\langle \left( \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0^{n+1}} \right)^2 \right\rangle}. \quad (2.9)$$

- If  $d > 0.95$ , so successive adjustments in  $\mathbf{u}_0$  are essentially aligned, then  $\epsilon$  is doubled. Otherwise, if  $d < -0.5$  (anti-alignment) or

$$\left\langle \left( \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0^{n+1}} \right)^2 \right\rangle > 4 \left\langle \left( \frac{\delta \mathcal{L}}{\delta \mathbf{u}_0^n} \right)^2 \right\rangle, \quad (2.10)$$

whereby the derivative has becomes large, then  $\epsilon$  is halved.

Close to apparent convergence, when  $\langle (\delta \mathcal{L} / \delta \mathbf{u}_0^n)^2 \rangle$  is very small, a constant  $\epsilon$  has sometimes been employed to prevent multiple looping with tiny updates each time.

The numerical code used here is based on the well-tested code described in Willis & Kerswell (2009). A Fourier decomposition is employed in the periodic directions and a finite difference approximation in the radial direction, so that a typical dependent variable is expanded as follows:

$$A(s_n, \phi, z, t) = \sum_{m=-MM}^{MM} \sum_{l=-LL}^{LL} A_{nml}(t) \exp(im\phi + i\alpha lz) \quad \text{for } n = 1, 2, \dots, NN, \quad (2.11)$$

where  $A$  is real, so only half the coefficients ( $m \geq 0$ ) need to be stored,  $\alpha = 2\pi/L$  is the longest wavelength allowed by the periodic axial boundary conditions, and  $s_n$  are the roots of a Chebyshev polynomial with finer resolution towards the wall. Typical resolutions used were  $(MM, NN, LL) = (23, 64, 11)$  for a  $\pi D/2$  (PK10) pipe and  $(23, 64, 37)$  for a  $5D$  pipe. Using finite differences in the radius is apt for parallelization, which has been implemented using MPI. Time integration is performed using a second-order predictor–corrector method.

A fast (parallel) numerical code for handling the Navier–Stokes equations and its adjoint is absolutely essential for successfully implementing this iterative approach to optimization. Each iteration requires integrating the Navier–Stokes equations forward from  $t = 0$  to  $T_{opt}$ , and the adjoint equations backwards from  $t = T_{opt}$  to 0, with typically  $O(10^3)$  iterations required to ensure convergence. There are also storage issues to circumvent, as the adjoint Navier–Stokes equations, although linear in  $\mathbf{v}$ , depend on  $\mathbf{u}$ . This either needs to be stored in totality (over the whole volume and time period), which is only practical for low-resolution short integrations, or must be recalculated piecemeal during the backward integration stage. This latter ‘check-pointing’ approach requires that  $\mathbf{u}$  is stored at regular intermediate points, e.g.  $t = T_i := iT_{opt}/n$  for  $i = 1, \dots, n - 1$ , during the forward integration stage. Then to integrate the adjoint equation backward over the time interval  $[T_i, T_{i+1}]$ ,  $\mathbf{u}$  is regenerated starting from the stored value at  $t = T_i$  by integrating the Navier–Stokes equations forward to  $T_{i+1}$  again. The extent of the check pointing is chosen such that the storage requirement for each subinterval is manageable. The extra overhead of this technique is to redo the forward integration for every backward integration, so approximately a 50% increase in CPU time, assuming forward and backward integrations take the same time. As memory restrictions may make full storage impossible, this is a small price to pay.



### 3. The nonlinear energy growth mechanism

The basic ingredient for the strategy being explored here is the solution of the variational problem (1.2) at a given initial perturbation energy  $E_0$ . PK10 demonstrated the feasibility of this and discovered that the nonlinearly adjusted linear optimal (LOP; see Schmid & Henningson 1994) is quickly outgrown by a completely new type of optimal (the NLOP) as  $E_0$  increases from 0 (e.g. see figures 1 and 5 in PK10). This NLOP exhibits both radial and azimuthal localization in a short  $\pi D/2$ -long pipe and would undoubtedly also localize in the axial direction if the geometry allowed. Localization allows the perturbation to retain velocities of sufficient amplitude in an adequate volume such that nonlinearity is important while permitting the global energy to be reduced. Without nonlinearity in the variational problem, any localized state could be decomposed into global linear optimals (e.g. by Fourier analysis), which would then evolve independently and, all except the LOP, sub-optimally. PK10 remarked that the NLOP had a two-phase evolution (e.g. see their figure 1) in which the initially three-dimensional optimal first delocalizes (slices *a* and *b* in their figure 2) followed by a second growth phase in which the flow becomes increasingly two-dimensional (streamwise independent). We now examine this evolution in more detail in order to understand how the NLOP is able to achieve more growth than the LOP.

A first inspection of the three-dimensional structure of the evolving NLOP actually reveals three distinct phases of development. Figure 1 shows how the axial structure of the NLOP evolves in time by plotting isocontours of streamwise perturbation velocity along the pipe (isocontours of streamwise vorticity show the same qualitative behaviour). Initially the contours are tightly layered and backward-facing, i.e. inclined into the shear: see the  $t = 0$  contour plot in figure 1 near  $(x, y) + (-0.5, -0.5)$  where the contours are slanted away from the pipe wall in the anti-streamwise direction (downwards). By  $t = 0.4$ , these contour layers have been tilted into the mean shear direction (i.e. the angle they make with the pipe wall has increased) by the shear and unpacked or separated slightly. This kinematic deformation of the perturbation by the underlying shear is the inviscid Orr mechanism (Orr 1907) and gives an initial spurt of energy growth. By  $t = 1$  the flow is then dominated by helical waves growing – the ‘oblique’ phase – before the flow becomes essentially although not completely two-dimensional by  $t = 10$  during the ‘lift-up’ phase. This evolution, consisting of the Orr, oblique and lift-up phases in sequence, is also apparently seen for the critical disturbance found by Monokrousos *et al.* (2011) (D. S. Henningson, private communication).

To clarify the oblique and lift-up phases, we reduce the considerable degrees of freedom of the fully resolved NLOP evolution down to those that really matter. As the Fourier–Fourier basis functions  $\exp(im\phi + i\alpha lz)$  in the velocity representation naturally partition the linearized problem, we considered the optimal growth calculation at  $Re = 1750$ ,  $T_{opt} = T_{lin} \approx 0.0122Re = 21.3D/U$  in a pipe length of  $\pi D/2$  (so PK10 settings) at  $E_0 = 1.8 \times 10^{-5}$  using the full resolution  $(MM, LL) = (23, 11)$  and reduced resolutions  $(7, 1)$ ,  $(2, 1)$  and  $(1, 1)$  (full  $NN = 64$  radial resolution was used for all). Figure 2 shows how the growth evolves as a function of time for each optimal initial condition. All the calculations bar that for  $(1, 1)$  show the distinctive ‘shoulder’ in the growth centred at  $t \approx 3D/U$  (the time units  $D/U$  will be suppressed hereafter), which signifies the end of the (what PK10 called the ‘first’) delocalization phase and the start of the next phase. The  $(1, 1)$  calculation fails to capture the NLOP at all, so the optimal that emerges is just the nonlinear version of the LOP. It is also striking that the  $(7, 1)$  optimal is quantitatively so similar to the full  $(23, 11)$ -optimal indicating that axial wavenumbers beyond the lowest are not important for this short






FIGURE 1. (Colour online available at [journals.cambridge.org/film](https://journals.cambridge.org/film)) The NLOP at  $Re = 1750$  and  $E_0 = 1.8 \times 10^{-5}$  calculated using resolution  $(MM, NN, LL) = (23, 64, 11)$  at (a)  $t = 0$ , (b)  $t = 0.4$ , (c)  $t = 1$ , (d)  $t = 2.5$ , (e)  $t = 5$  and (f)  $t = 10D/U$ . Isocontours are shown of the streamwise perturbation velocity: light/yellow, 50 % of the maximum; dark/red, 50 % of the minimum. Mean flow in each pipe section is from bottom to top.

pipe calculation. Drastically reducing the azimuthal resolution to just  $(2, 1)$  has a noticeable quantitative effect but still manages to preserve the qualitative features of the NLOP. In particular, the  $(2, 1)$ -optimal (figure 2c) captures the essential structure of the  $(7, 1)$ -optimal (figure 2b), which is itself almost identical to the  $(23, 11)$ -optimal (figure 4a) (although no attempt has been made to match phases of the solutions along the pipe).

The temporal evolutions of the modal kinetic energies  $E_{ml}(t)$  (defined as the kinetic energy associated with the Fourier–Fourier wavenumbers  $m, l$ ) for the  $(7, 1)$  calculation are shown in figure 3 (the equivalent plot for the  $(2, 1)$  calculation is qualitatively similar but not shown). The modal energy for streamwise independent velocities,  $E_{m0}$ , is further split into that associated with the streamwise velocity,  $E_{m0}^w$ , and that associated with the cross-plane velocities  $u$  and  $v$ ,  $E_{m0}^{uv}$ . Each modal energy can change because of three effects: input from the underlying basic state due to the non-normality of the linearized operator, loss due to viscous dissipation and either loss or gain through nonlinear mixing with the other modes. Generally, it is difficult to distinguish between these effects without explicitly monitoring the various terms in the Navier–Stokes equations. However, for streamwise independent modes, the cross-plane energy  $E_{m0}^{uv}$  cannot grow by non-normal effects so any energy gain must be the result




FIGURE 2. (Colour online) (a) Energy growth  $E(t)/E(0)$  against time  $t$  for the NLOP of PK10 at  $Re = 1750$  and  $E_0 = 1.8 \times 10^{-5}$  using full resolution  $(MM, NN, LL) = (23, 64, 11)$  (solid black line) and reduced resolutions  $(7, 64, 1)$  (dot-dashed red line) and  $(2, 64, 1)$  (dashed blue line):  $MM$  value used as a label. Removing the helical mode effect by using resolution  $(1, 64, 1)$  (dotted green line) destroys the NLOP to leave the LOP (notice the absence of a ‘shoulder’ in the curve at  $T \approx 3D/U$ ). (b,c) NLOPs for  $(7, 64, 1)$  (b) and  $(2, 64, 1)$  (c). Contours indicate streamwise velocity perturbation (total velocity with the laminar state of equivalent mass flux subtracted off) using the same levels in both plots (10 across  $-0.005$  and  $0.005$ ). Arrows indicate cross-sectional velocities (same scale used for both,  $\max = 0.009$ ). Note that the slices are taken at the same point in the pipe but there may be a phase difference between the solutions as they are calculated using two different calculations. Also compare the slice for  $(7, 64, 1)$  with that for  $(23, 64, 11)$  given in figure 4(a).

of nonlinear input alone. This observation is crucial for interpreting figure 3, which shows that after the Orr mechanism has played out, the NLOP evolution is dominated by the non-normal energy growth of helical modes ( $m, l \neq 0$ ) in the second phase ( $0.4 \lesssim t \lesssim 2.5$ ). As these modes grow quickly, they feed energy via their nonlinear interactions into the streamwise independent modes as evidenced in the increase in  $E_{m0}^{uv}$  over the interval  $2 \lesssim t \lesssim 4$ . When the non-normal energy growth of the helical modes runs out of steam (at  $\approx 2.5$ ), they decline quickly through the combined effect of this nonlinear energy drain and viscous dissipation. Thereafter, the evolution is dominated by each streamwise independent mode experiencing slow but sustained non-normal growth as the secularly decaying streamwise rolls advect the mean shear to produce streaks – the well-known lift-up process. The uniform decay rates of the streamwise rolls indicates that there is minimal nonlinear energy mixing at this point, at least in the cross-plane velocities. This is because they are insensitive to axial advection (e.g.  $w \partial u / \partial z = 0$ ) and the cross-plane velocities are so small. Figure 2 shows that these two non-normal growth processes together with the Orr mechanism cooperate nonlinearly to produce a larger overall growth at  $T_{opt}$  than separately. After the initial rotation and unpacking by the Orr mechanism, the helical modes grow but quickly run out of steam.




FIGURE 3. (Colour online) Modal energies for the NLOP run with resolution (7, 64, 1).  $E$  is the total perturbation energy;  $E_{ml}^w$  indicates the energy in the  $w$  (streamwise) velocity component with azimuthal and axial wavenumbers  $m$  and  $l$  respectively, and  $E_{ml}^{uv}$  indicates the energy in the  $u$  and  $v$  (cross-stream) velocity components with azimuthal and axial wavenumbers  $m$  and  $l$  respectively. So  $E_{m0}^w$  is the (streamwise) streak energy,  $E_{m0}^{uv}$  is the (streamwise) roll energy and  $E_{m1}$  is a helical mode energy.

They then dump their energy into the third streamwise-rolls-driving-streaks process, which is subsequently boosted to reach higher growth factors than otherwise.

The fact that helical (or more generically ‘oblique’) waves grow best over short times and streamwise independent flows grow larger but over longer times is well known (e.g. figure 8 of Farrell & Ioannou 1993, figure 4 of Schmid & Henningson 1994 and figure 5 of Meseguer & Trefethen 2003). Furthermore, the scenario of oblique waves growing transiently, feeding their energy into streamwise rolls that then drive streamwise streaks (which then become unstable) has also been proposed before as an efficient bypass mechanism in Reddy *et al.* (1998) (called the ‘oblique wave scenario’). This general picture, or at least the first stages of it, appear to be confirmed here in the nonlinear growth problem (see also Duguet *et al.* 2010a). However, the initial localization of the perturbation and how it ‘unwraps’ to give a final, large, predominantly streamwise independent flow is a new feature born out of a need to cheat the starting (global) energy constraint. Figures 4 and 5 show how the structure of the NLOP across one (fixed) slice of the pipe evolves in time. The initial slice shown (in figure 4a and again in figure 4b but rescaled) has a peak cross-plane speed of  $\approx 0.02U$  concentrated in a tight vortex pair near the pipe wall and peak axial speed of  $\approx 0.012U$ . The delocalization or ‘unwrapping’ is effectively completed by  $t \approx 1$  when the peak cross-plane speed is essentially unchanged, whereas the peak axial speed has grown to  $0.06U$ . Figure 5 (a is a rescaled version of figure 4f) shows that both the cross-plane and axial speeds grow considerably in the interval  $1 \lesssim t \lesssim 2.5$  (peak cross-plane speed increases from  $0.02U$  to  $0.08U$  and peak axial speed from  $0.06U$  to  $0.13U$ ). In fact, by  $t = 2.5$  the initial energy has experienced most of its




FIGURE 4. (Colour online) The NLOP at  $Re = 1750$  and  $E_0 = 1.8 \times 10^{-5}$  calculated using resolution  $(MM, NN, LL) = (23, 64, 11)$  at  $(a, b)$   $t = 0$  (the same velocity field is shown but with different contour levels),  $(c)$   $t = 0.2$ ,  $(d)$   $t = 0.6$ ,  $(e)$   $t = 0.8$  and  $(f)$   $t = 1D/U$ . Contours indicate the streamwise velocity perturbation (total velocity with the laminar state of equivalent mass flux subtracted off) and the arrows indicate the cross-stream velocity at a fixed slice in the pipe. All plots except  $(a)$  have 10 contour levels between the extremes of the streamwise velocity perturbation at  $t = 1$  and cross-stream velocities similarly scaled. Plot  $(a)$  uses 10 contours between the extremes of the streamwise velocity perturbation at  $t = 0$  with arrows automatically scaled.

growth (a factor of  $\approx 50$ ) and only a further magnification by  $\approx 7$  follows in the next  $\approx 20D/U$ . In this latter period the cross-plane velocities manoeuvre the streak structure into place and then die away, so that even by  $t = 10$ , the predominantly streamwise independent and axial flow has been established (peak cross-plane speed is  $0.012U$  and peak axial speed  $0.34U$  now). It is worth stressing that even at this point the flow does not match the LOP (see plot  $c'$  in figure 2 of PK10), which depends solely on the Fourier–Fourier basis function  $\exp(i\phi)$  and is strictly two-dimensional, being streamwise invariant.

#### 4. Optimals for moderate times in a $\pi D/2$ pipe

As discussed in § 1, PK10 demonstrated how nonlinearity can qualitatively change the form of the optimal disturbance of a given initial energy that achieves the most energy growth over a fixed period. The new NLOP could not, however, be followed up to the initial energy level at which turbulence was triggered. We now revisit this situation armed with a more efficient and parallel code that allows use of higher resolution and more carefully refined steps in  $\mathbf{u}_0$ .

In PK10, a short  $\pi D/2$  periodic pipe was adopted to minimize the axial resolution needed and the relatively short time period was taken equal to  $T_{lin}$ , the time for maximum energy growth in the linearized Navier–Stokes equations, in order to




FIGURE 5. (Colour online) The NLOP at  $Re = 1750$  and  $E_0 = 1.8 \times 10^{-5}$  calculated using resolution  $(MM, NN, LL) = (23, 64, 11)$  at (a)  $t = 1$ , (b)  $t = 2.5$ , (c)  $t = 4$  and (d)  $t = 10, D/U$  and the representation is as in figure 4. All plots have 10 contour levels between the extremes of the streamwise velocity perturbation at  $t = 10$  and cross-stream velocities similarly scaled and set by those at  $t = 2.5$ . Note that plot (a) indicates exactly the same velocity field as figure 4(f) but using revised contour and arrow levels.

highlight the effect of nonlinearity. Working at  $Re = 1750$ , PK10 report failure to converge for  $E_0 > 2 \times 10^{-5}$ . Their best estimate for  $E_c$  was  $\mathcal{E}_c = 6 \times 10^{-5}$ , the energy required to trigger turbulence when using a perturbation of the form  $A\mathbf{u}_{3d}(\mathbf{x}; E_0 = 2 \times 10^{-5}, Re = 1750)$ . With the new code using a resolution  $(MM, LL) = (23, 11)$  and  $NN = 64$  finite difference radial points, as opposed to PK10's  $(MM, LL) = (14, 5)$  and a 25 Chebyshev polynomial expansion radially, we were able to confirm PK10's results for  $E_0 \leq 2 \times 10^{-5}$  as well as continuing convergence up to  $E_0 = 2.52 \times 10^{-5}$ . Above this point, the amount of growth grows sharply compared with the amount of growth produced by rescaling arguments (figure 6).

Examining how the residual  $\langle (\delta\mathcal{L}/\delta\mathbf{u}_0)^2 \rangle$  decreases as the algorithm proceeds indicates that the code has converged at  $E_0 = 2.52 \times 10^{-5}$ : see figure 7 (outer). The time evolution of the optimal solution is relatively smooth, relaminarizing after the initial transient growth. Surprisingly, however, one of the velocity fields it iterates through (marked with a black dot in figure 7, outer) does lead to a turbulent episode.




FIGURE 6. (Colour online) Reproduction of figure 5 from Pringle & Kerswell (2010) showing growth as a function of initial energy. The (red) crosses correspond to the nonlinear optimal perturbations previously calculated in PK10, while the uppermost (green) line represents the optimals calculated with the new code presented in this paper (PWK11). The middle (magenta) line shows the amount of growth produced by simply rescaling the nonlinear optimal for  $E_0 = 2 \times 10^{-5}$  and therefore lower-bounds the true (green) optimal growth curve. The lowest flat (blue) line is the growth provided by the nonlinearly modified two-dimensional optimal (LOP).

A comparison of the two evolutions confirms that the optimization procedure has worked properly: the optimal produces more growth than the other initial condition despite the fact it does not lead to turbulence (figure 7, inner). This observation seems to go against our assertion that the optimization algorithm will latch onto a turbulence-triggering state and then fail to converge. There are two important lessons to be learned from this apparent pathology. The first, most obvious one is that the turbulence-triggering initial condition has not had enough time to reach the turbulent state by the end of the (short) period  $T_{lin}$ , so

- (i)  $T_{opt}$  needs to be large enough.

Cherubini *et al.* (2010, 2011) work in this short  $T$  regime, which is why they are able to converge onto an optimal disturbance ‘above’ the edge (it triggers turbulence but much later than their target time  $T$ ). Secondly, figure 7 shows that the energy level reached by the optimal a little after  $T_{lin}$  is actually higher than that typically associated with the turbulent state at this (low)  $Re$  in this (tight) pipe geometry. This situation is fatal for the approach being advocated here, which relies on the turbulent state producing the highest values of the energy growth (or whatever functional is being considered) in comparison to non-turbulent states. By choosing  $T$  large enough to avoid initial transients, this requirement is essentially the condition that the attracting energy plateau on the edge is lower than turbulent energy levels. The recent success of edge tracking (e.g. Mellibovsky *et al.* 2009; Schneider & Eckhardt 2009; Duguet, Willis & Kerswell 2010b; Schneider, Marinc & Eckhardt 2010) indicates that this






FIGURE 7. (Colour online) Convergence of procedure at  $E_0 = 2.52 \times 10^{-5}$ . The algorithm smoothly converges to a growth of 2030 (upper right line), while the residual decays to  $O(10^{-8})$  (lower right line). The algorithm was continued for over 800 iterations in total to ensure that there was no further change. *Inset*. The time evolutions of two disturbances found by this sequence of iteration. The lower right (green) line is the converged optimal that smoothly relaminarizes, while the upper right (red) line is the disturbance corresponding to the black dot leading to a turbulent episode. The vertical dashed line shows the target time  $T_{opt} = T_{lin} \approx 21.3D/U$  from the optimization procedure.

is routinely achieved provided that the flow is not too tightly constrained (i.e. the geometry is not too small as it evidently is in a  $\pi D/2$  long pipe) or too close (in  $Re$ ) to the first appearance of the turbulent state. For example, Schneider & Eckhardt (2009, figure 7) show how the edge state and turbulent energies quickly separate with increasing  $Re$  for a  $5D$  pipe (see also figure 3 of Mellibovsky *et al.* 2009 for a  $50D$  pipe). Therefore, the second lesson is that

- (ii) the optimization strategy can only be used to calculate the minimal seed sufficiently far (in  $Re$ ) from the first appearance of turbulence and/or for flows in non-small domains.

In hindsight then, the geometry and  $Re$  value chosen in PK10 is not suited to determining  $E_c$  there using this optimization approach. We therefore switch to a longer  $5D$  periodic pipe (theoretically popular since the work of Eggels *et al.* 1994) and a higher  $Re = 2400$  where the edge shows typical behaviour (Duguet *et al.* 2008) and the turbulent state is clearly energetically separated from the edge state (Schneider & Eckhardt 2009).

## 5. Optimals for long times in a $5D$ pipe

In order to assess the twin conjectures discussed in § 1, a practical decision needs to be made as to what constitutes a ‘asymptotically long’ optimization time. Figure 7




FIGURE 8. (Colour online) The axial distribution of energy for the NLOP at  $E_0 = 7.058 \times 10^{-6}$  (dashed blue lines) and for the turbulent seed (solid red lines). In both cases it is clear that although the energy has a clear peak it never fully decays away. The amplitude in the cross-stream direction changes by less than an order of magnitude across the domain.

shows that an initial condition is capable of growing through several orders of magnitude into a turbulent episode within  $\sim 50D/U$ . We therefore chose  $T_{opt} = 75D/U$ , which should be large enough to capture this behaviour, especially in a larger  $5D$  domain although  $Re = 2400$  is higher (so  $T_{opt} > 2.5T_{lin}$  at this  $Re$ ). It is worth remarking, though, that this finite choice will limit the accuracy to which we can determine the energy threshold. The algorithm senses initial conditions that have reached the turbulent state by the end of the observational window. This sets a lower limit on how close they can be to the edge, as the time for a turbulence-triggering initial condition to reach turbulence becomes arbitrarily large as it is taken closer to the edge. This said, our choice of  $T_{opt}$  gives acceptable accuracy, yet the way to improve this is clearly by integrating for longer.

The results of the energy growth optimization in a  $5D$  pipe at  $Re = 2400$  as a function of  $E_0$  (note that the non-dimensionalization of energy is dependent on the size of the flow domain being considered, and so equivalent absolute energies will appear smaller after non-dimensionalizing in this longer pipe) exactly mimics the situation uncovered in PK10. For  $E_0$  small enough, the linear (streamwise independent) optimal (LOP) is selected, albeit with slight nonlinear modification, which suppresses the growth of the two-dimensional optimal as  $E_0$  increases. Then there is a finite value (PK10 refer to this as  $E_{3d}$ ) when a new three-dimensional optimal (NLOP) is preferentially selected, which shows localization in the azimuthal and radial directions (recall that the ‘NLOP’ refers to this optimal, which, not being smoothly connected to the LOP as  $E_0$  increases, has distinctly different structure). There is also some localization in the streamwise direction, but the domain is by no means long enough for us to observe truly localized optimals as opposed to spatially periodic disturbances (see figure 8).




FIGURE 9. (Colour online) The amount of growth produced by successive choices for  $\mathbf{u}_0$  in the iterative scheme. For  $E_0 = 7.058 \times 10^{-6}$  the growth plateaus out at 483 (lower red line). For  $E_0 = 7.124 \times 10^{-6}$ , the growth briefly plateaus at  $\sim 520$  before rapidly rising to excess of 1000 (upper blue line). The circle plotted shows a growth of 14480 and has reached the turbulent attractor before the algorithm steps momentarily back to a region below the edge. The turbulent seed and the optimal for  $E_0 = 7.058 \times 10^{-6}$  are both plotted in figure 19. Before this has been reached, however, the growth being produced is already non-smooth between steps due to the lack of smoothness in the hypersurface  $\mathcal{L}$ , as predicted by Conjecture 1. We also include the choice of  $E_0 = 7.077 \times 10^{-6}$  for illustrative purposes (middle green line). This iterative run has not been fully converged and it is not clear whether it will converge to an optimal or depart to the turbulent state.

As  $E_0$  is increased further, there comes a point at which the algorithm struggles to converge properly. Successive bisection indicates that this value,  $E_{fail}$ , is bracketed by the initial energy values of  $E_0 = 7.058 \times 10^{-6}$ , which converges smoothly to the NLOP, and  $7.124 \times 10^{-6}$ , which clearly fails due to the occurrence of a turbulence-triggering initial condition: see figures 9–11.

An attempt to improve this bracketing by taking  $E_0 = 7.077 \times 10^{-6}$  appears to show convergence, yet there still remains some doubt even after running the algorithm for nearly 1600 iterations and 50 000 CPU hours ( $\approx 6$  years). Figure 9 indicates convergence yet at a much higher level compared to that reached by the ‘nearby’ initial energy of  $E_0 = 7.058 \times 10^{-6}$ . Moreover the fact that there is a jump up to this higher level after  $\approx 1000$  iterations is mildly disconcerting. This adjustment is reflected in the evolution of the residual (see figure 10) which, after 200 iterations, appears to show convergence for the next 500 iterations before being followed by a rapid transition that ends after 1150 iterations. It is not clear whether the algorithm has now finally converged or whether it will subsequently encounter a turbulence-inducing initial condition. This example demonstrates that it is clearly very important to take care when deciding whether the procedure has converged (e.g. stopping the algorithm after 600 iterations would appear to indicate clear convergence). It is probable that the algorithm is struggling to distinguish between turbulence-inducing initial conditions




FIGURE 10. (Colour online) The residuals corresponding to the three iterative runs described in figure 9: flat (red) line,  $7.058 \times 10^{-6}$ ; (green) line with spike at 1100,  $7.077 \times 10^{-6}$ ; upper (blue) line,  $7.124 \times 10^{-6}$ . The sudden adjustment centred on 1100 iterations is a warning that deciding upon convergence can be a subtle affair.




FIGURE 11. (Colour online) Evolution of the final states produced by the iterative scheme for  $E_0 = 7.058 \times 10^{-6}$  (lowest red line) and  $7.124 \times 10^{-6}$  (uppermost blue line). One clearly leads to a turbulent episode while the other simply relaminarizes after the initial transient growth. Also shown is the evolution of the initial condition arrived at after 1600 iterations for  $E_0 = 7.077 \times 10^{-6}$  (middle green line) which again relaminarizes. The vertical line marks the optimization time.




FIGURE 12. (Colour online) Snapshots showing isocontours of streamwise perturbation velocity during the evolution of the final states produced by the iterative scheme for (a)  $E_0 = 7.058 \times 10^{-6}$  and (b)  $7.124 \times 10^{-6}$ . The isocontours in each plot 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 snapshots correspond to times  $t = 0, 0.5, 5, 10, 20, 40$  and  $75D/U$ . In both cases the energy is initially localized in the streamwise direction and the disturbance quickly spreads. By  $t = 10$  both disturbances have created streamwise streaks, but only for the lower energy do they become streamwise independent. The larger amplitudes of the higher energy streaks are subject to a turbulence-triggering instability.

and the NLOP because the time to reach turbulence is comparable to, or exceeds,  $T_{opt}$ . Consequently, the estimate that  $7.058 \times 10^{-6} < E_{fail} < 7.124 \times 10^{-6}$  is the best we can hope for working with  $T_{opt} = 75D/U$  and the fate of  $E_0 = 7.077 \times 10^{-6}$  could be decided by taking a longer  $T_{opt}$  (not pursued here).

The physical evolution of the two disturbances is shown in figure 12. Initially the disturbances look streamwise localized because of the contouring, but they do in fact occupy the full length of the domain. Both subsequently develop into coherent domain-length streaks. In the  $E_0 = 7.058 \times 10^{-6}$  case, these streaks continue to evolve yet remain stable, becoming almost totally streamwise invariant before ultimately decaying. In the  $E_0 = 7.124 \times 10^{-6}$  case, the streaks have higher amplitude and a streak instability clearly occurs, leading to turbulence. The nature of this instability is shown in figure 13, which plots the streamwise dependent and independent components




FIGURE 13. (Colour online) The energy associated with the  $m = 1$  axial component of the disturbances calculated for  $E_0 = 7.058 \times 10^{-6}$  (solid lines) and  $E_0 = 7.124 \times 10^{-6}$  (dashed lines). Each energy is split into streamwise independent (dark/red uppermost lines at  $t = 25D/U$ ) and streamwise dependent (light/green lowermost lines at  $t = 25D/U$ ) parts. The former measures the streaks created by the disturbance while the latter shows the instability of these streaks.

of the energy of the axial velocity field with (lead) azimuthal wavenumber  $m = 1$ . For the relaminarizing disturbance, the three-dimensional part of the energy decays monotonically from around  $15D/U$  onwards. For the more energetic disturbance the decay is abated after  $20D/U$ , at which point an instability of the streaks occurs, eventually leading to turbulence. It is worth remarking that the final plot for  $E_0 = 7.058 \times 10^{-6}$  in figure 12 resembles more the two streaks produced by the linear optimal, rather than the three-streak field generated in PK10 (cf. figure 5d in this paper). Whether this is due to the increased target time or the lengthened flow domain is unclear.

Conjecture 1 claims  $E_c = E_{fail}$ . As there is no evidence of turbulence-triggering initial conditions at  $E_0 = 7.058 \times 10^{-6}$  but there is at  $7.124 \times 10^{-6}$ , we have  $7.058 \times 10^{-6} < E_c < 7.124 \times 10^{-6}$ . To this level of accuracy we have found that  $E_c = E_{fail}$ . That the optimization scheme will fail if turbulent seeds exist within the  $E_0$ -hypersurface seems clear, provided the iterative scheme can find them. Establishing this is very difficult if not impossible, but a weaker practical alternative is to demonstrate that the procedure is not dependent on the initial starting guess  $\mathbf{u}_0$ . To do this we have compared six very different choices for the initial seed for both  $E_0 < E_{fail}$  and  $E_0 > E_{fail}$ , and plotted their evolution between iterations on a two-dimensional projection of energy in the axisymmetric part of the perturbation against energy in the streamwise independent part (figures 14 and 15). The scatter of the initial crosses illustrates the variety of initial conditions used, which range from turbulent velocity fields to known travelling wave solutions. In both cases, irrespective of where the scheme begins, the eventual (iterative) evolution brings it to the same trajectory in






FIGURE 14. (Colour online) Convergence of six different initial conditions (crosses) towards the same optimal (circle) for  $E_0 = 7.058 \times 10^{-6}$ : axes are normalized perturbation energies associated with the streamwise independent part (abscissa) and the axisymmetric part (ordinate). The initial conditions chosen correspond to the various combinations of turbulent flow fields, travelling wave solutions and the nonlinear optimal from § 4. The red (thick) line corresponds to this final choice and is the iterative scheme shown in figures 9 and 10. Arrows are included to indicate the direction of convergence.

this ‘phase space’. This provides some evidence that the algorithm does sample the  $E_0$ -hypersurface well and that Conjecture 1 indeed holds true.

In order to assess Conjecture 2, we now consider the behaviour of the optimal solution close to the edge. Figures 14 and 15 already provide some evidence that the NLOP at  $7.058 \times 10^{-6}$  and the turbulent seed found at  $7.124 \times 10^{-6}$  are similar at least in terms of their axisymmetric and streamwise independent energy fractions. In order to probe the accuracy of Conjecture 2 further, we look at the one initial condition identified for  $E_0 = 7.124 \times 10^{-6}$  by our algorithm in § 5 that was a turbulent seed,  $\mathbf{u}_s$  (indicated by the circle in figure 9). The fact that only the one condition was found above the edge of chaos suggests that in this region there is only a very small set of turbulence-triggering initial conditions. We attempt to quantify this by considering the evolution of initial conditions of the form

$$\mathbf{u}_{ic}(A, d) := A[(1 + d)\mathbf{u}_{l=0} + (1 - d)\mathbf{u}_{l \neq 0}], \quad (5.1)$$

where  $\mathbf{u}_{l=0}$  and  $\mathbf{u}_{l \neq 0}$  are the streamwise independent and dependent parts of  $\mathbf{u}_s$  and  $A$  is adjusted to give the required value of  $E_0$ . The amount of growth after  $75D/U$  is shown as a function of  $E_0$  and  $d$  in figure 16. The jump in growth from  $O(10^3)$  to  $O(10^4)$  clearly demarcates where the edge of chaos is crossed. The narrowness of the peak for  $E_0 = 7.124 \times 10^{-6}$ , and the observation that a mere  $\sim 0.3\%$  reduction in amplitude is enough to dip beneath the edge, indicate that  $\mathbf{u}_s$  is very close to a local minimum of the edge. The precise energy at which  $\mathbf{u}_{ic}(A, d = 0)$  crosses the edge is plotted in figure 17. Here two bracketing cases are shown:  $E_0 = 7.121011 \times 10^{-6}$ , which ultimately relaminarizes, and  $E_0 = 7.121019 \times 10^{-6}$ , which leads to turbulence.




FIGURE 15. (Colour online) The result of using the same six initial conditions (crosses) for  $E_0 = 7.124 \times 10^{-6}$ . The red line corresponds to the iterative progression shown in figures 9 and 10. Clearly the procedure is independent of the starting guess. Note also how similar the progression is here to that for  $E_0 = 7.058 \times 10^{-6}$ . The red (thick) line indicates the initial path taken towards the turbulent seed (see figure 9). The path after entering the chaotic region is not included.




FIGURE 16. (Colour online) Growth factors in the neighbourhood of the seed state  $\mathbf{u}_i$  that triggers turbulence at  $E_0 = 7.124 \times 10^{-6}$ . The neighbourhood is defined by  $d$ , where  $\mathbf{u}_{ic}(\mathbf{x}; d) := A[(1 + d)\mathbf{u}_{l=0} + (1 - d)\mathbf{u}_{l \neq 0}]$ ,  $\mathbf{u}_{l=0}$  is the streamwise independent part of  $\mathbf{u}_s$  and  $\mathbf{u}_{l \neq 0} := \mathbf{u}_s - \mathbf{u}_{l=0}$ . The amplitude  $A$  is used to rescale the initial state to ensure the correct starting energy  $E_0$ . The figure shows that  $\mathbf{u}_s$  only just triggers transition at  $7.124 \times 10^{-6}$ , but for higher energies there exists an ever-increasing neighbourhood of initial conditions surrounding  $\mathbf{u}_s$  that trigger transition (indicated by the jump in  $G$ ).




FIGURE 17. (Colour online) Trajectories close to the laminar–turbulent boundary, or ‘edge’. Nearby initial conditions are  $\mathbf{u}_{ic} = A\mathbf{u}_s$ , with  $A$  selected to give the indicated energies. The case  $E_0 = 7.121011 \times 10^{-6}$  relaminarizes, but not before passing close by the edge attractor (marked by the dot, which is the energy of the  $S1$  travelling wave embedded in it). The case  $E_0 = 7.121019 \times 10^{-6}$  tracks the edge before leading to turbulence. The vertical line is  $T = T_{opt} = 75D/U$ .

The closeness of these energies means that both evolutions track the edge to  $T \approx T_{opt}$  before going their separate ways. This emphasizes that to improve the estimate of  $E_c$  discussed above,  $T_{opt}$  has to be increased.

Coincidentally, while following the relaminarizing case, the flow was found to transiently resemble the asymmetric travelling wave believed to be embedded in the edge state (Pringle & Kerswell 2007, later named  $S1$  in Pringle *et al.* 2009) at  $t \approx 100D/U$ . This was verified by calculating the two correlation functions,  $I_{tot}$  and  $(I_{tot} + I_{uv})/2$ , introduced in Kerswell & Tutty (2007, definitions (2.3) and (2.5)). Figure 18 shows that both these correlations simultaneously exceed 0.75 at  $t \approx 100D/U$ , clearly indicating a very close ‘visit’ (0.6 was deemed good enough to indicate a ‘close’ visit by Kerswell & Tutty 2007). The fact that this visit takes place is not a surprise but more a check of consistency: the edge state is believed to be unique and therefore a global attractor on the edge at this  $Re$  and pipe length (Schneider *et al.* 2007). It is worthy of note, however, that it takes a comparatively long time of  $\approx 100D/U$  for a flow trajectory starting at the lowest energy point on the edge to reach the  $S1$  state.

Examining the NLOP for  $E_0 = 7.058 \times 10^{-6}$  and  $\mathbf{u}_s$ , it seems reasonable to conclude that the minimal seed of turbulence is ‘sandwiched’ in between. The form of the two solutions is shown in figure 19 (along with the last iterate calculated at  $7.077 \times 10^{-6}$ ) and it is clear that they do not alter much as the edge is crossed. This supports Conjecture 2.

## 6. Discussion

We first summarize what has been done in this paper. An exploratory nonlinear energy growth calculation in PK10 showed that the form of the optimal initial




FIGURE 18. (Colour online) Correlation function data that measure how close the instantaneous velocity field is to the asymmetric travelling wave  $S1$  for the trajectory that relaminarizes in figure 17 with  $E_0 = 7.121011 \times 10^{-6}$ . The fact that  $I_{tot}$  and  $(I_{tot} + I_{uv})/2$  exceed 0.75 at  $t = 100D/U$  indicates a very close visit (Kerswell & Tutty 2007). The importance of this visit is that  $S1$  is believed to be embedded in the chaotic edge state for the pipe length and  $Re$ .




FIGURE 19. (Colour online) The nonlinear optimal calculated in § 5 for  $E_0 = 7.058 \times 10^{-6}$  (a) and the turbulent seed  $\mathbf{u}_s$  (c). The close similarity between the two solutions is striking, with the turbulent seed having stronger streaks but an otherwise comparable structure. For comparison we have also included the final state found for  $E_0 = 7.077 \times 10^{-6}$  (b). This state appears to be an intermediary between the two other states. (Note that no effort has been made to match the phases along the pipe length.)

disturbance changes suddenly at a small (pre-threshold) but finite initial energy level  $E_{3d}$  from a global linear optimal (weakly modified by nonlinearity) to a localized strongly nonlinear optimal. This has been confirmed at higher spatial and temporal resolution. The physical processes responsible for the enhanced energy growth of the new nonlinear optimal (NLOP) have been identified as three known linear growth mechanisms – the Orr mechanism, oblique wave transient growth and the lift-up effect – acting sequentially and coupled together via the nonlinearity of the Navier–Stokes

equations. These mechanisms operate on differing time scales, yet appear able to pass on their growth to the next (slower) process so that the total growth outweighs any of their individual contributions. The NLOP is also localized, which is an inherently nonlinear feature designed to cheat the global initial energy constraint. The calculations here have also managed to converge this nonlinear optimal state beyond the threshold energy level at which turbulence could be triggered. This highlighted two issues for the optimization strategy to identify this energy threshold: (i) the optimization time needs to be large enough for turbulence-triggering initial conditions to have time to reach the turbulent state; and (ii) the energy levels of the turbulent state need to be above those for laminar flows so that the optimization algorithm will naturally seek them out. This indicates that the optimization strategy discussed here is better suited to larger flow domains and more supercritical (higher  $Re$ ) regimes. Ironically, it is now clear that doing exploratory (cheap) calculations in a small domain at low  $Re$  was a natural but bad choice in PK10.

Calculations in a longer pipe  $5D$  at higher  $Re$  ( $=2400$ ) over a larger time period found the same general situation as in the  $\pi D/2$  pipe of PK10. At low initial energies, the optimal is the (global) linear optimal weakly modified by nonlinearity. At a certain small but finite initial energy  $E_{3d}$ , a new localized three-dimensional optimal (NLOP) is preferred, which stays the optimal until the algorithm fails to converge at  $E_{fail}$ . The only significant difference is that in the longer  $5D$  pipe the NLOP is starting to streamwise-localize (figure 8), in contrast to the  $\pi D/2$  NLOP, where the shortness of the domain prevents this. Above  $E_{fail}$ , initial conditions that trigger turbulence are found to exist on the energy hypersurface and to the accuracy available,  $E_{fail} = E_c$ . This supports Conjecture 1, which presupposes that the optimization algorithm will find any turbulent-triggering states if they exist on the energy hypersurface and then fail to converge as a result. As way of confirming this, the algorithm was tested with a variety of very different starting conditions with the same optimal emerging, indicating that the optimization algorithm is able to explore the energy hypersurface.

Intriguingly, good evidence was also found that  $NLOP \rightarrow$  minimal seed as  $E_0 \rightarrow E_c^-$  in support of the stronger Conjecture 2, at least for this flow, geometry and  $Re$ . Pictorially, this means that the NLOP for  $E_c$  and the minimal seed actually coincide in figure 20, rather than the more general situation shown where the two differ. It was also argued that, with enough computational power, the threshold energy  $E_c$  and the minimal seed could be calculated to arbitrary accuracy by increasing the spatial and temporal resolution as well as  $T_{opt}$ , which improves the algorithm's ability to distinguish between trajectories that become turbulent and those that relaminarize.

The strategy advocated here for determining the minimal finite amplitude disturbance to trigger transition to turbulence in shear flows involves construction and iterative solution of a variational problem. The objective functional must be selected such that it identifies turbulent velocity fields by taking enhanced values compared to those for laminar fields. This is then maximized via searching over all incompressible disturbances of fixed amplitude that respect the boundary conditions over an asymptotically long time period constrained by the full Navier–Stokes equations. All of the results discussed here were obtained using the perturbation energy growth over a given period as the key functional. This certainly takes enhanced values for turbulent velocity fields in the  $5D$  pipe at  $Re = 2400$ , and Conjecture 1 seems to hold true. However, other choices should also work equally well, e.g. the total dissipation (Monokrousos *et al.* 2011), provided that they share this crucial property. If Conjecture 1 is indeed true, then the optimization strategy discussed here will identify the threshold energy level  $E_c$ . However, more is available too, albeit




FIGURE 20. Sketch to illustrate Conjectures 1 and 2. Conjecture 1 asserts that the optimization approach will yield a well-defined nonlinear optimal (NLOP) up to  $E_0 = E_c$ , whereupon turbulence can be triggered by the minimal seed. Conjecture 2 asserts that the NLOP converges as  $E_0 \rightarrow E_c^-$  to the minimal seed (the converse situation where the limiting NLOP state and the minimal seed are different is shown for clarity). Note that the LOP is streamwise independent in pipe flow and hence traces the  $E_{l=0}$  axis.

indirectly, as the minimal seed should be the unique initial condition that triggers turbulence as  $E_0 \rightarrow E_c^+$ . Given this, the status of Conjecture 2, although conceptually fascinating, seems less important practically. Whether or not Conjecture 2 holds for energy growth (and, admittedly, we only have one supportive data analysis here), it can of course be restated for any functional. Then the question becomes: Is there a universal functional that, when optimized, always identifies the minimal seed as  $E \rightarrow E_c^-$  for a class of flows (e.g. wall-bounded shear flows)? This seems unlikely to be exactly true but may nevertheless be approximately true for some subset of functionals. Then any of these could give acceptable predictions depending on how the results are to be used subsequently (e.g. designing disturbances in the laboratory). Certainly this would seem to be the case using the energy growth functional given the comparison in figure 19.

The variational approach espoused here is, of course, incredibly flexible. Changing the key functional is straightforward, as is the initial (norm) constraint on the competitor initial fields. Although the discussion above has concentrated on the initial perturbation energy  $E_0$ , it should be clear that other norm choices can be made. Provided that the functional under consideration jumps to large values for turbulent flows, the optimization algorithm should converge up to the first point (as the norm hypersurface ‘expands’ away from the laminar state) at which the edge penetrates the



hypersurface (as per Conjecture 1). Furthermore, the turbulent state does not have to be the only target of the approach. Identifying the peak instantaneous pressure in a transitional flow is a key concern for pipeline structural integrity. One could easily imagine formulating an optimization problem to maximize the pressure after time  $T_{opt}$  over all disturbances of initial energy  $E_0$  where  $T_{opt}$  is also part of the optimization procedure.

The long-term objective of this theoretical work is to design better (lower-energy) ways to trigger turbulence with a view to informing control techniques. Further calculations clearly need to be carried out in more realistic geometries to see, for example, if universal localized minimal seeds emerge. Even now, though, this work is in a position to stimulate new experiments. The NLOP identified here indicates that structures that initially point into the shear will outgrow the equivalent structure directed across the shear (the Orr mechanism). This suggests a modification of the recent experiments of Peixinho & Mullin (2007) that were designed to generate oblique rolls by blowing and sucking directly across the shear. The calculations performed here indicate that their threshold scaling exponent (a non-trivial flux  $\sim Re^{-1.5}$ ) for transition may possibly be further reduced if the blowing and sucking is inclined upstream to take advantage of the Orr mechanism.

## 7. Glossary

$E_0$	Initial energy of a perturbation.
$E_T$	Final energy of a perturbation after time $T$ .
$E_{3d}$	Initial energy at which the NLOP first emerges as the new optimal.
$E_{fail}$	The minimum energy for which the optimization routine fails to converge.
$E_c$	The critical energy corresponding the minimum energy of the edge.
$\mathcal{E}_c(\mathbf{u})$	The minimum energy of a perturbation of the form $A\mathbf{u}$ required to trigger turbulence.
$T_{opt}$	The target time in an optimization procedure.
$T_{lin}$	The time for which transient growth is maximized in the linear problem.
$T_{turb}$	Typical time period required for the onset of turbulence.
$\mathbf{u}_{2D}(\mathbf{x}; Re, E, L, T)$	The two-dimensional optimal for the Reynolds number, energy, domain length and optimization time $T$ specified.
$\mathbf{u}_{3D}(\mathbf{x}; Re, E, L, T)$	The three-dimensional optimal for the Reynolds number, energy, domain length and optimization time $T$ specified.

## Acknowledgements

The authors acknowledge insightful conversations with B. Farrell, P. Ioannou and D. Henningson. The calculations in this paper were carried out at the Advanced Computing Research Centre, University of Bristol.

## REFERENCES

- ANDERSSON, P., BERGGREN, M. & HENNINGSON, D. S. 1999 Optimal disturbances and bypass transition in boundary layers. *Phys. Fluids* **11**, 134–150.
- BUTLER, K. M. & FARRELL, B. F. 1992 Three-dimensional optimal perturbations in viscous shear flow. *Phys. Fluids* **4**, 1637–1650.

- CHERUBINI, S., DE PALMA, P., ROBINET, J.-C. & BOTTARO, A. 2010 Rapid path to transition via nonlinear localized optimal perturbations in a boundary-layer flow. *Phys. Rev. E* **82**, 066302.
- CHERUBINI, S., DE PALMA, P., ROBINET, J.-C. & BOTTARO, A. 2011 The minimal seed of turbulent transition in the boundary layer. *J. Fluid Mech.* **689**, 221–253.
- CORBETT, P. & BOTTARO, A. 2000 Optimal perturbations for boundary layers subject to streamwise pressure gradient. *Phys. Fluids* **12**, 120–130.
- DUGUET, Y., BRANDT, L. & LARSSON, B. R. J. 2010a Towards minimal perturbations in transitional plane Couette flow. *Phys. Rev. E* **82**, 026316.
- DUGUET, Y., WILLIS, A. P. & KERSWELL, R. R. 2010b Slug genesis in cylindrical pipe flow. *J. Fluid Mech.* **663**, 180–208.
- DUGUET, Y., WILLIS, A. P. & KERSWELL, R. R. 2008 Transition in pipe flow: the saddle structure on the boundary of turbulence. *J. Fluid Mech.* **613**, 255–274.
- EGGELS, J. G. M., UNGER, F., WEISS, M. H., WESTERWEE, J., ADRIAN, R. J., FRIEDRICH, R. & NIEUWSTADT, F. T. M. 1994 Fully developed turbulence pipe flow: a comparison between direct numerical simulation and experiment. *J. Fluid Mech.* **268**, 175–209.
- FARRELL, B. F. & IOANNOU, P. J. 1993 Optimal excitation of three-dimensional perturbations in viscous constant shear flow. *Phys. Fluids* **5**, 1390–1400.
- GUSTAVSSON, L. H. 1991 Energy growth of three-dimensional disturbances in plane Poiseuille flow. *J. Fluid Mech.* **224**, 241–260.
- ITANO, T. & TOH, S. 2001 The dynamics of bursting process in wall turbulence. *J. Phys. Soc. Japan* **70**, 703–716.
- KERSWELL, R. R. & TUTTY, O. R. 2007 Recurrence of travelling waves in transitional pipe flow. *J. Fluid Mech.* **584**, 69–102.
- LUCHINI, P. 2000 Reynolds-number-independent instability of the boundary layer over a flat surface: optimal perturbations. *J. Fluid Mech.* **404**, 289–309.
- LUCHINI, P. & BOTTARO, A. 1998 Görtler vortices: a backward-in-time approach to the receptivity problem. *J. Fluid Mech.* **363**, 1–23.
- MELLIBOVSKY, F., MESEGUER, A., SCHNEIDER, T. M. & ECKHARDT, B. 2009 Transition in localized pipe flow turbulence. *Phys. Rev. Lett.* **103**, 054502.
- MESEGUER, A. & TREFETHEN, L. N. 2003 Linearized pipe flow to Reynolds number  $10^7$ . *J. Comput. Phys.* **186**, 178–197.
- MONOKROUSOS, A., BOTTARO, A., BRANDT, L., DI VITA, A. & HENNINGSON, D. S. 2011 Nonequilibrium thermodynamics and the optimal path to turbulence in shear flows. *Phys. Rev. Lett.* **106**, 134502.
- ORR, W. M. F. 1907 The stability or instability of the steady motions of a perfect liquid and of a viscous liquid. Part I. A perfect liquid. Part II. A viscous liquid. *Proc. R. Irish Acad. A* **27**, 9–138.
- PEIXINHO, J. & MULLIN, T. 2007 Finite-amplitude thresholds for transition in pipe flow. *J. Fluid Mech.* **582**, 169–178.
- PRINGLE, C. C. T., DUGUET, Y. & KERSWELL, R. R. 2009 Highly symmetric travelling waves in pipe flow. *Phil. Trans. R. Soc. A* **367**, 457–472.
- PRINGLE, C. C. T. & KERSWELL, R. R. 2007 Asymmetric, helical and mirror-symmetric travelling waves in pipe flow. *Phys. Rev. Lett.* **99**, 074502.
- 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 (referred to as PK10 in the text).
- REDDY, S. C. & HENNINGSON, D. S. 1993 Energy growth in viscous channel flows. *J. Fluid Mech.* **252**, 209–238.
- REDDY, S. C., SCHMID, P. J., BAGGETT, J. S. & HENNINGSON, D. S. 1998 On the stability of streamwise streaks and transition thresholds in plane channel flows. *J. Fluid Mech.* **365**, 269–303.
- SCHMID, P. J. & HENNINGSON, D. S. 1994 Optimal energy growth in Hagen–Poiseuille flow. *J. Fluid Mech.* **277**, 197.
- SCHNEIDER, T. M. & ECKHARDT, B. 2009 Edge states intermediate between laminar and turbulent dynamics in pipe flow. *Phil. Trans. R. Soc. A* **367**, 577–587.

- SCHNEIDER, T. M., ECKHARDT, B. & YORKE, J. A. 2007 Turbulence transition and the edge of chaos in pipe flow. *Phys. Rev. Lett.* **99**, 034502.
- SCHNEIDER, T. M., MARINC, D. & ECKHARDT, B. 2010 Localized edge states nucleate turbulence in extended plane Couette cells. *J. Fluid Mech.* **646**, 441–451.
- SKUFCA, J. D., YORKE, J. A. & ECKHARDT, B. 2006 Edge of chaos in a parallel shear flow. *Phys. Rev. Lett.* **96**, 174101.
- TREFETHEN, L. N., TREFETHEN, A. E. & REDDY, S. C. 1993 Hydrodynamic stability without eigenvalues. *Science* **261**, 578–584.
- VISWANATH, D. & CVITANOVIC, P. 2009 Stable manifolds and the transition to turbulence in pipe flow. *J. Fluid Mech.* **617**, 215–233.
- WILLIS, A. P. & KERSWELL, R. R. 2009 Turbulent dynamics of pipe flow captured in a reduced model: puff relaminarization and localized ‘edge’ states. *J. Fluid Mech.* **619**, 213–233.
- ZUCCHER, S., LUCHINI, P. & BOTTARO, A. 2004 Algebraic growth in a Blasius boundary layer: optimal and robust control by mean suction in the nonlinear regime. *J. Fluid Mech.* **513**, 135–160.