CrossMark

J. Fluid Mech. (2016), *vol.* 795, *pp.* 278–312. © Cambridge University Press 2016
doi:10.1017/jfm.2016.203

An adjoint-based approach for finding invariant solutions of Navier–Stokes equations

M. Farazmand†

Center for Nonlinear Science, School of Physics, Georgia Institute of Technology, Atlanta, GA 30332, USA

(Received 26 August 2015; revised 20 December 2015; accepted 14 March 2016; first published online 14 April 2016)

We consider the incompressible Navier–Stokes equations with periodic boundary conditions and time-independent forcing. For this type of flow, we derive adjoint equations whose trajectories converge asymptotically to the equilibrium and travelling-wave solutions of the Navier–Stokes equations. Using the adjoint equations, arbitrary initial conditions evolve to the vicinity of a (relative) equilibrium at which point a few Newton-type iterations yield the desired (relative) equilibrium solution. We apply this adjoint-based method to a chaotic two-dimensional Kolmogorov flow. A convergence rate of 100 % is observed, leading to the discovery of 21 new steady-state and travelling-wave solutions at Reynolds number $Re = 40$. Some of the new invariant solutions have spatially localized structures that were previously believed to exist only on domains with large aspect ratios. We show that one of the newly found steady-state solutions underpins the temporal intermittencies, i.e. high energy dissipation episodes of the flow. More precisely, it is shown that each intermittent episode of a generic turbulent trajectory corresponds to its close passage to this equilibrium solution.

Key words: chaos, intermittency, transition to turbulence

1. Introduction

The idea that turbulent fluid flow can be studied in terms of the invariant solutions of Navier–Stokes equations dates back to the 1940s (Hopf 1948). Examples of such invariant solutions (also called ‘exact coherent states’ or ‘recurrent flows’ in the fluid dynamics literature) are steady states (or equilibria), travelling waves (or relative equilibria), periodic orbits, relative periodic orbits and partially hyperbolic tori. A turbulent flow is then viewed as a walk from the neighbourhood of one invariant solution to the other. If all these solutions are unstable, the turbulent trajectory never settles down and its itinerary becomes desperately complex.

Initially a mathematical endeavour, this view has been put to practice over the last 15 years in many experimental and numerical studies, providing insights that are beyond the reach of purely statistical descriptions of turbulence – see Kawahara, Uhlmann & van Veen (2012) and Cvitanović (2013), for reviews.

The invariant solutions of the Navier–Stokes equations often exhibit complex spatiotemporal behaviours, and hence analytic expressions for them are unavailable.

† Email address for correspondence: mfarazmand6@gatech.edu

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

Their numerical computation was first explored in the pioneering works of Marcus & Tuckerman (1987) and Nagata (1990). Nagata found the first non-trivial equilibrium and travelling-wave solutions of the plane Couette flow (Nagata 1990, 1997). Nagata's approach to finding these solutions is the following. Through a finite truncation, the invariant solutions are formulated as the zeros of a large system of nonlinear algebraic equations $\boldsymbol{F}(\boldsymbol{a}) = \boldsymbol{0}$. The solution $\boldsymbol{a}$ is then found through Newton–Raphson iterations. Depending on the type of the discretization, the finite vector $\boldsymbol{a}$ contains either the Galerkin coefficients or the collocation values.

Nagata's approach underlies the state-of-the-art methods for finding invariant solutions of fluid flow (see e.g. Waleffe 1998; Kawahara & Kida 2001; Viswanath 2007; Gibson, Halcrow & Cvitanović 2009; Willis, Cvitanović & Avila 2013). Such Newton-type approaches suffer from two major drawbacks:

(i) Obtaining the Newton direction $\delta\boldsymbol{a}$ requires solving the linear system of equations $\boldsymbol{\nabla F}(\boldsymbol{a})\delta\boldsymbol{a} = -\boldsymbol{F}(\boldsymbol{a})$. Even for moderate Reynolds numbers, this linear equation is too large to be solved accurately (within numerical precision) in reasonable computation time.

(ii) The convergence of the Newton iterations is not generally guaranteed. The iterations converge only if very good initial guesses are provided.

As regards item (i), the linear system is typically too large for the matrix $\boldsymbol{\nabla F}(\boldsymbol{a})$ to be even explicitly formed (due to memory considerations). To address this issue, the generalized minimal residual (GMRES) method is used to approximate the solution of the linear system (Saad & Schultz 1986). The GMRES method replaces the exact Newton direction with its least-square approximation within a Krylov subspace. Forming this subspace only requires the action of the matrix $\boldsymbol{\nabla F}(\boldsymbol{a})$ on certain vectors, hence avoiding the formation of the matrix itself. The resulting method is often referred to as Newton–GMRES iterations. Although computationally feasible, Newton–GMRES iterations are still a formidable numerical undertaking in terms of implementation and computation time.

As mentioned in item (ii) above, even when the exact Newton descent direction $\delta\boldsymbol{a}$ is known, the iterations are not guaranteed to converge. Newton's domain of convergence in the context of Navier–Stokes equations is typically small, demanding a very good initial guess for the iterations to converge (Waleffe 1997). The Newton iterations, therefore, are very effective for bifurcation analysis where the invariant solutions bifurcate from known solutions (Tuckerman *et al.* 2000; Waleffe 2003; Kreilos & Eckhardt 2012). In this case, the known solution (corresponding to the old bifurcation parameter value) is used as the initial guess for the Newton iterations. Even then, the method might fail to determine solutions far from the bifurcation point.

In practice, good initial guesses for Newton iterations are obtained through heuristic methods, such as close recurrences of a turbulent trajectory (Auerbach *et al.* 1987; Kawahara, Kida & Van Veen 2006; Viswanath 2007; Cvitanović & Gibson 2010). Such heuristics lack a solid mathematical basis. As a result, many important invariant solutions may and will remain undiscovered (see § 5).

Viswanath (2007) used the hook step together with the Newton–GMRES method, which significantly improved the convergence of the iterations. The underlying idea of the hook step is to approximate the Newton direction within the Krylov subspace with the constraint that $\delta\boldsymbol{a}$ belongs to a 'trust region' such that $|\delta\boldsymbol{a}|$ is smaller than some prescribed value ϵ. The constraint $|\delta\boldsymbol{a}| < \epsilon$ ensures that the linearization $\boldsymbol{\nabla F}(\boldsymbol{a})\delta\boldsymbol{a}$ involved in Newton iterations is a valid approximation (Dennis & Schnabel 1996).

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

The resulting method, i.e. the Newton–GMRES–hook iterations, is state of the art in computing invariant solutions in the context of fluid flow. While improving the convergence chance of the Newton–GMRES iterations, the hook step still requires a good initial guess.

Such drawbacks have slowed down research progress in transitional and moderate-Reynolds-number turbulence. As the study of fluid turbulence through its invariant solutions reaches a mature state, it is time to revisit the methods by which these invariant solutions are found. Better methods should ideally scale linearly (in the sense of computational complexity) with the number of degrees of freedom and be universally convergent. The term 'universally convergent' implies that every initial guess converges to some solution. Universally convergent methods, however, do not immediately guarantee the identification of all possible solutions from a finite set of initial guesses.

The main goal of the present paper is to propose one such new direction. We develop an adjoint-based method for computing steady-state and travelling-wave solutions of incompressible Navier–Stokes equations with periodic boundary conditions. In particular, adjoint partial differential equations (PDEs) are derived whose trajectories converge to (relative) equilibria of the Navier–Stokes equations.

Adjoint equations appear naturally in optimal control, where the governing equations act as constraints – see Gunzburger (2002) for a survey of applications in fluid dynamics. Pringle & Kerswell (2010) and Monokrousos *et al.* (2011) used adjoint-based optimization to find the optimal route to turbulence in transitional shear flows – see also Kerswell, Pringle & Willis (2014) for a review. To investigate the regularity of the Navier–Stokes equations, Ayala & Protas (2014) use adjoint-based optimization to find the initial conditions that result in maximal palinstrophy growth.

To the best of our knowledge, adjoint-based methods have not been used to find invariant solutions of the Navier–Stokes equations. They have, however, been utilized in the context of nonlinear wave equations. For instance, Ambrose & Wilkening (2010*b*) formulate the time-periodic solutions of a unidirectional water wave equation as the minima of an appropriate functional. The minima are found iteratively by a 'steepest descent' method. At each iteration, the descent direction is obtained as the solution of a backward-time adjoint PDE (see also Ambrose & Wilkening 2010*a*). Yang & Lakoba (2007) realized that, for equilibrium solutions, the descent direction can be expressed explicitly. Namely, they show that the solitary solutions of nonlinear wave equations coincide with the asymptotically stable equilibria of an adjoint PDE.

The case of Navier–Stokes equations is more involved due to the presence of non-local pressure gradients enforcing incompressibility. This calls for a special treatment as presented here. We also extend the adjoint-based approach to finding travelling-wave solutions with *a priori* unknown wave speeds.

Recently, Chandler & Kerswell (2013) and Lucas & Kerswell (2014) carried out the most exhaustive search for invariant solutions of a chaotic two-dimensional Kolmogorov flow using Newton–GMRES–hook iterations. Here, we apply our adjoint-based method to the same Kolmogorov flow and discover several new steady-state and travelling-wave solutions. We show that some of these new solutions underpin the temporal intermittencies associated with short-lived, rapid growth of energy dissipation.

This paper is organized as follows. The preliminary material is reviewed in § 2, in a rather general setting. The explicit form of the adjoint equations for Navier–Stokes equations are presented in § 3. We test the adjoint-based approach in § 4 on a two-dimensional Kolmogorov flow and also carry out a thorough comparison with Newton–GMRES–hook iterations. In § 5, the temporal intermittency of the Kolmogorov flow is

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

studied in terms of its invariant solutions. Finally, § 6 contains our concluding remarks and an outline of future research directions.

2. Preliminaries: Newton descent versus adjoint descent

In this section, we review Newton and adjoint descent methods in a general framework. We find this exposition helpful for understanding the Navier–Stokes adjoint equations presented in § 3.

Consider the PDE

$$\partial_t \boldsymbol{u} = \boldsymbol{F}(\boldsymbol{u}), \tag{2.1}$$

where the real-valued vector field $\boldsymbol{u}(\boldsymbol{x}, t)$ is a function of space $\boldsymbol{x}$ and time t, and $\boldsymbol{F}$ is a nonlinear differential operator acting on an inner product function space $\mathscr{H}$. We seek equilibrium (or steady-state) solutions of this PDE, i.e. time-independent vector fields $\boldsymbol{u} = \boldsymbol{u}(\boldsymbol{x})$ such that $\boldsymbol{F}(\boldsymbol{u}) = \boldsymbol{0}$.

The finite-dimensional counterpart of (2.1) is the system of ordinary differential equations (ODEs) $\mathrm{d}\boldsymbol{a}/\mathrm{d}t = \boldsymbol{F}(\boldsymbol{a})$, where $\boldsymbol{a} \in \mathbb{R}^d$ and $\boldsymbol{F} : \mathbb{R}^d \to \mathbb{R}^d$. While our derivations are in the infinite-dimensional setting, we will occasionally invoke this finite-dimensional analogue to elaborate the ideas.

We define the 'weighted' inner product

$$\langle \boldsymbol{u}, \boldsymbol{u}' \rangle_{\mathscr{A}} = \langle \boldsymbol{u}, \mathscr{A}\boldsymbol{u}' \rangle_{L^2}, \tag{2.2}$$

for any two functions $\boldsymbol{u}, \boldsymbol{u}' \in \mathscr{H}$, where $\mathscr{A}$ denotes a real-valued, positive-definite, self-adjoint (with respect to the L^2 inner product $\langle \cdot, \cdot \rangle_{L^2}$) operator. The prime signs should not be confused with derivatives; we use them here to distinguish different functions. The inner product (2.2) defines a natural norm given by

$$\|\boldsymbol{u}\|_{\mathscr{A}} = \sqrt{\langle \boldsymbol{u}, \boldsymbol{u} \rangle_{\mathscr{A}}}, \tag{2.3}$$

for $\boldsymbol{u} \in \mathscr{H}$. The choice of $\mathscr{A}$ is somewhat arbitrary and should be judiciously chosen for a specific problem. For the Navier–Stokes equation, for instance, we will use an operator that is closely related to the inverse of the Laplacian (see § 3). For this exposition, however, we set $\mathscr{A}$ to the identity, which renders (2.2) as the usual L^2 inner product.

With this setting, we seek a second PDE

$$\partial_\tau \boldsymbol{u} = \boldsymbol{G}(\boldsymbol{u}), \tag{2.4}$$

such that its solutions converge to the equilibrium solutions of (2.1) as the fictitious time τ tends to infinity. To ensure this convergence, the differential operator $\boldsymbol{G}$ is designed in a specific way such that

$$\|\boldsymbol{F}(\boldsymbol{u}(\tau))\|_{L^2} \to 0, \quad \text{as } \tau \to \infty, \tag{2.5}$$

where $\boldsymbol{u}(\tau) = \boldsymbol{u}(\cdot, \tau)$ denotes a solution of (2.4). Newton and adjoint descents correspond to two different choices of the operator $\boldsymbol{G}$.

The finite-dimensional analogue of (2.4) is the fictitious-time ODE $\mathrm{d}\boldsymbol{a}/\mathrm{d}\tau = \boldsymbol{G}(\boldsymbol{a})$ with $\boldsymbol{G} : \mathbb{R}^d \to \mathbb{R}^d$.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

2.1. *Newton descent*

The main observation here is that discrete Newton iterations are an explicit Euler approximation of a continuous fictitious-time dynamical system (Saupe 1988). This well-known fact is rarely mentioned in the literature, prompting Smale (1981) to write 'This geometry is based on an idea which is known, but not usually explicated. Namely, Newton's method for solving $f(z) = 0$ is an Euler approximation to the ordinary differential equation ... '. Here, we first derive the continuous fictitious-time Newton method in the infinite-dimensional setting and then show how the discrete Newton iterations follow from there.

The norm $\|\boldsymbol{F}(\boldsymbol{u})\|_{L^2}$ evolves along the trajectories of (2.4) according to

$$\partial_\tau \|\boldsymbol{F}(\boldsymbol{u})\|_{L^2}^2 = 2\langle \mathscr{L}(\boldsymbol{u}; \boldsymbol{G}(\boldsymbol{u})), \boldsymbol{F}(\boldsymbol{u})\rangle_{L^2}, \tag{2.6}$$

where the linear operator $\mathscr{L}(\boldsymbol{u}; \cdot)$ is the Gâteaux derivative of $\boldsymbol{F}$ evaluated at the state $\boldsymbol{u}$, and is defined by

$$\mathscr{L}(\boldsymbol{u}; \boldsymbol{u}') := \lim_{\epsilon \to 0} \frac{\boldsymbol{F}(\boldsymbol{u} + \epsilon \boldsymbol{u}') - \boldsymbol{F}(\boldsymbol{u})}{\epsilon}, \quad \forall \boldsymbol{u}' \in \mathscr{H}. \tag{2.7}$$

Requiring $\boldsymbol{G}$ to satisfy

$$\mathscr{L}(\boldsymbol{u}; \boldsymbol{G}(\boldsymbol{u})) = -\boldsymbol{F}(\boldsymbol{u}), \tag{2.8}$$

one obtains the continuous fictitious-time Newton descent in infinite dimensions. Substituting this expression in (2.6), we obtain

$$\partial_\tau \|\boldsymbol{F}(\boldsymbol{u}(\tau))\|_{L^2}^2 = -2\|\boldsymbol{F}(\boldsymbol{u}(\tau))\|_{L^2}^2, \tag{2.9}$$

which has the exact solution $\|\boldsymbol{F}(\boldsymbol{u}(\tau))\|_{L^2} = \|\boldsymbol{F}(\boldsymbol{u}(0))\|_{L^2} \exp(-\tau)$. This shows that, as long as (2.8) has a solution, $\|\boldsymbol{F}(\boldsymbol{u}(\tau))\|_{L^2}$ decays to zero exponentially fast along the solutions of (2.4). In other words, the continuous fictitious-time Newton method converges to an equilibrium solution of (2.1) for almost all initial conditions (Saupe 1988; Cvitanović & Lan 2003; Lan & Cvitanović 2004).

For a given state $\boldsymbol{u}$, (2.8) is a PDE that can in principle be solved for $\boldsymbol{G}$. In practice, it is approximated through some finite truncation to take the form $\boldsymbol{\nabla F}(\boldsymbol{a})\boldsymbol{G}(\boldsymbol{a}) = -\boldsymbol{F}(\boldsymbol{a})$, with $\boldsymbol{a}$ being the Galerkin coefficients. The solution $\boldsymbol{G}$ of this large, but finite-dimensional, linear system is often approximated by some variant of GMRES iterations (Trefethen & Bau 1997).

Furthermore, the PDE (2.4) is discretized in time to yield the explicit Euler step

$$\boldsymbol{u}_{i+1} = \boldsymbol{u}_i + \delta\tau \boldsymbol{G}(\boldsymbol{u}_i), \tag{2.10}$$

with $0 < \delta\tau \leqslant 1$. The standard Newton iterations correspond to $\delta\tau = 1$, while damped Newton iterations adjust $\delta\tau$ in order to ensure that the L^2 norm decreases (Boyd & Vandenberghe 2004).

The above discrete iterations do not guarantee the global convergence that the continuous fictitious-time Newton descent does (Saupe 1988). As mentioned in the introduction, the convergence of Newton iterations is only guaranteed if the initial guess is sufficiently close to an equilibrium. In fact, if such a close initial guess is available, the convergence to the equilibrium is super-exponential (Deuflhard 2011).

Anecdotal evidence suggests that, in the context of fluid flow, the basin of attraction of Newton iterations is rather small (Waleffe 1997; Viswanath 2007). This domain can be enlarged by choosing a higher-order scheme for temporal discretization of (2.4). This is, however, computationally too demanding for large-dimensional systems, such as turbulent fluid flow.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

2.2. *Adjoint descent*

The adjoint descent corresponds to an alternative choice of $\boldsymbol{G}$ that may be expressed analytically, thus avoiding the approximation involved in the Newton–GMRES iterations. Moreover, the adjoint direction can be evaluated at a relatively low computational cost, rendering a higher-order time discretization of (2.4) feasible.

To express the adjoint direction, note that (2.6) can be written as

$$\partial_\tau \|\boldsymbol{F}(\boldsymbol{u})\|_{L^2}^2 = 2\langle \boldsymbol{G}(\boldsymbol{u}), \mathscr{L}^\dagger(\boldsymbol{u}; \boldsymbol{F}(\boldsymbol{u}))\rangle_{L^2}, \tag{2.11}$$

where the adjoint $\mathscr{L}^\dagger(\boldsymbol{u}; \cdot)$ is the linear operator satisfying

$$\langle \mathscr{L}(\boldsymbol{u}; \boldsymbol{u}'), \boldsymbol{u}''\rangle_{L^2} = \langle \boldsymbol{u}', \mathscr{L}^\dagger(\boldsymbol{u}; \boldsymbol{u}'')\rangle_{L^2}, \quad \forall \boldsymbol{u}, \boldsymbol{u}', \boldsymbol{u}'' \in L^2. \tag{2.12}$$

Requiring

$$\boldsymbol{G}(\boldsymbol{u}) = -\mathscr{L}^\dagger(\boldsymbol{u}; \boldsymbol{F}(\boldsymbol{u})) \tag{2.13}$$

ensures the descent of the norm $\|\boldsymbol{F}(\boldsymbol{u}(\tau))\|_{L^2}$ along the trajectories of (2.4), since

$$\partial_\tau \|\boldsymbol{F}(\boldsymbol{u})\|_{L^2}^2 = -2\|\mathscr{L}^\dagger(\boldsymbol{u}; \boldsymbol{F}(\boldsymbol{u}))\|_{L^2}^2 \leqslant 0. \tag{2.14}$$

We refer to $\boldsymbol{G}$ given by (2.13) as the adjoint direction. The finite-dimensional analogue of this adjoint direction is $\boldsymbol{G}(\boldsymbol{a}) = -[\nabla \boldsymbol{F}(\boldsymbol{a})]^{\mathrm{T}} \boldsymbol{F}(\boldsymbol{a})$, where T denotes the usual matrix transposition.

The downside of using the adjoint descent is that, unlike the continuous fictitious-time Newton descent, (2.14) does not guarantee an exponential decay of the L^2 norm. In fact, we observe a relatively slow decay in § 4.3, and propose a hybrid adjoint–Newton algorithm to rectify this drawback.

In § 3, we derive the explicit form of the adjoint operator for the incompressible Navier–Stokes equations.

3. Adjoint descent for the Navier–Stokes equations

Consider the incompressible Navier–Stokes equation in non-dimensional variables,

$$\partial_t \boldsymbol{u} = -\boldsymbol{u} \cdot \nabla \boldsymbol{u} - \nabla p + \nu \Delta \boldsymbol{u} + \boldsymbol{f}, \tag{3.1a}$$

$$\nabla \cdot \boldsymbol{u} = 0, \tag{3.1b}$$

defined on the spatial domain $\boldsymbol{x} = (x_1, x_2, x_3) \in \mathscr{D} = [0, L_1] \times [0, L_2] \times [0, L_3]$ with periodic boundary conditions. The parameter ν is the inverse of the Reynolds number, $\nu = Re^{-1}$. We assume that the time-independent forcing term $\boldsymbol{f} = \boldsymbol{f}(\boldsymbol{x})$ is divergence-free, $\nabla \cdot \boldsymbol{f} = 0$.

The goal is to find an adjoint PDE such that, along its solutions $(\boldsymbol{u}(\tau), p(\tau))$, the norm of the right-hand side of (3.1*a*) decays to zero monotonically, while ensuring that $\nabla \cdot \boldsymbol{u} = 0$ for all τ. We derive the adjoint equations for the general norms of the type (2.3). In § 3.3, we make a specific choice for the norm that is used in the subsequent computations.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

3.1. *The adjoint descent equation for equilibria*

Define

$$\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}) = -\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u} - \boldsymbol{\nabla}p + \nu\Delta\boldsymbol{u} + \boldsymbol{f}, \tag{3.2}$$

which is the right-hand side of (3.1*a*). The reason for using the subscript $\boldsymbol{0}$ in $\boldsymbol{F}_{\boldsymbol{0}}$ becomes clear in § 3.2, where we derive the adjoint equation for travelling waves. We seek steady states $\boldsymbol{u}$ such that $\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}) = \boldsymbol{0}$.

To this end, we derive an adjoint PDE, such that, along its trajectories $\boldsymbol{u}(\tau)$, the norm $\|\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}(\tau))\|_{\mathscr{A}}$ decreases monotonically and converges to zero asymptotically. Our choice of the norm does not compromise the accuracy of the resulting equilibrium solutions of the Navier–Stokes equations: owing to the positive-definiteness of the operator $\mathscr{A}$, $\|\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u})\|_{\mathscr{A}} = 0$ if and only if $\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}) = \boldsymbol{0}$

Leaving the derivation to appendix A, the adjoint descent equations read

$$\partial_\tau \boldsymbol{u} = -\{[\boldsymbol{\nabla}\tilde{\boldsymbol{u}}'' + (\boldsymbol{\nabla}\tilde{\boldsymbol{u}}'')^{\mathrm{T}}]\boldsymbol{u} - \boldsymbol{\nabla}p'' + \nu\Delta\tilde{\boldsymbol{u}}''\}, \tag{3.3a}$$

$$\boldsymbol{\nabla}\cdot\boldsymbol{u}'' = 0, \quad \boldsymbol{\nabla}\cdot\boldsymbol{u} = 0, \tag{3.3b,c}$$

where $\boldsymbol{u}'' = \boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u})$, i.e.

$$\boldsymbol{u}'' = -\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u} - \boldsymbol{\nabla}p + \nu\Delta\boldsymbol{u} + \boldsymbol{f}, \tag{3.4}$$

and we use the shorthand notation $\tilde{\boldsymbol{u}}'' := \mathscr{A}\boldsymbol{u}''$.

Two pressure-type terms p and p'' appear in the adjoint equation. As in the case of the Navier–Stokes equation, the pressure terms play the role of the Lagrange multipliers, enforcing the divergence-free constraints on $\boldsymbol{u}$ and $\boldsymbol{u}''$ (cf. appendix A). Taking the divergence of (3.4) we get

$$\Delta p = -\boldsymbol{\nabla}\cdot[\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u}], \tag{3.5}$$

enforcing $\boldsymbol{\nabla}\cdot\boldsymbol{u}'' = 0$. Similarly, taking the divergence of (3.3*a*) yields

$$\Delta p'' = \boldsymbol{\nabla}\cdot\{[\boldsymbol{\nabla}\tilde{\boldsymbol{u}}'' + (\boldsymbol{\nabla}\tilde{\boldsymbol{u}}'')^{\mathrm{T}}]\boldsymbol{u}\}, \tag{3.6}$$

ensuring $\boldsymbol{\nabla}\cdot\boldsymbol{u} = 0$.

Owing to the periodic boundary conditions, these divergence-free constraints can be easily enforced by a projection operator in the spectral space. In appendix C, we derive the Fourier-space representation of the adjoint equations (3.3). This spectral representation closely resembles that of the Navier–Stokes equations. Therefore, an existing pseudo-spectral code for the Navier–Stokes equations can also be used, with straightforward modifications, to solve the adjoint equations. Furthermore, the adjoint equations (3.3) enjoy the same degree of parallelizability as the Navier–Stokes equations.

3.2. *Adjoint descent for the travelling waves*

In the absence of forcing, the Navier–Stokes equation (3.1) is invariant under Galilean translations. This symmetry allows for the existence of travelling-wave (or relative equilibrium) solutions. A travelling wave has the general form $\boldsymbol{u}(\boldsymbol{x}, t) = \boldsymbol{u}(\boldsymbol{x} - \boldsymbol{c}t)$, where $\boldsymbol{c} = (c_1, c_2, c_3) \in \mathbb{R}^3$ is a constant wave velocity. The forcing term $\boldsymbol{f}$ may break

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

the translational symmetry in one or more directions, in which case, the wave speed c_i corresponding to the symmetry-broken direction x_i is identically zero.

Substituting the travelling-wave solution $\boldsymbol{u}(\boldsymbol{x}-\boldsymbol{c}t)$ into the Navier–Stokes equation (3.1) gives the time-independent equations

$$-\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u}-\boldsymbol{\nabla}p+\nu\Delta\boldsymbol{u}+\boldsymbol{f}+\boldsymbol{c}\cdot\boldsymbol{\nabla}\boldsymbol{u}=\boldsymbol{0},\quad \boldsymbol{\nabla}\cdot\boldsymbol{u}=0. \tag{3.7a,b}$$

Therefore, finding travelling-wave solutions to the Navier–Stokes equation amounts to finding the kernel of the operator

$$\boldsymbol{F}_{\boldsymbol{c}}(\boldsymbol{u})=-\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u}-\boldsymbol{\nabla}p+\nu\Delta\boldsymbol{u}+\boldsymbol{f}+\boldsymbol{c}\cdot\boldsymbol{\nabla}\boldsymbol{u}. \tag{3.8}$$

If the wave velocity $\boldsymbol{c}$ is known, the solutions to $\boldsymbol{F}_{\boldsymbol{c}}(\boldsymbol{u})=\boldsymbol{0}$ can be found through an adjoint descent equation similar to (3.3). In general, however, the wave velocity $\boldsymbol{c}$ is unknown. To address this general case, we allow $\boldsymbol{c}$ to be a function of the fictitious time τ, and enforce its derivative $\dot{\boldsymbol{c}}=\mathrm{d}\boldsymbol{c}/\mathrm{d}\tau$ to change in such a way that $\|\boldsymbol{F}_{\boldsymbol{c}(\tau)}(\boldsymbol{u}(\tau))\|_{\mathscr{A}}$ decreases monotonically to zero as τ increases.

Leaving the details to appendix B, we obtain the set of adjoint descent equations

$$\partial_\tau\boldsymbol{u}=-\{[\boldsymbol{\nabla}\tilde{\boldsymbol{u}}''+(\boldsymbol{\nabla}\tilde{\boldsymbol{u}}'')^{\mathrm{T}}]\boldsymbol{u}-\boldsymbol{\nabla}p''+\nu\Delta\tilde{\boldsymbol{u}}''\}+\boldsymbol{c}\cdot\boldsymbol{\nabla}\tilde{\boldsymbol{u}}'', \tag{3.9a}$$

$$\frac{\mathrm{d}\boldsymbol{c}}{\mathrm{d}\tau}=-\int_{\mathscr{D}}(\boldsymbol{\nabla}\boldsymbol{u})^{\mathrm{T}}\tilde{\boldsymbol{u}}''\,\mathrm{d}\boldsymbol{x}, \tag{3.9b}$$

$$\boldsymbol{\nabla}\cdot\boldsymbol{u}''=0,\quad \boldsymbol{\nabla}\cdot\boldsymbol{u}=0, \tag{3.9c,d}$$

where $\boldsymbol{u}''=\boldsymbol{F}_{\boldsymbol{c}}(\boldsymbol{u})$, i.e.

$$\boldsymbol{u}''=-\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u}-\boldsymbol{\nabla}p+\nu\Delta\boldsymbol{u}+\boldsymbol{f}+\boldsymbol{c}\cdot\boldsymbol{\nabla}\boldsymbol{u} \tag{3.10}$$

and $\tilde{\boldsymbol{u}}''=\mathscr{A}\boldsymbol{u}''$. The pressure-type functions p and p'' satisfy (3.5) and (3.6), respectively, with $\boldsymbol{u}''$ given by (3.10). The spectral representation of the adjoint equations (3.9) is given in appendix C.

The adjoint descent for travelling waves includes the evolution equation (3.9*b*) for the unknown wave velocity $\boldsymbol{c}$. If the symmetry in the x_i direction is broken due to the forcing term $\boldsymbol{f}$, the corresponding wave speed c_i is identically zero and the equation for $\dot{c}_i$ is eliminated from (3.9*b*).

3.3. *Choice of the norm*

The conventional choice for the norm $\|\cdot\|_{\mathscr{A}}$ is the L^2 norm, corresponding to $\mathscr{A}$ being the identity operation. We find, however, that the resulting equations are stiff, requiring very small time steps for their numerical integration. This is in line with a similar observation made by Yang & Lakoba (2007) in the context of nonlinear wave equations. Here we choose the H^{-1} norm that renders the adjoint PDEs less stiff by damping the higher Fourier modes.

This choice is somewhat arbitrary, but we find it to work very well for the Kolmogorov flow presented in § 4.1. Moreover, the implementation of the resulting equations in the Fourier space is straightforward.

The H^{-1} can be computed as follows. Let $\boldsymbol{q}=(\boldsymbol{u},p)$ be the velocity–pressure pair and define the operator $\mathscr{A}$ through its action in the Fourier space,

$$\widehat{\mathscr{A}\boldsymbol{q}}(\boldsymbol{k})=\frac{\widehat{\boldsymbol{q}}(\boldsymbol{k})}{1+|\boldsymbol{k}|^2}, \tag{3.11}$$

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

where $\widehat{\boldsymbol{q}}$ denotes the Fourier transform of $\boldsymbol{q}$ and $\boldsymbol{k}=(2\pi k_1/L_1, 2\pi k_2/L_2, 2\pi k_3/L_3)$, with $k_i\in\mathbb{Z}$ being the wavenumber. Note that the operator $\mathscr{A}$ commutes with the divergence operation, i.e. $\mathscr{A}(\boldsymbol{\nabla}\cdot\boldsymbol{u})=\boldsymbol{\nabla}\cdot\mathscr{A}\boldsymbol{u}=0$.

With this choice and using Parseval's identity, the inner product (2.2) can be written explicitly in the Fourier space as

$$\langle \boldsymbol{q}, \boldsymbol{q}'\rangle_{H^{-1}} = \sum_{\boldsymbol{k}} \frac{\widehat{\boldsymbol{q}}(\boldsymbol{k})\cdot\widehat{\boldsymbol{q}}'(\boldsymbol{k})^*}{1+|\boldsymbol{k}|^2}, \tag{3.12}$$

where * denotes complex conjugation. This inner product induces the H^{-1} norm $\|\boldsymbol{q}\|^2_{H^{-1}}=\langle \boldsymbol{q}, \boldsymbol{q}\rangle_{H^{-1}}$.

The operator $\mathscr{A}$ appears in the adjoint equations (3.3) and (3.9) through $\tilde{\boldsymbol{u}}''=\mathscr{A}\boldsymbol{u}''$, which can be readily computed in the Fourier space through (3.11).

4. Equilibria and travelling waves for Kolmogorov flow

We test the performance of the adjoint descent equations on a two-dimensional Kolmogorov flow. This flow (with identical geometry and parameters as the ones chosen here) was studied thoroughly by Chandler & Kerswell (2013), providing a basis for comparison with our results.

In § 4.3, we introduce the hybrid adjoint–Newton method, which requires the Newton–GMRES–hook (or NGh, for short) iterations. The NGh iterations are explained in detail by Chandler & Kerswell (2013). Our implementation and choice of parameters are identical to theirs with the exception that we use the velocity formulations of the Kolmogorov flow, while Chandler & Kerswell (2013) use the vorticity formulation. We chose to use the velocity formulation since the adjoint equations (3.3) and (3.9) are derived for the general three-dimensional flow, and hence are in the velocity–pressure form.

4.1. *Kolmogorov flow*

Considered the forced Navier–Stokes equation (3.1) on the two-dimensional torus $\boldsymbol{x}=(x_1, x_2)\in\mathbb{T}^2=[0, 2\pi]\times[0, 2\pi]$ with the forcing $\boldsymbol{f}(\boldsymbol{x})=\sin(nx_2)\boldsymbol{e}_1$, where $\boldsymbol{e}_1=(1, 0)^{\mathrm{T}}$ and n is a positive integer. This choice of forcing yields the two-dimensional Kolmogorov equation,

$$\partial_t\boldsymbol{u} = -\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u} - \boldsymbol{\nabla}p + \nu\Delta\boldsymbol{u} + \sin(nx_2)\boldsymbol{e}_1, \tag{4.1a}$$

$$\boldsymbol{\nabla}\cdot\boldsymbol{u}=0, \tag{4.1b}$$

where $\boldsymbol{u}=(u_1, u_2)$ and $\nu=1/Re$ with Re being the Reynolds number.

Here, we review some of the relevant properties of this equation that can also be found in Platt, Sirovich & Fitzmaurice (1991) and Chandler & Kerswell (2013). The Kolmogorov equation has an equilibrium solution that can be expressed analytically as

$$u_1(x_1, x_2)=\frac{Re}{n^2}\sin(nx_2), \quad u_2(x_1, x_2)=0, \tag{4.2a,b}$$

at any Reynolds number Re. Following Chandler & Kerswell (2013), we refer to this solution as the laminar state E_0.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

The instantaneous energy E, energy dissipation D and energy input I of a state $\boldsymbol{u}$ are defined, respectively, by

$$E(t) = \frac{1}{2(2\pi)^2} \iint_{\mathbb{T}^2} |\boldsymbol{u}(\boldsymbol{x}, t)|^2 \,\mathrm{d}\boldsymbol{x}, \tag{4.3a}$$

$$D(t) = \frac{1}{(2\pi)^2 Re} \iint_{\mathbb{T}^2} |\nabla \boldsymbol{u}(\boldsymbol{x}, t)|^2 \,\mathrm{d}\boldsymbol{x} = \frac{1}{(2\pi)^2 Re} \iint_{\mathbb{T}^2} |\omega(\boldsymbol{x}, t)|^2 \,\mathrm{d}\boldsymbol{x}, \tag{4.3b}$$

$$I(t) = \frac{1}{(2\pi)^2} \iint_{\mathbb{T}^2} u_1(x_1, x_2, t) \sin(n x_2) \,\mathrm{d}\boldsymbol{x}, \tag{4.3c}$$

where $|\cdot|$ denotes the usual Euclidean norm and ω is the non-zero component of the vorticity field $\nabla \times \boldsymbol{u}$. One can readily show that

$$\frac{\mathrm{d}E}{\mathrm{d}t} = I - D, \tag{4.4}$$

along an arbitrary solution of the Navier–Stokes equation. For equilibrium and travelling-wave solutions, the energy is time-independent and therefore $I = D = \text{const}$. For the laminar solution E_0, for instance, we have

$$E_{lam} = \frac{Re^2}{4n^4}, \quad D_{lam} = I_{lam} = \frac{Re}{2n^2}. \tag{4.5a,b}$$

For $n = 1$, the laminar state is the global attractor for all Reynolds numbers, and therefore all solutions converge to it asymptotically (see e.g. Marchioro 1986; Foias *et al.* 2001, § 3.1). For $n > 1$ and large enough Reynolds numbers, the laminar state becomes unstable. The numerical study of Platt *et al.* (1991) suggests that, for $n = 4$ and large enough Reynolds numbers, all invariant solutions of the Kolmogorov equation become unstable, resulting in a chaotic attractor. More recently, Chandler & Kerswell (2013) confirmed this observation. To be in agreement with these studies, we also choose the forcing wavenumber $n = 4$. Our numerical results are carried out at $Re = 40$, unless otherwise mentioned explicitly. Chandler & Kerswell (2013) also carry out their most exhaustive search for invariant solutions at $Re = 40$.

4.2. *Temporal and spatial discretization*

To evaluate the right-hand side of the Navier–Stokes equation (4.1) and the adjoint equations (3.3) and (3.9), we use a standard pseudo-spectral method with 2/3 dealiasing. Chandler & Kerswell (2013) consider $Re = 40$, 60, 80 and 100 using uniform 256×256 Fourier modes for all Reynolds numbers. For $Re = 40$, this resolution is unnecessarily high. Instead, we use 128×128 Fourier modes for $Re = 40$. The resulting energy spectra (cf. figure 1), and the fact that we have been able to reproduce the (relative) equilibria found by Chandler & Kerswell (2013), reassures that 128^2 modes are sufficient.

For the time integration of the adjoint equations (3.3) and (3.9), one can take an explicit adaptive Euler time step. Denoting the right-hand side of either adjoint equation by $\boldsymbol{G}$, the ith Euler step reads

$$\boldsymbol{u}_{i+1} = \boldsymbol{u}_i + \delta\tau_i \boldsymbol{G}(\boldsymbol{u}_i), \tag{4.6}$$

where the length $\delta\tau_i$ of the time step is chosen small enough such that the residue $\|\boldsymbol{F}(\boldsymbol{u}_{i+1})\|_{\mathscr{A}}$ is less than $\|\boldsymbol{F}(\boldsymbol{u}_i)\|_{\mathscr{A}}$.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

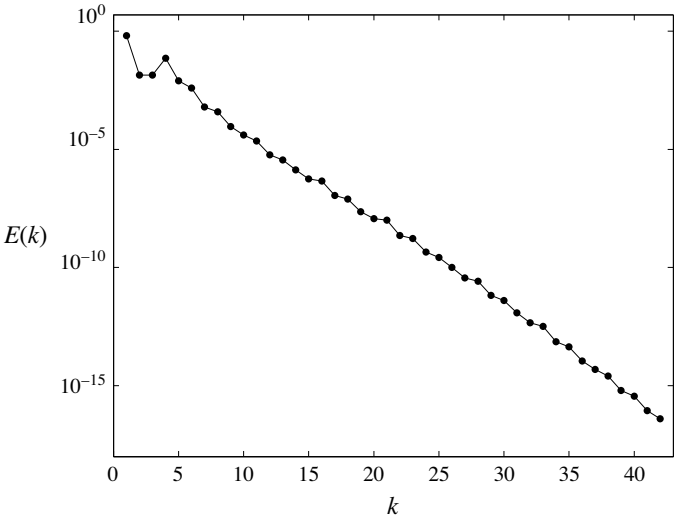


FIGURE 1. Typical energy spectrum of a turbulent trajectory at $Re = 40$.

Yang & Lakoba (2007) derive an upper bound for the admissible time step $\delta\tau_i$. Computing this upper bound for the Navier–Stokes equation is, however, not straightforward. In practice, therefore, one starts with a large value $\delta\tau_i$ and decreases it incrementally until the residue $\|\boldsymbol{F}(\boldsymbol{u}_{i+1})\|_{\mathscr{A}}$ reduces compared to the previous iteration.

To gain accuracy, however, we use a higher-order numerical scheme for our time integrations. This is computationally feasible due to the low cost of evaluating the right-hand sides of (3.3) and (3.9) by the pseudo-spectral method. More specifically, we use the embedded Runge–Kutta scheme `RK5(4)` of Dormand & Prince (1980). This scheme allows for an automatic adaptive time step size. Roughly speaking, `RK5(4)` takes fourth- and fifth-order Runge–Kutta (RK) steps. The fifth order is eventually used for the time stepping. The fourth-order prediction is used to adjust the step size as follows. Let `err` denote the absolute difference between the fourth- and fifth-order predictions. Then the step size is chosen such that $\texttt{err} < \texttt{atol} + |\boldsymbol{u}_i|\texttt{rtol}$. The prescribed constants `atol` and `rtol` are the absolute and relative errors, respectively. We refer the reader to Press *et al.* (2007, § 17.2) for further details and an implementation of the `RK5(4)` scheme. This integrator is also implemented in the MATLAB function `ode45`. For integrating the adjoint equations, we choose $\texttt{atol} = \texttt{rtol} = 10^{-10}$. Even with this conservative choice, step sizes as large as 10 fictitious-time units were taken by `RK5(4)`.

We also use `RK5(4)` for the temporal discretization of the Navier–Stokes equation (4.1). However, as the time step sizes were much smaller for integrating this equation, we used the less stringent error tolerances $\texttt{atol} = \texttt{rtol} = 10^{-5}$.

4.3. *Hybrid adjoint–Newton iterations*

We find that the adjoint equation (3.3) does take arbitrary initial conditions to Navier–Stokes equilibria. The convergence is, however, rather slow. To demonstrate this, we take the initial condition $\boldsymbol{u}(\boldsymbol{x}, 0) = (\cos(2x_2), \cos(x_1))$ (see figure 2*a*) and evolve it under the adjoint equation (3.3) to time $\tau = 500$. The result is shown in

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

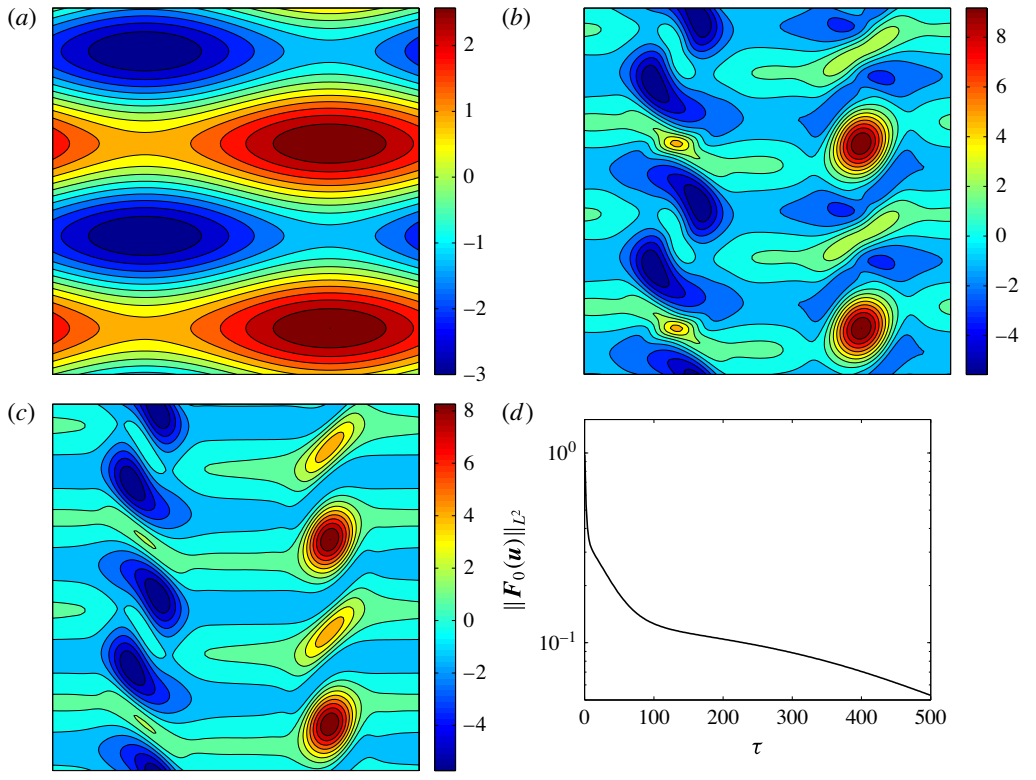


FIGURE 2. (Colour online) (*a*) The vorticity field for the initial condition $\boldsymbol{u}(\boldsymbol{x}, 0) = (\cos(2x_2), \cos(x_1))^{\mathrm{T}}$. (*b*) The initial condition evolved to time $\tau = 500$ under the adjoint equation (3.3) to obtain $\boldsymbol{u}(\boldsymbol{x}, 500)$. (*c*) The result $\boldsymbol{u}(\boldsymbol{x}, 500)$ shown in panel (*b*) is used as the initial guess for the Newton–GMRES–hook iterations, which converged after seven iterations with residue 5.98×10^{-11}. This equilibrium solution is labelled in table 1 as E_4. (*d*) The evolution of the L^2 residue $\|\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}(\cdot, \tau))\|_{L^2}$ as the initial condition in panel (*a*) evolves under the adjoint equation from $\tau = 0$ to $\tau = 500$.

figure 2(*b*). This integration took 54 s (on an iMac with a single processor: Intel Core i5, 2.9 GHz). The L^2 norm of the residue $\boldsymbol{F}_{\boldsymbol{0}}$, however, decreases to $\simeq 5 \times 10^{-2}$ approximately, indicating that much longer integration times are required to reduce the residue sufficiently enough, say to 10^{-10}.

This slow convergence is due to nearly marginal, *a priori* unknown eigenmodes of the adjoint operator (Lakoba & Yang 2007). For particular wave equations, *ad hoc* methods have been proposed to eliminate these modes, and hence speed up the convergence (Yang & Lakoba 2007). Owing to the complexity of the Navier–Stokes equations, it is unclear how such mode elimination techniques could be employed here.

Using the state $\boldsymbol{u}(\boldsymbol{x}, 0) = (\cos(2x_2), \cos(x_1))$ directly as the initial guess for the NGh iterations does not converge to an equilibrium either: after the first 20 iterations, the residue plateaued around 2×10^{-2} and remained so for the 75 iterations that were carried out.

Instead, when we use the result of the adjoint integration, i.e. $\boldsymbol{u}(\boldsymbol{x}, \tau = 500)$, as the initial guess for NGh iterations, it converges after seven iterations with residue 5.98×10^{-11} (figure 2*c*). These seven iterations took 156.84 s.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

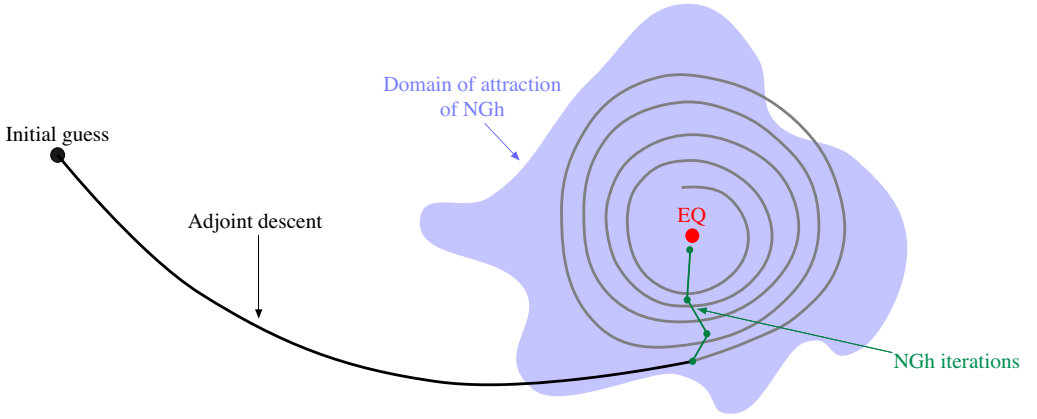


FIGURE 3. (Colour online) State space depiction of the hybrid adjoint–Newton method. An initial guess may be too far away from the desired equilibrium solution EQ for the Newton–GMRES–hook iterations (NGh) to converge to it. The adjoint descent may eventually converge to EQ, but alone it would take a long integration time to do so. In the hybrid adjoint–Newton method the adjoint descent takes the initial guess to the domain of attraction of NGh. Once there, NGh converges to EQ in a few iterations.

This suggests that, although the convergence of the adjoint equation to the equilibrium is slow, it takes generic initial guesses to the vicinity of the equilibrium at a relatively low computational cost. This can be seen by comparing panels (*b*) and (*c*) of figure 2, which shows that the state $\boldsymbol{u}(\boldsymbol{x}, \tau = 500)$ is quite similar to the equilibrium E_4, and thus likely to lie in its NGh domain of attraction. Indeed, switching to NGh took this state to the equilibrium solution within a few iterations.

Based on this observation, we propose the following hybrid adjoint–Newton iterations, sketched in figure 3. We take a prescribed positive real number τ_0 and a positive integer `N`. An initial condition $\boldsymbol{u}(\boldsymbol{x}, 0)$ is then integrated under the adjoint equation (3.3) to obtain $\boldsymbol{u}(\boldsymbol{x}, \tau_0)$. This state is then fed into the NGh algorithm and `N` iterations of NGh are performed. The output is then fed back into the adjoint equation as the initial condition. This hybrid loop is repeated until the residue $\|\boldsymbol{F_0}\|_{L^2}$ decreases below some prescribed tolerance `tol`. This procedure is summarized in Algorithm 1. The hybrid adjoint–Newton iterations for finding travelling waves are similar except that the adjoint equations (3.9) are solved at each iteration (see Algorithm 2).

For the computations reported below, we set $\tau_0 = 100$, $\texttt{tol} = 10^{-10}$ and $\texttt{N} = 1$. Since NGh steps are relatively expensive, we only take one NGh step ($\texttt{N} = 1$) per iteration of adjoint–Newton. At Reynolds number $Re = 40$, for instance, the integration of the adjoint equations to $\tau = 100$ took approximately 10 s on average while each NGh step took approximately 50 s on average.

4.4. *Equilibrium solutions*

In this section, we report the equilibria found by the hybrid adjoint–Newton iterations (Algorithm 1). We also compare its performance with that of pure NGh iterations without the adjoint step.

To this end, we consider the set of initial guesses

$$\boldsymbol{u}_0^{(m_1,m_2)} = (\cos(m_2 x_2), \cos(m_1 x_1)), \tag{4.7}$$

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

Algorithm 1: Hybrid adjoint–Newton algorithm for finding equilibrium solutions of the Navier–Stokes equation.

Input: $\tau_0, \mathtt{tol} \in \mathbb{R}^+$, $\mathtt{N} \in \mathbb{N}$, state $\boldsymbol{u}_0$
while $\|\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}_0)\|_{L^2} > \mathtt{tol}$ **do**
  Integrate the adjoint equation (3.3) for τ_0 fictitious-time units with the initial condition $\boldsymbol{u}(\cdot, 0) = \boldsymbol{u}_0$
  $\boldsymbol{u}_0 \longleftarrow \boldsymbol{u}(\cdot, \tau_0)$
  if $\|\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}_0)\|_{L^2} < \mathtt{tol}$ **then**
    STOP
  for $(k = 1$ **to** $\mathtt{N})$ **do**
    Take one Newton–GMRES–hook step to get NGh($\boldsymbol{u}_0$)
    $\boldsymbol{u}_0 \longleftarrow$ NGh($\boldsymbol{u}_0$)
    if $\|\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{u}_0)\|_{L^2} < \mathtt{tol}$ **then**
      STOP
Output: $\boldsymbol{u}_0$

Algorithm 2: Hybrid adjoint–Newton algorithm for finding travelling-wave solutions of the Navier–Stokes equation.

Input: $\tau_0, \mathtt{tol} \in \mathbb{R}^+$, $\mathtt{N} \in \mathbb{N}$, state $\boldsymbol{u}_0$, wave speed $\boldsymbol{c}_0$
while $\|\boldsymbol{F}_{c_0}(\boldsymbol{u}_0)\|_{L^2} > \mathtt{tol}$ **do**
  Integrate the adjoint equation (3.9) for τ_0 fictitious-time units with the initial conditions $\boldsymbol{u}(\cdot, 0) = \boldsymbol{u}_0$ and $\boldsymbol{c}(0) = \boldsymbol{c}_0$
  $\boldsymbol{u}_0 \longleftarrow \boldsymbol{u}(\cdot, \tau_0)$
  $\boldsymbol{c}_0 \longleftarrow \boldsymbol{c}(\tau_0)$
  if $\|\boldsymbol{F}_{c_0}(\boldsymbol{u}_0)\|_{L^2} < \mathtt{tol}$ **then**
    STOP
  for $(k = 1$ **to** $\mathtt{N})$ **do**
    Take one Newton–GMRES–hook step to get NGh($\boldsymbol{u}_0, \boldsymbol{c}_0$)
    $(\boldsymbol{u}_0, \boldsymbol{c}_0) \longleftarrow$ NGh($\boldsymbol{u}_0, \boldsymbol{c}_0$)
    if $\|\boldsymbol{F}_{c_0}(\boldsymbol{u}_0)\|_{L^2} < \mathtt{tol}$ **then**
      STOP
Output: $\boldsymbol{u}_0, \boldsymbol{c}_0$

for a range of integers m_1 and m_2. We find that for $m_2 > 4$, both Newton–GMRES–hook and the hybrid adjoint–Newton iterations converge to the laminar equilibrium E_0. Therefore, we restrict the range of the integers to $1 \leqslant m_1, m_2 \leqslant 4$. This leads to 16 distinct initial guesses.

These initial conditions are divergence-free by construction, consistent with our incompressible Navier–Stokes setting. They are also explicit, rendering the following results reproducible. Furthermore, the generic nature of the initial conditions illustrates the main advantage of the adjoint descent method, namely that its convergence does not require good initial guesses. For brevity, we shall refer to the states $\boldsymbol{u}_0^{(m_1,m_2)}$ by their indices (m_1, m_2).

Using the hybrid adjoint–Newton iterations for equilibria (see Algorithm 1), all 16 runs converged, resulting in 10 distinct equilibria. These equilibria are listed as E_0 to E_9 in table 1, where E_0 is the laminar state (4.2). All initial guesses converged

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

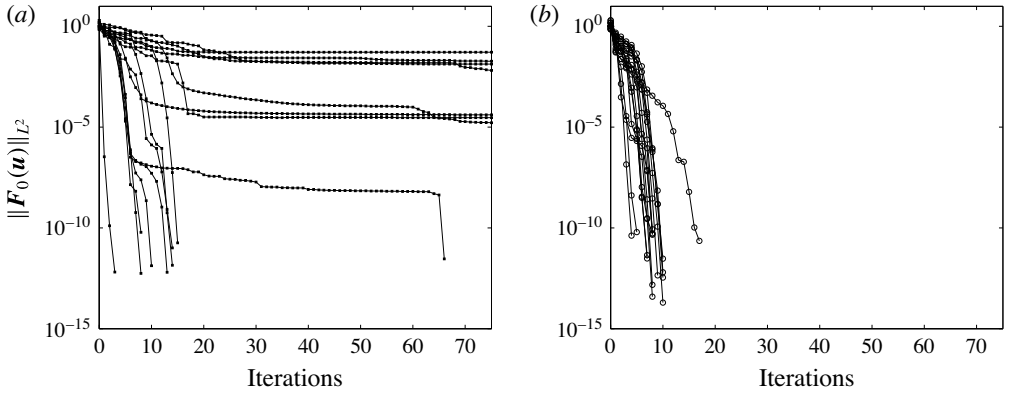


FIGURE 4. The decay of the residue (i.e. the L^2-norm of the right-hand side) using Newton–GMRES–hook iterations (*a*) and the hybrid adjoint–Newton iterations (*b*).

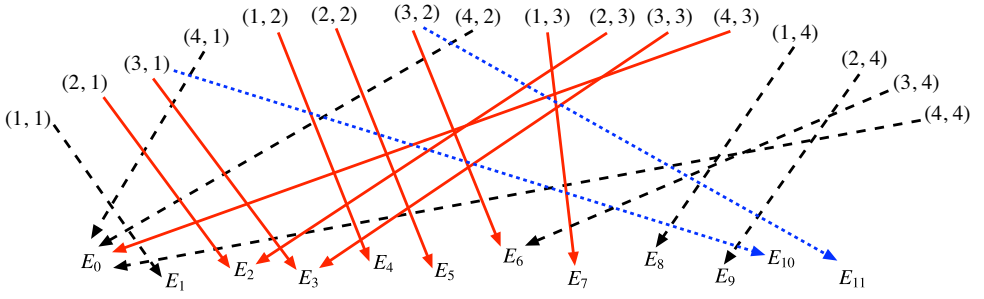


FIGURE 5. (Colour online) The convergence diagram. Double indices (m_1, m_2) label the initial guesses (4.7). A dashed black arrow indicates that both the hybrid adjoint–Newton and the NGh iterations converge to the same equilibrium. A dotted blue arrow indicates that the NGh iterations converge to a different equilibrium from the one reached by the hybrid adjoint–Newton (red arrow). A solid red arrow (with no blue arrow originating from the same initial guess) indicates that only the hybrid adjoint–Newton converged to the equilibrium, i.e. the NGh iterations alone failed to converge.

to an equilibrium within the first 10 iterations of hybrid adjoint–Newton iterations, except initial guess (2, 1) which took 17 iterations (see figure 4). Figure 5 shows the convergence diagram, connecting each initial guess (m_1, m_2) to the resulting equilibrium solution.

Using NGh alone, only nine out of 16 runs converged. They converged to seven distinct equilibria: E_0, E_1, E_6, E_8, E_9, E_{10} and E_{11} (see figure 5). This comes initially as a surprise since the exhaustive search carried out by Chandler & Kerswell (2013) only returned a single equilibrium (i.e. E_1). This can be explained, however, in terms of the method used for initiating the NGh iterations. Chandler & Kerswell (2013) use recurrences to find initial guesses for their searches, as opposed to the generic initial guesses used here. Recurrences only happen in a subset of the state space where a generic turbulent trajectory spends most of its lifetime. As a result, using recurrences for initiating the NGh searches might preferentially yield the equilibria 'close' to this region. We will return to this subject in § 5, where the temporal intermittency is discussed.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

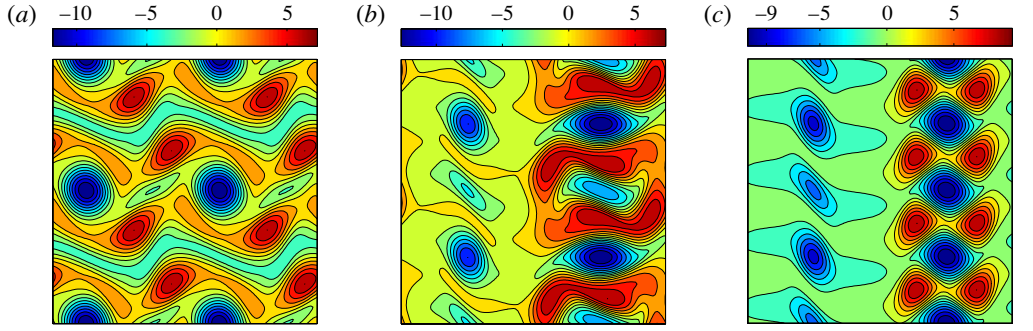


FIGURE 6. (Colour online) Vorticity fields of the equilibrium solutions E_5 (*a*), E_{12} (*b*) and E_{15} (*c*). All panels show the entire domain $[0, 2\pi] \times [0, 2\pi]$.

Equilibrium	$I = D$	E	μ_1	ω_1	dim E^u
E_0	1.25000	1.56250	2.35340	0.0	38
E_1	0.12732	0.76108	0.24929	2.56010	9
E_2	0.15406	0.47313	0.66722	0.0	13
E_3	0.26287	0.45173	0.84744	0.0	27
E_4	0.08433	0.57317	0.62697	0.0	5
E_5	0.26661	0.55183	0.30119	0.0	19
E_6	0.26227	0.45574	0.86924	0.0	22
E_7	0.07530	0.58104	0.58318	0.0	5
E_8	0.17452	0.53493	0.61189	0.00675	17
E_9	0.15315	0.47396	0.67006	0.0	17
E_{10}	0.61437	0.91740	0.85655	1.75030	30
E_{11}	0.48020	0.71284	0.81175	0.0	29
E_{12}	0.31049	0.89032	0.68928	0.0	13
E_{13}	0.26151	0.50921	0.42190	1.49660	16
E_{14}	0.31152	0.54066	0.71722	0.0	21
E_{15}	0.27070	0.76570	0.87760	0.0	14
E_{16}	0.34954	0.61168	0.83092	0.02382	24

TABLE 1. List of equilibrium solutions at $Re = 40$. Energy E, energy dissipation D and energy input I are defined in (4.3). The leading unstable eigenvalue of the equilibrium is $\mu_1 + \mathrm{i}\omega_1$. The dimension of the linear unstable manifold of the equilibrium is denoted by dim E^u.

We also sought further equilibria using states of the form $(\sin(m_2 x_2), \cos(m_1 x_1))$ as the initial guesses for the hybrid adjoint–Newton iterations. This resulted in five more equilibria: E_{12} to E_{16} in table 1. While all of these additional searches did converge, most of them re-converged to previously discovered equilibria, including E_{10} and E_{11} that were only found by NGh iterations when initial guesses (4.7) were used. Figure 6 shows the vorticity field of three select equilibrium solutions.

Some of the equilibria (e.g. E_{15} shown in figure 6) exhibit vertical bands of localized vorticity that are separated by an almost zero-vorticity background. Such spatially localized equilibria of the Kolmogorov flow were only observed previously on domains with small aspect ratio $\alpha = L_2/L_1$ (Lucas & Kerswell 2014). The fact that they also exist on a domain with aspect ratio $\alpha = 1$ comes as a surprise.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

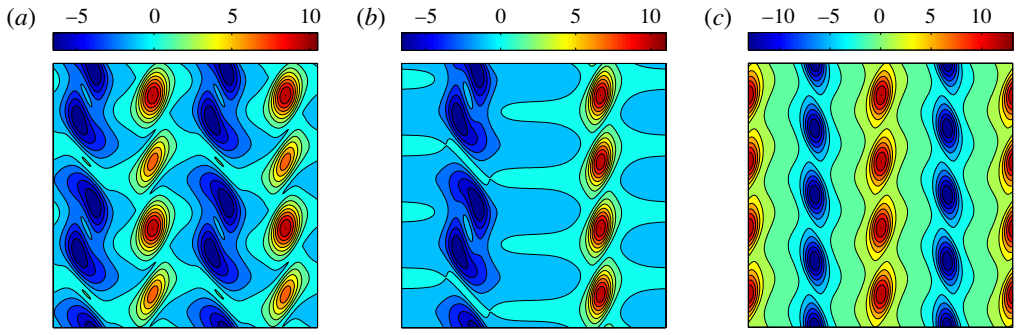


FIGURE 7. (Colour online) Vorticity fields of the equilibrium solutions at Reynolds number $Re = 60$ (a), $Re = 80$ (b) and $Re = 100$ (c). All panels show the entire $[0, 2\pi] \times [0, 2\pi]$ domain.

Although our focus here is on Reynolds number $Re = 40$, figure 7 showcases select equilibria at $Re = 60$, 80 and 100. These equilibria are computed using the higher resolution of 256×256 Fourier modes. They are found by our hybrid adjoint–Newton method while the previous studies using Newton–GMRES–hook iterations had not been able to discover them (Chandler & Kerswell 2013).

4.5. *Travelling-wave solutions*

The forcing term in the Kolmogorov equation (4.1) breaks the continuous symmetry in the x_2 direction. Therefore, only travelling-wave solutions of the type $\boldsymbol{u}(\boldsymbol{x} - \boldsymbol{c}t)$ with $\boldsymbol{c} = (c, 0)$ are permitted. This reduces the wave velocity equations (3.9b) to the scalar equation

$$\frac{\mathrm{d}c}{\mathrm{d}\tau} = -\int_{\mathbb{T}^2} \frac{\partial \boldsymbol{u}}{\partial x_1} \cdot \tilde{\boldsymbol{u}}'' \, \mathrm{d}^2\boldsymbol{x}. \tag{4.8}$$

Similarly, the term $\boldsymbol{c} \cdot \nabla \tilde{\boldsymbol{u}}''$ in (3.9a) reduces to $c\, \partial_{x_1} \tilde{\boldsymbol{u}}''$ and the term $\boldsymbol{c} \cdot \nabla \boldsymbol{u}$ in (3.10) reduces to $c\, \partial_{x_1} \boldsymbol{u}$.

We search for travelling waves using Algorithm 2 and the generic initial guesses discussed in the previous section. For the initial wave speed, we used $c(0) = 1$. Other values of $c(0)$ yielded similar results.

Some of our searches for travelling waves converged to equilibrium solutions. This is to be expected, as equilibria are degenerate travelling waves with wave speed $c = 0$. In fact, the adjoint equation (3.9) admits such solutions. This can be readily verified by letting $\boldsymbol{u}(\boldsymbol{x})$ be an equilibrium solution of the Navier–Stokes equation and setting $\boldsymbol{c} \equiv 0$. Then $\boldsymbol{u}''$ in (3.10) is zero, resulting in vanishing right-hand sides in (3.9).

Nonetheless, our hybrid adjoint–Newton searches led to nine distinct travelling waves listed in table 2. Only travelling waves T_1 and T_2 had been discovered previously (Chandler & Kerswell 2013). Figure 8 shows the vorticity field for three select travelling waves. As in the case of equilibria, we find that some of the travelling waves (e.g. T_8 in figure 8) exhibit localized spatial structures, although the domain aspect ratio is one.

5. Temporal intermittency in Kolmogorov flow

In this section, we explore the significance of the invariant solutions, found in the previous section, on the global dynamics of the Kolmogorov flow.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

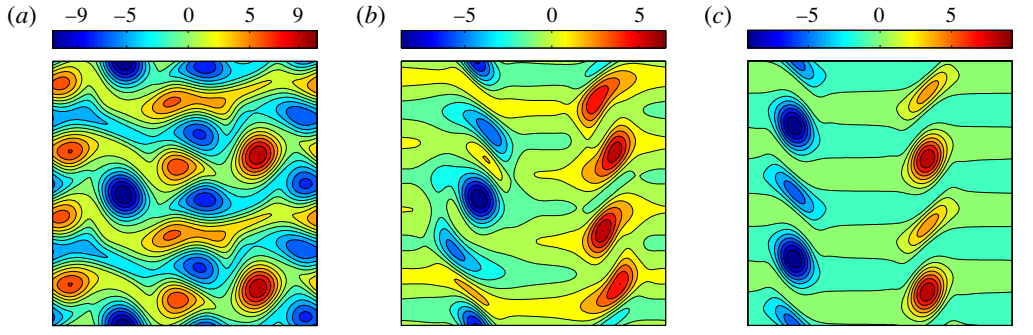


FIGURE 8. (Colour online) Vorticity field for the travelling-wave solutions T_4 (*a*), T_7 (*b*) and T_8 (*c*). All panels show the entire domain $[0, 2\pi] \times [0, 2\pi]$.

Travelling wave	c	$I = D$	E	μ_1	ω_1	$\dim E^u$
T_1	0.01978	0.08873	0.69747	0.06815	0.35451	4
T_2	0.00944	0.08680	0.63969	0.45288	0.02134	4
T_3	0.01826	0.13432	0.38056	0.49301	0.0	10
T_4	0.03062	0.31453	0.52793	0.55183	0.0	21
T_5	0.05266	0.40027	0.61651	0.82112	0.0	17
T_6	0.04223	0.32063	0.52963	0.54292	0.0	18
T_7	0.00208	0.08058	0.61200	0.50598	0.0	2
T_8	0.00156	0.09641	0.59883	0.63357	0.0	7
T_9	0.00642	0.08867	0.61764	0.57681	0.0	4

TABLE 2. List of travelling-wave solutions at $Re = 40$. The constant c denotes the wave speed. Energy E, energy dissipation D and energy input I are defined in (4.3). The leading unstable eigenvalue of the travelling wave is $\mu_1 + \mathrm{i}\omega_1$. The dimension of the linear unstable manifold of the travelling wave is denoted by $\dim E^u$.

Figure 9 shows the energy input I versus the energy dissipation D for a generic turbulent trajectory computed for 10^3 time units and recorded every 0.1 time units. The energy input and dissipation of each equilibrium and travelling wave are marked by circles and squares, respectively. As mentioned earlier, the energy input and dissipation coincide for these invariant solutions, locating them on the diagonal $I = D$.

The equilibria and travelling waves assume a wide range of energy input and dissipation. The turbulent trajectory, on the other hand, mostly resides in a relatively low energy input–dissipation regime. In particular, approximately 85 % of this trajectory belongs to the $I/I_{lam} < 0.12$ and $D/D_{lam} < 0.12$ regime, marked by the green square in figure 9. For the lack of a better term, we refer to this regime as the 'ergodic sea'.

At the same time, the turbulent trajectory also experiences sporadic episodes of high energy input and dissipation. Such rare, extreme events are usually referred to as temporal intermittency and are ubiquitous in turbulent fluid flow (see e.g. Batchelor & Townsend 1949; Sreenivasan & Antonia 1997). Non-Gaussian probability distributions of turbulent quantities are a footprint of intermittency that produces the heavy tails of the distribution functions (Frisch 1996; Mininni & Pouquet 2010).

The short lifetime of the intermittent bursts is better seen in figure 10(*a*) where the normalized energy input I/I_{lam} is shown as a function of time. The time series for

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

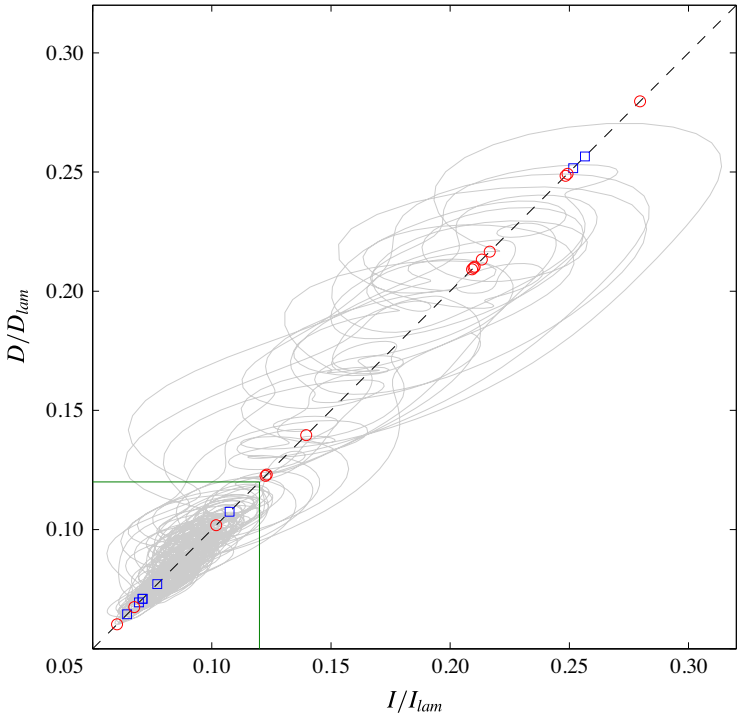


FIGURE 9. (Colour online) Results for $Re = 40$. Grey curve: turbulent trajectory spanning 10^3 time units. Red circles: equilibria. Blue squares: travelling waves. The green square marks the region where $I/I_{lam} < 0.12$ and $D/D_{lam} < 0.12$. The turbulent trajectory spends 86.62 % of the total 10^3 time units inside this region. The diagonal $I = D$ is marked by the dashed black line. Equilibria and travelling waves with $I/I_{lam} = D/D_{lam} > 0.32$ are not shown.

the energy dissipation (not shown here) is very similar, except that the intermittent bursts of the dissipation occur with a short delay of 0.5 to 1.5 time units relative to the energy input bursts. This suggests that, once in a while, the turbulent velocity field mostly aligns with the external forcing $\sin(nx_2)\boldsymbol{e}_1$, resulting in the growth of the energy input, which kicks the trajectory out of its ergodic sea. After a short time delay, the energy dissipation also increases, bringing the trajectory back to the ergodic regime.

Spatial intermittency is also a characteristic of turbulent fluid flow, which refers to unusually large velocity (or vorticity) amplitudes occurring in a relatively small subset of the physical domain (see e.g. Kuo & Corrsin 1971; Schneider, Farge & Kevlahan 2004; Ruppert-Felsot, Farge & Petitjeans 2009). Although temporal and spatial intermittencies are sometimes conflated in the literature, the relation between the two is not well understood (Gibbon & Doering 2003). For our turbulent trajectory, in fact, an appreciable correlation between them was not found. Figure 10(*b*), for instance, shows that the normalized maximum vorticity amplitude oscillates rapidly throughout the simulation time, exhibiting no clear correlation with the energy input.

Focusing on the temporal intermittency, we first review the dynamical systems perspective of this phenomena.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

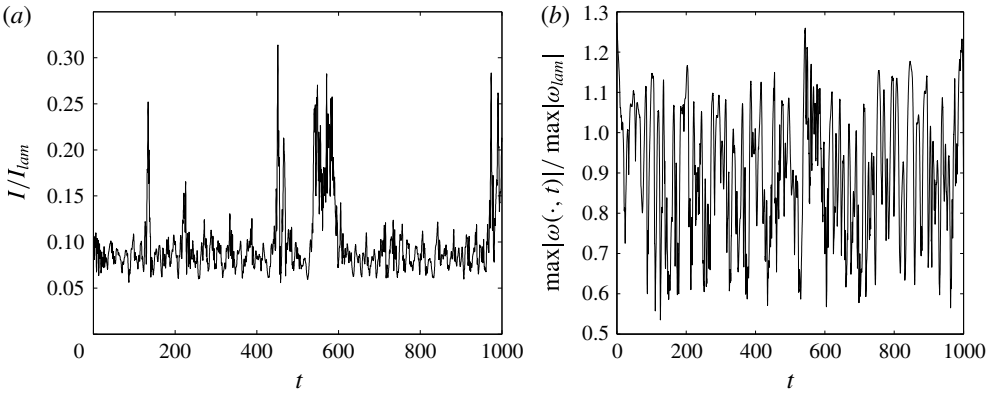


FIGURE 10. (*a*) The energy input I as a function of time normalized by the energy input of the laminar state $I_{lam} = 1.25$. (*b*) The spatial maximum of vorticity magnitude, i.e. $\max_{\boldsymbol{x}\in\mathbb{T}^2} |\omega(\boldsymbol{x}, t)|$, normalized by its value corresponding to the laminar state, i.e. $\max_{\boldsymbol{x}\in\mathbb{T}^2} |\omega_{lam}(\boldsymbol{x})| = 10$.

5.1. *A dynamical systems perspective on temporal intermittency*

Although the Navier–Stokes equations generate an infinite-dimensional dynamical system, it is believed that, owing to the dissipative term $\nu\Delta\boldsymbol{u}$, its solutions converge exponentially fast to a finite-dimensional, invariant set, usually referred to as the inertial manifold (see Constantin *et al.* (1989) for the rigorous definition). The existence of the inertial manifold for the Navier–Stokes equation, in its most general form, is an open mathematical problem. In practice, however, its existence is often assumed. In fact, this assumption underlies the finite Galerkin truncations used in computations (Foias *et al.* 1988).

Some (relative) equilibria and (relative) periodic orbits and portions of their stable and unstable manifolds belong to the inertial manifold. For large enough Reynolds numbers, these invariant solutions are typically unstable. A generic turbulent trajectory visits the neighbourhood of an invariant solution for a finite time before it is repelled along its unstable manifold towards the neighbourhood of another invariant solution (Ruelle 1991; Halcrow *et al.* 2009). This process continues indefinitely in a somewhat unpredictable fashion, thereby causing the complex temporal behaviour of turbulent trajectories (see figure 11 for an illustration).

From this perspective, intermittent episodes are viewed as close passages of the turbulent trajectory to invariant solutions that reside in a 'less accessible' region of the inertial manifold or, more precisely, the attractor (Holmes 1993). As depicted in figure 11, such passages are viable along the heteroclinic connections between the invariant solutions (Holmes & Stone 1992; Holmes *et al.* 2012).

To characterize the less accessible regions of the attractor, one naturally needs to answer the following question: how frequently is an invariant solution visited by generic turbulent trajectories? There is no straightforward, *a priori* answer to this question (Cvitanović *et al.* 2014, ch. 23). There are, however, some characteristics of the invariant solutions that are relevant. In tables 1 and 2, for instance, we report the dimension of the linear unstable manifold (i.e. $\dim E^u$) of each invariant solution. The invariant solutions E_1, E_4, E_7, T_1, T_2, T_3, T_7, T_8 and T_9 that reside close to the ergodic sea have at most 10 linearly unstable eigenmodes. The remaining invariant

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

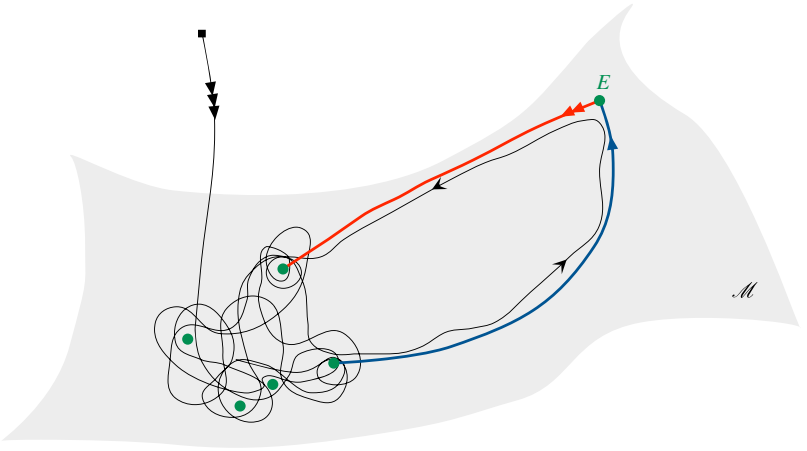


FIGURE 11. (Colour online) An initial condition (black square) decays rapidly to the inertial manifold $\mathscr{M}$ where the dynamics is governed by the (relative) equilibria, (relative) periodic orbits and portions of their stable/unstable manifolds that lie within the inertial manifold. The green dots represent equilibria. Highly unstable invariant solutions (e.g. the equilibrium E) are rarely visited by a generic trajectory.

solutions have at least 13 unstable eigenmodes and seem to reside further away from the ergodic sea.

In tables 1 and 2, we also report the stability exponent $\mu_1 + \mathrm{i}\omega_1$ of the most unstable eigenmode of each invariant solution. It is tempting to assert that the invariant solutions with larger μ_1 are less likely to be visited by a generic turbulent trajectory. This is, however, not the case. For instance, we have $\mu_1 = 0.62697$ for equilibrium E_4 and $\mu_1 = 0.61189$ for equilibrium E_8. Equilibrium E_4 is located in the heart of the ergodic sea and is, in fact, visited by the turbulent trajectory quite often. Equilibrium E_8, in spite of having a similar stability exponent, is rarely visited by the turbulent trajectory. The crucial difference between these two equilibria is the dimension of the linear unstable manifold, which is $\dim E^u = 5$ for E_4 and $\dim E^u = 17$ for E_8.

In retrospect, the lack of correlation between the stability exponent and the frequency of visitations by the turbulent trajectory is to be expected. The stability exponent of an invariant solution is a local quantity. As such, for it to be meaningful, the turbulent trajectory should already be in the vicinity of the invariant solution. Once there, the stability exponent μ_1 determines how quickly the trajectory will leave the neighbourhood.

At any rate, neither the dimension of the unstable manifold nor the stability exponents of an invariant solution decisively determine the frequency at which its neighbourhood is visited by a turbulent trajectory. Therefore, we take a more direct approach to quantify intermittency. Namely, we measure the L^2 distance between the turbulent trajectory and the computed invariant solutions.

The Kolmogorov equation is equivariant under a one-parameter family of continuous symmetries and a number of discrete symmetries (Sirovich 1987). This implies that each solution $\boldsymbol{u}(t)$ has infinitely many equivalent copies. Therefore, when measuring the 'distance' between two states, one needs to make an informed choice among the equivalent copies of each state. This necessitates a discussion on the symmetries of the Kolmogorov flow.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

5.2. *Symmetries of Kolmogorov flow*

The Kolmogorov equation (4.1) is equivariant with respect to $4n$ discrete symmetries and a continuous translational symmetry (Sirovich 1987). We denote the complete set of such symmetries by

$$(\mathscr{T}_\ell \boldsymbol{u})(x_1, x_2) = \boldsymbol{u}(x_1 + \ell, x_2), \tag{5.1a}$$

$$(\mathscr{R}\boldsymbol{u})(x_1, x_2) = -\boldsymbol{u}(-x_1, -x_2), \tag{5.1b}$$

$$(\mathscr{S}^m \boldsymbol{u})(x_1, x_2) = \begin{pmatrix} (-1)^m u_1((-1)^m x_1, x_2 + m\pi/n) \\ u_2((-1)^m x_1, x_2 + m\pi/n) \end{pmatrix}, \tag{5.1c}$$

where $m \in \{0, 1, \ldots, 2n-1\}$. Here, $\mathscr{T}_\ell$ denotes ℓ shift in the x_1 direction, $\mathscr{R}$ denotes rotation through π and $\mathscr{S}$ denotes a simultaneous (π/n) shift in the x_2 direction and a reflection in the x_1 direction.

One can readily verify that the above symmetry operations act on the vorticity field according to the following rules:

$$(\mathscr{T}_\ell \omega)(x_1, x_2) = \omega(x_1 + \ell, x_2), \tag{5.2a}$$

$$(\mathscr{R}\omega)(x_1, x_2) = \omega(-x_1, -x_2), \tag{5.2b}$$

$$(\mathscr{S}^m \omega)(x_1, x_2) = (-1)^m \omega((-1)^m x_1, x_2 + m\pi/n). \tag{5.2c}$$

The glide reflection (or shift–reflect operation) $\mathscr{S}$ generates a cyclic group of order $2n$,

$$C_{2n} = \{e, \mathscr{S}, \mathscr{S}^2, \ldots, \mathscr{S}^{2n-1}\}, \tag{5.3}$$

where e denotes the identity $e = \mathscr{S}^0$. The rotation-through-π operation $\mathscr{R}$ generates a cyclic group of order two, $R_2 = \{e, \mathscr{R}\}$. The complete set of discrete symmetries of the Kolmogorov equation, therefore, is the dihedral group of order $4n$, i.e.

$$D_{4n} = R_2 \ltimes C_{2n} = \{e, \mathscr{S}, \ldots, \mathscr{S}^{2n-1}, \mathscr{R}, \mathscr{R}\mathscr{S}, \ldots, \mathscr{R}\mathscr{S}^{2n-1}\}. \tag{5.4}$$

Note that the operations $\mathscr{R}$ and $\mathscr{S}$ do not commute, $\mathscr{R}\mathscr{S} \neq \mathscr{S}\mathscr{R}$. Instead, we have $\mathscr{S}\mathscr{R}\mathscr{S} = \mathscr{R}$.

Therefore, the solutions of the Kolmogorov equation have up to $4n$ distinct but equivalent copies due to its equivariance under the discrete symmetries D_{4n}. They also have infinitely many equivalent copies due to the continuous symmetry $\mathscr{T}_\ell$ for any $\ell \in [0, 2\pi]$.

An invariant solution may itself have some, all or none of the symmetries of the equations. The laminar solution E_0, for instance, has the complete set of symmetries, i.e. $gE_0 = E_0$ for all $g \in D_{4n}$ and $\mathscr{T}_\ell E_0 = E_0$ for all $\ell \in [0, 2\pi]$. The laminar state, therefore, has only one copy. The travelling wave T_7 (see figure 8), on the other hand, has no symmetries and therefore possesses infinitely many equivalent copies. Incidentally, travelling wave T_7 happens to have the lowest-dimensional unstable manifold, $\dim E^u = 2$, among the solutions found here.

These symmetry-related copies greatly complicate the analysis of the state space of the Kolmogorov flow. When comparing the L^2 distance between two states $\boldsymbol{u}^1$ and $\boldsymbol{u}^2$, one needs to take the minimum L^2 distance between $\boldsymbol{u}^1$ and $\boldsymbol{u}^2$ and all their symmetry-related copies. For example, let $\boldsymbol{u}^2(t)$ be a symmetry copy of a solution $\boldsymbol{u}^1(t)$ such that

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

$\boldsymbol{u}^2(t) = (\mathscr{T}_\ell \boldsymbol{u}^1)(t)$ for some $\ell \in (0, 2\pi)$. These two states are equivalent and both solve the Kolmogorov equation. However, the norm $\|\mathscr{T}_\ell \boldsymbol{u}^1 - \boldsymbol{u}^1\|_{L^2}$ is generally non-zero. Therefore, an appropriate norm on the state space of the Kolmogorov flow is

$$\min \|\boldsymbol{u}^1 - g\mathscr{T}_\ell \boldsymbol{u}^2\|_{L^2}, \tag{5.5}$$

where the minimum is taken over all $\ell \in [0, 2\pi]$ and $g \in D_{4n}$. Evaluating this norm can be somewhat cumbersome.

Another approach is to map each state into a symmetry-invariant polynomial basis. For low-dimensional dynamical systems with simple discrete symmetries, such coordinates are available analytically (Gilmore & Letellier 2007). As the dimension of the system (and/or the order of the group) increases, the determination of the invariant coordinates becomes quickly prohibitive (Siminos & Cvitanović 2011). As a result, and to the best of our knowledge, symmetry-invariant polynomial coordinates for the Kolmogorov flow are not known.

Here, we take an alternative approach, which also proves to be insightful in analysis of the temporal intermittency. We define the projection operator

$$\mathscr{P}\boldsymbol{u} = \frac{1}{4n}\sum_{m=0}^{2n-1} \mathscr{S}^m(\boldsymbol{u} + \mathscr{R}\boldsymbol{u}), \tag{5.6}$$

which is the average over all copies of $\boldsymbol{u}$ given by the discrete symmetries D_{4n} (Cvitanović *et al.* 2014). The projection $\mathscr{P}\boldsymbol{u}$ is invariant under all discrete symmetries $g \in D_{4n}$ and therefore we refer to it as the symmetric part of the state $\boldsymbol{u}$. The symmetric part of vorticity ω is defined analogously.

All $4n$ symmetry copies of the state $\boldsymbol{u}$ have a unique projection $\mathscr{P}\boldsymbol{u}$. Working with the symmetric part of the states, therefore, eliminates the complications arising from the discrete symmetries.

Considering this symmetric part also has a physical motivation. The energy E, dissipation D and energy input I defined in (4.3) are invariant under symmetry operations. Denoting the energy input of a state $\boldsymbol{u}$ by $I[\boldsymbol{u}]$, we have $I[\boldsymbol{u}] = I[g\boldsymbol{u}]$ for g being any symmetry of the Kolmogorov flow. The same holds for energy E and dissipation D. The particular linear form of the energy input $I[\boldsymbol{u}]$ implies that it is also invariant under the projection operation $\mathscr{P}$, i.e. $I[\boldsymbol{u}] = I[\mathscr{P}\boldsymbol{u}]$.

Furthermore, we have

$$I[\boldsymbol{u} - \mathscr{P}\boldsymbol{u}] = 0, \tag{5.7}$$

that is, the remainder $\boldsymbol{u} - \mathscr{P}\boldsymbol{u}$ does not contribute to the energy input. As discussed in § 5.1, the intermittent episodes of the flow are triggered by high energy inputs. Therefore, to study the temporal intermittency of the Kolmogorov flow, it is sufficient to consider the symmetric part $\mathscr{P}\boldsymbol{u}$. Note, however, that, owing to the quadratic form of E and D, the energy and dissipation of the remainder $\boldsymbol{u} - \mathscr{P}\boldsymbol{u}$ are generally non-zero.

The remaining continuous symmetry $\mathscr{T}_\ell$ is handled here by the method of slices (Cartan 1935; Field 1980; Rowley *et al.* 2003). This method replaces all continuous symmetry copies of a state with a copy that belongs to a given hypersurface called the slice. This hypersurface is such that each group orbit $\mathscr{T}_\ell \boldsymbol{u}$ in a neighbourhood of a given state intersects the slice transversely at a unique point. Each group orbit $\mathscr{T}_\ell \boldsymbol{u}$ is then replaced by its unique intersection with the slice. The method of slices has only recently been used in the context of fluid dynamics (Willis *et al.* 2013; Willis, Short & Cvitanović 2015). We use the first-Fourier-mode implementation of this method developed by Budanur *et al.* (2015) (cf. appendix D for more detail).

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

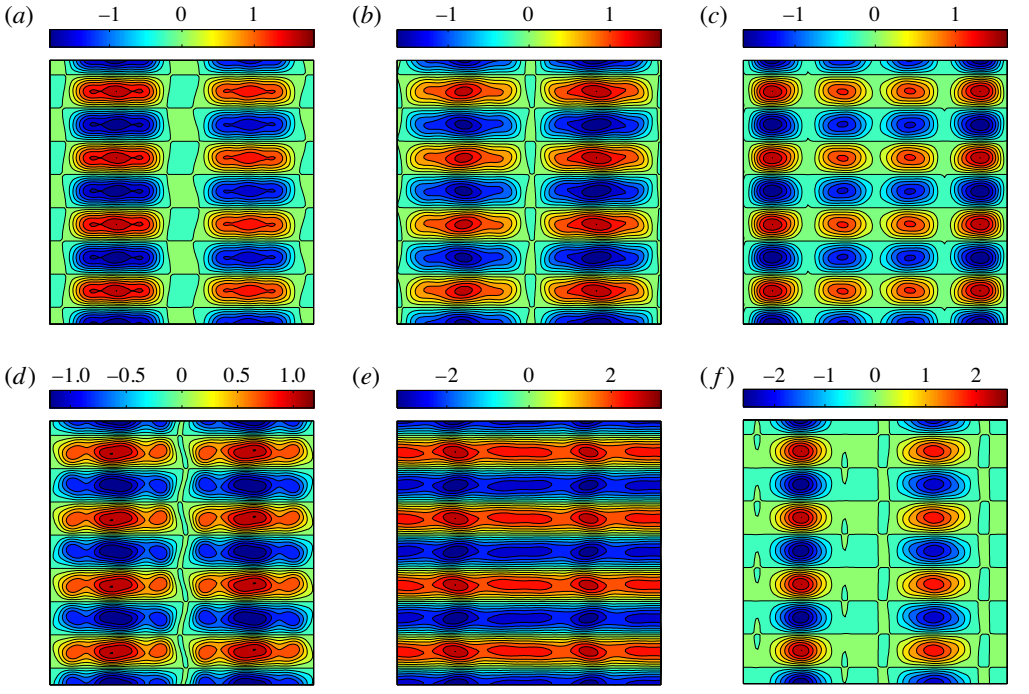


FIGURE 12. (Colour online) The symmetric part of the vorticity field $\mathscr{P}\omega$ of the turbulent trajectory at $t=84$ (*a*), $t=94$ (*b*), $t=104$ (*c*), $t=124.2$ (*d*), $t=134.2$ (*e*) and $t=144.2$ (*f*). The energy input–dissipation pairs (I, D) for the states are $(0.109, 0.107)$, $(0.115, 0.114)$, $(0.104, 0.104)$, $(0.095, 0.094)$, $(0.314, 0.269)$ and $(0.099, 0.093)$, respectively.

5.3. *Temporal intermittency*

Figure 12 shows the symmetric part of the vorticity field $\mathscr{P}\omega$ for six select times along the turbulent trajectory. Since the forcing wavenumber is $n=4$, there are $4n=16$ discrete symmetries. As a result, the symmetric part of each state exhibits recurring patterns that are related through discrete symmetry operations $\mathscr{R}$ and $\mathscr{S}$. In other words, knowing the symmetric part $\mathscr{P}\omega$ on one-sixteenth of the domain $[0, 2\pi] \times [0, 2\pi]$, one can reproduce $\mathscr{P}\omega$ on the entire domain.

Figure 12(*e*), showing the state at time $t=134.2$, corresponds to the first intermittent peak in figure 10(*a*). A distinct change of topology occurs in the symmetric part of the vorticity field as the turbulent trajectory undergoes an intermittent episode. Before and after the episode, $\mathscr{P}\omega$ has at least two distinct co-rotating vortices in each positive (or negative) vorticity band. As the trajectory gets closer to the intermittent episode, the co-rotating vortices merge, resulting in horizontal bands of alternating positive and negative vorticity. After the episode (see figure 12*f*), these bands split again into two distinct co-rotating vortices. This sharp distinction is not immediate from the full vorticity field ω (cf. figure 13). The same trend (i.e. the merger of co-rotating vortices) was observed during the intermittent episodes occurring at later times.

We found by inspection that the symmetric part of equilibrium solution E_{13} (see figure 14) is strikingly similar to that of the turbulent trajectory as it undergoes intermittency. This is visually appreciable from comparing figure 12(*e*) with figure 14(*b*). This observation suggests that close passages to equilibrium E_{13} trigger the intermittent behaviour.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

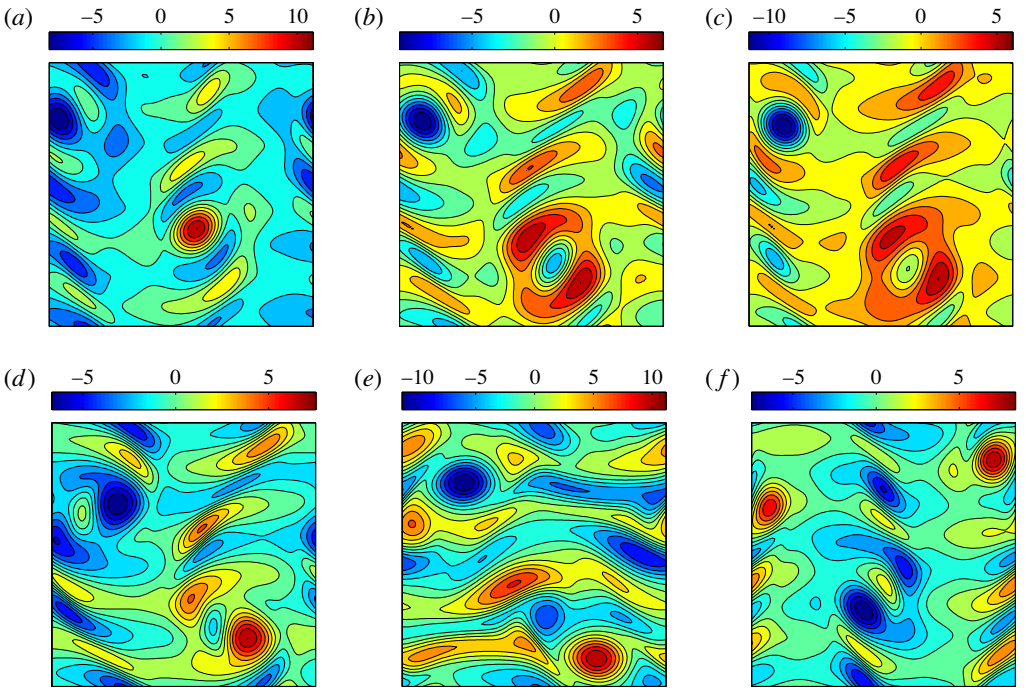


FIGURE 13. (Colour online) The vorticity field ω for the snapshots of the turbulent trajectory shown in figure 12. The panels correspond to times $t = 84$ (a), $t = 94$ (b), $t = 104$ (c), $t = 124.2$ (d), $t = 134.2$ (e) and $t = 144.2$ (f).

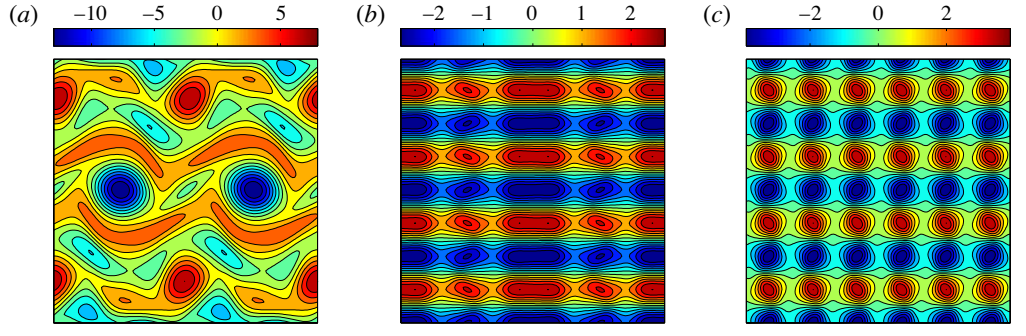


FIGURE 14. (Colour online) (a) The vorticity field ω for the equilibrium solution E_{13}. (b) The symmetric part of the vorticity $\mathscr{P}\omega$ for the equilibrium solution E_{13}. (c) The symmetric part of the vorticity $\mathscr{P}\omega$ for the equilibrium solution E_6.

In figure 14(c), we also show the symmetric part $\mathscr{P}\omega$ for equilibrium E_6. While the energy input–dissipation of equilibria E_6 and E_{13} are very close (cf. table 1), their vorticity fields are very different. This demonstrates the fact that closeness of the energy input–dissipation of two states does not imply their closeness in the state space.

To quantitatively examine the role of equilibrium E_{13} on intermittency, we consider the L^2 distance between the symmetric part of each turbulent state and the symmetric part of equilibrium E_{13}, i.e. $\|\mathscr{P}\omega(t) - \mathscr{P}E_{13}\|_{L^2}$. To account for the continuous

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

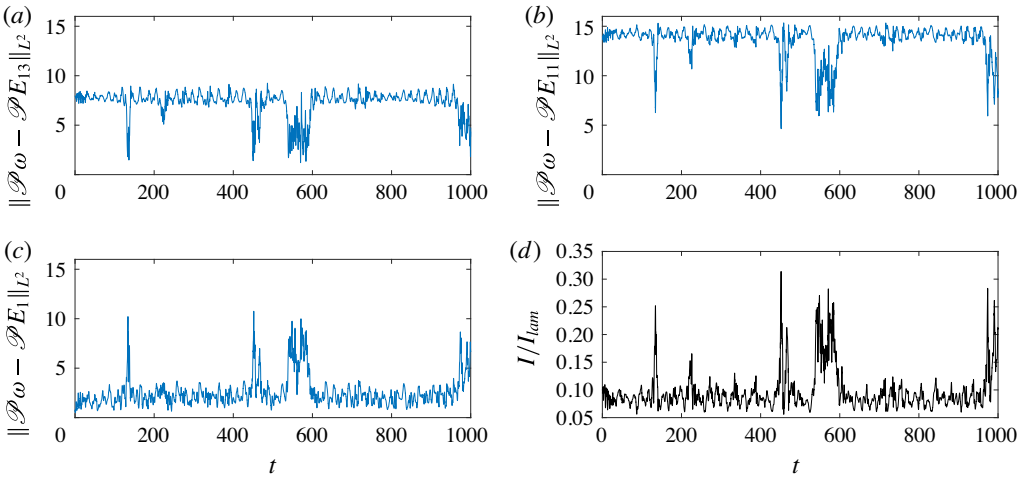


FIGURE 15. (Colour online) The L^2 distance $\|\mathscr{P}\omega(t) - \mathscr{P}E\|_{L^2}$ between the symmetric part of the turbulent trajectory $\mathscr{P}\omega$ and the symmetric parts of the equilibria E_{13} (*a*), E_{11} (*b*) and E_1 (*c*). (*d*) The normalized energy input I/I_{lam} of the turbulent trajectory, shown for reference (same as figure 10*a*). Each distinguished peak of the energy input (i.e. intermittency) coincides with a significant decrease in the distance from the equilibrium E_{13}.

translational symmetry, each symmetric part is first brought to the first-Fourier-mode slice, as explained in appendix D, before the distance is computed.

Figure 15 shows the evolution of the L^2 distance as a function of time. The evolution of the energy input I is also shown in the same figure. Every intermittent episode, i.e. significant peaks in the energy input, is concurrent with a significant dip in the L^2 distance. This demonstrates that the intermittent episodes correspond to close passages of the turbulent trajectory to equilibrium E_{13}.

The distance between the turbulent trajectory and some of the other invariant solutions (e.g. E_{11} and T_4) also decreases during the intermittent episodes. Their decrease is, however, not as dramatic as the one observed for equilibrium E_{13}. For instance, the evolution of the distance from E_{11} is also shown in figure 15. The distance from E_{11} fluctuates mostly around 14.0, which is significantly larger than the average distance from E_{13}, which fluctuates around 8.0. During the intermittent episodes, the distance from both equilibria decreases. The distance from E_{11} reaches 4.64 at its minimum, which is approximately one-third of its average distance of 14.0. The decrease in the distance from E_{13} is more significant: it attains a minimum as low as 1.21, which is approximately 6.6 times smaller than its average distance of 8.0. In contrast, the distance from equilibria residing in the ergodic sea increases during each intermittent episode. The distance from the equilibrium E_1 is shown in figure 15 as an example.

If the Kolmogorov flow at Reynolds number $Re = 40$ possesses an ergodic attractor, as the numerical evidence suggests (see e.g. Platt *et al.* 1991), every invariant solution embedded in the attractor will eventually be visited by a generic turbulent trajectory. Nonetheless, for the finite-time (and relatively short) turbulent trajectory computed here, equilibrium E_{13} appears to explain its observed intermittencies.

Finding the invariant equilibrium and travelling-wave solutions is only the first step. A complete understanding of the Kolmogorov flow in terms of its invariant solutions

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

will require a detailed state space analysis of the type carried out by Gibson, Halcrow & Cvitanović (2008) for the plane Couette flow. Such a thorough study deserves its own treatment, which will be presented elsewhere.

6. Conclusions and perspectives

We have proposed and developed here a new method for finding the equilibrium and travelling-wave (relative equilibrium) solutions of the forced Navier–Stokes equations with periodic boundary conditions. Namely, adjoint PDEs were derived whose (relative) equilibria include those of Navier–Stokes equations. Furthermore, the (relative) equilibria of the adjoint equations are asymptotically stable, and therefore their trajectories converge to desired invariant solutions.

Applying this method to the Kolmogorov flow led to the discovery of several new equilibrium and travelling-wave solutions. Specifically, for Reynolds number $Re = 40$, we found a total of 24 non-trivial equilibrium and travelling-wave solutions, where only three of them were known previously (Chandler & Kerswell 2013).

Some of these new solutions exhibit highly localized spatial structures (cf. figures 6 and 8). Such localized structures were previously believed to exist only in rectangular domains with large aspect ratios (Schneider, Gibson & Burke 2010; Lucas & Kerswell 2014).

One of the equilibrium solutions appears responsible for the observed temporal intermittencies in the Kolmogorov flow. Such intermittencies manifest themselves as sudden and short-lived bursts in the energy dissipation (and energy input) along a generic turbulent trajectory. We showed that these bursts correspond to close passages of the trajectory to equilibrium E_{13}. This supports the dynamical systems view that such rare, extreme events are the result of heteroclinic orbits carrying the trajectory to less accessible regions of the state space (Holmes 1993).

Owing to the periodic boundary conditions, the adjoint PDEs can be easily integrated numerically using a pseudo-spectral method. The spectral representation of the adjoint equations, presented in appendix C, exhibits a close resemblance to that of the Navier–Stokes equations. Therefore, existing pseudo-spectral codes can be easily adapted to solve the adjoint equations.

We found the rate of convergence of the adjoint method to be rather slow. More precisely, the trajectories of the adjoint PDEs converge rapidly to the vicinity of an invariant solution, but further convergence takes place at a slow rate. To achieve better convergence, we proposed a hybrid adjoint–Newton algorithm, consisting of two steps. First, the adjoint equations are integrated from some initial condition in order to reach the neighbourhood of an invariant solution. Once in the neighbourhood of the invariant solution, the standard Newton–GMRES–hook iterations (Viswanath 2007) were used to converge further to the solution. This hybrid algorithm yielded a 100 % converge rate from generic initial conditions (4.7), obviating the preprocessing step for finding 'good' initial guesses (Viswanath 2007; Chandler & Kerswell 2013).

Newton–GMRES–hook iterations are relatively expensive. Hence, a modification of our adjoint-based method that would eliminate the Newton-type steps altogether is highly desirable (see Lakoba & Yang (2007), in the context of nonlinear wave equations).

While we only considered periodic boundary conditions, the general adjoint-based approach applies to wall-bounded turbulence such as plane Couette and pipe flows. Our preliminary analysis (not presented here) shows that, in the presence of boundaries, the resulting adjoint equations require more boundary conditions

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

than the corresponding Navier–Stokes equations. This is to be expected, as the adjoint equations have higher-order spatial derivatives. This calls for a special numerical treatment of the adjoint equations for wall-bounded flows. Our results on the Kolmogorov flow, however, show that the gain is worth the pain.

Finally, we point out that our adjoint-based method does not immediately apply to the computation of periodic and relative periodic orbits. They are currently found through Newton–GMRES–hook iterations (see e.g. Kawahara *et al.* 2006; Viswanath 2007; Willis *et al.* 2013; Lucas & Kerswell 2015) with the drawbacks discussed in the introduction. Alternatives include the variational method of Lan & Cvitanović (2004), which shares the universally convergent property of our adjoint-based method. Its computational complexity is, however, formidable (Fazendeiro *et al.* 2010). The more recent method of Yang (2015) has proved promising for unidirectional wave equations, but its feasibility for Navier–Stokes equations remains to be explored.

Acknowledgements

I am grateful to P. Cvitanović for his support, helpful discussions and a careful reading of the manuscript that greatly improved the presentation of this work. I am thankful to A. P. Willis for his help with the implementation of Newton–GMRES–hook iterations and his comments on an earlier version of the manuscript. I would like to acknowledge fruitful discussions with R. de la Llave, F. Fedele, J. F. Gibson, R. R. Kerswell, B. Protas A. Souza and L. S. Tuckerman. I thank three anonymous referees for their comments and suggestions that improved the presentation of the results. Finally, I would like to thank the family of G. Robinson, Jr. for J. Ford Fellowship support.

Supplementary material

A tutorial, with the accompanying MATLAB code, is available at https://farazmand.wordpress.com/2015/12/18/adjoint/.

Appendix A. Derivation of the adjoint equations for equilibria

We start by redefining the vector $\boldsymbol{F_0}$ by adding the term $\boldsymbol{\nabla \cdot u}$ as a new element; that is

$$\boldsymbol{F_0}(\boldsymbol{q}) = \begin{pmatrix} -\boldsymbol{u}\cdot\boldsymbol{\nabla u} - \boldsymbol{\nabla} p + \nu\Delta\boldsymbol{u} + \boldsymbol{f} \\ \boldsymbol{\nabla\cdot u} \end{pmatrix}, \tag{A 1}$$

where $\boldsymbol{q} = (\boldsymbol{u}, p)$ and with the understanding that, for divergence-free vector fields, the last row of $\boldsymbol{F_0}$ is identically zero. This twist in the notation proves useful below, where we compute the adjoint of the Gâteaux derivative within the space of divergence-free vector fields. The Gâteaux derivative of $\boldsymbol{F_0}$ is given by

$$\mathscr{L}_{\boldsymbol{0}}(\boldsymbol{q};\boldsymbol{q}') = \begin{pmatrix} -\boldsymbol{u}'\cdot\boldsymbol{\nabla u} - \boldsymbol{u}\cdot\boldsymbol{\nabla u}' - \boldsymbol{\nabla} p' + \nu\Delta\boldsymbol{u}' \\ \boldsymbol{\nabla\cdot u}' \end{pmatrix}, \tag{A 2}$$

where $\boldsymbol{q}' = (\boldsymbol{u}', p')$ with $\boldsymbol{\nabla\cdot u}' = 0$.

The key part of deriving the adjoint equation (3.3) is to find the adjoint operator $\mathscr{L}_{\boldsymbol{0}}^{\dagger}$. Its derivation is standard and can be found in the literature on adjoint-based optimal control (see e.g. Gunzburger (2002); our notation is closer to Farazmand, Kevlahan & Protas (2011)). The difference is that, in optimal control, one seeks to minimize a

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

cost functional with the constraint that the Navier–Stokes equations are satisfied. Here, instead, the only constraint is the divergence-free condition and we seek to find the states $\boldsymbol{q} = (\boldsymbol{u}, p)$ such that $\|\boldsymbol{F}_{\boldsymbol{0}}(\boldsymbol{q})\|_{\mathscr{A}} = 0$. We derive the adjoint with respect to the L^2 inner product first. The adjoint with respect to the more general inner product (2.2) follows easily.

Let the function space $\mathscr{H}$ be the space of square integrable functions $\boldsymbol{q} = (\boldsymbol{u}, p)$ such that the $\boldsymbol{u}$ component is divergence-free. More precisely,

$$\mathscr{H} = \{\boldsymbol{q} = (\boldsymbol{u}, p) \in L^2(\mathscr{D}) \mid \boldsymbol{\nabla} \cdot \boldsymbol{u} = 0\}. \tag{A 3}$$

We define the usual L^2 inner product on $\mathscr{H}$, i.e. for any $\boldsymbol{q}, \boldsymbol{q}' \in \mathscr{H}$ we define

$$\langle \boldsymbol{q}, \boldsymbol{q}' \rangle_{L^2} = \int_{\mathscr{D}} [\boldsymbol{u}(\boldsymbol{x}, t) \cdot \boldsymbol{u}'(\boldsymbol{x}, t) + p(\boldsymbol{x}, t) p'(\boldsymbol{x}, t)] \, \mathrm{d}\boldsymbol{x}. \tag{A 4}$$

The adjoint of the linear operator (A 2) with respect to the inner product (A 4) on the function space $\mathscr{H}$ (A 3) is given by

$$\mathscr{L}_{\boldsymbol{0}}^{\dagger}(\boldsymbol{q}; \boldsymbol{q}'') = \begin{pmatrix} [\boldsymbol{\nabla} \boldsymbol{u}'' + (\boldsymbol{\nabla} \boldsymbol{u}'')^{\mathrm{T}}] \boldsymbol{u} - \boldsymbol{\nabla} p'' + \nu \Delta \boldsymbol{u}'' \\ \boldsymbol{\nabla} \cdot \boldsymbol{u}'' \end{pmatrix}, \quad \forall \boldsymbol{q}, \boldsymbol{q}'' \in \mathscr{H}, \tag{A 5}$$

where $^{\mathrm{T}}$ denotes matrix transposition. As we are restricted to the space of divergence-free vector fields, the last component of the adjoint (A 5) is identically zero, that is, $\boldsymbol{\nabla} \cdot \boldsymbol{u}'' = 0$.

This expression follows directly from definition (2.12) of the adjoint and (A 2). Namely, substituting (A 2) into (2.12), we have

$$\begin{aligned} \langle \mathscr{L}_{\boldsymbol{0}}(\boldsymbol{q}; \boldsymbol{q}'), \boldsymbol{q}'' \rangle_{L^2} &= \int_{\mathscr{D}} \begin{pmatrix} -\boldsymbol{u}' \cdot \boldsymbol{\nabla} \boldsymbol{u} - \boldsymbol{u} \cdot \boldsymbol{\nabla} \boldsymbol{u}' - \boldsymbol{\nabla} p' + \nu \Delta \boldsymbol{u}' \\ \boldsymbol{\nabla} \cdot \boldsymbol{u}' \end{pmatrix} \cdot \begin{pmatrix} \boldsymbol{u}'' \\ p'' \end{pmatrix} \mathrm{d}\boldsymbol{x}, \\ &= \int_{\mathscr{D}} [((\boldsymbol{\nabla} \boldsymbol{u}'' + \boldsymbol{\nabla} \boldsymbol{u}''^{\mathrm{T}}) \boldsymbol{u} - \boldsymbol{\nabla} p'' + \nu \Delta \boldsymbol{u}'') \cdot \boldsymbol{u}' + (\boldsymbol{\nabla} \cdot \boldsymbol{u}'') p'] \, \mathrm{d}\boldsymbol{x}, \\ &= \int_{\mathscr{D}} \begin{pmatrix} (\boldsymbol{\nabla} \boldsymbol{u}'' + \boldsymbol{\nabla} \boldsymbol{u}''^{\mathrm{T}}) \boldsymbol{u} - \boldsymbol{\nabla} p'' + \nu \Delta \boldsymbol{u}'' \\ \boldsymbol{\nabla} \cdot \boldsymbol{u}'' \end{pmatrix} \cdot \begin{pmatrix} \boldsymbol{u}' \\ p' \end{pmatrix} \mathrm{d}\boldsymbol{x}, \\ &= \langle \boldsymbol{q}', \mathscr{L}_{\boldsymbol{0}}^{\dagger}(\boldsymbol{q}; \boldsymbol{q}'') \rangle_{L^2}, \end{aligned} \tag{A 6}$$

where the first line is the definition of the L^2 inner product and the second line follows from integration by parts. Note that the boundary terms vanish due to the periodic boundary conditions. Only one of these integration by parts is not straightforward, which we outline below:

$$\begin{aligned} \int_{\mathscr{D}} (-\boldsymbol{u}' \cdot \boldsymbol{\nabla} \boldsymbol{u} - \boldsymbol{u} \cdot \boldsymbol{\nabla} \boldsymbol{u}') \cdot \boldsymbol{u}'' \, \mathrm{d}\boldsymbol{x} &= \int_{\mathscr{D}} (-u_i'' u_j' \partial_j u_i - u_i'' u_j \partial_j u_i') \, \mathrm{d}\boldsymbol{x}, \\ &= \int_{\mathscr{D}} (u_i \partial_j (u_i'' u_j') + u_i' \partial_j (u_i'' u_j)) \, \mathrm{d}\boldsymbol{x}, \\ &= \int_{\mathscr{D}} (u_i u_j' \partial_j u_i'' + u_i' u_j \partial_j u_i'') \, \mathrm{d}\boldsymbol{x}, \\ &= \int_{\mathscr{D}} (u_j u_i' \partial_i u_j'' + u_i' u_j \partial_j u_i'') \, \mathrm{d}\boldsymbol{x}, \\ &= \int_{\mathscr{D}} [(\boldsymbol{\nabla} \boldsymbol{u}'' + \boldsymbol{\nabla} \boldsymbol{u}''^{\mathrm{T}}) \boldsymbol{u}] \cdot \boldsymbol{u}' \, \mathrm{d}\boldsymbol{x}, \end{aligned} \tag{A 7}$$

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

where we used the summation notation on repeated indices and the fact that $\boldsymbol{u}$ and $\boldsymbol{u}'$ are divergence-free: $\partial_j u_j = \partial_j u'_j = 0$.

Identity (A 6) holds for general $\boldsymbol{q}, \boldsymbol{q}', \boldsymbol{q}'' \in \mathscr{H}$. Along the trajectories of the adjoint descent PDE, we have $\boldsymbol{q} = (\boldsymbol{u}, p)$, $\boldsymbol{q}' = (\partial_\tau \boldsymbol{u}, \partial_\tau p)$ and $\boldsymbol{q}'' := (\boldsymbol{u}'', p'') = \boldsymbol{F_0}(\boldsymbol{q})$ (cf. (2.11)). This yields

$$\langle \boldsymbol{q}', \mathscr{L}_{\boldsymbol{0}}^{\dagger}(\boldsymbol{q}; \boldsymbol{q}'')\rangle_{L^2} = \int_{\mathscr{D}} \{\partial_\tau \boldsymbol{u} \cdot ([\boldsymbol{\nabla u}'' + (\boldsymbol{\nabla u}'')^{\mathrm{T}}]\boldsymbol{u} - \boldsymbol{\nabla} p'' + \nu \Delta \boldsymbol{u}'') + (\boldsymbol{\nabla} \cdot \boldsymbol{u}'')\partial_\tau p\}\, \mathrm{d}\boldsymbol{x}. \tag{A 8}$$

Therefore, (3.3*a*) ensures that the above inner product is negative-semidefinite. This together with identity (2.11) ensures that the norm $\|\boldsymbol{F_0}(\boldsymbol{q}(\tau))\|_{L^2}$ decreases along the trajectories of the adjoint PDE.

Owing to the divergence-free condition, the term $(\boldsymbol{\nabla} \cdot \boldsymbol{u}'')\partial_\tau p$ vanishes. Therefore, no independent evolution equation is obtained for the pressure p. As in the case of the Navier–Stokes equation, the pressure terms simply enforce the divergence-free conditions (3.3*b*,*c*).

The adjoint descent with respect to the more general inner product (2.2) is obtained similarly, the only difference being that $\boldsymbol{q}''$ is replaced by $\mathscr{A}\boldsymbol{q}''$.

Appendix B. Derivation of the adjoint equation for travelling waves

As in appendix A, we redefine $\boldsymbol{F}_c$ as

$$\boldsymbol{F}_c(\boldsymbol{q}) = \begin{pmatrix} -\boldsymbol{u} \cdot \boldsymbol{\nabla u} - \boldsymbol{\nabla} p + \nu \Delta \boldsymbol{u} + \boldsymbol{f} + \boldsymbol{c} \cdot \boldsymbol{\nabla u} \\ \boldsymbol{\nabla} \cdot \boldsymbol{u} \end{pmatrix}, \tag{B 1}$$

where $\boldsymbol{q} = (\boldsymbol{u}, p)$ and

$$\boldsymbol{F}_c(\boldsymbol{q}) = \boldsymbol{F_0}(\boldsymbol{q}) + \begin{pmatrix} \boldsymbol{c} \cdot \boldsymbol{\nabla u} \\ 0 \end{pmatrix}, \tag{B 2}$$

with $\boldsymbol{F_0}$ defined in (A 1). Owing to the linearity of the second term, the Gâteaux derivative of $\boldsymbol{F}_c$ is

$$\mathscr{L}_c(\boldsymbol{q}; \boldsymbol{q}') = \mathscr{L}_{\boldsymbol{0}}(\boldsymbol{q}; \boldsymbol{q}') + \begin{pmatrix} \boldsymbol{c} \cdot \boldsymbol{\nabla u}' \\ 0 \end{pmatrix}, \tag{B 3}$$

where $\boldsymbol{q}' = (\boldsymbol{u}', p')$ and $\mathscr{L}_{\boldsymbol{0}}$ is defined by (A 5).

Analogous to the analysis of appendix A, the adjoint $\mathscr{L}_c^{\dagger}$ can be shown to be given by

$$\mathscr{L}_c^{\dagger}(\boldsymbol{q}; \boldsymbol{q}'') = \mathscr{L}_{\boldsymbol{0}}^{\dagger}(\boldsymbol{q}; \boldsymbol{q}'') - \begin{pmatrix} \boldsymbol{c} \cdot \boldsymbol{\nabla u}'' \\ 0 \end{pmatrix}, \tag{B 4}$$

where $\boldsymbol{q}'' = (\boldsymbol{u}'', p'')$.

We would like to define the adjoint descent equation in such a way that, along its trajectories $(\boldsymbol{q}(\tau), \boldsymbol{c}(\tau))$, the residue $\|\boldsymbol{F}_{\boldsymbol{c}(\tau)}(\boldsymbol{q}(\tau))\|_{\mathscr{A}}$ decreases monotonically. Taking the derivative with respect to the fictitious time τ, we obtain

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

$$\begin{aligned}\partial_\tau \|\boldsymbol{F}_c(\boldsymbol{q})\|^2_{\mathscr{A}} &= \langle \mathscr{L}_c(\boldsymbol{q};\boldsymbol{q}'), \mathscr{A}\boldsymbol{F}_c(\boldsymbol{q})\rangle_{L^2} + \left\langle \begin{pmatrix} \dot{\boldsymbol{c}}\cdot\nabla\boldsymbol{u} \\ 0 \end{pmatrix}, \mathscr{A}\boldsymbol{F}_c(\boldsymbol{q}) \right\rangle_{L^2} \\ &= \langle \boldsymbol{q}', \mathscr{L}_c^\dagger(\boldsymbol{q}; \mathscr{A}\boldsymbol{F}_c(\boldsymbol{q}))\rangle_{L^2} + \langle \dot{\boldsymbol{c}}\cdot\nabla\boldsymbol{u}, \tilde{\boldsymbol{u}}''\rangle_{L^2} \\ &= \langle \boldsymbol{q}', \mathscr{L}_c^\dagger(\boldsymbol{q}; \mathscr{A}\boldsymbol{F}_c(\boldsymbol{q}))\rangle_{L^2} + \dot{\boldsymbol{c}}\cdot\int_{\mathscr{D}} (\nabla\boldsymbol{u})^{\mathrm{T}}\tilde{\boldsymbol{u}}''\,\mathrm{d}\boldsymbol{x}, \end{aligned} \tag{B 5}$$

where $\boldsymbol{q}' = \partial_\tau \boldsymbol{q}$ and $\boldsymbol{u}''$ is given by (3.10). The adjoint set of equations (3.9) is designed in such a way that the right-hand side of the above equation is always negative, resulting in monotonic decrease in the residue along its solutions $(\boldsymbol{u}(\tau), p(\tau))$.

Appendix C. The adjoint equations in the Fourier space

In this appendix, we derive the spectral representation of the adjoint descent equations (3.3) and (3.9). As shorthand notation, we define

$$\boldsymbol{N} := [\nabla\tilde{\boldsymbol{u}}'' + (\nabla\tilde{\boldsymbol{u}}'')^{\mathrm{T}}]\boldsymbol{u}, \tag{C 1}$$

and denote Fourier transforms with a hat sign. Then one can write (3.3*a*) in the Fourier space as

$$\partial_\tau \widehat{\boldsymbol{u}}(\boldsymbol{k}) = -\left(\boldsymbol{I} - \frac{\boldsymbol{k}\otimes\boldsymbol{k}}{|\boldsymbol{k}|^2}\right)\widehat{\boldsymbol{N}}(\boldsymbol{k}) + \nu|\boldsymbol{k}|^2\widehat{\tilde{\boldsymbol{u}}''}(\boldsymbol{k}), \tag{C 2}$$

where $\boldsymbol{I}$ is the identity matrix and $\boldsymbol{k}\otimes\boldsymbol{k}$ denotes the tensor product. Note that the pressure p'' is eliminated using identity (3.6), which implies

$$\widehat{p''}(\boldsymbol{k}) = \frac{-\mathrm{i}\boldsymbol{k}\cdot\widehat{\boldsymbol{N}}(\boldsymbol{k})}{|\boldsymbol{k}|^2}. \tag{C 3}$$

For $\widehat{\tilde{\boldsymbol{u}}''}$, using definition (3.11), we have

$$\widehat{\tilde{\boldsymbol{u}}''}(\boldsymbol{k}) = \frac{\widehat{\boldsymbol{u}}''}{1+|\boldsymbol{k}|^2}. \tag{C 4}$$

The term $\widehat{\boldsymbol{u}}''$ is in turn computed from identity (3.4) as

$$\widehat{\boldsymbol{u}}''(\boldsymbol{k}) = -\left(\boldsymbol{I} - \frac{\boldsymbol{k}\otimes\boldsymbol{k}}{|\boldsymbol{k}|^2}\right)\widehat{\boldsymbol{R}}(\boldsymbol{k}) - \nu|\boldsymbol{k}|^2\widehat{\boldsymbol{u}}(\boldsymbol{k}) + \widehat{\boldsymbol{f}}(\boldsymbol{k}), \tag{C 5}$$

where $\boldsymbol{R}$ is the shorthand notation for the nonlinear term

$$\boldsymbol{R} := \boldsymbol{u}\cdot\nabla\boldsymbol{u}. \tag{C 6}$$

The nonlinear terms $\boldsymbol{R}$ and $\boldsymbol{N}$ are computed by the pseudo-spectral method, i.e. the differentiations are carried out in the Fourier space and the products are computed in the physical space.

Note that the spectral representation (C 2) of the adjoint equation closely resembles that of the Navier–Stokes equation (C 5). Therefore, existing pseudo-spectral codes for Navier–Stokes can easily be adapted to solve the adjoint PDEs.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

The adjoint descent for the travelling waves is computed similarly, except that the term $\mathrm{i}(\boldsymbol{c}\cdot\boldsymbol{k})\widehat{\tilde{\boldsymbol{u}}''}(\boldsymbol{k})$ is added to the right-hand side of (C 2) accounting for the term $\boldsymbol{c}\cdot\nabla\tilde{\boldsymbol{u}}''$ in (3.9). More precisely, the adjoint descent equation for the travelling waves in the Fourier space reads

$$\partial_\tau\widehat{\boldsymbol{u}}(\boldsymbol{k}) = -\left(\boldsymbol{I}-\frac{\boldsymbol{k}\otimes\boldsymbol{k}}{|\boldsymbol{k}|^2}\right)\widehat{\boldsymbol{N}}(\boldsymbol{k}) + \nu|\boldsymbol{k}|^2\widehat{\tilde{\boldsymbol{u}}''}(\boldsymbol{k}) + \mathrm{i}(\boldsymbol{c}\cdot\boldsymbol{k})\widehat{\tilde{\boldsymbol{u}}''}(\boldsymbol{k}), \tag{C 7}$$

with $\tilde{\boldsymbol{u}}'' = \mathscr{A}\boldsymbol{u}''$ and

$$\widehat{\boldsymbol{u}}''(\boldsymbol{k}) = -\left(\boldsymbol{I}-\frac{\boldsymbol{k}\otimes\boldsymbol{k}}{|\boldsymbol{k}|^2}\right)\widehat{\boldsymbol{R}}(\boldsymbol{k}) - \nu|\boldsymbol{k}|^2\widehat{\boldsymbol{u}}(\boldsymbol{k}) + \widehat{\boldsymbol{f}}(\boldsymbol{k}) + \mathrm{i}(\boldsymbol{c}\cdot\boldsymbol{k})\widehat{\boldsymbol{u}}(\boldsymbol{k}). \tag{C 8}$$

Appendix D. Symmetry actions in the Fourier space

The symmetry operations (5.1) can be readily implemented in the Fourier space. The following transformations follow directly from the definition of the Fourier transform. Denoting the Fourier modes of the velocity field $\boldsymbol{u}$ by $\widehat{\boldsymbol{u}}(\boldsymbol{k})$ with $\boldsymbol{k} = (k_1, k_2) \in \mathbb{Z}^2$, we have

$$\boldsymbol{u}(x_1, x_2) = \sum_{k_1\in\mathbb{Z}}\sum_{k_2\in\mathbb{Z}}\widehat{\boldsymbol{u}}(k_1, k_2)\mathrm{e}^{\mathrm{i}k_1x_1}\mathrm{e}^{\mathrm{i}k_2x_2}. \tag{D 1}$$

Therefore, the shift operation $\mathscr{T}_\ell$ satisfies

$$(\mathscr{T}_\ell\boldsymbol{u})(x_1, x_2) = \boldsymbol{u}(x_1+\ell, x_2) = \sum_{k_1\in\mathbb{Z}}\sum_{k_2\in\mathbb{Z}}\widehat{\boldsymbol{u}}(k_1, k_2)\mathrm{e}^{\mathrm{i}k_1\ell}\mathrm{e}^{\mathrm{i}k_1x_1}\mathrm{e}^{\mathrm{i}k_2x_2}, \tag{D 2}$$

implying that the action of the continuous symmetry $\mathscr{T}_\ell$ on the Fourier mode $\widehat{\boldsymbol{u}}(k_1, k_2)$ is multiplicative, such that

$$\widehat{\mathscr{T}_\ell\boldsymbol{u}}(k_1, k_2) = \mathrm{e}^{\mathrm{i}k_1\ell}\widehat{\boldsymbol{u}}(k_1, k_2). \tag{D 3}$$

Similarly, for the rotation-through-π symmetry $\mathscr{R}$, we have

$$\widehat{\mathscr{R}\boldsymbol{u}}(k_1, k_2) = -\widehat{\boldsymbol{u}}(-k_1, -k_2) = -[\widehat{\boldsymbol{u}}(k_1, k_2)]^*, \tag{D 4}$$

where the * denotes complex conjugation and the last identity follows from the fact that $\boldsymbol{u}$ is real-valued.

Finally, the shift–reflect symmetry $\mathscr{S}$ acts on the Fourier modes through

$$\widehat{\mathscr{S}^m\boldsymbol{u}}(k_1, k_2) = \mathrm{e}^{\mathrm{i}k_2(m\pi/n)}\begin{pmatrix}(-1)^m\widehat{u}_1((-1)^mk_1, k_2)\\ \widehat{u}_2((-1)^mk_1, k_2)\end{pmatrix}. \tag{D 5}$$

To reduce the continuous symmetry, we use the first-Fourier-mode slice of Budanur *et al.* (2015). It follows from (D 3) that the shift $\mathscr{T}_\ell$ acts on the Fourier mode $(k_1, k_2) = (1, 0)$ of the vorticity through

$$\widehat{\omega}(1, 0) \mapsto \widehat{\omega}(1, 0)\mathrm{e}^{\mathrm{i}\ell} = |\widehat{\omega}(1, 0)|\mathrm{e}^{\mathrm{i}(\phi(1,0)+\ell)}, \tag{D 6}$$

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

where $\phi(k_1, k_2) \in (-\pi, \pi]$ denotes the principal value of the phase of mode (k_1, k_2). To bring the vorticity field to the first-Fourier-mode slice, the shift value ℓ is chosen such that $\phi(1, 0) + \ell = 0$. More precisely, a given vorticity field ω is replaced with its symmetry-equivalent copy through the transformation

$$\widehat{\omega}(k_1, k_2) \mapsto \mathrm{e}^{-\mathrm{i}k_1\phi(1,0)}\widehat{\omega}(k_1, k_2). \tag{D 7}$$

As a result, the mode $\widehat{\omega}(1, 0)$ of the symmetry-reduced vorticity field has vanishing imaginary part.

REFERENCES

AMBROSE, D. M. & WILKENING, J. 2010*a* Computation of symmetric, time-periodic solutions of the vortex sheet with surface tension. *Proc. Natl Acad. Sci. USA* **107** (8), 3361–3366.

AMBROSE, D. M. & WILKENING, J. 2010*b* Computation of time-periodic solutions of the Benjamin–Ono equation. *J. Nonlinear Sci.* **20** (3), 277–308.

AUERBACH, D., CVITANOVIĆ, P., ECKMANN, J.-P., GUNARATNE, G. & PROCACCIA, I. 1987 Exploring chaotic motion through periodic orbits. *Phys. Rev. Lett.* **58**, 23.

AYALA, D. & PROTAS, B. 2014 Maximum palinstrophy growth in 2D incompressible flows. *J. Fluid Mech.* **742**, 340–367.

BATCHELOR, G. K. & TOWNSEND, A. A. 1949 The nature of turbulent motion at large wave-numbers. *Proc. R. Soc. Lond.* A **199**, 238–255.

BOYD, S. & VANDENBERGHE, L. 2004 *Convex Optimization*. Cambridge University Press.

BUDANUR, N. B., CVITANOVIĆ, P., DAVIDCHACK, R. L. & SIMINOS, E. 2015 Reduction of SO(2) symmetry for spatially extended dynamical systems. *Phys. Rev. Lett.* **114**, 084102.

CARTAN, E. 1935 *La Méthode du Repère Mobile, la Théorie des Groupes Continues et les Espaces Généralisés. (Exposés de Géométrie)*, vol. 5. Hermann.

CHANDLER, G. J. & KERSWELL, R. R. 2013 Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. *J. Fluid Mech.* **722**, 554–595.

CONSTANTIN, P., FOIAS, C., NICOLAENKO, B. & TEMAM, R. 1989 *Integral Manifolds and Inertial Manifolds for Dissipative Partial Differential Equations*. Springer.

CVITANOVIĆ, P. 2013 Recurrent flows: the clockwork behind turbulence. *J. Fluid Mech.* **726**, 1–4.

CVITANOVIĆ, P., ARTUSO, R., MAINIERI, R., TANNER, G. & VATTAY, G. 2014 *Chaos: Classical and Quantum*. Niels Bohr Institute.

CVITANOVIĆ, P. & GIBSON, J. F. 2010 Geometry of turbulence in wall-bounded shear flows: periodic orbits. *Phys. Scr.* T **142**, 014007.

CVITANOVIĆ, P. & LAN, Y. 2003 Turbulent fields and their recurrences. In *Proceedings of 10th International Workshop on Multiparticle Production: Correlations and Fluctuations in QCD* (ed. N. Antoniou), World Scientific.

DENNIS, J. E. JR & SCHNABEL, R. B. 1996 *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM.

DEUFLHARD, P. 2011 *Newton Methods for Nonlinear Problems: Affine Invariance and Adaptive Algorithms,* vol. 35. Springer.

DORMAND, J. R. & PRINCE, P. J. 1980 A family of embedded Runge–Kutta formulae. *J. Comput. Appl. Maths* **6** (1), 19–26.

FARAZMAND, M., KEVLAHAN, N. K.-R. & PROTAS, B. 2011 Controlling the dual cascade of two-dimensional turbulence. *J. Fluid Mech.* **668**, 202–222.

FAZENDEIRO, L., BOGHOSIAN, B. M., COVENEY, P. V. & LÄTT, J. 2010 Unstable periodic orbits in weak turbulence. *J. Comput. Sci.* **1**, 13–23.

FIELD, M. J. 1980 Equivariant dynamical systems. *Trans. Am. Math. Soc.* **259** (1), 185–205.

FOIAS, C., JOLLY, M. S., KEVREKIDIS, I. G., SELL, G. R. & TITI, E. S. 1988 On the computation of inertial manifolds. *Phys. Lett.* A **131**, 433–436.

FOIAS, C., MANLEY, O., ROSA, R. & TEMAM, R. 2001 *Navier–Stokes Equations and Turbulence*. Cambridge University Press.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

FRISCH, U. 1996 *Turbulence*. Cambridge University Press.

GIBBON, J. D. & DOERING, C. R. 2003 Intermittency in solutions of the three-dimensional Navier–Stokes equations. *J. Fluid Mech.* **478**, 227–235.

GIBSON, J. F., HALCROW, J. & CVITANOVIĆ, P. 2008 Visualizing the geometry of state-space in plane Couette flow. *J. Fluid Mech.* **611**, 107–130.

GIBSON, J. F., HALCROW, J. & CVITANOVIĆ, P. 2009 Equilibrium and traveling-wave solutions of plane Couette flow. *J. Fluid Mech.* **638**, 243–266.

GILMORE, R. & LETELLIER, C. 2007 *The Symmetry of Chaos*. Oxford University Press.

GUNZBURGER, M. D. 2002 *Perspectives in Flow Control and Optimization*. SIAM.

HALCROW, J., GIBSON, J. F., CVITANOVIĆ, P. & VISWANATH, D. 2009 Heteroclinic connections in plane Couette flow. *J. Fluid Mech.* **621**, 365–376.

HOLMES, P. 1993 Symmetries, heteroclinic cycles and intermittency in fluid flow. In *Turbulence in Fluid Flows*, pp. 49–58. Springer.

HOLMES, P., LUMLEY, J. L., BERKOOZ, G. & ROWLEY, C. W. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd edn. Cambridge University Press.

HOLMES, P. & STONE, E. 1992 Heteroclinic cycles, exponential tails, and intermittency in turbulence production. In *Studies in Turbulence* (ed. T. B. Gatski, C. G. Speziale & S. Sarkar), pp. 179–189. Springer.

HOPF, E. 1948 A mathematical example displaying features of turbulence. *Commun. Appl. Maths* **1**, 303–322.

KAWAHARA, G. & KIDA, S. 2001 Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. *J. Fluid Mech.* **449**, 291–300.

KAWAHARA, G., KIDA, S. & VAN VEEN, L. 2006 Unstable periodic motion in turbulent flows. *Nonlinear Process. Geophys.* **13** (5), 499–507.

KAWAHARA, G., UHLMANN, M. & VAN VEEN, L. 2012 The significance of simple invariant solutions in turbulent flows. *Annu. Rev. Fluid Mech.* **44**, 203–225.

KERSWELL, R. R., PRINGLE, C. C. T. & WILLIS, A. P. 2014 An optimization approach for analysing nonlinear stability with transition to turbulence in fluids as an exemplar. *Rep. Prog. Phys.* **77** (8), 085901.

KREILOS, T. & ECKHARDT, B. 2012 Periodic orbits near onset of chaos in plane Couette flow. *Chaos* **22** (4), 047505.

KUO, A. Y.-S. & CORRSIN, S. 1971 Experiments on internal intermittency and fine-structure distribution functions in fully turbulent fluid. *J. Fluid Mech.* **50**, 285–319.

LAKOBA, T. I. & YANG, J. 2007 A mode elimination technique to improve convergence of iteration methods for finding solitary waves. *J. Comput. Phys.* **226** (2), 1693–1709.

LAN, Y. & CVITANOVIĆ, P. 2004 Variational method for finding periodic orbits in a general flow. *Phys. Rev.* E **69**, 016217.

LUCAS, D. & KERSWELL, R. 2014 Spatiotemporal dynamics in two-dimensional Kolmogorov flow over large domains. *J. Fluid Mech.* **750**, 518–554.

LUCAS, D. & KERSWELL, R. R. 2015 Recurrent flow analysis in spatiotemporally chaotic 2-dimensional Kolmogorov flow. *Phys. Fluids* **27**, 518–554.

MARCHIORO, C. 1986 An example of absence of turbulence for any Reynolds number. *Commun. Math. Phys.* **105**, 99–106.

MARCUS, P. S. & TUCKERMAN, L. S. 1987 Simulation of flow between concentric rotating spheres. Part 1. Steady states. *J. Fluid Mech.* **185**, 1–30.

MININNI, P. D. & POUQUET, A. 2010 Rotating helical turbulence. II. Intermittency, scale invariance, and structures. *Phys. Fluids* **22** (3), 035106.

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** (13), 134502.

NAGATA, M. 1990 Three-dimensional finite-amplitude solutions in plane Couette flow: bifurcation from infinity. *J. Fluid Mech.* **217**, 519–527.

NAGATA, M. 1997 Three-dimensional traveling-wave solutions in plane Couette flow. *Phys. Rev.* E **55**, 2023–2025.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press

PLATT, N., SIROVICH, L. & FITZMAURICE, N. 1991 An investigation of chaotic Kolmogorov flows. *Phys. Fluids* A **3**, 681–696.

PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T. & FLANNERY, B. P. 2007 *Numerical Recipes: The Art of Scientific Computing*, 3rd edn. Cambridge University Press.

PRINGLE, C. C. T. & KERSWELL, R. R. 2010 Using nonlinear transient growth to construct the minimal seed for shear flow turbulence. *Phys. Rev. Lett.* **105**, 154502.

ROWLEY, C. W., KEVREKIDIS, I. G., MARSDEN, J. E. & LUST, K. 2003 Reduction and reconstruction for self-similar dynamical systems. *Nonlinearity* **16**, 1257.

RUELLE, D. 1991 The turbulent fluid as a dynamical system. In *New Perspectives in Turbulence* (ed. L. Sirovich), pp. 123–138. Springer.

RUPPERT-FELSOT, J., FARGE, M. & PETITJEANS, P. 2009 Wavelet tools to study intermittency: application to vortex bursting. *J. Fluid Mech.* **636**, 427–453.

SAAD, Y. & SCHULTZ, M. H. 1986 GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.* **7**, 856–869.

SAUPE, D. 1988 Discrete versus continuous Newton's method: a case study. *Acta Appl. Maths* **13**, 59–80.

SCHNEIDER, K., FARGE, M. & KEVLAHAN, N. K.-R. 2004 Spatial intermittency in two-dimensional turbulence: a wavelet approach. In *Woods Hole Mathematics. Perspectives in Mathematics and Physics,* vol. 34, pp. 302–328.

SCHNEIDER, T. M., GIBSON, J. F. & BURKE, J. 2010 Snakes and ladders: localized solutions of plane Couette flow. *Phys. Rev. Lett.* **104** (10), 104501.

SIMINOS, E. & CVITANOVIĆ, P. 2011 Continuous symmetry reduction and return maps for high-dimensional flows. *Physica* D **240**, 187–198.

SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures. Part II: symmetries and transformations. *Q. Appl. Maths* **45** (3), 573–582.

SMALE, S. 1981 The fundamental theorem of algebra and complexity theory. *Bull. Am. Math. Soc. (N.S.)* **4** (1), 1–36.

SREENIVASAN, K. R. & ANTONIA, R. A. 1997 The phenomenology of small-scale turbulence. *Annu. Rev. Fluid Mech.* **29** (1), 435–472.

TREFETHEN, L. N. & BAU, D. 1997 *Numerical Linear Algebra.* SIAM.

TUCKERMAN, L., BERTAGNOLIO, F., DAUBE, O., QUÉRÉ, P. L. & BARKLEY, D. 2000 Stokes preconditioning for the inverse Arnoldi method. In *Continuation Methods for Fluid Dynamics* (ed. D. Henry & A. Bergeon), Notes on Numerical Fluid Mechanics, vol. 74, pp. 241–255. Springer.

VISWANATH, D. 2007 Recurrent motions within plane Couette turbulence. *J. Fluid Mech.* **580**, 339–358.

WALEFFE, F. 1997 On a self-sustaining process in shear flows. *Phys. Fluids* **9**, 883–900.

WALEFFE, F. 1998 Three-dimensional coherent states in plane shear flows. *Phys. Rev. Lett.* **81**, 4140–4143.

WALEFFE, F. 2003 Homotopy of exact coherent structures in plane shear flows. *Phys. Fluids* **15**, 1517–1543.

WILLIS, A. P., CVITANOVIĆ, P. & AVILA, M. 2013 Revealing the state space of turbulent pipe flow by symmetry reduction. *J. Fluid Mech.* **721**, 514–540.

WILLIS, A. P., SHORT, K. Y. & CVITANOVIĆ, P. 2015 Relative periodic orbits form the backbone of turbulent pipe flow. *Phys. Rev.* E **93**, 022204.

YANG, J. 2015 A numerical method for computing time-periodic solutions in dissipative wave systems. *Stud. Appl. Maths* **134** (4), 420–455.

YANG, J. & LAKOBA, T. I. 2007 Universally-convergent squared-operator iteration methods for solitary waves in general nonlinear wave equations. *Stud. Appl. Maths* **118** (2), 153–197.

https://doi.org/10.1017/jfm.2016.203 Published online by Cambridge University Press