

Fully localised nonlinear energy growth optimals in pipe flow

Chris C. T. Pringle, Ashley P. Willis, and Rich R. Kerswell

Citation: [Physics of Fluids](#) **27**, 064102 (2015); doi: 10.1063/1.4922183

View online: <http://dx.doi.org/10.1063/1.4922183>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/pof2/27/6?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Turbulent-laminar patterns in plane Poiseuille flow](#)

Phys. Fluids **26**, 114103 (2014); 10.1063/1.4900874

[Laminar and turbulent nozzle-jet flows and their acoustic near-field](#)

Phys. Fluids **26**, 086103 (2014); 10.1063/1.4890493

[Minimal transition thresholds in plane Couette flow](#)

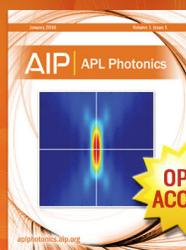
Phys. Fluids **25**, 084103 (2013); 10.1063/1.4817328

[A self-adjusting flow dependent formulation for the classical Smagorinsky model coefficient](#)

Phys. Fluids **25**, 055102 (2013); 10.1063/1.4804393

[Pipe-flow relaminarization by temporal acceleration](#)

Phys. Fluids **11**, 3478 (1999); 10.1063/1.870205



Launching in 2016!

The future of applied photonics research is here

AIP | APL
Photonics

Fully localised nonlinear energy growth optimals in pipe flow

Chris C. T. Pringle,¹ Ashley P. Willis,² and Rich R. Kerswell³

¹*Applied Mathematics Research Centre, Department of Mathematics and Physics, Coventry University, Coventry, United Kingdom*

²*School of Mathematics and Statistics, University of Sheffield, Sheffield S3 7RH, United Kingdom*

³*Department of Mathematics, University of Bristol, Bristol BS8 1TW, United Kingdom*

(Received 3 March 2015; accepted 15 May 2015; published online 5 June 2015)

A new, fully localised, energy growth optimal is found over large times and in long pipe domains at a given mass flow rate. This optimal emerges at a threshold disturbance energy below which a nonlinear version of the known (streamwise-independent) linear optimal [P. J. Schmid and D. S. Henningson, “Optimal energy density growth in Hagen-Poiseuille flow,” *J. Fluid Mech.* **277**, 192–225 (1994)] is selected and appears to remain the optimal up until the critical energy at which transition is triggered. The form of this optimal is similar to that found in short pipes [Pringle *et al.*, “Minimal seeds for shear flow turbulence: Using nonlinear transient growth to touch the edge of chaos,” *J. Fluid Mech.* **702**, 415–443 (2012)], but now with full localisation in the streamwise direction. This fully localised optimal perturbation represents the best approximation yet of the *minimal seed* (the smallest perturbation which is arbitrarily close to states capable of triggering a turbulent episode) for “real” (laboratory) pipe flows. Dependence of the optimal with respect to several parameters has been computed and establishes that the structure is robust. © 2015 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4922183>]

I. INTRODUCTION

In wall-bounded shear flows such as pipe flow, transition to turbulence remains a problem of great theoretical and practical importance. The transition is typically abrupt but occurs at flow rates for which the underlying base flow is stable. This implies that any disturbance to the laminar state can lead to either relaminarisation or a turbulent episode. In fact, there is also a third possibility. On the cusp between the two former outcomes are perturbations which neither go off to turbulence nor relaminarise and instead remain at an intermediate threshold. This is known as the laminar-turbulent boundary or “edge of chaos” and represents a hypersurface in phase space, separating those initial conditions which relaminarise from those that seed turbulence.

For any given initial disturbance, the end state is expected to depend on the amplitude of the disturbance, with smaller ones more likely to decay away. This immediately raises the question as to what the critical amplitude is, that is required to trigger turbulence. The amplitude will depend on the precise form of the disturbance imposed but is expected to be much smaller than the ensuing turbulent state.^{1,2} A more precise question asks—what is the smallest perturbation of any shape capable of triggering turbulence at a given flow rate? Such a perturbation would be expected to undergo significant growth and so early attempts to answer this centered on exploiting the non-normality of the linearised Navier-Stokes equations to find those perturbations that undergo the most growth.^{3–13} This work identified a number of processes by which a disturbance can grow in energy despite the flow being linearly stable before it has to eventually decay.¹⁴ Ultimately, however, the linear nature of this work means that these results cannot explain the full transition.

Reference 15 presented the argument that the *minimal seed* of turbulence—the smallest amplitude perturbation that is arbitrarily close to states which lead to turbulence—is the perturbation within the edge that has the most energy growth after a large time (“large” meaning large enough for the flow

to reach the attractor in the edge). This is a fully nonlinear calculation, and recently, the approach of finding optimal perturbations that maximise growth over a finite time has been extended to include the necessary nonlinearity.^{16–18} Computationally, this is a very intensive procedure and so far only small computational domains^{15,16,18,19} or short integration times^{17,20,21} have been used to demonstrate feasibility of the approach. Nonetheless, it serves as a basis for a proposed procedure for finding the minimal seed.^{15,22}

Constraining the choice of initial conditions for the transient growth calculation to within the edge is impractical. Instead, a two stage process has been proposed¹⁵ in which the maximum energy growth across a given time horizon is sought over all disturbances of a given initial energy (the “nonlinear transient growth problem”). The initial energy is then slowly increased and the calculation rerun until a rapid increase in the energy growth is identified. For very large target times, this growth increase tends to a discontinuous jump, and the first optimal initial condition (as the initial energy increases) to achieve heightened growth, and thereby be outside the laminar state’s basin of attraction, is then an estimate of the minimal seed (this estimate converges to the minimal seed as the target time becomes infinite: see the review in Ref. 22). For more moderate optimisation times, the jump in energy growth is smoothed, and there is a window in initial energy where a new nonlinear optimal is more efficient than the linear problem’s optimal, but where transition cannot be triggered. Provided the time is still larger than the transition time, the amplitude of the minimal seed is also the maximum amplitude of initial perturbation for which convergence of the optimisation algorithm is possible. If much shorter times are chosen, it is possible to converge at energies well above this critical amplitude,²¹ but these optimals are not necessarily expected to be related to the minimal seed.

When sufficiently long target times are used, short-pipe-domain calculations^{15,16} suggest that the minimal seed should be a fully localised disturbance and therefore of immediate interest to experimentalists. In the present work, we perform the above procedure within experimentally relevant, long pipe domains (up to 10 times longer than in Ref. 15) to (a) try to confirm this and (b) to see if any new minimal seed emerges when the optimal is allowed to fully localise in the streamwise direction. We focus on the case of maximising energy growth (the ratio of energy at target time T to initial energy, $G := E(T)/E(0)$) for an intermediate choice of time. In the 25 diameters ($25D$) long pipe considered mostly here, the computational demands of this are already heavy. Nonetheless, the results are shown to be independent of the domains size taken once this length has been reached and give a good predictor of the true minimal seed.

The paper is split into five further sections: Sec. II formulates the problem, Sec. III presents and describes the new fully localised nonlinear optimal, Sec. IV explores the sensitivity of the new optimal to changes in either T or L to ensure that the new optimal is genuinely localised and does not change as the large optimisation times required to calculate minimum seeds are used, Sec. V explores the natural symmetries that appear within the optimal in very short domains and how they can effect the optimal, and Sec. VI contains a discussion of the results.

II. FORMULATION

We consider the problem of a Newtonian fluid flowing through a straight pipe of length L and circular cross section (diameter D). For localised disturbances in the limit of infinitely long pipes, forcing the flow by imposing constant mass flux or constant pressure gradient is equivalent. We take the former to remain consistent with previously published work.^{15,16} Nondimensionalising by the pipe radius ($D/2$) and the mean axial velocity (U), the governing equations of motion are

$$\partial_t \mathbf{u} + \mathcal{U} \partial_z \mathbf{u} + u \mathcal{U}' \hat{\mathbf{z}} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + Re^{-1} \nabla^2 \mathbf{u}, \quad (1)$$

where $\mathcal{U} \hat{\mathbf{z}} = 2(1 - s^2) \hat{\mathbf{z}}$ is the underlying laminar flow to which $\mathbf{u} = (u, v, w)$ is the not-necessarily small perturbation in cylindrical coordinates (s, ϕ, z) and $Re := UD/\nu$ is the Reynolds number. Periodic boundary conditions are imposed across the ends of the pipe (i.e., in z) and no slip conditions on the pipe wall. We wish to identify the perturbation with initial energy $E_0 := E(0)$ that will undergo the most energy growth over a given period of time (T), and to this end, we employ the usual variational approach.²² Defining

$$\langle \cdots \rangle := \int_0^{2\pi/\alpha} \int_0^{2\pi} \int_0^1 \cdots s ds d\phi dz, \quad (2)$$

we numerically maximise the functional

$$\begin{aligned} \mathcal{L} := \langle \tfrac{1}{2} \mathbf{u}(\mathbf{x}, T)^2 \rangle - \lambda \left[\langle \tfrac{1}{2} \mathbf{u}(\mathbf{x}, 0)^2 \rangle - E_0 \right] - \int_0^T \langle \mathbf{v} \cdot \left[\partial_t \mathbf{u} + \mathcal{O}u \partial_z \mathbf{u} + u \mathcal{O}u' \hat{\mathbf{z}} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p - Re^{-1} \nabla^2 \mathbf{u} \right] \rangle dt \\ - \int_0^T \langle \Pi \nabla \cdot \mathbf{u} \rangle dt - \int_0^T \Gamma \langle \mathbf{u} \cdot \hat{\mathbf{z}} \rangle dt, \end{aligned} \quad (3)$$

in the manner laid out in Ref. 15. Throughout this paper, we will consider the two quantities,

$$e(z, t) := \int_0^{2\pi} \int_0^1 \tfrac{1}{2} \mathbf{u}^2 s ds d\phi \quad \text{and} \quad E(t) := \langle \tfrac{1}{2} \mathbf{u}^2 \rangle = \int_0^{2\pi/\alpha} e(z, t) dz. \quad (4)$$

$E(t)$ is the total energy of the perturbation, while $e(z, t)$ is the energy per unit length of the perturbation at a given axial position along the pipe at a given time. We also consider subdivisions of the energy into the roll and streak energies

$$E_{uv}(t) := \langle \tfrac{1}{2} (u^2 + v^2) \rangle \quad \text{and} \quad E_w(t) := \langle \tfrac{1}{2} w^2 \rangle, \quad (5)$$

with the equivalent versions $e_{uv}(z, t)$ and $e_w(z, t)$ defined in the expected way. In places, we also use E_{ml} , E_{ml}^{uv} , and E_{ml}^w to denote the total energy, the cross-stream energy, and the streamwise energy, respectively, in the Fourier-Fourier mode with azimuthal wavenumber m and axial wavenumber l . If m or l is given specific values then this is the wavenumber being considered; otherwise, it is summed over all wavenumbers.

All calculations are performed with 64 finite difference points in s (concentrated near the boundary) and azimuthal Fourier modes running from -23 to 23 . For a pipe of length $L = 25D$, we use axial Fourier modes between -128 and 128 , while for other pipe lengths, this is adjusted to keep the same physical resolution in z (e.g., the $16\pi D$ run in Figure 6 includes Fourier modes from -256 and 256). Throughout we use $Re = 2400$ and, except where indicated otherwise, we take the optimisation time to be $T = T_{lin} = 29.35 D/U$ —the time that maximises the linear optimal growth at this Reynolds number.

III. LOCALISED NONLINEAR OPTIMAL

In order to find the localised nonlinear optimal, we performed the transient growth calculation outlined in Sec. II in an $8\pi \approx 25$ diameter long pipe. As with shorter domains, for initial energies below a certain threshold (here, $E = 1.12 \times 10^{-4}$), a streamwise-independent minor variation of the linear optimal (Quasi-Linear Optimal Perturbation, abbreviated to QLOP) is found (green curve in Fig. 1, left). At energies larger than this threshold, a new three dimensional optimal (NonLinear Optimal Perturbation, or NLOP) emerges (red curve in Fig. 1, left). The growth produced by the new NLOP quickly dwarfs the energy growth of the corresponding QLOP as the initial energy is increased further, until convergence ceases to be possible and we begin to find turbulent seeds—perturbations that lead to a turbulent end state by $t = T$. The energy level at which convergence fails, E_{fail} , is in the interval $1.7 \times 10^{-4} < E_{fail} < 1.8 \times 10^{-4}$.

Unlike for shorter periodic domains, the QLOP remains a local maximum within the optimisation procedure even for energies where the NLOP produces more growth. This makes it harder to assess whether the NLOP is genuinely the global maximum. In Fig. 2, we plot the path in phase space that the maximisation algorithm iterates through for a variety of initial guesses. These states were formed by taking linear combinations of the QLOP, the NLOP, and turbulent flow fields. Despite trying a range of initial conditions, we were unable to find any that led to any optimals other than either NLOP (the circle) or QLOP (this has no energy associated with streamwise-dependent flow and so would be at $y \rightarrow -\infty$ on the plot).

For the nonlinear optimal corresponding to $E_0 = 1.6 \times 10^{-4}$ (marked as a red dot in Fig. 1, left), we plot cross sections of the perturbation during its development (Fig. 1, right). The sequence shown is very similar to that observed in shorter pipes.^{15,16} Like the previously known optimals, the initial

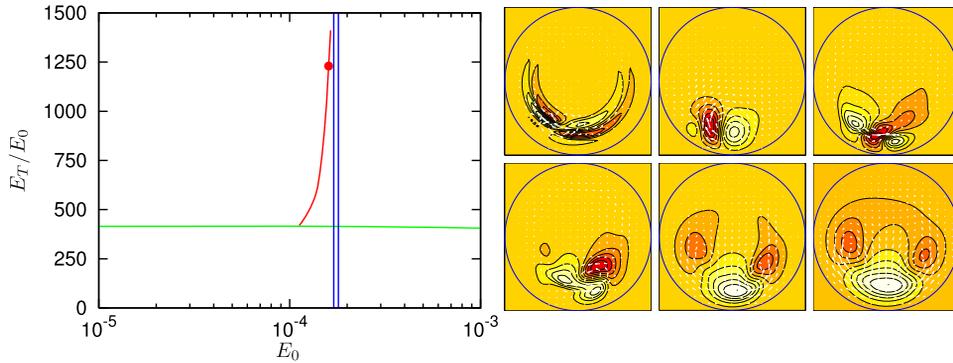


FIG. 1. **Left:** Growth $:= E(T)/E(0)$ as a function of initial energy $E(0)$ for $T = T_{lin}$. The green (almost flat) line is the result of a streamwise-independent, nonlinear calculation where the optimal (QLOP) is found. The red (steeply climbing) line shows the new 3D optimal NLOP, while the vertical blue lines represent the interval within which turbulent seeds begin to appear. The solid dot indicates the nonlinear optimal at $E_0 = 1.6 \times 10^{-4}$, used as the exemplar localised optimal throughout. **Right:** The evolution of the NLOP in a $25D$ long pipe with $E_0 = 1.6 \times 10^{-4}$. The slices are taken at z corresponding to the maximum value of $e(z, t)$ at times $t = 0, 1, 2.5, 10, 20$ and T_{opt} (all in D/U) (left to right, top to bottom). Streak contour levels are varied between slices to show the structure of the growing disturbance.

disturbance is strongly localised in the cross-sectional plane and unpacks through a complicated procedure¹⁵ to produce two larger rolls straddling three streaks. Unlike those previously reported, however, the optimal found here is also strongly localised in the streamwise direction with 99% of the energy contained within a $7D$ section of the pipe—shown in Fig. 3. The rolls shown as isosurfaces weave their way along one side of the pipe, threading through the streak contours shown at discrete cross sections along the pipe. These structures are tightly layered and inclined back into the oncoming flow. This mirrors the structure of the optimal in shorter pipes¹⁵ where the initial growth is driven by the Orr-mechanism in which the layers are tilted up into the underlying shear.

Localisation is present throughout the energy window in which convergence to a nonlinear optimal is possible (Fig. 4). As the initial energy is varied, the streak structure remains essentially unchanged. The roll structure separates slightly in the axial direction as the initial energy is increased leading to two slightly distinct peaks for higher energies.

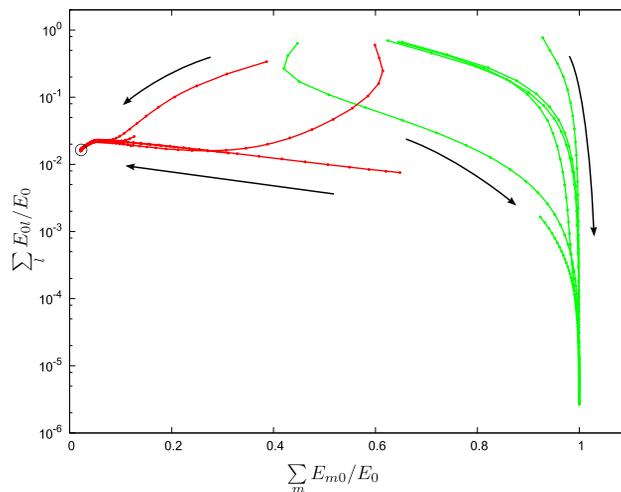


FIG. 2. The paths in phase space of thirteen different initial conditions which either converge to the NLOP (circle at extreme left) or the QLOP (at $(1, 0)$). The axes are normalised perturbation energies associated with the streamwise-independent part of the flow field (x-axis) and the streamwise-dependent part of the flow field (y-axis). The initial conditions were constructed from different linear combinations of the QLOP, NLOP, and 3 turbulent states, each scaled to $E_0 = 1.6 \times 10^{-4}$. Dots indicate individual iterates during the optimisation process. In this projection, some of the initial conditions are slightly obscured.

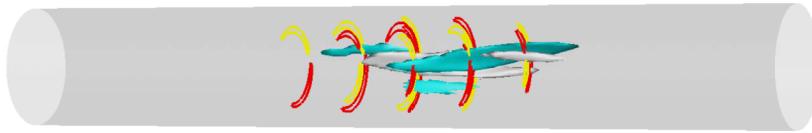


FIG. 3. A $7D$ section of the NLOP at $t=0$ from a $25D$ pipe for $E_0=1.6\times 10^{-4}$. The white (cyan) surface is an isosurface where the vorticity is 30% (-30%) of the maximum vorticity in the pipe. The yellow (red) lines are contours on cross-sectional surfaces of positive (negative) streamwise velocity. Flow is from left to right.

IV. PARAMETER SENSITIVITY

For the optimal described in Sec. III to be a good approximation to the localised minimal seed, it must be insensitive to changes in the chosen periodic length (L) and to increases in the target time (T). If changes to L alter the optimal, then the evolution of the optimal is dependent on the periodicity of the domain and so is not truly localised. Using nonlinear transient growth to find the minimal seed only formally works if asymptotically large target times are used. We are seeking to approximate the minimal seed with a shorter time optimal, and so this will only be reasonable if increasing the target time leads to no notable changes in the form of the optimal.

In order to capture a localised optimal, we must make sure that not only is the optimal initially localised but also that it remains localised throughout its evolution. The perturbation is expected to swell as it grows in energy, and the pipe must be long enough that as it expands it does not begin to interact with itself through the periodic boundary conditions. As an initial test of this, we plot how the perturbation energy spreads through the pipe as it evolves in Figure 5. The lines represent contours of $e(z,t)/E(t)$ with the thick lines representing 0.25, 0.5, and 0.75 levels and the thinner lines the other octiles (the plot is in the comoving frame). The optimal seed gradually unpacks in the first $10D/U$ before the lift-mechanism takes over causing streaks to develop and elongate (Figure 1 gives flow snapshots of the same evolution).

This plot is highly suggestive that the perturbation remains localised; however, the degree of self-interaction was tested further by repeating the nonlinear transient growth calculation for pipes of lengths $L=2\pi D$, $4\pi D$, and $16\pi D \approx 50D$ at $E_0=1.6\times 10^{-4}$, close to that required to trigger turbulence. In Fig. 6, we plot the energy evolution of these optimals (each with $E_0=1.6\times 10^{-4}$) along with that of the benchmark $8\pi D \approx 25D$ optimal. In all four cases, the initial evolution is indistinguishable, but as the optimals begin to unfurl along the pipe, the evolutions begin to diverge. Unsurprisingly, this self-interaction has the greatest effect upon the optimal in the shortest periodic domain. For the $4\pi D$ optimal, it is only the very final part of the evolution that is affected, while the $25D$ and $50D$ optimals are indistinguishable. The inset of Fig. 6 shows the initial axial distribution of the energy within these optimals. The central portion of each of these optimals align very closely with the distributions only diverging as each optimal reaches the ends of its periodic domain. Taken

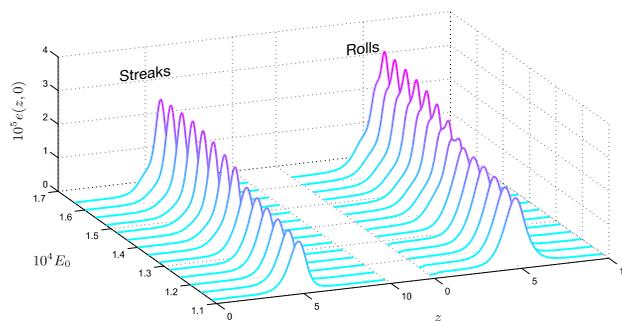


FIG. 4. The axial distribution of the initial energy in the streaks ($e_w(z,0)$) and the rolls ($e_{uv}(z,0)$) of the NLOP in a $25D$ pipe, and how it changes as a function of E_0 . z is in units of D so only a $10D$ section of the pipe is represented. The streamwise-independent QLOP would be represented as a flat line of amplitude E_0/L on the roll plot and zero across the streak plot.

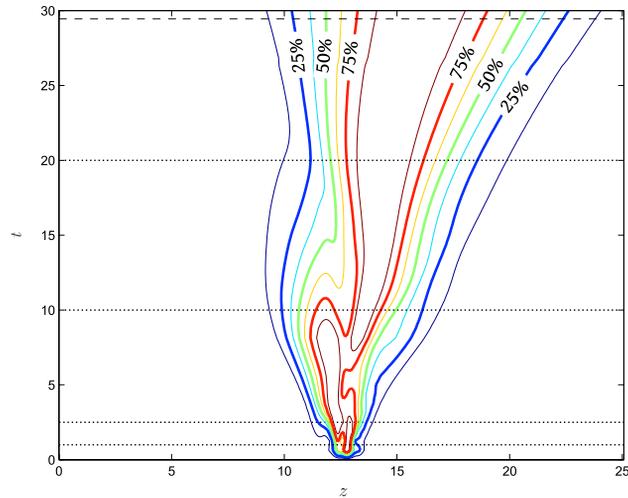


FIG. 5. For each time, the thick lines show where the local energy is 25% (blue), 50% (green), and 75% (red) of the instantaneous maximum energy with the thin lines giving the remaining octiles: 12.5% (blue), 37.5% (light blue), 62.5% (orange), and 87.5% (brown). The horizontal dotted lines, along with the $t = 0$ axis and the horizontal dashed line, correspond to the times for which the cross sections are plotted in Fig. 1.

altogether, these results suggest that calculations performed in much smaller domains are not only able to capture the same mechanisms and qualitative results as those observed in domains large enough to capture localised dynamics but also that the optimals found are in fact precisely the same. The only apparent difference is that for longer domains, there is more space for the energy levels to drop off. This energy drop-off is passive and does not directly influence the form of the perturbation.

As well as domain size and initial energy, the nonlinear transient growth calculation depends upon the target optimisation time. It has previously been reported that it is possible to converge at high initial amplitudes for which turbulence can be triggered, but this is only if short times are considered.^{15,21} What is not clear, however, is whether the *form* of the optimal found in the nonlinear calculation also depends upon the choice of target time.

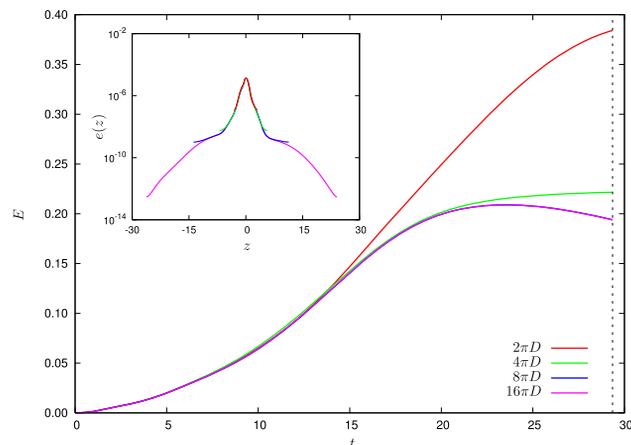


FIG. 6. **Outer:** Evolutions of the NLOPs found in different length pipes at $E_0 = 1.6 \times 10^{-4}$. All four optimals initially have indistinguishable evolutions, before they begin to separate. The shortest ($2\pi D$) domain separates first generating the most growth followed by the $4\pi D$ domain. For the two longest domains, they remain inseparable throughout the evolution period (the blue $8\pi D$ line is underneath the magenta $16\pi D$ curve). **Inner:** The axial distribution of the initial energy of the same NLOPs. The central structure of the differing optimals closely match, diverging only as the ends of the periodic domain are reached.

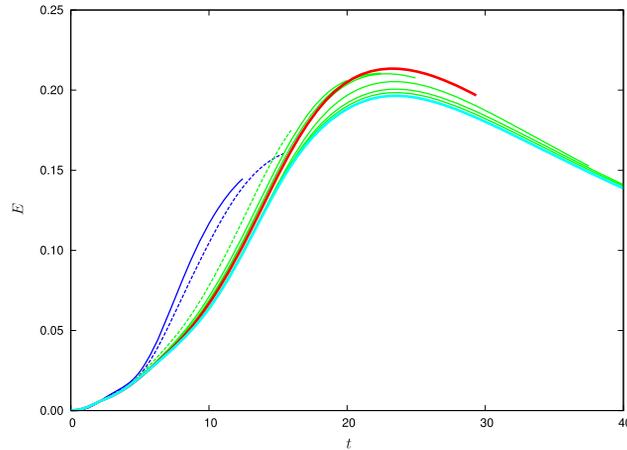


FIG. 7. Evolution of optimals with differing T . The two blue (dark thin) lines have very short optimisation times ($T = 12.5 D/U$ and $15.5 D/U$). These two optimals evolve in a notably different manner to all the other optimals ($T \geq 16 D/U$) being considered. The optimal corresponding to $T = T_{lim}$ is shown in (thick red), while the largest optimisation time considered ($T = 100D/U$) is shown in cyan (only the first $40 D/U$ is plotted). The two dashed lines correspond to $T = 15.5 D/U$ (blue/dark) and $T = 16 D/U$ (green/light) which bracket the abrupt change in evolution of their respective optimals.

To this end, we performed the nonlinear transient growth calculation for a range of target times, again all with $E_0 = 1.6 \times 10^{-4}$. The evolutions of the optimals found through this are shown in Fig. 7. Two separate forms of evolution are observed. For values of $T \gtrsim 16D/U$, the optimals all evolved in similar manners, reaching a peak energy level at $T \approx 20D/U$ before decaying away. Smaller values of T give an optimal that undergoes an accelerated evolution. The eventual maximum energy levels obtained (were the perturbation allowed to evolve indefinitely) are lower than before. The transition between these two optimals is not abrupt. For $13.75D/U \lesssim T \lesssim 16D/U$, both these states are local optimals and either can be found depending on the initial condition selected.

Unless drastically shorter times are taken, the optimal observed is relatively insensitive to the value of T chosen. Previous work^{15,22} conjectured, and provided evidence that, if large optimisation times are used, then the nonlinear transient growth algorithm can be used to identify both the minimum amplitude of disturbance required to trigger turbulence and the minimum seed that this equates to. Due to the computationally demanding nature of the computation even in short domains, in this work, we have considered intermediate optimisation times. Nonetheless, it appears clear that this is sufficient to produce a good approximation of the minimal seed.

V. SYMMETRIES

Although the optimal is robust to even quite significant changes in pipe length, extremely short domains reveal a different optimal. The new optimal that emerges seems fundamentally different to our previous optimal in that it possesses both mirror symmetry

$$\mathcal{Z} : (u, v, w)(s, \pi/2 + \phi, z) \rightarrow (u, -v, w)(s, \pi/2 - \phi, z) \quad (6)$$

and shift-and-reflect symmetry

$$\mathcal{S} : (u, v, w)(s, \phi, z) \rightarrow (u, -v, w)(s, -\phi, z + L/2). \quad (7)$$

In Fig. 8, the optimal for $E_0 = 0.8 \times 10^{-4}$ (chosen to ensure convergence to a nonlinear optimal in all lengths of periodic domain considered) is found for varying choices of L . At $L \approx 3.5D$, the solid red line corresponding to symmetry-free nonlinear optimals is superseded by the blue line of optimals exhibiting both \mathcal{S} and \mathcal{Z} symmetries. It is worth emphasizing that these symmetries are *not* enforced but chosen naturally by the optimisation procedure. However, to extend the latter line past the size of domain for which the non-symmetric optimal is the global optimal, we enforced

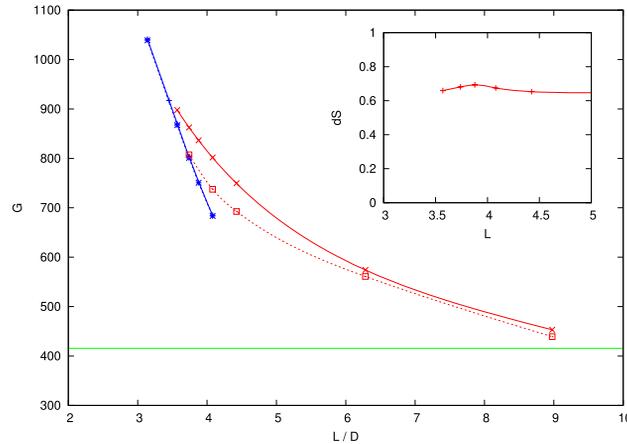


FIG. 8. Growth against length of pipe L for the streamwise independent optimal, QLOP (horizontal green), the nonlinear optimal, NLOP (red with crosses), and the nonlinear optimal with shift-and-reflect symmetry enforced within the domain (blue with stars). The switch in optimal type can clearly be seen at $L \approx 3.5D$ with the shift-and-reflect optimal being the global optimal for shorter domains than this. The red dashed line shows the optimal when only mirror-symmetry is enforced. The inset demonstrates that the switchover is a crossing over of distinct maximums as the energy in the symmetry-breaking part of the large L optimal does not vanish at the switchover length. The dashed red line represents an optimal where only mirror symmetry is enforced. All results are for fixed global energy $E_0 = 0.8 \times 10^{-4}$ so as L increases, the growth decreases (consistent with Fig. 1).

S symmetry within the optimisation procedure. For all the parameters tried, whenever S symmetry was enforced, the optimal found also naturally possessed \mathcal{Z} symmetry. Although the two lines seem to come together smoothly, the two optimals are not smooth continuations of each other—there is a step change at a critical length. This is revealed by examining the quantity

$$dS := \sqrt{\frac{\langle (\mathbf{u} - S\mathbf{u})^2 \rangle}{\langle \mathbf{u}^2 \rangle}}, \quad (8)$$

which measures how far from S symmetry each state is. For the symmetry-free nonlinear optimal, this quantity does not vanish as the critical length is approached (see Fig. 8 inset).

If \mathcal{Z} symmetry is explicitly enforced, although the S and \mathcal{Z} -symmetric optimal persists to slightly longer pipes, the S symmetry is soon broken and an optimal that only possesses \mathcal{Z} -symmetry appears. This is not unexpected as unlike shift-and-reflect symmetry, mirror symmetry is supportable by localised disturbances. Both the S and \mathcal{Z} -symmetric and the \mathcal{Z} -symmetric optimals in a $L = 4.1D$ pipe are plotted in Figure 9. The contours for these two optimals are the same and clearly indicate that although the configuration is very similar, the amplitudes are greater for the \mathcal{Z} optimal—this optimal is already beginning to unevenly distribute its energy in the axial direction. The evolution of the S and \mathcal{Z} optimal is still the same pattern of Orr-mechanism, then oblique waves, and finally lift up as previously observed. This can be seen from Figure 9 (bottom) where the energy in different Fourier modes is plotted. For a full interpretation, comparison should be made with Ref. 15, Figure 3.

It is also possible to find localised, \mathcal{Z} symmetric optimals in much longer domains. In Figure 9, we show cross sections of the optimal found in a $12.5D$ pipe with \mathcal{Z} symmetry enforced at times $t = 0$ (top, third from left) and T_{opt} (top, right). As is now expected, comparing with the symmetry-free optimal (Figure 1), we see a similarity. The structure is still dominated by tightly layered streaks interweaved with rolls, but now they are stretched out over the full azimuthal extent of the cross section to allow for mirror-symmetry.

VI. DISCUSSION

We have found the first energy growth optimal, which despite the underlying equations supporting strictly periodic domain-filling flows, is fully localised. Reassuringly, our examinations of

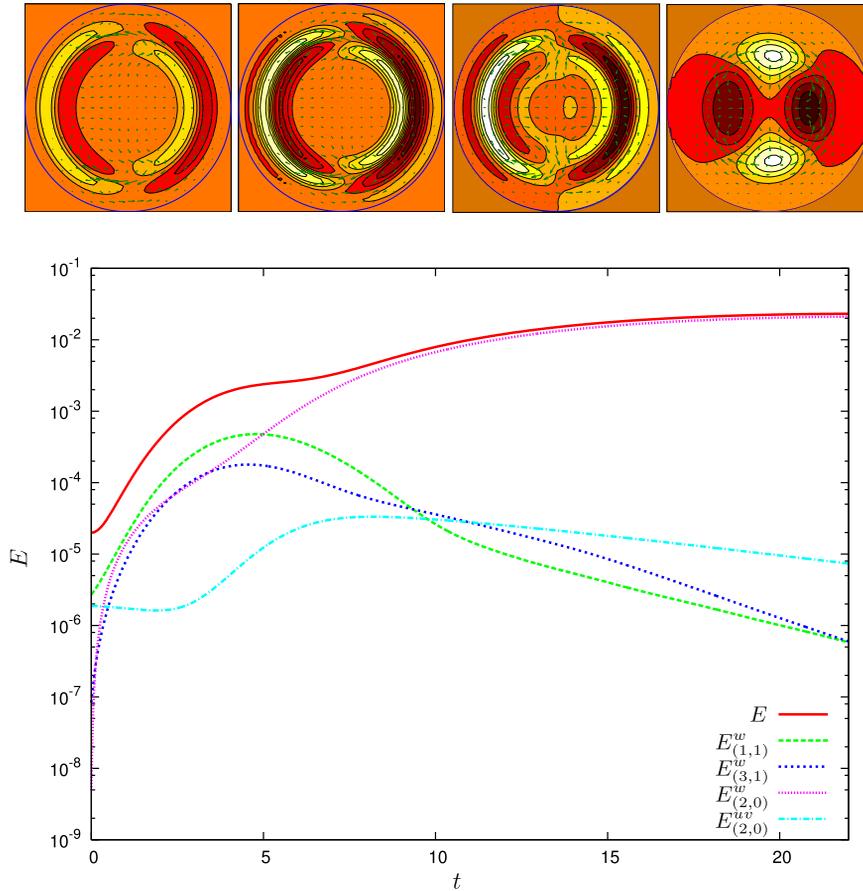


FIG. 9. **Across the top, from left to right:** The \mathcal{S} and \mathcal{Z} symmetric optimal at $L = 4.08D$; the \mathcal{Z} symmetric optimal at $L = 4.08D$; the \mathcal{Z} symmetric optimal at $L = 12.5D$; the state the $12.5D$ optimal evolves into after T_{opt} . The broad structure of the optimal—tightly layered streaks with interweaving rolls—is the same as for the previously found optimals. **Below:** The breakdown of energy into individual Fourier modes for the \mathcal{S} and \mathcal{Z} symmetric optimal at $L = 4.08D$ as it evolves in time. This reveals the same evolution process of the Orr-mechanism, followed by oblique waves and lift up.

convergence show that the optimal is robust and fits well with previous results. The central structure of the optimal is strikingly similar to the optimal found in much shorter domains and appears to be essentially the same optimal but with an extended region of (spatial) exponential decay. From this it seems clear that the energy growth mechanisms in the localised optimal are essentially the same as for the optimal found in short pipes,¹⁵ though now the optimal also expands along the domain as it grows in amplitude. Less clear is what sets the rate of energy drop off in z —a strong scaling appears to be present, but it is not the result of a simple energy balance as the optimal is not a steady solution of the Navier-Stokes equations.

The importance of T has also been illuminated. In order to estimate the minimum amplitude of the edge (E_c) and the minimal seed to high accuracy, large optimisation target times are required. Despite this, the form of the minimal seed is accurately revealed by more intermediate choices. Further, with this choice of T , we were able to find reasonable energy bounds, $1.7 \times 10^{-4} < E_c < 1.8 \times 10^{-4}$. The upper bound is firm as it comes from finding turbulent seeds at this energy. The lower bound is less definite from this calculation alone but is confirmed by performing larger T calculations, for which $E_c = E_{fail}$ —see Refs. 15 and 22. This estimate for E_c is consistent with the $5D$ result found much more precisely¹⁵ as the slightly lower (since the perturbation can self-interact) value of $E_c = 1.5 \times 10^{-4}$. It is only when we consider much lower choices of T or very short L that any substantial differences appear. In the case of reducing T , a second optimal emerges that prioritises fast unsustainable growth. Over a small interval of choices of T , both of these optimals are local maximums.

A similar change takes place for very short choices of L , where the optimal switches to one exhibiting both shift-and-reflect symmetry and mirror symmetry appears. Outside of these two extreme cases, we have shown that using more computationally viable parameter regimes than those previously stipulated¹⁵ still allows us to ascertain insight into the minimal seed and the corresponding critical minimum amplitude of turbulence. One immediate observation is that the minimal seed has 99% of its energy concentrated in just $7D$ of the pipe length at $Re = 2400$. This resonates with the observation in experimental work^{23,24} that once disturbances generated by jets become more than $\approx 6D$ long, the ensuing dynamics is largely independent of the disturbance length. Recently discovered, localised relative periodic orbits in pipe flow also share this length scale of $5\text{--}10D$ as do turbulent puffs: see Figure 2 of Ref. 25 and Figure 4 of Ref. 26.

In terms of future work, the way is now clear to map out the threshold energy E_c for transition as a function of Re just as has been recently done in small-box plane Couette flow²⁷ (the imposed streamwise periodicity in Ref. 27 is equivalent to $L = 2\pi D$ here). Our results indicate that using small-to-intermediate periodic domains (at least in the streamwise direction) can still yield useful results. Experimentally, of course, only a small subset Σ of all possible disturbances considered theoretically can actually be generated. To move the theory closer to this reality just requires that the optimisation be performed over Σ , which means simply projecting the variational derivative of the energy growth with respect to the initial perturbation down onto Σ . The greater theoretical challenge is actually to accurately model the disturbances routinely generated in the laboratory by injecting or removing fluid through small holes.^{23,24} Adding an artificial body force temporarily to the Navier–Stokes equations, however, seems to work well.²⁸

Another direction to take this work is into control. Here, the aim could be to increase E_c by manipulating some aspect of the flow. A first promising step along these lines has already been made in plane Couette flow by oscillating the boundaries in their plane and perpendicular to the shearing direction.²⁹

ACKNOWLEDGMENTS

The calculations in this paper were carried out at the Advanced Computing Research Centre, University of Bristol.

- ¹ A. G. Darbyshire and T. Mullin, “Transition to turbulence in constant-mass-flux pipe-flow,” *J. Fluid Mech.* **289**, 83–114 (1995).
- ² H. Faisst and B. Eckhardt, “Sensitive dependence on initial conditions in transition to turbulence in pipe flow,” *J. Fluid Mech.* **504**, 343–352 (2004).
- ³ B. F. Farrell, “Optimal excitation of perturbations in viscous shear flow,” *Phys. Fluids* **31**, 2093–2101 (1988).
- ⁴ B. F. Farrell, “Optimal excitation of baroclinic waves,” *J. Atmos. Sci.* **46**, 1193–1206 (1989).
- ⁵ L. H. Gustavsson, “Energy growth of 3-dimensional disturbances in plane Poiseuille flow,” *J. Fluid Mech.* **224**, 241–260 (1991).
- ⁶ K. M. Butler and B. F. Farrell, “3-dimensional optimal perturbations in viscous shear-flow,” *Phys. Fluids* **4**, 1637–1650 (1992).
- ⁷ P. Schmid and D. Henningson, “A new mechanism for rapid transition involving a pair of oblique waves,” *Phys. Fluids* **4**, 1986–1989 (1992).
- ⁸ L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, “Hydrodynamic stability without eigenvalues,” *Science* **261**, 578–584 (1993).
- ⁹ L. Bergström, “Optimal-growth of small disturbances in pipe Poiseuille flow,” *Phys. Fluids* **5**, 2710–2720 (1993).
- ¹⁰ S. C. Reddy and D. S. Henningson, “Energy growth in viscous flows,” *J. Fluid Mech.* **365**, 209–238 (1993).
- ¹¹ P. J. Schmid and D. S. Henningson, “Optimal energy density growth in Hagen-Poiseuille flow,” *J. Fluid Mech.* **277**, 192–225 (1994).
- ¹² O. Y. Zikanov, “On the instability of pipe Poiseuille flow,” *Phys. Fluids* **8**, 2923–2932 (1996).
- ¹³ S. C. Reddy, P. J. Schmid, J. S. Baggett, and D. S. Henningson, “On stability of streamwise streaks and transition thresholds in plane channel flows,” *J. Fluid Mech.* **365**, 269–303 (1998).
- ¹⁴ P. J. Schmid, “Nonmodal stability theory,” *Annu. Rev. Fluid Mech.* **39**, 129–162 (2007).
- ¹⁵ C. C. T. Pringle, A. P. Willis, and R. R. Kerswell, “Minimal seeds for shear flow turbulence: Using nonlinear transient growth to touch the edge of chaos,” *J. Fluid Mech.* **702**, 415–443 (2012).
- ¹⁶ C. C. T. Pringle and R. R. Kerswell, “Using nonlinear transient growth to construct the minimal seed for shear flow turbulence,” *Phys. Rev. Lett.* **105**, 154502 (2010).
- ¹⁷ S. Cherubini, P. de Palma, J.-C. Robinet, and A. Bottaro, “Rapid path to transition via nonlinear localised optimal perturbations in a boundary layer flow,” *Phys. Rev. E* **82**, 066302 (2010).
- ¹⁸ A. Monokrousos, A. Bottaro, L. Brandt, A. Di Vita, and D. S. Henningson, “Nonequilibrium thermodynamics and the optimal path to turbulence in shear flows,” *Phys. Rev. Lett.* **106**, 134502 (2011).

- ¹⁹ S. M. E. Rabin, C. P. Caulfield, and R. R. Kerswell, "Triggering turbulence efficiently in plane Couette flow," *J. Fluid Mech.* **712**, 244–272 (2012).
- ²⁰ S. Cherubini, P. de Palma, J.-C. Robinet, and A. Bottaro, "The minimal seed of turbulent transition in the boundary layer," *J. Fluid Mech.* **689**, 221–253 (2011).
- ²¹ S. Cherubini and P. de Palma, "Nonlinear optimal perturbations in a Couette flow: Bursting and transition," *J. Fluid Mech.* **724**, 251–279 (2013).
- ²² R. R. Kerswell, C. C. T. Pringle, and A. P. Willis, "An optimisation approach for analysing nonlinear stability with transition to turbulence in fluids as an exemplar," *Rep. Prog. Phys.* **77**, 085901 (2014).
- ²³ B. Hof, A. Juel, and T. Mullin, "Scaling of the turbulence transition in a pipe," *Phys. Rev. Lett.* **91**, 244502 (2003).
- ²⁴ J. Peixinho and T. Mullin, "Finite amplitude thresholds for transition in pipe flow," *J. Fluid Mech.* **582**, 169–178 (2007).
- ²⁵ M. Avila, F. Mellibovsky, N. Roland, and B. Hof, "Streamwise-localised solutions at the onset of turbulence in pipe flow," *Phys. Rev. Lett.* **110**, 224502 (2013).
- ²⁶ M. Chantry, A. P. Willis, and R. R. Kerswell, "Genesis of streamwise-localized solutions from globally periodic travelling waves in pipe flow," *Phys. Rev. Lett.* **112**, 164501 (2014).
- ²⁷ Y. Duguet, A. Monokrousos, L. Brandt, and D. S. Henningson, "Minimal transition thresholds in plane Couette flow," *Phys. Fluids* **25**, 084103 (2013).
- ²⁸ F. Mellibovsky and A. Meseguer, "Critical threshold in pipe flow transition," *Philos. Trans. R. Soc., A* **367**, 545–560 (2009).
- ²⁹ S. M. E. Rabin, C. P. Caulfield, and R. R. Kerswell, "Designing a more nonlinearly stable laminar flow via boundary manipulation," *J. Fluid Mech.* **738**, R1 (2014).