

Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow

Gary J. Chandler and Rich R. Kerswell†

School of Mathematics, University of Bristol, University walk, Bristol, UK

(Received 17 July 2012; revised 19 February 2013; accepted 24 February 2013;
first published online 28 March 2013)

We consider long-time simulations of two-dimensional turbulence body forced by $\sin 4y\hat{x}$ on the torus $(x, y) \in [0, 2\pi]^2$ with the purpose of extracting simple invariant sets or ‘exact recurrent flows’ embedded in this turbulence. Each recurrent flow represents a sustained closed cycle of dynamical processes which underpins the turbulence. These are used to reconstruct the turbulence statistics using periodic orbit theory. The approach is found to be reasonably successful at a low value of the forcing where the flow is close to but not fully in its asymptotic (strongly) turbulent regime. Here, a total of 50 recurrent flows are found with the majority buried in the part of phase space most populated by the turbulence giving rise to a good reproduction of the energy and dissipation p.d.f. However, at higher forcing amplitudes now in the asymptotic turbulent regime, the generated turbulence data set proves insufficiently long to yield enough recurrent flows to make viable predictions. Despite this, the general approach seems promising providing enough simulation data is available since it is open to extensive automation and naturally generates dynamically important exact solutions for the flow.

Key words: chaos, nonlinear dynamical systems, turbulence theory

1. Introduction

Ideas from dynamical systems have recently provided fresh insight into transitional and weakly turbulent flows where the system size is smaller than the spatial correlation length. Viewing such flows as a trajectory through a phase space littered with invariant (‘exact’) solutions and their stable and unstable manifolds has proved a fruitful way of understanding such flows (Eckhardt *et al.* 2002; Kerswell 2005; Eckhardt *et al.* 2007; Gibson, Halcrow & Cvitanović 2008; Cvitanović & Gibson 2010; Kawahara, Uhlmann & van Veen 2012). It is therefore natural to ask whether any ideas attempting to rationalize chaos may have something to say about developed turbulence. This is not to presuppose the two phenomena are simply related, that they are not has surely been appreciated for over 30 years, but merely that an approach found useful in one may provide some insight into the other. One promising line of thinking in low-dimensional, hyperbolic dynamical systems stands out as a possibility: periodic orbit theory.

† Email address for correspondence: R.R.Kerswell@bristol.ac.uk

The study of periodic orbits as a tool to understand chaos has been a longstanding theme in dynamical systems dating back to Poincaré's original work on the three body problem in the 1880s (Poincaré 1892; Ruelle 1978; MacKay & Miess 1987; Gutzwiller 1990). The fact that chaotic solutions can fleetingly, but also recurrently, resemble different periodic flows over time has always suggested that the statistics of the former may be expressible as a weighted sum of properties of the latter. However, this has generally remained a vague hope except for a special subclass of dynamical system where periodic orbit theory has formalized this link (see Auerbach *et al.* 1987; Cvitanović 1988; Artuso, Aurell & Cvitanović 1990a and, for a recent review, Lan 2010). For these systems, very low-dimensional, fully hyperbolic invariant sets in which periodic orbits are dense ('axiom A' attractors), there have been some notable successes (e.g. Artuso, Aurell & Cvitanović 1990b; Cvitanović 1992 and later papers in the same journal issue, the evolving webbook Cvitanović *et al.* 2013). Here the invariant measure across the attractor can be expressed in terms of the periodic orbits which are dense within it so that ergodic averages can be determined from suitably weighted sums across the periodic orbits. Central to applying the approach is identifying a symbolic dynamics which can catalogue and order the infinite number of periodic orbits present in a chaotic attractor to give convergent expressions.

Extending periodic orbit theory to higher dimensional dynamical systems, most notably spatiotemporal systems, would obviously be highly desirable but represents a very considerable challenge. However, there are encouraging signs that something approaching this could be possible in fluid turbulence. The fact that a turbulent flow fleetingly yet recurrently resembles a series of smoother coherent structures or spatiotemporal patterns is a familiar observation perhaps first recorded by Leonardo da Vinci in his famous drawings. Mathematically, Hopf (Hopf 1948; see also Robinson 1991, Holmes, Lumley & Berkooz 1996 and Panton 1997 for overviews of subsequent work) made the key step forward in this direction by viewing the evolving solution of the Navier–Stokes equations as a point moving in an infinite-dimensional state space and observing that viscosity would lead to a contraction of the dynamics onto a finite-dimensional (now known as 'inertial') manifold. Within this manifold, it is then natural to view turbulence as a phase space flow transiently visiting the neighbourhoods of unstable, simple invariant (spatiotemporal) solutions of the governing equations (e.g. equilibria, periodic orbits, tori, etc.; hereafter also referred to as 'recurrent flows'). However, an attempt to build a prediction of turbulence statistics from the recurrent flows present is fraught with difficulties. Not only is there the daunting problem of initially identifying enough of them in such high-dimensional systems (typically 10^4 – 10^5 degrees of freedom (d.o.f.s)) to make such a prediction seem feasible, but there is the problem of understanding how each should be weighted in any expansion. Finally, in the very likely eventuality that there is no symbolic dynamics for turbulence, it is difficult to know if important recurrent flows have been missed thereby compromising any prediction.

The situation although very daunting, promises much and is not without hope. Efforts to extend the ideas of periodic orbit theory to higher-dimensional systems have focused on 1-space and 1-time partial differential equations, most notably the one-dimensional Kuramoto–Sivashinsky system (Christiansen, Cvitanović & Putkaradze 1997; Zoldi & Greenside 1998; Lan & Cvitanović 2008; Cvitanović, Davidchack & Siminos 2010) and the complex Ginzburg–Landau equation (Lopez *et al.* 2005). The emphasis in this work has mostly been to establish the feasibility of extracting recurrent flows directly from the 'turbulent' dynamics although some predictions were made (Christiansen *et al.* 1997; Lopez *et al.* 2005). The first attempt to extract

a recurrent flow from three-dimensional Navier–Stokes turbulence was made in a landmark calculation by Kawahara & Kida (2001). In this work they found one periodic orbit embedded in the turbulent attractor in a 15 422-d.o.f. simulation of small box plane Couette flow. This immediately raised the ‘bar’ of what had been thought possible and interestingly, they found that this one orbit was a very good proxy for their turbulence statistics. van Veen, Kawahara & Kida (2006) drew a similar conclusion albeit after discarding all but one of the few orbits they found when studying highly symmetric three-dimensional body-forced box turbulence. Subsequent work in plane Couette flow by Viswanath (2007) essentially confirmed the existence of Kawahara & Kida (2001) periodic orbit (using 180 670 d.o.f.), found another one and identified four new relative periodic orbits (see also Lopez *et al.* 2005). These are periodic orbits where the flow repeats in time but drifts spatially in directions where the system has a continuous translational symmetry. Cvitanović & Gibson (2010) report (using 61 506 d.o.f.s) having identified 40 periodic solutions, 15 relative periodic solutions with streamwise shifts and one relative periodic orbit with a small spanwise shift in low-Reynolds-number and small box plane Couette flow. More recent work has focused on finding recurrent flows in symmetric subspaces at transitional Reynolds numbers either in small box plane Couette flow (Kreilos & Eckhardt 2012) or in short pipe flow (Willis, Cvitanović & Avila 2013).

The state of the field is then that recurrent flows can be found in three-dimensional Navier–Stokes turbulence calculations requiring up to $O(10^5)$ d.o.f. (low Reynolds number turbulence) but understanding how many can be found in a reasonable (tolerable) time and then identifying how dynamically important they are, remain outstanding issues. As a result, making useful predictions with any confidence using the set of recurrent motions found seems some way off. With this background, our objective here is to make some contribution to this effort by mounting a systematic investigation of these issues in the simpler context of two-dimensional Navier–Stokes turbulence.

It’s worth emphasizing that even if a ‘turbulence’ version of periodic orbit theory ultimately proves beyond our grasp, the procedure of identifying recurrent flows buried within a turbulent solution has considerable value in its own right. This is because each recurrent flow can be thought of as a sustainable dynamical process which helps underpin the turbulent state. Since they are ‘closed’ (recur exactly), their spatial and temporal structure can be dissected to reveal the fundamental physics involved. Just such an approach has helped uncover the ‘self-sustaining process’ (Waleffe 1997) – streamwise vortices generate streaks which are unstable to streamwise-dependent flows which subsequently invigorate the streamwise vortices – in wall-bounded shear flows following the discovery of a quasi-cycle in highly constrained plane Couette flow by Hamilton, Kim & Waleffe (1995). Beautifully, this quasi-cycle turned out to indicate the presence of families of exact (unstable) travelling wave (TW) solutions to the Navier–Stokes equations (Waleffe 1998, 2001, 2003), the existence of which have revolutionized our thinking about transitional and weakly turbulent shear flows (see the reviews by Kerswell (2005), Eckhardt *et al.* (2007) and Kawahara *et al.* (2012)).

The specific framework investigated here is two-dimensional Kolmogorov flow on a $[0, 2\pi]^2$ torus (efficiently simulated using spectral methods) where the flow is forced monochromatically and steadily at a large length scale. This flow has been extensively studied since Kolmogorov introduced the model in 1959 (Arnol’d & Meshalkin 1960) as a simple example of linear instability which could be studied analytically (Meshalkin & Sinai 1961). The flow has many possible variations: torus aspect ratio (e.g. Marchioro 1986; Okamoto & Shoji 1993; Sarris *et al.* 2007), forcing wavelength

(She 1988; Platt, Sirovich & Fitzmaurice 1991; Armbruster *et al.* 1996), forcing form (e.g. Gotoh & Yamada 1986; Kim & Okamoto 2003) and three-dimensionalization (e.g. Borue & Orszag 1996; Shebalin & Woodruff 1997; Sarris *et al.* 2007). It has been experimentally realized using magnetohydrodynamic forcing (e.g. Bondarenko, Gak & Dolzhanskii 1979; Obukhov 1983; Sommeria 1986) and latterly in soap films (e.g. Burgess *et al.* 1999). With an additional Coriolis term, Kolmogorov flow can also be used as a barotropic ocean model on the β -plane (e.g. Kazantsev 1998, 2001; Tsang & Young 2008).

The work by Kazantsev (1998, 2001) is particularly relevant for this study as this was the first attempt to use periodic orbits to reproduce properties of the chaotic attractor in a 211-d.o.f. discretization of a two-dimensional Kolmogorov-like flow (differences include the addition of non-periodic boundary conditions, rotation and bottom friction). The work is most notable for his use of a minimization procedure to identify periodic orbits (59 found) as well as a good survey of relevant atmospheric literature. More recent work by Fazendeiro *et al.* (2010) (see also Boghosian *et al.* 2011) has started to study triply periodic body-forced turbulence using lattice-Boltzmann computations. Their focus was on developing another variational approach for identifying periodic orbits based upon the idea of Lan & Cvitanović (2004) and they describe convergence evidence for two periodic orbits. The approach starts with a closed orbit that does not satisfy the Navier–Stokes equations and uses a variational method to adjust the orbit until it does. This requires manipulating the whole orbit at once and requires massive computations which are facilitated by the inherent parallelism of the lattice-Boltzmann approach. In contrast, the approach adopted here is to start with a flow trajectory which does satisfy the Navier–Stokes equations but is not closed and to adjust the start of the trajectory until it does. This boils down to a Newton–Raphson root search in very high dimensions and iterative methods have to be employed to make things feasible. We adopt a Newton–GMRES–hook-step procedure developed by Viswanath (2007, 2009) and subsequently used with success by Cvitanović & Gibson (2010) (see Duguet, Pringle & Kerswell 2008 for a slight variation which replaces the ‘hook step’ with the ‘double dogleg’ step; Dennis & Schnabel 1996).

The plan of the paper is as follows. Section 2 describes two-dimensional Kolmogorov flow in detail, discusses its symmetries (§ 2.1) and makes connections with some previous direct numerical simulations (DNSs) (§ 2.2). Key flow measures to be used subsequently are listed in § 2.3. Section 3 describes the methodology used, starting with the time stepping code in § 3.1, how initial guesses for recurrent flows are identified in § 3.2, and then the Newton–GMRES–hook-step algorithm in § 3.3 (this draws its inspiration from Viswanath (2009)). Section 3.4 discusses how the algorithms were tested. Section 4 describes the results, first giving a flow orientation in § 4.1, then reporting on how recurrent flows were actually extracted, before giving details of the recurrent flows found in § 4.3. Section 5 describes an attempt to reproduce properties of two-dimensional Kolmogorov turbulence before § 6 discusses the results and the outlook for future work.

2. Formulation

The incompressible Navier–Stokes equation with what is called ‘Kolmogorov forcing’ is

$$\frac{\partial \mathbf{u}^*}{\partial t^*} + \mathbf{u}^* \cdot \nabla^* \mathbf{u}^* + \frac{1}{\rho} \nabla^* p^* = \nu \nabla^{*2} \mathbf{u}^* + \chi \sin(2\pi n y^*/L_y) \hat{\mathbf{x}}, \quad (2.1)$$

$$\nabla^* \cdot \mathbf{u}^* = 0 \quad (2.2)$$

where ρ is the density, ν the kinematic viscosity, n an integer describing the scale of the (monochromatic) Kolmogorov forcing and χ is the forcing amplitude per unit mass of fluid over a doubly periodic domain $[0, L_x] \times [0, L_y]$ (and in this section only $*$ indicates a dimensional quantity). The system is non-dimensionalized by the length scale $L_y/2\pi$ and time scale $\sqrt{L_y/2\pi\chi}$ so that the equations become

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p = \frac{1}{Re} \nabla^2 \mathbf{u} + \sin(ny) \hat{\mathbf{x}}, \tag{2.3}$$

$$\nabla \cdot \mathbf{u} = 0 \tag{2.4}$$

where the Reynolds number is

$$Re := \frac{\sqrt{\chi}}{\nu} \left(\frac{L_y}{2\pi} \right)^{3/2} \tag{2.5}$$

to be solved over the domain $[0, 2\pi/\alpha] \times [0, 2\pi]$ ($\alpha := L_y/L_x$). Given the doubly periodic boundary conditions, dealing with the cross-plane vorticity equation is more natural and reduces simply to the scalar equation

$$\frac{\partial \omega}{\partial t} = \hat{\mathbf{z}} \cdot \nabla \times (\mathbf{u} \times \omega \hat{\mathbf{z}}) + \frac{1}{Re} \nabla^2 \omega - n \cos(ny) \tag{2.6}$$

where $\omega \hat{\mathbf{z}} := \nabla \times \mathbf{u}$. (The form of the nonlinearity on the right-hand side is convenient for computation but can be further reduced to simply $-\mathbf{u} \cdot \nabla \omega$ as the vortex stretching term $\boldsymbol{\omega} \cdot \nabla \mathbf{u} = 0$ is, of course, absent in two dimensions.) Dealing with this equation is analogous to working with the streamfunction $\mathbf{u} = \nabla \times \psi(x, y) \hat{\mathbf{z}}$ since spatially constant velocity and vorticity fields are not present so $\psi = \nabla^{-2} \omega$.

2.1. Symmetries

There is a shift-and-reflect symmetry

$$\mathcal{S} : [u, v, \omega](x, y) \rightarrow [-u, v, -\omega] \left(-x, y + \frac{\pi}{n} \right). \tag{2.7}$$

which shifts half a wavelength of the forcing function in y and reflects in x ($\mathbf{u} := u\hat{\mathbf{x}} + v\hat{\mathbf{y}}$ and $\omega := \partial v/\partial x - \partial u/\partial y$). Since there are n wavelengths in the domain, this transformation forms a cyclic group of order $2n - 1$. There is also a rotation-through- π symmetry

$$\mathcal{R} : [u, v, \omega](x, y) \rightarrow [-u, -v, \omega](-x, -y) \tag{2.8}$$

and the continuous group of translations

$$\mathcal{T}_l : [u, v, \omega](x, y) \rightarrow [u, v, \omega](x + l, y) \quad \text{for } 0 \leq l < \frac{2\pi}{\alpha}. \tag{2.9}$$

The focus here is (unusually) not to take advantage of these, that is, the flow is allowed to fully explore phase space.

2.2. Past literature

Of all the previous work on two-dimensional Kolmogorov flow, Platt *et al.* (1991) seem to have carried out the most detailed study with $n = 4$ over the non-dimensional domain $[0, 2\pi] \times [0, 2\pi]$. The same choices $n = 4$ and $\alpha = 1$ were therefore made throughout the calculations reported here. With this, $Re = 8\sqrt{Re^{latt}}$ and so the critical Reynolds number for linear instability is $Re_c = 8\sqrt[4]{2}$ ($Re_c^{latt} = \sqrt{2}$). Platt *et al.*

(1991) looked at the flow regime $Re/Re_c \leq 3.54$ over a 32×32 spatial grid so that $9.51 \leq Re \leq 33.6$. Here we consider a 256×256 grid and look at $9.5 \leq Re \leq 100$ ($Re/Re_c \leq 10.5$). Unfortunately, we were only able to confirm the detailed dynamics reported by Platt *et al.* (1991) if we reduced our resolution down to theirs.

The next closest study was She's (1988) which took $n = 8$, a 64×64 grid and examined $\sqrt{2} \leq Re_{She} \leq 30$ ($26.9 \leq Re \leq 123.9$ as $Re = 8^{3/2} \sqrt{Re_{She}}$) which corresponds to $Re/Re_c \leq 4.6$. More recently, Sarris *et al.* (2007) considered three-dimensional Kolmogorov flow over a variety of box aspect ratios considering $65 \leq Re^{Sarris} \leq 180$ including the three-dimensionalization of the flow considered here (then $Re = 8Re^{Sarris}$). Typically they use 128 mesh points per wavelength of the forcing. At the time of writing, the world record for resolution when simulating doubly periodic body-forced turbulence seems to be $32, 768^2$ (Boffetta & Musacchio 2010).

2.3. Key measures of the flow

Key measures of the flow which will aid the subsequent discussion are as follows: the mean flow,

$$\bar{u}(y) := \langle \mathbf{u} \cdot \hat{\mathbf{x}} \rangle_{t,x}, \tag{2.10}$$

(initial conditions are such that $\langle \mathbf{u}(\mathbf{x}, 0) \cdot \hat{\mathbf{y}} \rangle_x = 0$ so that $\langle \mathbf{u}(\mathbf{x}, t) \cdot \hat{\mathbf{y}} \rangle_x = 0$ for all time); the bulk mean square of the fluctuations around the mean,

$$\hat{u}_{rms}^2(t) := \langle (u - \bar{u})^2 \rangle_V, \quad \hat{v}_{rms}^2(t) := \langle v^2 \rangle_V, \tag{2.11}$$

and root mean square (r.m.s.) of the fluctuations as a function of y ,

$$u_{rms}(y) := \sqrt{\langle (u - \bar{u})^2 \rangle_{t,x}}, \quad v_{rms}(y) := \sqrt{\langle v^2 \rangle_{t,x}}; \tag{2.12}$$

the total kinetic energy and the kinetic energy of the fluctuation field

$$E(t) := \frac{1}{2} \langle \mathbf{u}^2 \rangle_V, \quad E_t(t) := \frac{1}{2} \langle (\mathbf{u} - \langle \mathbf{u} \rangle_{t,x})^2 \rangle_V; \tag{2.13}$$

the total dissipation rate and the instantaneous power input

$$D(t) := \frac{1}{Re} \langle |\nabla \mathbf{u}|^2 \rangle_V, \quad I(t) := \langle u \sin(ny) \rangle_V, \tag{2.14}$$

with finally the laminar state, bulk laminar kinetic energy and bulk dissipation rate

$$\mathbf{u}_{lam} := \frac{Re}{n^2} \sin ny \hat{\mathbf{x}}, \quad E_{lam} := \frac{Re^2}{4n^4}, \quad D_{lam} := \frac{Re}{2n^2}, \tag{2.15}$$

where the various averagings are defined as

$$\langle \rangle_V := \frac{\alpha}{4\pi^2} \int_0^{2\pi/\alpha} \int_0^{2\pi} dx dy, \tag{2.16a}$$

$$\langle \rangle_x := \frac{\alpha}{2\pi} \int_0^{2\pi/\alpha} dx, \tag{2.16b}$$

$$\langle \rangle_t := \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt. \tag{2.16c}$$

3. Methodology

3.1. Time stepping code

A two-dimensional fully de-aliased pseudo-spectral code was used as developed in Bartello & Warn (1996). The original Leapfrog+filter approach was replaced by the Crank–Nicolson method for the viscous terms and Heun’s method (Euler predictor method) for the nonlinear and forcing terms so that only one state vector was required to accurately restart the code. This together with a constant time step size (except for the last step) means that the discretized flow is a dynamical system which closely matches the Navier–Stokes flow. Specifically, if $\Omega(\mathbf{k}) = \text{fft}(\omega(\mathbf{x}))$ is the Fourier transform of ω with $\mathbf{k} = (k_x, k_y)$, the vorticity equation (2.6) in spectral space is

$$\frac{\partial \Omega}{\partial t} = -G\Omega + f(\Omega) \tag{3.1}$$

with

$$G(k_x, k_y) := \frac{k_x^2 + k_y^2}{Re} = \frac{|\mathbf{k}|^2}{Re} \tag{3.2}$$

and

$$f(\Omega) := -i(k_x \text{fft}[u\omega] + k_y \text{fft}[v\omega]) - \frac{n}{2} \delta_{k_x} \delta_{(|k_y|-n)}. \tag{3.3}$$

Here δ_i is the Kronecker delta function and takes the value 1 when $i = 0$ and 0 otherwise. A time step is performed by solving

$$\frac{\tilde{\Omega}^{i+1} - \Omega^i}{\Delta t} = -\frac{G}{2} (\tilde{\Omega}^{i+1} + \Omega^i) + f(\Omega^i) \tag{3.4}$$

followed by solving

$$\frac{\Omega^{i+1} - \Omega^i}{\Delta t} = -\underbrace{\frac{G}{2} (\Omega^{i+1} + \Omega^i)}_{C-N} + \underbrace{\frac{1}{2} (f(\tilde{\Omega}^{i+1}) + f(\Omega^i))}_{Heun} \tag{3.5}$$

where the superscript is a time step index. With de-aliasing, a resolution of $N_x \times N_y$ corresponds in practice to the vorticity representation

$$\omega(x, y, t) = \sum_{j=-N_x/3}^{N_x/3-1} \sum_{l=-N_y/3}^{N_y/3-1} \Omega_{jl}(t) e^{i(\alpha jx + ly)} \tag{3.6}$$

where $\mathbf{k} = (\alpha j, l)$, $\Omega_{00} = 0$ and a mask is employed so that $\Omega_{jl} = 0$ for wavenumbers outside a specified domain Σ . Calculations reported here have $\alpha = 1$, $n = 4$ and $N_x = N_y$, so $\Sigma := \{(j, l) : j^2 + l^2 \leq (N_x/3)^2\}$ is used. The number of active (real) d.o.f.s is therefore $\approx \pi N_x^2/9$ which is $\approx 22\,800$ (or exactly 22 428) for the $N_x = 256$ used here ($0 \leq j \leq 85$ and $-85 \leq l \leq 85$ since $\Omega(-j, -l) = \Omega^*(j, l)$).

3.2. Near recurrences

The key idea pursued here is to extract recurrent flows directly from the turbulent DNS data with the implication that they are clearly dynamically important. With this in mind, the time stepping code was run for 10^5 time units starting from random initial conditions and ‘near recurrences’ of the flow field searched for. These near recurrences

were defined as episodes where

$$\mathcal{I}_s \mathcal{S}^{2m} \omega(x, y, t + T) := \omega(x + s, y + \frac{1}{2}\pi m, t + T) = \omega(x, y, t) \tag{3.7}$$

‘approximately’ holds for some choice of the continuous shift $0 \leq s < 2\pi$, the discrete shift $m \in 0, 1, 2, \dots, n - 1$ and $T > 0$ over $0 \leq x, y < 2\pi$. Periodic orbits correspond to $s = m = 0$ and some period $T > 0$, TWs to $m = 0$ and $s = cT$ with $T > 0$ free where c is the phase speed, equilibria have $s = m = 0$ and T free and relative periodic orbits have one or both of s and m not equal to zero with period $T > 0$. The existence of relative periodic orbits, permitted by the inclusion of two free parameters (s and m) here, is the signature feature of systems with discrete and/or continuous translational symmetries. The presence of the further symmetry (2.8) means we could also have searched separately for relative periodic orbits with reflection

$$\mathcal{R} \mathcal{I}_s \mathcal{S}^{2m} \omega(x, y, t + T) = \omega(-x - s, -y - \frac{1}{2}\pi m, t + T) = \omega(x, y, t). \tag{3.8}$$

This extra search was not done for reasons of expediency since these flows are ‘preperiodic’ (Cvitanović *et al.* 2010) to relative periodic orbits of period $2T$ and are therefore picked up by checking (3.7). This was also the rationale behind not extending the recurrence search to include

$$\mathcal{I}_s \mathcal{S}^{2m+1} \omega(x, y, t + T) := -\omega(-x + s, y + \frac{1}{4}\pi(2m + 1), t + T) = \omega(x, y, t). \tag{3.9}$$

Again, these recurrent flows are captured as relative periodic orbits by (3.7) albeit after integrating over two periods (presuming that they are not too unstable).

The key to this search is to understand how approximately (3.7) should hold to signify the presence of a recurrent flow structure nearby. The only way to answer this seems to be to perform computations and experiments. The search for near recurrences was done most efficiently by calculating every, say $t = 0.1$ or 0.2 steps, the normalized difference between states in wavenumber space suitably minimized over continuous shifts in x and discrete shifts in y as follows:

$$R(t, T) := \min_{0 \leq s < 2\pi} \min_{m \in 0, 1, 2, \dots, n-1} \frac{\sum_j \sum_l |\Omega_{jl}(t) e^{i\alpha js + 2iml\pi/n} - \Omega_{jl}(t - T)|^2}{\sum_j \sum_l |\Omega_{jl}(t)|^2} \tag{3.10}$$

where $\sum_j \sum_l |\Omega_{jl}|^2 = \alpha / (4\pi)^2 \int_0^{2\pi/x} \int_0^{2\pi} \omega^2 dx dy$ and $0 < T_{thres} \approx 0.5 < T < 100$. Since $R(t, 0) = 0$ and $dR(t, 0)/dT > 0$, the offset T_{thres} is defined adaptively as the first time at which $dR(t, T_{thres})/dT < 0$. Figure 1 is a typical example of how $R(t, T)$ looks as a function of t and T during a recurrent episode. The nine black dots are the guesses identified by the code ($R < R_{thres} = 0.3$) over this time interval. All except one (the last dot at $t \approx 171$) subsequently converged to an exactly recurrent solution (the four dots for $t < 130$ to a periodic orbit (P1 in table 2) with period 5.3807 and the next four dots with $t \in [130, 160]$ to a TW (T1 in table 2) with phase speed $c = 0.0198$). The threshold R_{thres} was chosen judiciously to give enough good quality guesses.

3.3. UPO extraction method: Newton-GMRES-hook step

Once a near recurrence has been found by the above-stated criterion, we then attempted to find whether an exact recurrent flow was lurking nearby in phase space. This required a high-dimensional root finding algorithm acting on a state vector which




FIGURE 1. (Colour online) Example of $R(t, T)$ at $Re = 40$ contoured over $t \in [95, 175]$ (x -axis) and $T \in [thres, 50]$ (y -axis). The nine blue dots are the guesses identified by the code ($R < R_{thres}$) over this time interval. All except one (the last dot at $t \approx 171$) subsequently converged to exact solutions: the four dots for $t < 130$ to a periodic orbit with period 5.3807 ($P1$ in table 2) and the next four dots with $t \in [130, 160]$ to a TW ($T1$ in table 2) with phase speed $c = 0.0198$. The R values above 0.55 are not drawn/coloured for clarity.

completely specifies the velocity field

$$X = \begin{bmatrix} \Omega \\ s \\ T \end{bmatrix} \tag{3.11}$$

and contains information about the potential recurrence (Ω is a vector containing the scalars Ω_{jl} arranged in some fashion). The shift s is included since it can be adjusted continuously whereas the discrete shift m cannot and therefore is preset. To set up the Newton–Raphson algorithm (and we follow the excellent description by Viswanath 2009), it is convenient to define the infinitesimal generators T_x and T_y of translations in x and y , respectively,

$$T_x \omega(x, y, t) \rightarrow \frac{\partial \omega}{\partial x} = \sum_{-N_y/3}^{N_y/3-1} \sum_{-N_x/3}^{N_x/3-1} i\alpha_j \Omega_{jl}(t) e^{i(\alpha_j x + l y)}, \tag{3.12a}$$

$$T_y \omega(x, y, t) \rightarrow \frac{\partial \omega}{\partial y} = \sum_{-N_y/3}^{N_y/3-1} \sum_{-N_x/3}^{N_x/3-1} i l \Omega_{jl}(t) e^{i(\alpha_j x + l y)} \tag{3.12b}$$

as they act in spectral space

$$T_x \Omega \rightarrow \Omega_x \quad \text{and} \quad T_y \Omega \rightarrow \Omega_y \tag{3.13}$$

where each element Ω_{jl} of Ω is mapped to $i\alpha_j\Omega_{jl}$ in Ω_x and $il\Omega_{jl}$ in Ω_y . Then, in spectral space the recurrence condition (3.7) becomes

$$F(\Omega_0, s, T; m) := \exp(sT_x + \frac{1}{2}\pi m T_y)\hat{\Omega}(\Omega_0, T) - \Omega_0 = \mathbf{0} \tag{3.14}$$

where $\Omega_0 = \Omega(t)$ and $\hat{\Omega} = \Omega(t + T)$. If $X_0 = (\Omega_0, s_0, T_0)^T$ is an initial guess for a solution, then a better (next) guess $X_0 + \delta X_0 = (\Omega_0 + \delta\Omega, s_0 + \delta s, T_0 + \delta T)^T$ is given by

$$\frac{\partial F}{\partial \Omega_0} \delta\Omega + \frac{\partial F}{\partial s} \delta s + \frac{\partial F}{\partial T} \delta T = -F(\Omega_0, s_0, T_0; m) \tag{3.15}$$

These are $\dim(\Omega)$ equations for $\dim(\Omega) + 2$ unknowns. The extra two equations come from removing the degeneracy associated with these translational symmetries (the system is invariant under $(x, t) \rightarrow (x + s, t + T)$). This can be done by imposing that $\delta\Omega$, has no component which shifts the solution infinitesimally in the x -direction or the t -direction (i.e. just redefines the time origin of the flow). The Newton–Raphson problem is then to solve

$$\begin{bmatrix} \ddots & & & & & \\ & \frac{\partial \hat{\Omega}_s}{\partial \Omega_0} - \mathbf{I} & & & & \\ & & \ddots & & & \\ \dots & (T_x \Omega_0)^T & \dots & 0 & 0 & \\ \dots & \frac{\partial \Omega_0^T}{\partial t} & \dots & 0 & 0 & \end{bmatrix} \begin{bmatrix} \delta\Omega \\ \delta s \\ \delta T \end{bmatrix} = - \begin{bmatrix} F(\Omega_0, s_0, T_0; m) \\ 0 \\ 0 \end{bmatrix} \tag{3.16}$$

where $\hat{\Omega}_s := \exp(sT_x + (1/2)\pi m T_y)\hat{\Omega}$ is the ‘back-shifted’ final state and \mathbf{I} is the $\dim(\Omega) \times \dim(\Omega)$ identity matrix. This is now in the standard form $A\delta X = b$ with only the Jacobian matrix $\partial \hat{\Omega}_s / \partial \Omega$ not straightforward to evaluate ($\partial \hat{\Omega}_s / \partial T$ and $\partial \Omega_0 / \partial t$ are found by substituting $\hat{\Omega}_s$ or Ω_0 into the Navier–Stokes equations).

Typically, the size of the matrix A is too large to store explicitly let alone attempt to solve $A\delta X = b$ directly. As a result, the only way to proceed is iteratively and GMRES (Saad & Schultz 1986) is convenient (see the excellent description by Trefethen & Bau (1997)). Here only the effect of A on an arbitrary vector is needed. The effect of the troublesome Jacobian can be handled easily by a forward difference approach since

$$\frac{\partial \hat{\Omega}_s}{\partial \Omega_0} \mathbf{y} \approx \frac{\hat{\Omega}_s(\Omega_0 + \epsilon \mathbf{y}) - \hat{\Omega}_s(\Omega_0)}{\epsilon} \tag{3.17}$$

where ϵ is chosen such that $\|\epsilon \mathbf{y}\| = 10^{-7} \|\Omega_0\|$ which balances truncation error with round-off error using double precision arithmetic and $\|\cdot\|$ is the Euclidean norm (using a more physically orientated norm is clearly an interesting direction awaiting exploration).

Straight Newton–GMRES is typically not good enough as guesses are usually not in the region where linearization holds sufficiently well and divergence to infinity is commonplace. Instead it proves useful to modify the approach to incorporate a trust

region. Following Viswanath (2007, 2009), we use the ‘hook-step’ method (Dennis & Schnabel 1996, §6.4.1) which can be easily built on top of the GMRES process. Exactly how the approach is implemented can vary and we adopt what looks to be a slightly different algorithm to Viswanath (2007, 2009) in which GMRES is used first to derive an approximate solution to $A\delta X = b$ before this ‘solution’ δX is moved into the trust region. The advantage of this is there is a clear convergence criterion that can be imposed to terminate the initial GMRES algorithm. Before stating this, it is worth first briefly describing the GMRES algorithm itself which is based upon a simple idea. The GMRES algorithm for solving $A\delta X = b$ at iteration n approximates δX by the vector δX_n in the Krylov space $\mathcal{K}_n := \langle b, Ab, A^2b, \dots, A^{n-1}b \rangle$ that minimizes the norm of the residual

$$\|A\delta X_n - b\| \tag{3.18}$$

(Trefethen & Bau 1997). For numerical stability, an orthonormal basis for \mathcal{K}_n is constructed using a Gram–Schmidt-style iteration as follows

$$q_1 = \frac{b}{\|b\|}, \quad \tilde{q}_{i+1} = Aq_i - \sum_{j=1}^i q_j \frac{q_j^T Aq_i}{\|q_j\|^2} \quad q_{i+1} = \tilde{q}_{i+1} / \|\tilde{q}_{i+1}\| \quad i \in 1 \dots, n - 1 \tag{3.19}$$

so that if Q_n is the matrix with columns q_1, q_2, \dots, q_n , then $AQ_n = Q_{n+1}H_{n+1,n}$ where $H_{n+1,n}$ is the upper $(n + 1) \times n$ left section of an upper Hessenberg matrix generated by the basis orthonormalization (Trefethen & Bau 1997, p. 252). With this basis the solution $\delta X_n = Q_n y_n$ (where y_n is an n -vector) minimizes $\|Q_{n+1}H_{n+1,n}y_n - b\|$ (N equations and n unknowns) or equivalently $\|H_{n+1,n}y - \|b\|\hat{e}_1\|$ ($n + 1$ equations and n unknowns) since the only non-zero entry in $Q_{n+1}^T b$ is the first entry. This can be accomplished by a singular value decomposition (SVD) of $H_{n+1,n}$ into $U_{n+1}DV_n^T$ (where U_{n+1} and V_n are orthonormal matrices and D is an $(n + 1) \times n$ diagonal matrix with a zeroed bottom row) through straightforwardly solving the first n equations $Dz_n = p_{n+1} := \|b\|U_{n+1}^T \hat{e}_1$ followed by $V_n^T y_n = z_n$ and then $\delta X_n = Q_n y_n$. The modulus of the remaining unbalanced component $p_{n+1}(n + 1)$ then gives the minimum value or residual. The iterations are continued until

$$\frac{\|A\delta X_n - b\|}{\|b\|} = \frac{\|Dz_n - p_{n+1}\|}{\|p_{n+1}\|} = \frac{|p_{n+1}(n + 1)|}{\|p_{n+1}\|} \leq tol \tag{3.20}$$

where tol is a small number typically chosen in the range 10^{-4} to 10^{-2} (the majority of the computations reported here were obtained using a value of 10^{-3}). If $\|F(X + \delta X_n)\|$ is not smaller than $\|F(X)\|$ or more specifically not well predicted by the linearization around X , i.e. $\|F(X) + A\delta X_n\|$, then the approximate solution of the linearized problem is transformed back to a smaller trust region where the linearized problem is valid. This is done by adding the constraint $\|\delta X_n\| \leq \Delta$ or equivalently $\|y_n\| \leq \Delta$ to the GMRES minimization (3.18): this is the hook step. The beauty of this adjustment is that it is a very natural modification of the GMRES approximate solution since the (innermost) problem for z_n is then

$$\min \|Dz_n - p_{n+1}\| \quad \text{s.t.} \quad \|y_n\| = \|z_n\| \leq \Delta. \tag{3.21}$$

Constructing the Lagrangian

$$\mathcal{L} := (Dz_n - p_{n+1})^2 + \mu(z_n^2 + \beta^2 - \Delta^2) \tag{3.22}$$

where μ is a Lagrange multiplier imposing the trust region constraint, leads to the minimization equations

$$2d(i)(d(i)z_n(i) - p_{n+1}(i)) + 2\mu z_n(i) = 0 \quad (3.23)$$

$$2\mu\beta = 0 \quad (3.24)$$

$$z_n^2 + \beta^2 - \Delta^2 = 0 \quad (3.25)$$

where d_i is the i th diagonal element of \mathbf{D} . The solution to this is

$$z_n(i) := \frac{p_{n+1}(i)d_i}{d_i^2 + \mu} \quad 1 \leq i \leq n \quad (3.26)$$

with either $\mu = 0$ as $\|z_n\| < \Delta$ (the original GMRES solution) or $\mu \neq 0$ chosen so that $\|z_n\| = \Delta$ (in practice μ is just increased until $\|z_n\| < \Delta$). An acceptable solution $\delta\mathbf{X}_n$ is signalled by

$$\mathbf{F}(\mathbf{X} + \delta\mathbf{X}_n) \leq \mathbf{F}(\mathbf{X}) + c\mathbf{A}\delta\mathbf{X}_n \quad (3.27)$$

where some value of $c \in (0, 0.5)$ is chosen (see Dennis & Schnabel 1996, (6.4.14): we took the least demanding value of $c = 0$). If this does not hold, Δ is decreased and the hook step repeated until it is. Depending on how easily this improvement condition is met, the trust region may be relaxed (e.g. if linearization holds well, $\mathbf{F}(\mathbf{X} + \delta\mathbf{X}_n) \approx \mathbf{F}(\mathbf{X}) + \mathbf{A}\delta\mathbf{X}_n$) or not for subsequent Newton steps.

This algorithm can be readily extended to perform solution branch continuation: see Appendix. Furthermore, since we know how to calculate the action of the Jacobian on any vector (see (3.17)), the linear stability of an exactly recurrent flow can also be readily found using the Arnoldi technique (e.g. using ARPACK to extract extremal eigenvalues).

3.4. Testing

The modified (Crank–Nicholson+ Heun) time stepping code was thoroughly validated against the well-tested Leapfrog+filter code developed by Bartello & Warn (1996). The Newton–GMRES–hook-step algorithm developed on top of this was tested by attempting to converge onto a known periodic orbit. This orbit was originally found by tracing bifurcations up from the basic state. For $n = 4$, the one-dimensional basic state becomes linearly unstable at $Re = 9.9669$ for disturbances 2π -periodic in x giving rise to a steady two-dimensional state which is \mathcal{R} -symmetric. This state loses stability to a stable periodic orbit within the \mathcal{R} -symmetric subspace for $30 < Re < 31$ before this orbit becomes unstable at $Re \gtrsim 32$ through a torus bifurcation. The periodic orbit at $Re = 31$ was easily found by time stepping within the \mathcal{R} -symmetric subspace yet is unstable to \mathcal{R} -asymmetric disturbances in the full unrestricted space. Having such an orbit to experiment with was invaluable for building up confidence in the code and some feel for how the tolerances of the algorithm should be set (e.g. tol in (3.20)).

4. Results

4.1. Flow orientation

Two-dimensional Kolmogorov flow is linearly unstable at a comparatively low Re which depends strongly on the imposed periodicity in the forcing direction: see figure 2. For the domain studied here ($\alpha = 1$), disturbances to the base flow (2.15) fail to decay monotonically at $Re_E = 6.8297$ and then start to grow exponentially at $Re_{lin} = 9.9669$. Figure 3 shows that this initial bifurcation is to a steady




FIGURE 2. (Colour online) The energy stability Reynolds number Re_E (dashed line, shown in red online) and the linear instability Reynolds number Re_{lin} (solid line, shown in blue online) for $\sin 4y\hat{x}$ forcing over the torus $[0, 2\pi/\alpha] \times [0, 2\pi]$. Note that $Re_{lin} \rightarrow \infty$ as $\alpha \rightarrow 4$ so the domain is squeezed down to $[0, \pi/2] \times [0, 2\pi]$. The dotted line at $\alpha = 1$ is the present case ($Re_E = 6.8297$ and $Re_{lin} = 9.9669$). Here $Re_E \rightarrow 0$ and $Re_{lin} \rightarrow 8\sqrt[4]{2} \approx 9.5137$ as $\alpha \rightarrow 0$.

flow ($D/D_{lam} < 1$ and $E_t/E = 0$) until $Re \approx 15$ whereupon time dependence appears (note instabilities tend to decrease the dissipation rate because the flow is body-forced). For $15 \lesssim Re \lesssim 23$, some metastability is noticed which is illustrated in figure 4 at $Re = 22$ for two different initial conditions. One leads to a chaotic-looking dissipation signal across the time interval $[500, 1500]$ whereas the other drops out of this chaotic state at just over $t = 1000$ to converge on a stable TW solution (later named T1). Beyond $Re \approx 23$, the chaotic state presumably becomes an attractor or the probability of dropping out of this state becomes so small that it is not picked up over the time windows studied (10^3 units here and 10^5 later). Finally, an asymptotic regime is approached for $Re \gtrsim 50$. The preliminary calculations performed here for $100 < Re < 200$ tentatively support the asymptotic scaling laws $D \sim Re^{-1/2}$ and $U := \sqrt{2E} \sim Re^{2/5}$ although the noisy data clearly warrants much longer time averaging to confirm this.

Given this general flow behaviour, we chose to concentrate on analysing the flow at $Re = 40$ (approaching the asymptotic regime), and three values, $Re = 60, 80$ and 100 , which get deeper into the asymptotic regime. Figures 4 and 5 give an idea of the temporal and spatial scales in the flow at the two extremes, $Re = 40$ and $Re = 100$, of our study. Both indicate a hierarchy of temporal and spatial scales (which broaden with Re) indicative of two-dimensional turbulence. Figure 6 confirms that the flows studied for $Re \leq 100$ are well-resolved: there is 10 orders of drop off in the enstrophy spectrum in the most demanding case ($Re = 100$) for the 256^2 resolution used throughout this work.

4.2. Finding recurrent structures

A standard hunt for recurrent flows involved integrating the flow from random initial conditions for a period of 10^5 time units. Initially, R_{thres} was set at 0.15 for three runs




FIGURE 3. (Colour online) The normalized dissipation D/D_{lam} (solid line with dots, shown in blue online) and fractional fluctuation kinetic energy E_t/E (dotted line connecting dots, shown in red online) to illustrate how the flow changes with Re . The first bifurcation at $Re_{in} = 9.9669$ is steady so D/D_{lam} decreases below 1 but E_t/E remains 0 until time-dependent flows appear at around $Re = 15$. Curves were generated by using a random initial condition, running for 500 time units and then calculating averages over the next 1000 time units. This produces a unique average except over the interval $15 \lesssim Re \lesssim 23$ where there is some metastability as shown. Inset: $Re^{1/2}D$ (upper line, shown in blue online) and $3Re^{-2/5}E^{1/2}$ (lower thick line, shown in red online) against Re on a log scale showing the apparent start of the asymptotic regime $D \sim Re^{-1/2}$ and $U := \sqrt{2E} \sim Re^{2/5}$. The downward sloping lowest line shows $U/Re^{1/2}$ for comparison.

at $Re = 40$ (labelled a , b and c in table 1) which produced only 9, 7 and 13 guesses respectively. Relaxing R_{thres} to 0.3 (run d), however, produced 885. This threshold value proved adequate at $Re = 60$ (nearly 300 near recurrences detected over runs e , f and g) but had to be further relaxed to 0.35 at $Re = 80$ and 0.4 for $Re = 100$: see table 1. Unfortunately, it was noticed after these (series A) runs had been completed and the guesses tested for convergence that only $s = m = 0$ shifts had been searched over. So the runs were repeated (series B runs o , p , q and r) searching specifically for recurrences which selected either $s \neq 0$ and/or $m \neq 0$ to minimize R . This was done to indicate the frequency of observing strictly periodic near-recurrences and relative periodic near-recurrences.

The initial trawl for near recurrences took a few weeks (each case run on a Xeon X5670 processor) with the DNS code slowed considerably by the need to search for near-recurrences every 0.1 or 0.2 units in time (which is anything from 20 to 100 numerical time steps). The more time-consuming activity, however, was attempting to converge the near-recurrent guesses to exact solutions. Adopting fairly conservative limits for the Newton–GMRES–hook-step procedure (the maximum period considered was 100, maximum number of Newton, GMRES and hook steps were 75, 500 and 50, respectively) typically lead to run times of a couple of months for each of the $Re = 60$, 80 and 100 runs. The data for $Re = 40$ had to be subdivided 12 ways to make the process manageable. These numbers make it clear why a very efficient DNS code was important for this work.




FIGURE 4. (Colour online) The normalized dissipation $D(t)/D_{lam}$ versus $t \in [500, 1500]$ for $Re = 22$ (two different initial conditions: top solid line, shown in blue online, and dashed line, shown in green online), $Re = 40$ (second lowest solid line, shown in red online) and $Re = 100$ (lowest black line). Here $Re = 22$ shows chaotic saddle behaviour with one trajectory seemingly dropping out back randomly to the TW $T1$ (stable at $Re = 22$).

Table 1 also indicates the conversion rate of near-recurrence guesses to exactly recurrent solutions. There is considerable duplication of such solutions so that a much smaller set of distinct recurrent structures is obtained.

4.3. Recurrent structures found

4.3.1. $Re = 40$

Table 2 lists the recurrent structures found at $Re = 40$. The equilibrium flow $E1$ (see figure 7), which was found many times in the series A runs, corresponds to the \mathcal{R} -symmetric steady state which bifurcates off the basic solution at $Re = 9.9669$ as shown in figure 3. $E1$ loses stability at about $Re = 15$ to the TW $T1$ or later via $P1$ in the \mathcal{R} -symmetric subspace for a $Re \in (30, 31)$ ($P2$ which is \mathcal{R} -symmetric and $P3$ which is not bifurcate at yet higher Re from $E1$). All three flows, $E1$, $P1$ and $T1$ are found to be repeatedly visited by the (series A) DNS indicating the strong influence of the \mathcal{R} -symmetric subspace on the ‘turbulent’ dynamics despite them all being unstable (e.g. $E1$ has 9 unstable directions at $Re = 40$; see table 2). However, a further 47 recurrent flows were also identified from the DNS: another TW $T2$, two further periodic orbits $P2$ and $P3$, and 44 relative periodic orbits, $R1$ – $R6$ and $R18$ – $R55$ (note all have a non-zero shift s and some also a non-zero integer m). *A priori*, we expected to find mainly small period recurrent structures due to the method of extraction. Longer periods mean more time for the turbulent trajectory to diverge away from the unstable recurrent flow and hence a higher probability for: (a) the episode to escape detection as a nearly recurrent flow and; (b) even if detected, for GMRES to fail to converge due to the quality of the initial approximation. This seems borne out by the periodic orbits found but not for the relative period orbits where the majority have a period over 20 and some over 50 time units. That such long period structures exist and were ‘extractable’ from the DNS frankly was a surprise and begs




FIGURE 5. (Colour online) A time sequence of vorticity plots at $Re = 40$ (a) and $Re = 100$ (b) over the $2\pi \times 2\pi$ domain. Time difference between plots is 1 time unit and 10 contours are drawn between the maximum (light, shown in white online) and minimum values (dark, shown in red online) of the perturbation vorticity $-10.25 \leq \omega \leq 8.6$ at $Re = 40$ and $-18.4 \leq \omega \leq 20.3$ at $Re = 100$.

the question whether our ‘long’ runs of 10^5 time units (now known to be only a factor of $O(1000)$ longer than some recurrent flows) were actually really long enough to capture all of the structures possible. This issue will be raised again below.

With so many recurrent flows found, it becomes impractical to display and characterize each flow separately. Table 2 lists some key characteristics along with their stability information (all are unstable but none with more than 9 unstable directions out of 22,428 possible directions). One useful projection, however, is the ‘energy out ($D(t)$) versus energy in ($I(t)$)’ plot which is shown in figure 8 (both quantities normalized by D_{lam}). The line $D = I$ corresponds to dissipation exactly balancing energy input which has to be the case over all times for equilibria and TWs (which are just equilibria in an appropriate Galilean frame): these are therefore just points on this line in this plot. Figure 8 shows how a representative subset of these recurrent flows look when compared with the joint dissipation-input probability density function (p.d.f.) of the DNS. The darkest shading makes it clear that the DNS stays predominantly in the region $0.055 < I/D_{lam} < 0.115$, $0.06 < D/D_{lam} < 0.11$. The recurrent flows shown are also dominantly concentrated in this region although there are two relative periodic orbits shown, $R26$ and $R50$, which have large dissipation episodes (it is worth emphasizing that the basic state would be represented by the point $(1, 1)$ in this plot so the turbulent flow adopts a much reduced dissipative state). Since this D versus I plot is such a drastic projection of the dynamics, the fact that two flows look close there does not necessarily mean they are close in the full phase space (see Willis *et al.* 2013 for further discussion). However, because all




FIGURE 6. (Colour online) The enstrophy spectrum for the rightmost snapshots in figure 5 at $Re = 40$ (circles, shown in blue online) and $Re = 100$ (crosses, shown in red online) $D(n) := \sum_{n-(1/2) \leq \sqrt{j^2 + l^2} < n+(1/2)} |\Omega_{jl}|^2$ for $n = 1, 2, \dots, 85$: recall definition ((3.6) with $\alpha = 1$; e.g. there are 264 and 516 wavenumbers included for $n = 40$ and 84, respectively). The energy spectrum $E(n) \approx D(n)/n^2$ has a steeper drop off. The dotted line indicates the wavenumber of energy injection.

of the recurrent flows discussed here have been extracted from turbulent DNSs, this conclusion nevertheless seems reasonable.

In figure 9 we focus on one typical ‘embedded’ relative periodic orbit $R25$ which stays within the central region of the DNS joint p.d.f. There is a clear temporal cycle where the energy input increases (exceeding the dissipation) and then decreases (now exceeded by the dissipation). Plotting the associated vorticity fields over this cycle (figure 10) shows the character of the flow. At the dissipation low point (time 17 in figures 9 and 10), the vorticity is concentrated into weak y -aligned patches which are separated from each other whereas at the high dissipation point (time 8), the vorticity seems to be undergoing a shearing episode with only one stronger vortex recognizable. These two extremes bear more than a passing resemblance to either $T1$ ($D/D_{lam} = 0.071$) or $T2$ (0.071) and $E1$ (0.102) respectively suggesting that $R25$ is probably a closed trajectory linking their neighbourhoods. In contrast, $R50$ undergoes a large high dissipation excursion as shown more completely in figure 11. The associated vorticity fields (see figure 12) show similar structures to $R25$ when in the same part of (I, D) space (compare $t = 5$ for $R25$ with $t = 0$ for $R50$, and $t = 15$ for $R25$ and $t = 31$ for $R50$) but $R50$ exhibits intense shearing too and vortex break-up at times 5, 8 and 9. So $R50$ clearly reflects an important but infrequent aspect of the turbulent dynamics as indicated by the fact that the joint (D, I) p.d.f. of the DNS stretches to such high values of the dissipation. Whether we have extracted enough of such recurrent structures to capture this episodic behaviour is of course a key issue for this study and will be discussed in § 6.

Figure 13 is an attempt to show more of the recurrent structures found by zooming in on the central dashed box drawn in figure 8. This illustrates the intricacy of most of the flows found: many of the relative periodic orbits trace complicated D – I curves

	Run	R_{thres}	dt	Duration	Number of guesses	Number of convergences
Series A	$(s = m = 0)$					
$Re = 40$	<i>a</i>	0.15	0.005	10^5	9	5
	<i>b</i>	0.15	0.005	10^5	7	3
	<i>c</i>	0.15	0.005	10^5	13	5
	<i>d</i>	0.30	0.005	10^5	885	553
$Re = 60$	<i>e</i>	0.30	0.003	10^5	102	64
	<i>f</i>	0.30	0.003	10^5	104	67
	<i>g</i>	0.30	0.003	10^5	78	58
$Re = 80$	<i>h</i>	0.35	0.0025	10^5	53	31
	<i>i</i>	0.35	0.0025	10^5	60	37
	<i>j</i>	0.35	0.0025	10^5	41	25
$Re = 100$	<i>l</i>	0.4	0.002	10^5	75	34
	<i>m</i>	0.4	0.002	10^5	91	33
	<i>n</i>	0.4	0.002	10^5	93	42
	Series B	$(s \neq 0 \text{ or } m \neq 0)$				
$Re = 40$	<i>o</i>	0.30	0.005	10^5	1223	540
$Re = 60$	<i>p</i>	0.30	0.003	3×10^5	163	7
$Re = 80$	<i>q</i>	0.35	0.0025	3×10^5	66	15
$Re = 100$	<i>r</i>	0.4	0.002	3×10^5	84	12

TABLE 1. DNS data used to extract equilibria, TWs, periodic orbits and relative periodic orbits at $Re = 40, 60, 80$ and 100 . The value of R_{thres} cannot be set too ambitiously. Run *d* yielded all of the solutions thrown up by runs *a, b* and *c* combined.

whereas, in contrast, the periodic orbits are simple loops. Another key observation is that some relative periodic orbits look very similar, e.g. *R28* and *R29* (and other pairings not shown). This, of course, resonates with the mental picture one has of periodic orbits being dense in a chaotic attractor. In fact, the consecutive numbering of *R28* and *R29* indicates that these relative periodic orbits were found concurrently from the DNS confirming their proximity in phase space. Also it is clear that some relative periodic orbits look like merged versions of two shorter orbits (not shown) again consistent with low-dimensional dynamical systems thinking.

4.3.2. $Re = 60, 80$ and 100

At higher Re , we managed to extract far fewer recurrent flows from the DNS. There are certainly reasons to expect this, most notably that the recurrent flows present should become more unstable and it is therefore harder to find good guesses from the DNS. There is also the fact that the ‘turbulence’ should explore more of phase space and therefore close visits to simple invariant sets should become rarer. However, the sharp drop in the number of recurrent flows found (see table 3) was still a surprise. In keeping with the philosophy of this work, only recurrent flows extracted from the DNS at that Re are listed in table 3 under the relevant Re heading. This then says nothing about whether a certain recurrent flow found at one Re might not exist at another. To explore this a little, we carried out some branch continuation (see Appendix) on the recurrent flows extracted from the series A DNSs while the runs and analysis for the series B DNSs were progressing. The results are shown in figure 14 colour coded to group recurrent flows found at the same Re and with black dots indicating branches detected at a given Re (note the rescaled dissipation measure on the ordinate to make the plot clearer). For example, the *T3* branch is shown as a dashed line (second dashed

	UPO	Frequency	c	T	$-s$	$-m$	N	$\sum_{j=1}^N \text{Re}(\lambda_j) (\max \text{Re}(\lambda_j))$
$Re = 40$	E1	≥ 261					9	1.296(0.249)
	T1	127	0.0198				4	0.142(0.068)
	T2	1	0.0096				4	1.227(0.454)
	P1	143		5.380			7	0.570(0.191)
	P2	6		2.830			5	0.742(0.223)
	P3	2		2.917			7	0.992(0.236)
	R1	1		56.677	0.092	0	3	0.156(0.077)
	R2	1		25.401	0.199	0	5	0.254(0.123)
	R3	1		54.280	0.200	0	3	0.195(0.108)
	R4	1		6.720	0.106	0	8	0.870(0.343)
	R5	1		23.780	0.022	0	4	0.376(0.156)
	R6	4		20.808	0.060	0	3	0.258(0.172)
	R18	1		37.233	0.270	0	5	0.242(0.165)
	R19			12.207	0.243	0	2	0.141(0.070)
	R20			16.586	5.827	1	4	0.289(0.103)
	R21			17.470	5.765	3	5	0.348(0.143)
	R22			19.723	0.222	0	4	0.297(0.172)
	R23			19.762	0.513	0	4	0.302(0.127)
	R24			19.779	6.035	0	5	0.292(0.202)
	R25			20.201	5.898	3	6	0.380(0.138)
	R26			20.385	1.334	2	7	0.714(0.270)
	R27			20.632	5.871	3	4	0.365(0.127)
	R28			20.885	5.987	1	4	0.360(0.121)
	R29			20.909	0.306	1	5	0.380(0.124)
	R30			21.310	5.694	0	5	0.330(0.100)
	R31			21.725	5.799	0	3	0.319(0.133)
	R32			22.560	0.006	1	4	0.283(0.096)
	R33			22.617	5.660	0	5	0.478(0.156)
	R34			23.157	0.265	0	3	0.260(0.113)
	R35			23.417	5.936	3	4	0.489(0.183)
	R36			24.465	6.010	3	4	0.358(0.191)
	R37			25.870	0.182	0	3	0.272(0.122)
	R38			25.934	0.227	0	4	0.263(0.125)
	R39			27.138	6.248	0	4	0.391(0.107)
	R40			28.817	5.971	0	5	0.238(0.116)
R41			32.541	0.349	1	4	0.224(0.153)	
R42			34.316	5.886	0	5	0.163(0.120)	
R43			34.530	5.742	0	3	0.220(0.134)	
R44			34.917	-0.059	3	4	0.325(0.139)	
R45			36.549	6.027	3	3	0.183(0.118)	
R46			36.627	0.197	3	4	0.202(0.139)	
R47			36.812	5.648	0	3	0.155(0.074)	
R48			37.079	6.103	0	4	0.171(0.134)	
R49			37.233	0.270	0	3	0.241(0.165)	
R50			37.698	3.499	1	6	0.477(0.146)	

TABLE 2. (Continued on next page)

line up from the Re axis, shown in red online) as it was first found at $Re = 60$ and it bears three black dots marked at $Re = 60, 80$ and 100 as $T3$ was extracted from the DNSs at all three Re . This type of analysis can indicate the bifurcation structure, e.g. $T3$ clearly bifurcates off a recurrent flow found at $Re = 40$, but is very time-consuming to pursue through to completion as branches can become difficult to

UPO	Frequency	c	T	$-s$	$-m$	N	$\sum_{j=1}^N \text{Re}(\lambda_j)$ (max $\text{Re}(\lambda_j)$)
<i>R51</i>			39.368	6.070	0	5	0.192(0.098)
<i>R52</i>			39.619	0.380	0	5	0.242(0.067)
<i>R53</i>			41.400	5.806	0	4	0.176(0.081)
<i>R54</i>			49.645	6.054	3	4	0.160(0.065)
<i>R55</i>			53.073	6.031	0	4	0.189(0.105)

TABLE 2. All of the invariant sets found directly from turbulent DNS data (from series A above the separating blank line and from series B below). ‘Frequency’ is the number of times the solution was extracted (Series A runs). There is one steady Equilibrium, c is the phase speed of the Travelling waves found, T is the Period of periodic and Relative periodic orbits which also either have a shift s and/or shift m . Here N is the number of unstable directions and $\sum_{j=1}^N \text{Re}(\lambda_j)$ the sum of the real parts of all of the unstable eigenvalues.

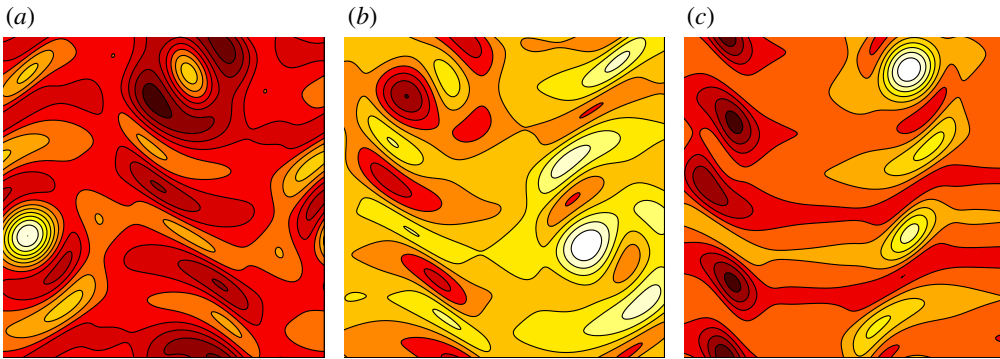


FIGURE 7. (Colour online) The steady solution $E1$ (a) and the TWs $T1$ (b) and $T2$ (c) at $Re = 40$. Vorticity is contoured using 15 contours between -7.2 (dark, shown in red online) and 12.1 (white); the range is $-6.5 \leq \omega \leq 12.1$ for $E1$, $-7.2 \leq \omega \leq 7.2$ for $T1$ and $-6.8 \leq \omega \leq 10.4$ for $T2$.

continue and interpret (note the number of open circles in figure 14 which indicate where the branch continuation procedure stagnated for some reason). This aside, the overriding impression is one of simple invariant sets proliferating with increasing Re . Notably, only two recurrent flows found at $Re = 40$ are also extracted from the $Re = 60$ DNSs: $E1$ (the highest line with a dot at $Re = 60$, shown in blue online) and $T1$. Here $E1$ seems to lose dynamical importance for yet higher Re but $T1$ is found for all four Re studied here.

Figure 15 shows the new TWs found and figures 16–18 show the $D-I$ plots where now all of the recurrent flows found at the respective Re are marked. Again most sit in the $D-I$ region where the DNS spends the majority of its time although as at $Re = 40$ there are some outliers (e.g. $E1$ at $Re = 60$, $R8$ at $Re = 80$, and $P4$ and $R14$ at $Re = 100$). That $R14$ actually appears outside the footprint of the DNS p.d.f. at first looks erroneous but is in fact merely an indication that when the ‘turbulence’ approached $R14$ in phase space, it maintained higher (global) dissipation and energy input than $R14$. This can occur when part of the domain resembles $R14$ while the rest does not and exhibits enhanced dissipation. A good example of this is

	UPO	Frequency	c	T	$-s$	$-m$	N	$\sum_{j=1}^N \text{Re}(\lambda_j)$ (max $\text{Re}(\lambda_j)$)
$Re = 60$	$E1$	4				0	14	5.053(0.858)
	$T1$	High	0.0019			0	4	0.139(0.064)
	$T3$	High	0.0124			0	17	3.377(0.684)
	$T4$	1	0.0082			0	3	0.257(0.178)
	$R7$	High		2.472	0.036	0	9	0.911(0.214)
	$R8$	1		1.638	0.022	0	14	2.903(0.681)
	$R56$			16.326	0.588	2	6	0.609(0.139)
	$R57$			17.909	5.802	0	7	0.805(0.169)
	$R58$			20.546	0.659	2	8	0.529(0.168)
$Re = 80$	$T1$	High	0.0115			0	6	0.360(0.105)
	$T3$	8	0.0154			0	21	5.588(0.958)
	$T5$	3	0.0831			0	20	4.183(0.658)
	$R7$	10		2.299	0.054	0	13	1.326(0.181)
	$R8$	2		1.705	0.028	0	18	4.318(1.105)
	$R9$	1		2.150	0.032	0	19	3.987(1.026)
	$R10$	1		1.280	0.020	0	20	5.310(0.878)
	$R11$	1		2.443	0.031	0	10	0.704(0.277)
	$R12$	1		2.095	0.034	0	11	3.083(1.004)
	$R13$	1		15.285	0.181	0	8	0.409(0.131)
	$R59$			15.667	0.397	1	11	0.949(0.176)
	$R60$			16.071	0.462	1	11	1.136(0.248)
	$Re = 100$	$T1$	High	0.0155			0	10
$T3$		7	0.0179			0	25	6.491(1.042)
$T4$		5	0.0118			0	3	0.684(0.370)
$T5$		6	0.0691			0	28	6.238(0.689)
$P4$		1		1.185		0	16	7.376(1.201)
$R7$		4		1.971	0.030	0	15	1.933(0.326)
$R11$		3		2.262	0.001	0	9	0.905(0.385)
$R12$		1		1.902	0.029	0	14	3.953(1.244)
$R14$		5		4.526	0.071	0	8	0.428(0.105)
$R15$		2		1.984	0.122	0	16	3.110(0.556)
$R16$		2		1.938	0.121	0	6	0.945(0.270)
$R17$		1		3.827	0.008	0	16	2.894(0.818)
$R61$		1		1.344	0.090	0	21	4.997(0.476)

TABLE 3. All of the invariant sets found directly from turbulent DNS data (from series A above the separating blank line and from series B below: none were found at $Re = 100$ in the series B runs despite the numerical gap between $R17$ and $R61$). ‘Frequency’ is the number of times the solution was extracted. Solutions listed under each Re indicate those actually extracted at that Re , hence multiple entries. ‘Frequency’ is the number of times the solution was extracted (series A runs). There is one steady Equilibrium, c is the phase speed of the Travelling waves found, T is the Period of periodic and Relative periodic orbits which also either have a shift s and/or shift m . Here N is the number of unstable directions and $\sum_{j=1}^N \text{Re}(\lambda_j)$ the sum of the real parts of all of the unstable eigenvalues.

shown in figure 19 which details the turbulent episode which signalled the presence of $P4$ (figure 19a) alongside the successfully converged periodic orbit $P4$ (figure 19b). Visually, the eye is drawn to the centre of the domain where in both columns an

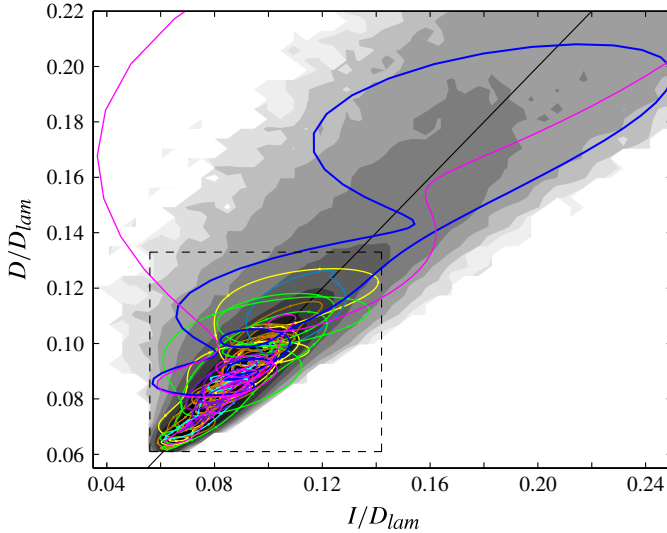


FIGURE 8. (Colour online) The normalized dissipation $D(t)/D_{lam}$ versus I/D_{lam} for a small collection of the recurrent flows found with the p.d.f. of the DNS turbulence plotted in the background (11 shades at levels 10^α where $\alpha = -5, -4.5, \dots, -0.5, 0$). Plotted are $E1$ ($D/D_{lam} = I/D_{lam} = 0.102$), $T1$ (0.071), $T2$ (0.070), $P1$, $P2$, $R1 - 6$, $R18$, $R26$ (large orbit, shown in blue online) and $R50$ (the even larger orbit, shown in magenta online). The dashed box is used to show more recurrent structures in figure 13.

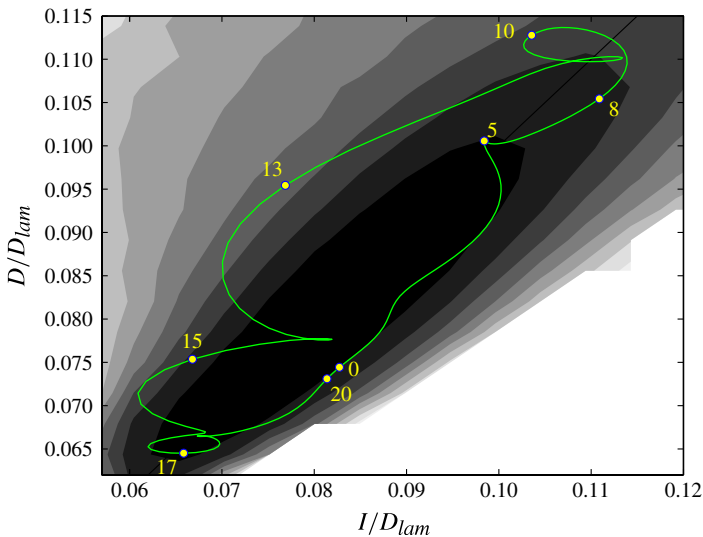


FIGURE 9. (Colour online) The normalized dissipation $D(t)/D_{lam}$ versus I/D_{lam} for the relative periodic orbit $R25$ at $Re = 40$ with the p.d.f. of the DNS turbulence plotted in the background (11 shades at levels 10^α where $\alpha = -5, -4.5, \dots, -0.5, 0$). The labels refer to times along the orbit at which snapshots are shown in figure 10.

isolated vortex is clearly seen rotating in a clockwise fashion. However, the corners are just as significant in that they also indicate an isolated vortex, yet this is stronger with




FIGURE 10. (Colour online) A time sequence of vorticity plots for $R25$ at $Re = 40$ at times (running left to right across the top and then bottom) $t = 0, 5, 8, 10, 13, 15, 17, 20$ (marked as dots on figure 9). The period is 20.2 so the flow in the bottom right is nearly the same as the top left except for shifts in x and y . For all, 15 contours are plotted from -12 to 12 .




FIGURE 11. (Colour online) The normalized dissipation $D(t)/D_{lam}$ versus I/D_{lam} for the relative periodic orbit $R50$ at $Re = 40$ with the p.d.f. of the DNS turbulence plotted in the background (11 shades at levels 10^α where $\alpha = -5, -4.5, \dots, -0.5, 0$). The labels refer to times along the orbit at which snapshots are shown in figure 12.

higher gradients (and, hence, larger dissipation) for $P4$ than the DNS signal. Plotting the two time sequences on a $D-I$ plot shows $P4$ as a closed loop much higher up the $D = I$ line than the DNS (not shown).




FIGURE 12. (Colour online) A time sequence of vorticity plots for $R50$ at $Re = 40$ at times (running left to right across the top and then bottom) $t = 0, 5, 8, 9, 10, 26, 31, 37$ (marked as dots on figure 11). The period is 37.7 so the flow in the bottom right is nearly the same as the top left except for shifts in x and y . In all plots, 15 contours are plotted from -12 to 12 .

5. Recurrent flows as a turbulent alphabet

Given the sets of recurrent flows extracted at each Re , the question is then how to use them to predict properties of the turbulence encountered. Periodic orbit theory advocates a weighted expansion of ‘pseudo-cycles’, sequences of ‘prime cycles’ (Cvitanović *et al.* 2013), such that

$$\Gamma_{prediction}^N := \frac{\sum_{i=1}^N w_i \Gamma_i}{\sum_{i=1}^N w_i} \quad (5.1)$$

where Γ is any property such as the mean dissipation rate, the mean profile or a p.d.f. and N is a finite but large number (to be discussed below). The weights w_i are not simply expressed but emerge from a recursive construction and depend on the type of periodic orbit averaging formula being used. Since the full set of Floquet multipliers corresponding to each recurrent orbit is very costly to obtain, we use the dynamical-zeta-function periodic orbit averaging formula which needs only information about the (typically much smaller number of) unstable Floquet multipliers. If $\Lambda_k^{(i)} := e^{\text{Re}(\lambda_k^{(i)})T_i}$ is the modulus of the k th Floquet multiplier of the linearized (Jacobian) operator around the i th recurrent flow of period T_i ($\lambda_k^{(i)}$ the complex growth rate), then the weight

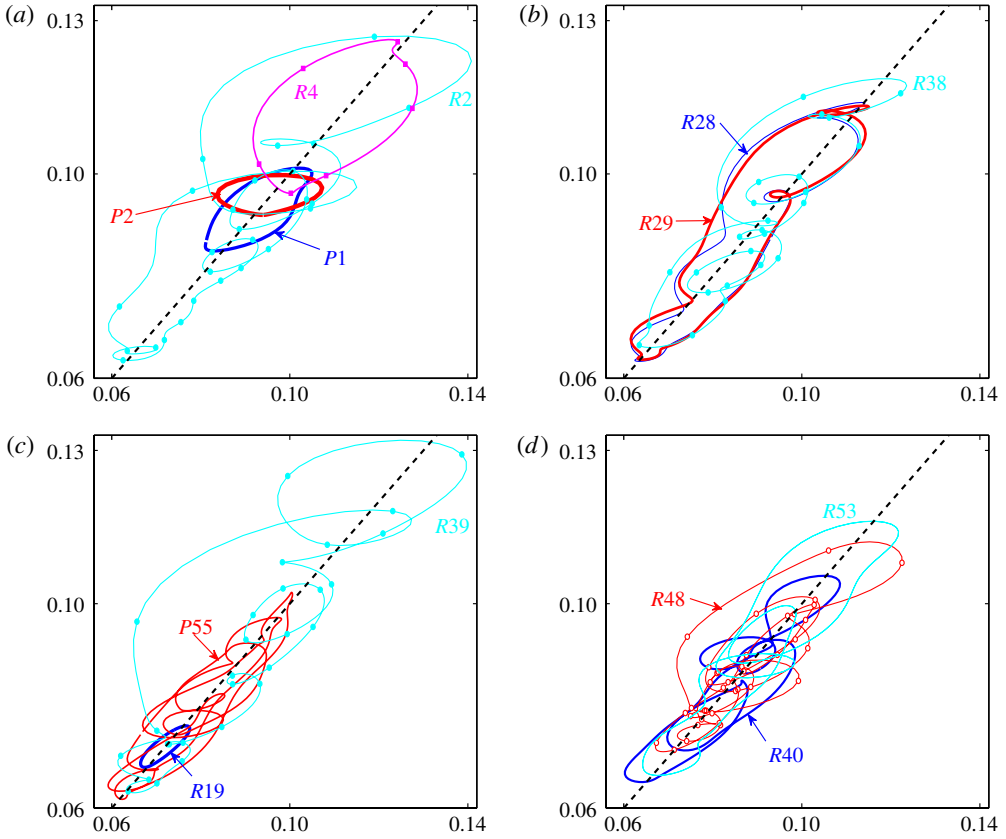


FIGURE 13. (Colour online) A sampling of the recurrent structures found at $Re = 40$ plotted over the zoom-in box shown in figure 8: (a) P1, P2, R2 and R4; (b) R28, R29 and R38; (c) R19, R39 and R55; and (d) R40, R48 and R53. In all cases, symbols added to lines are to assist in their distinction and are placed one time unit apart to indicate speed of flow.

associated with the i th recurrent flow depends on

$$A_i := \prod_{k, \Lambda_k^{(i)} > 1} \Lambda_k^{(i)} = \exp \left(\sum_{k \in \mathcal{K}_i} \text{Re}(\lambda_k^{(i)}) T_i \right). \tag{5.2}$$

(see, e.g., Gaspard 1997; Lan 2010, § 20 Cvitanović *et al.* 2013) where \mathcal{K}_i is the set of k such that $\text{Re}(\lambda_k^{(i)}) > 0$. For bounded flows (no trajectories escape; see Cvitanović *et al.* 2013, § 20.4.1), the zeta-function averaging formula takes the simple form

$$\Gamma_{prediction}^N := \frac{\langle \Gamma \rangle}{\langle T \rangle} \tag{5.3}$$

where

$$\langle \Gamma \rangle := \sum_{\pi} (-1)^{k+1} \frac{\sum_{i=1}^k T_{p_i} \Gamma_{p_i}}{\Lambda_{p_1} \Lambda_{p_2} \cdots \Lambda_{p_k}} \tag{5.4}$$




FIGURE 14. (Colour online) The scaled, time-averaged dissipation $Re \langle D/D_{lam} \rangle_t$ versus Re for recurrent flows discovered by the series A runs. Thin lines (shown in blue online) trace recurrent flows found at $Re = 40$, thin dashed lines (shown in red online) trace new recurrent flows found at $Re = 60$, thick lines (shown in cyan online) those found at $Re = 80$ and very thick lines (shown in magenta online) those at $Re = 100$. The black dots identify the subset of solutions which were identified by processing the DNS runs at the respective Re as opposed to just being continued up or down from other Re (e.g. there are six dots at $Re = 60$ corresponding to $E1$, $T1$, $T3$, $T4$, $R7$ and $R8$: $R56$ – $R58$ were discovered in the series B runs, and none of the dashed (red) lines join dots at $Re = 40$). Open circles indicate limits beyond which a solution branch could not be continued. The situation is clearly complicated with solutions seemingly dynamically important at some Re but not at others.




FIGURE 15. (Colour online) The TWs $T3$ (a), $T4$ (b) and $T5$ (c) at $Re = 100$. Vorticity is contoured using 15 contours between -20 (dark, shown in red online) and 15 (white); the range is $-18.9 \leq \omega \leq 8.03$ for $T3$, $-11.5 \leq \omega \leq 11.5$ for $T4$ and $-13.7 \leq \omega \leq 13.7$ for $T5$.




FIGURE 16. (Colour online) The normalized dissipation $D(t)/D_{lam}$ versus I/D_{lam} for the recurrent flows found at $Re = 60$ with the p.d.f. of the DNS turbulence plotted in the background (11 shades at levels 10^α where $\alpha = -5, -4.5, \dots, -0.5, 0$). Plotted are $E1$ ($D/D_{lam} = I/D_{lam} = 0.091$), $T1$ (0.033), $T3$ (0.068), $T4$ (0.034), $R7$, $R8$ and $R56$ – $R58$.




FIGURE 17. (Colour online) The normalized dissipation $D(t)/D_{lam}$ versus I/D_{lam} for the recurrent flows found at $Re = 80$ with the p.d.f. of the DNS turbulence plotted in the background (11 shades at levels 10^α where $\alpha = -5, -4.5, \dots, -0.5, 0$). Plotted are $T1$ ($D/D_{lam} = I/D_{lam} = 0.0196$), $T3$ (0.0514), $T5$ (0.0435, light dot, shown in yellow online), $R7$ – $R13$ and $R59$ – $R60$.

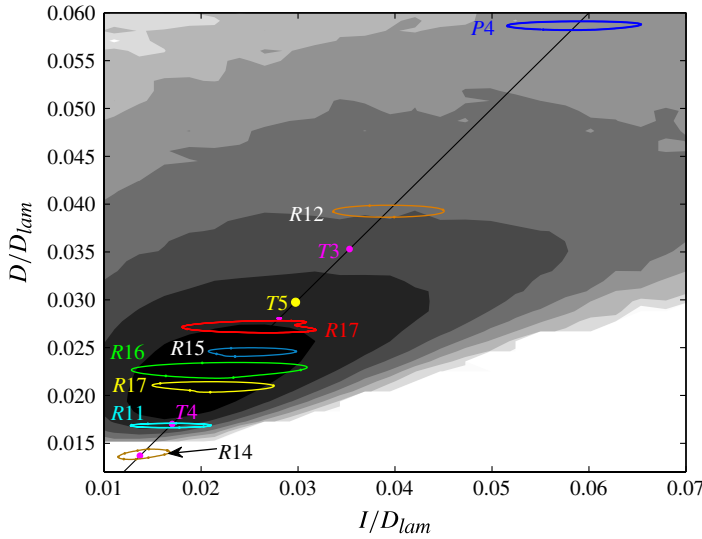


FIGURE 18. (Colour online) The normalized dissipation $D(t)/D_{lam}$ versus I/D_{lam} for the recurrent flows found at $Re = 100$ with the p.d.f. of the DNS turbulence plotted in the background (11 shades at levels 10^α where $\alpha = -5, -4.5, \dots, -0.5, 0$). Plotted are $T1$ ($D/D_{lam} = I/D_{lam} = 0.0137$), $T3$ (0.0353), $T4$ (0.0170), $T5$ (0.0297, light dot, shown in yellow online), $P4$, $R7$, $R11$, $R12$, $R14$ – $R17$ and $R61$.

and

$$\begin{aligned} \langle T \rangle &:= \sum_{\pi} (-1)^{k+1} \frac{\sum_{i=1}^k T_{p_i}}{\Lambda_{p_1} \Lambda_{p_2} \cdots \Lambda_{p_k}} \\ &= \sum_{i=1} T_i / \Lambda_i - \sum_{i=1} \sum_{j=i+1} (T_i + T_j) / (\Lambda_i \Lambda_j) + \sum_{i=1} \sum_{j=i+1} \sum_{k=j+1} (T_i + T_j + T_k) / (\Lambda_i \Lambda_j \Lambda_k) - \dots \end{aligned} \quad (5.5)$$

(see Cvitanović *et al.* 2013, §20.4.1). Here the subscript p_i refers to the p_i th prime cycle (here taken to be all of the recurrent flows identified), Γ_{p_i} is the temporal average of the quantity Γ over this cycle and \sum'_{π} represents a sum over all ($k = 1, 2, 3, \dots$) non-repeating, ordered combinations of prime cycles making up a pseudo-cycle (e.g. $\pi = (p_1, p_2, p_3, \dots, p_k)$ represents a pseudo-cycle of prime cycles p_1 to p_k concatenated to create a total period of $\sum_i T_{p_i}$). Very roughly, the geometrical meaning of a pseudo-cycle is that it is a sequence of shorter periodic orbits that shadow a longer periodic orbit along the segments p_1, p_2, \dots, p_k with the relative minus signs ensuring shadowing cancellations.

The truncation parameter N is now a function of which pseudo-cycles are to be included in the various sums. One strategy is to apply a total period cut-off in which only pseudo-cycles with $\sum_{i=1}^k T_{p_i} \leq T_{max}$ are included in the sum. A more appropriate choice is one based on stability (Dahlqvist & Russberg 1991; Dahlqvist 1994; Dettmann & Morriss 1997, §20.6; Cvitanović *et al.* 2013) in which only pseudo-cycles with

$$\Lambda_{p_1} \Lambda_{p_2} \cdots \Lambda_{p_k} \leq \Lambda_{max} \quad (5.6)$$

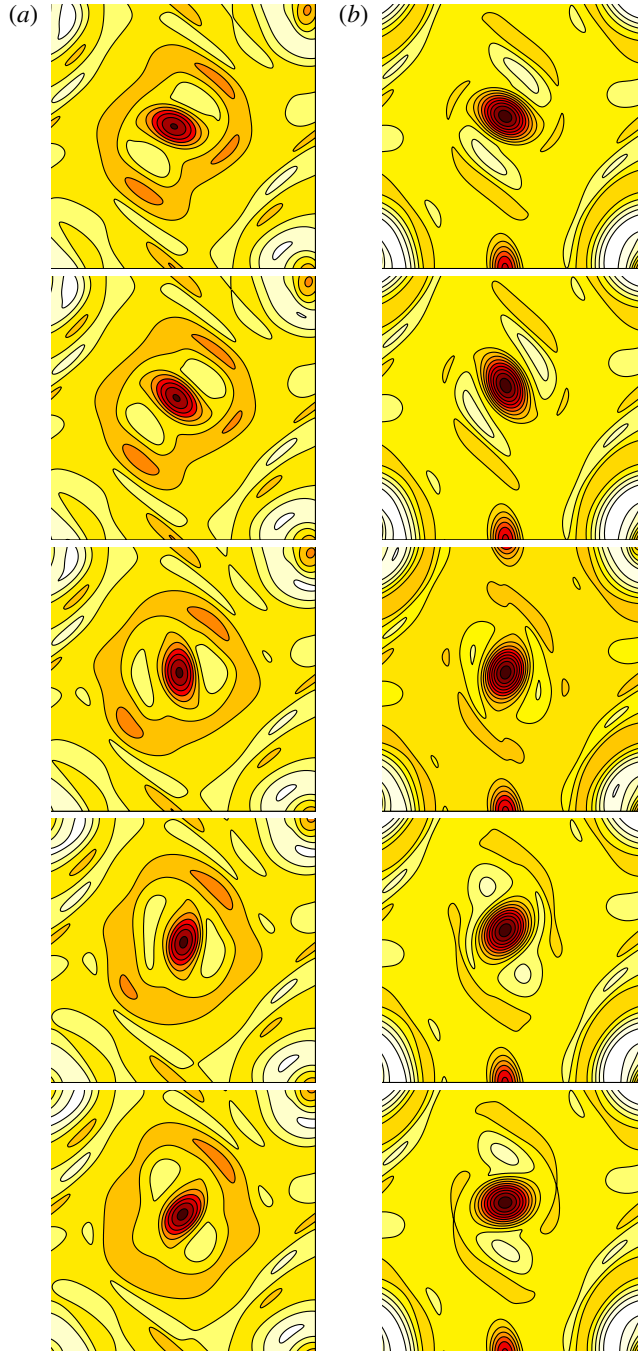


FIGURE 19. (Colour online) The DNS trajectory (a) synchronized with the subsequently converged periodic orbit $P4$ (b) at $Re = 100$. Time proceeds downwards with snapshots at t_0 , $t_0 + 0.2$, $t_0 + 0.6$, $t_0 + 0.8$ and $t_0 + 1.0$ (the period of $P4$ is 1.185). The vorticity scale ranges from -26 (dark, shown in red online) to 12 (white).

are included, where

$$\Lambda_{max} := \max_k \Lambda_k. \tag{5.7}$$

Figure 20 shows how Λ_i varies across the 47 periodic and relative periodic flows found at $Re = 40$ (they have been reordered in this plot by their stability so $\Lambda_{max} = \Lambda_{47} \approx 10^8$). The last two recurrent flows seem like outliers and there could be an argument to take $\Lambda_{max} = \Lambda_{45} \approx 10^5$ instead but this makes little difference to the results reported below.

The predictive formulae which emerge from periodic orbit theory are not straightforward to understand or their validity easy to assess. For example, why do TWs or equilibria not contribute? Why in constructing pseudo-cycles is it not necessary to include repeated traversals of prime cycles (this contribution is actually included but in a different way)? And how much is lost in working only with the unstable eigenvalues of a cycle rather than its full spectrum? The theory is derived under special conditions not necessarily satisfied by the Navier–Stokes equations (e.g. hyperbolicity) and requires the appropriate convergence and successful truncation of many singular-looking expressions during its derivation. Given this, it is tempting to also consider purely heuristic choices of weights based on the same general philosophy (i.e. an expansion over recurrent flows) as discussed by Zoldi & Greenside (1998) and Kazantsev (1998, 2001). These authors proposed and tested weights in expression (5.1) based solely on the unstable eigenvalues associated with the recurrent flow. Zoldi & Greenside (1998) used the ‘escape-time’ weighting

$$w_i := \frac{1}{\sum_{k \in \mathcal{K}_i} \text{Re}(\lambda_k^{(i)})} \quad \text{protocol 1} \tag{5.8}$$

(recall that \mathcal{K}_i is the set of k such that $\text{Re}(\lambda_k^{(i)}) > 0$) which they argue captures how unstable the recurrent flow is and inversely correlates this with how long the turbulent trajectory should spend in its vicinity. Kazantsev (1998, 2001) argued that this formula should be modified to reflect the fact that longer period orbits have a greater ‘presence’ in phase space than shorter period orbits and added the period $T^{(i)}$ to the numerator

$$w_i := \frac{T_i}{\sum_{k \in \mathcal{K}_i} \text{Re}(\lambda_k^{(i)})} \quad \text{protocol 2} \tag{5.9}$$

Significantly, this protocol suppresses any contribution from equilibria or TWs in common with periodic orbit theory. Finally we also consider a ‘control’ choice of ‘no weighting’ so just

$$w_i := 1 \quad \text{protocol 3.} \tag{5.10}$$

In what follows, we compare how periodic orbit theory (in the form of expression (5.3) and referred to hereafter as ‘POT’) and protocols 1, 2 and 3 perform using the sets of recurrent flows identified at each Re . The key measures we use to characterize the two-dimensional turbulence simulated here are p.d.f.s of $E(t)$ and $D(t)$ together with the profiles $\bar{u}(y)$, $u_{rms}(y)$ and $v_{rms}(y)$ (the p.d.f. of $I(t)$ was also considered but adds little information to that provided by the p.d.f. for $D(t)$).

5.1. $Re = 40$

The p.d.f.s of $E(t)/E_{lam}$ and $D(t)/D_{lam}$ at $Re = 40$ are shown in figure 21 along with the predictions using periodic orbit theory (5.3) and protocols 1–3 where the




FIGURE 20. (Colour online) The stabilities Λ of the 47 periodic and relatively periodic recurrent flows (reordered by stability from table 2) at $Re = 40$ (filled circles, shown in blue online). The upper (blue online) dashed line indicates Λ_{max} , the stability of the most unstable recurrent flow found (the cut-off stability value indicated by the lower (blue online) dashed line was also tried but produced little discernable difference in the periodic orbit theory predictions). The periods T of the recurrent flows (filled squares, shown red online) are also shown to indicate that there is a (expected) positive correlation with stability.




FIGURE 21. (Colour online) The p.d.f. for $E(t)/E_{lam}$ (a) and $D(t)/D_{lam}$ (b) from DNS (thick line, shown in blue online) and predictions using periodic orbit theory (thick line with squares, shown in magenta online) and weighting protocols 1 (medium line with circular dots, shown in red online), 2 (thin line, shown in green online) and 3 (black thick dashed line) at $Re = 40$. A total of 40 bins were used to calculate the p.d.f.s for the recurrent flows and 100 bins for the DNS due to its greater range. These choices gave the best balance of resolution with the data available.

individual p.d.f.s of each recurrent flow are weighted together appropriately to produce an overall p.d.f. as in (5.1). In the case of $D(t)/D_{lam}$ all three protocols give very good predictions of the central peak of the p.d.f. but each fails to capture the clear

shoulder at higher energies in the DNS albeit at a level of the p.d.f. a factor of ≈ 30 lower. The POT prediction is noticeably poorer but still captures the p.d.f. peak reasonably well. There is a similar story for the $E(t)/E_{lam}$ p.d.f. although in this case, the performance of protocols 1 and 2 seem marginally more effective than (the control) protocol 3 in capturing the DNS p.d.f. Again, the extremes of the DNS p.d.f. are not captured reflecting the fact, as commented earlier, that perhaps not enough recurrent flows with large or small dissipation excursions have been found. (The POT prediction seems especially sensitive to this.) Given the reduced value of the p.d.f. there, these are infrequently visited and thus require very long runs to realistically have a chance to extract them. As already mentioned, what initially looked like long runs of 10^5 time units were actually not long enough. As further independent evidence of this, plots of the mean profiles $\bar{u}(y)$ from each ‘long’ run at the same Re (see table 1) showed noticeable differences between each other and to the expected asymptotic state which respects all of the symmetries of the system: see figure 22. In particular, the obvious symmetry that the mean profile should be invariant under $\pi/2$ shifts in y was clearly violated. To ameliorate this, we decided to ‘symmetrize’ the DNS mean profile by extracting that part ($U^{SR}(y)$) from the signal (\bar{u}) which does satisfy all the symmetries listed in § 2.1. Explicitly

$$U^{SR}(y) := \frac{1}{2n} \sum_{m=0}^{2n-1} \mathcal{S}^{-m} U^R(\mathcal{S}^m y) \quad \text{with } U^R := \frac{1}{2} [\bar{u}(y) + \mathcal{R}^{-1} \bar{u}(\mathcal{R}y)] \quad (5.11)$$

(recall $n = 4$) and similarly for u_{rms}^{SR} and v_{rms}^{SR} . This process picks out the following Fourier coefficients

$$U^{SR}(y) := \sum_{m=0} a_m \sin(4(2m + 1)y) \quad (u_{rms}^{SR}, v_{rms}^{SR}) = \sum_{m=0} (b_m, c_m) \cos(8my)m, \quad (5.12)$$

from the complete Fourier series for \bar{u} , u_{rms} and v_{rms} . In particular, the symmetrized mean profile $U^{SR}(y)$ has the leading Fourier modal form of $\sin 4y$, which mimics the forcing, and a leading correction of $\sin 12y$. Such a profile needs only be plotted over $y \in [0, \pi/4]$ which is done in figure 23 along with the predictions. This comparison looks impressive with $a_0 = 0.232$ (cf. expression (5.12)) in the DNS, compared with 0.200 for POT, 0.214 (protocol 1), 0.215 (protocol 2) and 0.219 (protocol 3). For POT $a_1 = 1.1 \times 10^{-5}$ and $a_2 = -1.54 \times 10^{-7}$ whereas otherwise $a_1 = O(10^{-6})$ and $a_2 = O(10^{-8})$ (as way of comparison, the ‘raw’ mean flow \bar{u} has a leading non-symmetrized part given by $0.0280 \cos y + 0.0268 \sin y$, i.e. roughly 10% smaller than the symmetrized part: see figure 22). The explanation for why the (symmetrized) mean profile matches the forcing profile so well is currently unclear to us and it is tempting to speculate that actually $a_n \rightarrow 0$ ($n \geq 1$) with the period of averaging. Sarris *et al.* (2007) study the statistics of three-dimensional Kolmogorov flow for various computational domains and use two measures to signal whether their statistics have converged sufficiently over a period of time integration. The first

$$\gamma_2 := \frac{(\langle I \rangle_t - \langle D \rangle_t)^2}{\langle D \rangle_t^2} \quad (5.13)$$

(Sarris *et al.* 2007, equation (22)) assesses the extent to which the energy input into the flow matches the energy dissipated and is easily calculated from our output data: γ_2 is at most $O(10^{-8})$ for all our 10^5 time unit runs. Even with this small value, the mean profile is far from converged to what is expected (i.e. satisfies all of the




FIGURE 22. (Colour online) (a) The DNS mean flow (solid line, shown in blue online) compared with the symmetrized mean flow $U^{SR}(y)$ (dashed line, shown in red online) at $Re = 40$ from run d (see table 1). (b) DNS means from runs e, f and g (see table 1) at $Re = 60$ again showing a lack of convergence even after 10^5 time units.




FIGURE 23. (Colour online) The symmetrized mean flow $U^{SR}(y)$ (left) and symmetrized $u_{rms}^{RS}(y)$ (right plot left part) and $v_{rms}^{RS}(y)$ (right part of right plot) from DNS (thick line, shown in blue online) and predictions using periodic orbit theory (line with squares, shown in magenta online), weighting protocol 1 (thick line with dots, shown in red online), 2 (thin line, shown in green online) and 3 (black thick dashed line) at $Re = 40$.

symmetries of the problem) which emphasizes how easy it is to unwittingly collect unconverged statistics.

Figure 23 also shows the equivalent plot for u_{rms} and v_{rms} . For u_{rms} , $b_0 = 0.687$ for the DNS which clearly differs from all of the predictions: 0.115 (POT), 0.272 (protocol 1), 0.289 (protocol 2) and 0.276 (protocol 3). The comparison for v_{rms} , however, is much better: $c_0 = 0.933$ for the DNS versus 0.891 (POT), 0.887

(protocol 1), 0.891 (protocol 2) and 0.879 (protocol 3). One possible reason why the u_{rms} comparison is poor is that u_{rms} is calculated in the DNS ‘on the fly’ by subtracting the current best estimate of the mean \bar{u} from the current streamwise velocity (see (2.12)) rather than using the final mean profile to *a posteriori* calculate the streamwise fluctuation field. Given our realization now that the mean profile takes a long time to converge to its (symmetric) asymptotic state, there is likely to be a significant error (henceforth we consider only v_{rms} for $Re = 60, 80$ and 100).

An inescapable conclusion from these comparisons so far is that the ‘control’ protocol 3 of actually ‘no weighting’ performs almost as well as the other stability-motivated protocols 1 and 2 and better than POT. It is worthwhile at this point to clarify why. Figure 24 shows how the peak symmetrized mean value, $U^{SR}(\pi/8)$, varies for each recurrent flow and compares these values with the DNS and the predictions. From this it is clear that most of the recurrent flows are good predictors individually and so, however they are mixed together, the result is still reasonably good. The upper plot in figure 25 indicates how the weights of the recurrent flows vary for the three different protocols. Again, there is not that much variation over the majority (although note that $w_i = 0$ for $i = 1, 2, 3$ in protocol 2 since $T^{(i)} = 0$) which presumably reflects the fact that the values of $\sum_{k \in \mathcal{X}_i} \text{Re}(\lambda_k^{(i)})$ for the recurrent flows are all pretty similar. The lower plot of figure 25 shows the effective weights of the periodic and relative periodic flows which emerge from the POT prediction. These vary much more dramatically in amplitude than the heuristic formulae (protocols 1 and 2) because Λ_i rather than $\log \Lambda_i$ is used in the prediction (see (5.4) and (5.5)). In fact, the simple choice of weights

$$w_i^{approx} := \frac{\Lambda_i^{-1}}{\sum_i \Lambda_i^{-1}} \quad (5.14)$$

is quite a good approximation to the POT weights which would reflect only counting prime cycles rather than all the pseudo-cycles made up of multiple prime cycles (note the similarity to those of protocols 1 and 2 if $\sum_{k \in \mathcal{X}_i} \text{Re}(\lambda_k^{(i)})$ is replaced by $\Lambda_i = \exp(\sum_{k \in \mathcal{X}_i} \text{Re}(\lambda_k^{(i)}) T_i)$ up to factors of T_i). Figure 25 also indicates that one prime cycle dominates the prediction: $R19$ ($i = 11$). This is a simple relative periodic orbit (see figure 13) which does not explore much of the turbulent attractor and results in the lumpy form of the POT p.d.f. prediction. This highlights the importance of not only identifying lots of recurrent flows but also many which have similar stability characteristics (i.e. Λ).

5.2. $Re = 60, 80$ and 100

The smaller number of recurrent flows extracted for $Re > 40$ means that it is harder to generate reasonably smooth predictions for the p.d.f.s of the energy and dissipation. While the same number of 100 bins as at $Re = 40$ can be used to generate a smooth DNS p.d.f., only 60 bins could be used to sum the p.d.f.s of the recurrent flows. This number produced the best compromise of granularity across the range while ensuring that there is enough data in each bin for (reasonable) smoothness at least for $Re = 60$ and 80 (the sparse coverage of the dissipation range by the recurrent flows at $Re = 100$, see figure 18, prevented any useful plot from being generated). Figure 26 shows the result of this procedure for the dissipation p.d.f. at $Re = 60$ and 80 (plots not shown for the energy are similar). Here, protocol 2 offers the best partial fit both at $Re = 60$ and 80 and POT the worst (note both are only built upon five recurrent




FIGURE 24. (Colour online) The peak symmetrized mean $U_{max}^{RS}(\pi/8)$ for the DNS (thick horizontal line, shown in blue online), predictions using periodic orbit theory (line with squares, shown in magenta online), protocol 1 (dotted line, shown in red online), protocol 2 (thin line, shown in green online), and protocol 3 (dashed black line) and for each recurrent flow (crosses, shown in blue online) listed in the order given in table 2 at $Re = 40$ (so $i = 20$ is $R25$ and $i = 45$ is $R50$ marked with black dotted lines).

flows with $T_i \neq 0$). It is clear that much of the dissipation range extends higher than any of the recurrent flows found at both Re (also clear from figures 16 and 17) and that all of the ‘predictions’ are of limited quality. It is also worth remarking that the predictions are now more distinguishable at these higher Re which can be traced back to the separating stability characteristics of the recurrent flows. For example, $E1$ has $\sum \text{Re}(\lambda_j) = 5.053$ whereas this is just 0.139 for $T1$. Again the POT prediction struggles to smooth across the set of recurrent flows found because of the large differences in the sizes of the weights.

Figure 27 shows the symmetrized mean profile as calculated from the DNS and predicted by POT and the three protocols at $Re = 60$ and 80 . Again, somewhat paradoxically, the ‘control’ protocol does the best job in both cases. At $Re = 60$ $a_0 = 0.2277$, 0.2298 and 0.2296 (cf. expression (5.12)) across the three series A DNS runs listed in table 1 (with again $\approx 10\%$ non-symmetrized part in all cases). In comparison, $a_0 = 0.182$ for POT, $a_0 = 0.148$ (protocol 1), 0.175 (protocol 2) and 0.200 (protocol 3). For v_{rms} , $c_0 = 1.099$ in the DNS run to be compared with the predictions of 1.018 (POT), 1.064 (protocol 1), 1.057 (protocol 2) and 1.055 (protocol 3). At $Re = 80$, $a_0 = 0.2009$, 0.2016 and 0.1998 across the three series A DNS runs listed in table 1 with predictions $a_0 = 0.1323$ (POT), $a_0 = 0.1416$ (protocol 1), 0.1403 (protocol 2) and 0.1944 (protocol 3).

Finally, we note that at $Re = 100$, $a_0 = 0.1777$, 0.1782 and 0.1783 across the three series A DNS runs listed in table 1 (with now $\approx 20\%$ non-symmetrized part in all cases).

6. Discussion

In this study, we have considered two-dimensional turbulence on the torus $[0, 2\pi]^2$ forced monochromatically in one direction (Kolmogorov flow). By looking for near




FIGURE 25. (Colour online) (a) The weights for the three protocols (1, line with dots, shown in red online; 2, full line, shown in green online; 3, black horizontal line) plotted for all of the recurrent structures at $Re = 40$. Again the vertical (black) dotted lines indicate $R25$ and $R50$. (b) The (implicit) weights w_i which emerge from the periodic orbit theory (thick line with squares, shown in magenta online) for the recurrent flows as ordered in table 2. The thin black line with dots indicates the weights $\Lambda_i^{-1} / \sum_i \Lambda_i^{-1}$ which are strongly correlated with the weights w_i from periodic orbit theory (note the log scale).

recurrences of the flow in long DNS runs, sets of exactly closed flow solutions ‘embedded’ in this turbulence have been extracted at different forcing amplitudes (Re). We have then tried to use these sets of recurrent flows to reconstruct key statistics of the turbulence motivated by periodic orbit theory in low-dimensional chaos. The approach has been reasonably successful at $Re = 40$ (see figures 21 and 23), where 50 recurrent flows were found with the majority buried in the part of phase space most populated by the turbulence. In contrast, at $Re = 60, 80$ and 100 , the limited size of the recurrent flow sets found has made the approach largely impotent. Even at $Re = 40$, the success achieved seems more reliant on just extracting lots of similar-looking recurrent flows buried in the most popular part of phase space for the turbulence than on any sophisticated choice of weighting coefficients. Indeed, one is reminded of Kawahara & Kida’s (2001) conclusion that one judiciously chosen periodic orbit is ‘enough’ to be a valuable proxy of the turbulence. We can sympathize with this viewpoint but only if the comparison with the turbulence statistics is not too demanding. The key issue, of course, plaguing this investigation




FIGURE 26. (Colour online) The p.d.f. for $D(t)/D_{lam}$ from DNS (blue thick line, cyan and magenta lines) and predictions using periodic orbit theory (thick line with squares, shown in magenta online), weighting protocol 1 (thick line with dots, shown in red online), 2 (thin line with triangles, shown in green online) and 3 (black thick dashed with squares) at $Re = 60$ (a) and $Re = 80$ (b). A total of 60 bins were used to calculate the p.d.f.s for the recurrent flows and 100 bins for the DNS due to its greater range. These choices gave the best balance of resolution with the data available.

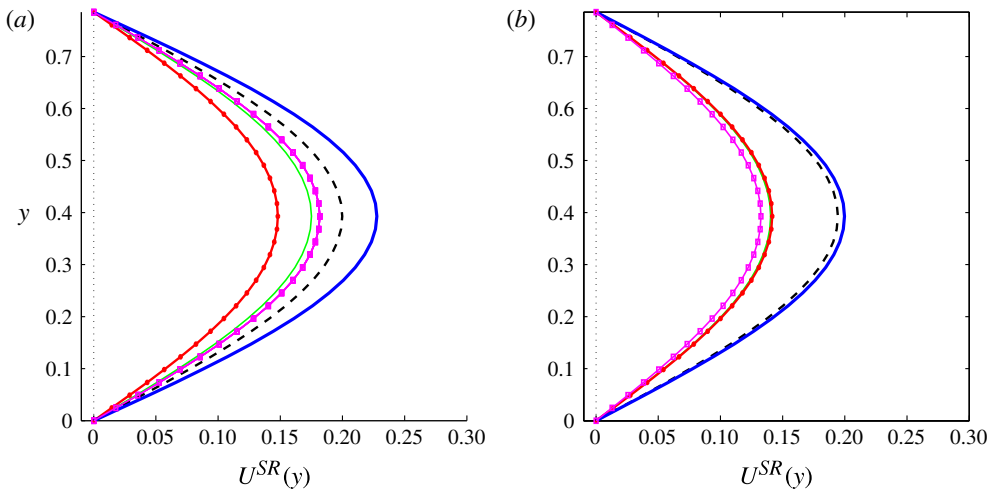


FIGURE 27. (Colour online) The symmetrized mean flow $U^{SR}(y)$ from DNS (thick line, shown in blue online) and predictions using periodic orbit theory (line with squares, shown in magenta online), weighting protocol 1 (thick line with dots, shown in red online), 2 (thin line, shown in green online) and 3 (black thick dashed line) for $Re = 60$ (a) and $Re = 80$ (b).

is the paucity of recurrent flows found from the finite DNS data generated. This is perhaps the main message to come out of this work: periodic orbit theory for fluid turbulence is a promising approach but only if enough (say, $O(100)$) recurrent flows of similar stability characteristics are gathered which requires very long turbulence data sequences. A time sequence of 10^5 time units seems marginally adequate for

$Re = 40$ but is maybe two orders of magnitude too short for $Re = 60$ and beyond. Unfortunately, without these large sets, it has been impossible to see periodic orbit theory, which has some rationale basis, outperform the other purely heuristic protocols tested which have none.

Operationally, the work described here has been time-consuming both computationally in generating near-recurrence episodes and attempting to converge them, as well as ‘manually’ because of all of the careful processing (e.g. calculating their stability) and archiving of the recurrent flows needed (e.g. does a new convergence from a DNS guess represent a new recurrent flow or a repeat of a previously extracted flow?). Fortunately, there is no reason why this process could not be automated with the objective being to ‘automatically’ generate a basis set of recurrent flows for each Re . Indeed, one could hope that such a set at given $Re = Re_1$ could be used to predict the turbulent statistics at another $Re = Re_2$. This would require each recurrent flow at Re_1 being continued to Re_2 and the fresh weights for an expansion being generated from the (new) stability information for each recurrent flow: again, painstaking work, but readily automated. One fly in the ointment is the possibility of bifurcations in the interval $[Re_1, Re_2]$, particularly saddle node bifurcations where two recurrent flows at Re_1 merge and annihilate before Re reaches Re_2 . Working with large enough recurrent flow sets would presumably smooth over this effect somewhat but will not eliminate it entirely.

Leaving aside these issues for a moment, it is worth re-emphasizing that any recurrent flow extracted from DNS data is a simple invariant solution ‘buried’ in the turbulence. As such, each represents a sustained sequence of dynamical processes which contributes to, if not underpins, the turbulence itself. Since they are closed in time, they can be analysed relatively easily in whatever detail is required to understand key dynamical relationships in the flow. This seems a very promising byproduct of the analysis whether one believes a periodic-orbit-theory-type expansion of turbulence is possible or not (pursuing this has not been the focus here due to the two-dimensionality of the flow).

Finally, the ever-improving computational resources available now have only recently made this type of study possible. Even with these, we have underestimated the demands of data collection in two-dimensional turbulence over the small torus $[0, 2\pi]^2$. Major challenges ahead include treating large aspect ratio domains (can we find localized recurrent flows?) and handling fully three-dimensional flows (with automated machinery, will the approach be practical?). There is plenty to explore.

Acknowledgements

We both would like to thank P. Bartello for generously sharing his DNS code. G.J.C. would like to thank I. Waugh for guidance on arclength continuation and R.R.K. is very grateful to C. Dettmann and particularly P. Cvitanović for numerous clarifying discussions concerning periodic orbit theory.

Appendix.

The Newton–GMRES–hook-step algorithm described in the main text is easily extended to continue solutions over parameter space such as Re or the domain geometry (e.g. α). We briefly describe this extension for solution branch continuation in Re which was used to generate figure 14. A simple strategy is to use the solution $X(Re)$ as an initial guess in the Newton–GMRES–hook-step algorithm with the hope of converging to $X(Re + \delta Re)$. This should work provided that δRe is ‘small enough’

Depending on how easily convergence is obtained, δr can be increased or decreased if the algorithm shows signs of divergence. A second-order approach to estimating $\partial X/\partial r$ was actually adopted for the predictive step but the first-order estimate proved sufficient for the constraint present in (A 6).

REFERENCES

- ARMBRUSTER, D., NICOLAENKO, B., SMAOUI, N. & CHOSSAT, P. 1996 Symmetries and dynamics of 2-D Navier–Stokes flow. *Physica D* **95**, 81–93.
- ARNOL'D, V. I. & MESHALKIN, L. D. 1960 The seminar of A.N. Kolmogorov on selected topics in analysis (1958–1959). *Usp. Mat. Nauk* **15**, 247–250.
- ARTUSO, R., AURELL, E. & CVITANOVIĆ, P. 1990a Recycling of strange sets: I cycle expansions. *Nonlinearity* **3**, 325–359.
- ARTUSO, R., AURELL, E. & CVITANOVIĆ, P. 1990b Recycling of strange sets: II applications. *Nonlinearity* **3**, 361–386.
- AUERBACH, D., CVITANOVIĆ, P., ECKMANN, J.-P., GUNARATNE, G. & PROCACCIA, I. 1987 Exploring chaotic motion through periodic orbits. *Phys. Rev. Lett.* **58**, 2387–2389.
- BARTELLO, P. & WARN, T. 1996 Self-similarity of decaying two-dimensional turbulence. *J. Fluid Mech.* **326**, 357–372.
- BOFFETTA, G. & MUSACCHIO, S. 2010 Evidence for a double cascade scenario in two-dimensional turbulence. *Phys. Rev. E* **82**, 016307.
- BOGHOSIAN, B. M., FAZENDEIRO, L. M., LÄTT, J., TANG, H. & COVENEY, P. V. 2011 New variational principles for locating periodic orbits of differential equations. *Phil. Trans. R. Soc. A* **369**, 2211–2218.
- BONDARENKO, N. F., GAK, M. Z. & DOLZHANSKII, F. V. 1979 Laboratory and theoretical models of plane periodic flow. *Izv. Acad. Sci. USSR Atmos. Ocean. Phys.* **15**, 711–716.
- BORUE, V. & ORSZAG, S. A. 1996 Numerical study of three-dimensional Kolmogorov flow at high Reynolds numbers. *J. Fluid Mech.* **306**, 293–323.
- BURGESS, J. M., BIZON, C., MCCORMICK, W. D., SWIFT, J. B. & SWINNEY, H. L. 1999 Instability of the Kolmogorov flow in a soap film. *Phys. Rev. E* **60**, 715–721.
- CHRISTIANSEN, F., CVITANOVIĆ, P. & PUTKARADZE, V. 1997 Spatiotemporal chaos in terms of unstable recurrent patterns. *Nonlinearity* **10**, 55–70.
- CVITANOVIĆ, P. 1988 Invariant measurement of strange sets in terms of cycles. *Phys. Rev. Lett.* **61**, 2729–2732.
- CVITANOVIĆ, P. 1992 Periodic orbit theory in classical and quantum mechanics. *Chaos* **2**, 1.
- CVITANOVIĆ, P. 1995 Dynamical averaging in terms of periodic orbits. *Physica D* **83**, 109–123.
- CVITANOVIĆ, P. 2007 Continuous symmetry reduced trace formulas
ChaosBook.org/~predrag/trace.pdf.
- CVITANOVIĆ, P. 2012 Continuous symmetry reduced trace formulas. Preprint.
- CVITANOVIĆ, P., ARTUSO, R., DAHLQVIST, P., MAINIERI, R., TANNER, G., VATTAY, G., WHELAN, N. & WIRZBA, A. 2013 Classical and Quantum Chaos webbook available at <http://chaosbook.org>.
- CVITANOVIĆ, P., DAVIDCHACK, R. & SIMINOS, E. 2010 On the state space geometry of the Kuramoto–Sivashinsky flow in a periodic domain. *SIAM J. Appl. Dyn. Sys.* **9**, 1–33.
- CVITANOVIĆ, P. & GIBSON, J. F. 2010 Geometry of turbulence in wall-bounded shear flows: periodic orbits. *Phys. Scr.* **142**, 014007.
- DAHLQVIST, P. 1994 Determination of resonance spectra for bound chaotic systems. *J. Phys. A: Math. Gen.* **27**, 763–785.
- DAHLQVIST, P. & RUSSBERG, G. 1991 Periodic orbit quantization of bound chaotic systems. *J. Phys. A* **24**, 4763–4778.
- DENNIS, J. E. & SCHNABEL, R. B. 1996 Numerical Methods for Unconstrained Optimisation and Nonlinear equations. In *SIAM Classics*. SIAM.
- DETMANN, C. P. & MORRIS, G. P. 1997 Stability ordering of cycle expansions. *Phys. Rev. Lett.* **78**, 4201–4204.

- DUGUET, Y., PRINGLE, C. C. T. & KERSWELL, R. R. 2008 Relative periodic orbits in transitional pipe flow. *Phys. Fluids* **20**, 114102.
- ECKHARDT, B., FAISST, H., SCHMIEGEL, A. & SCHUMACHER, J. 2002 Turbulence transition in shear flows. In *Advances in Turbulence IX: Proceedings 9th European Turbulence Conference (Southampton)* (ed. I. P. Castro *et al.*), p. 701, CISME.
- ECKHARDT, B., SCHNEIDER, T. M., HOF, B. & WESTERWEEL, J. 2007 Turbulence transition in pipe flow. *Annu. Rev. Fluid Mech.* **39**, 447–468.
- ECKMANN, P. & RUELLE, D. 1985 Ergodic theory of chaotic systems. *Rev. Mod. Phys.* **57**, 617–656.
- FAZENDEIRO, L., BOGHOSIAN, B. M., COVENEY, P. V. & LÄTT, J. 2010 Unstable periodic orbits in weak turbulence. *J. Comput. Sci.* **1**, 13–23.
- GASPARD, P. 1997 *Chaos Scattering and Statistical Mechanics*. CUP.
- 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 travelling-wave solutions of plane Couette flow. *J. Fluid Mech.* **638**, 243–266.
- GOTOH, K. & YAMADA, M. 1986 The instability of rhombic cell flows. *Fluid Dyn. Res.* **1**, 165–176.
- GUTZWILLER, 1990 *Chaos in Classical and Quantum Mechanics*. Springer.
- HALCROW, J., GIBSON, J. F., CVITANOVIĆ, P. & VISWANATH, D. 2009 Heteroclinic connections in plane Couette flow. *J. Fluid Mech.* **621**, 365–376.
- HAMILTON, J. M., KIM, J. & WALEFFE, F. 1995 Regeneration mechanism of near-wall turbulence structures. *J. Fluid Mech.* **287**, 317–348.
- HOF, B., VAN DOORNE, C. W. H., WESTERWEEL, J., NIEUWSTADT, F. T. M., FAISST, H., ECKHARDT, B., WEDIN, H., KERSWELL, R. R. & WALEFFE, F. 2004 Experimental observation of nonlinear travelling waves in turbulent pipe flow. *Science* **305**, 1594–1598.
- HOLMES, P., LUMLEY, J. L. & BERKOOZ, G. 1996 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press.
- 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: regenerative cycle and burst. *J. Fluid Mech.* **449**, 291–300.
- KAWAHARA, G., UHLMANN, M. & VAN VEEN, L. 2012 The significance of simple invariant solutions in turbulent flows. *Annu. Rev. Fluid Mech.* **44**, 203–225.
- KAZANTSEV, E. 1998 Unstable periodic orbits and attractor of the barotropic ocean model. *Nonlinear Process. Geophys.* **5**, 193–208.
- KAZANTSEV, E. 2001 Sensitivity of the attractor of the barotropic ocean model to external influences: approach by unstable periodic orbits. *Nonlinear Process. Geophys.* **8**, 281–300.
- KERSWELL, R. R. 2005 Recent Progress in understanding the transition to turbulence in a pipe. *Nonlinearity* **18**, R17–R44.
- KIM, S.-C. & OKAMOTO, H. 2003 Bifurcations and inviscid limit of rhombic Navier–Stokes flows in tori. *IMA J. Appl. Maths* **68**, 119–134.
- KREILOS, T. & ECKHARDT, B. 2012 Periodic orbits near onset of chaos in plane Couette flow. *Chaos* **22**, 047505.
- LAN, Y. 2010 Cycle expansions: from maps to turbulence. *Commun. Nonlinear Sci. Numer. Simul.* **15**, 502–526.
- LAN, Y. & CVITANOVIĆ, P. 2004 Variational method for finding periodic orbits in a general flow. *Phys. Rev. E* **69**, 016217.
- LAN, Y. & CVITANOVIĆ, P. 2008 Unstable recurrent patterns in Kuramoto–Sivashinsky dynamics. *Phys. Rev. E* **78**, 026208.
- LOPEZ, V., BOYLAND, P., HEATH, M. T. & MOSER, R. D. 2005 Relative periodic solutions of the complex Ginzburg–Landau equation. *SIAM J. Appl. Dyn. Syst.* **4**, 1042–1075.
- MACKAY, R. S. & MIESS, J. D. 1987 *Hamiltonian Dynamical Systems*. Adam Hilger.

- MARCHIORO, C. 1986 An example of absence of turbulence for any Reynolds number. *Commun. Math. Phys.* **105**, 99–106.
- MESHALKIN, L. D. & SINAI, YA. G. 1961 Investigation of stability of a steady-state solution of a system of equations for the plane motion of an incompressible viscous liquid. *Prikl. Mat. Mekh.* **25**, 1140–1143.
- OBUKHOV, A. M. 1983 Kolmogorov flow and laboratory simulation of it. *Usp. Mat. Nauk* **38**, 101–111.
- OKAMOTO, H. & SHOJI, M. 1993 Bifurcation diagrams in Kolmogorov's problem of viscous incompressible fluid on 2-D flat tori. *Japan. J. Indust. Appl. Math.* **10**, 191–218.
- PANTON, R.L. (Ed.) 1997 *Self-Sustaining Mechanisms of Wall Turbulence*. Computational Mechanics.
- PLATT, N., SIROVICH, L. & FITZMAURICE, N. 1991 An investigation of chaotic Kolmogorov flows. *Phys. Fluids* **3**, 681–696.
- POINCARÉ, H. 1892 *Les méthodes nouvelles de la mécanique céleste*. Guthier-Villars.
- ROBINSON, S. K. 1991 Coherent motions in the turbulent boundary layer. *Annu. Rev. Fluid Mech.* **23**, 601–639.
- RUELLE, D. 1978 *Statistical Mechanics, Thermodynamic Formalism*. Addison-Wesley.
- 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.
- SARRIS, I. E., JEANMART, H., CARATI, D. & WINCKELMANS, G. 2007 Box-size dependence and breaking of translational invariance in the velocity statistics computed from three-dimensional turbulent Kolmogorov flows. *Phys. Fluids* **19**, 095101.
- SHE, Z. S. 1988 Large-scale dynamics and transition to turbulence in the two-dimensional Kolmogorov flow. In *Proceedings on Current Trends in Turbulence Research* (ed. H. Branover, M. Mond & Y. Unger), vol. 117, pp. 374–396. American Institute of Aeronautics and Astronautics.
- SHEBALIN, J. V. & WOODRUFF, S. L. 1997 Kolmogorov flow in three dimensions. *Phys. Fluids* **9**, 164–170.
- SOMMERIA, J. 1986 Experimental study of the two-dimensional inverse energy cascade in a square box. *J. Fluid Mech.* **170**, 139–168.
- TREFETHEN, L. N. & BAU, D. 1997 *Numerical Linear Algebra*. SIAM.
- TSANG, Y.-K. & YOUNG, W. R. 2008 Energy-ensrophy stability of beta-plane Kolmogorov flow with drag. *Phys. Fluids* **20**, 084102.
- VAN VEEN, L., KAWAHARA, G. & KIDA, S. 2006 Periodic motion representing isotropic turbulence. *Fluid Dyn. Res.* **38**, 19–46.
- VISWANATH, D. 2007 Recurrent motions within plane Couette turbulence. *J. Fluid Mech.* **580**, 339–358.
- VISWANATH, D. 2009 The critical layer in pipe flow at high Reynolds number. *Phil. Trans R. Soc. A* **367**, 561–576.
- WALEFFE, F. 1997 On a self-sustaining process in shear flows. *Phys. Fluids* **9**, 884–900.
- WALEFFE, F. 1998 Three-dimensional coherent states in plane shear flows. *Phys. Rev. Lett.* **81**, 4140–4143.
- WALEFFE, F. 2001 Exact coherent structures in channel flow. *J. Fluid Mech.* **435**, 93–102.
- WALEFFE, F. 2003 Homotopy of exact coherent structures in plane shear flows. *Phys. Fluids* **15**, 1517–1534.
- WILLIS, A. P., CVITANOVIĆ, P. & AVILA, M. 2013 Revealing the state space of turbulent pipe flow by symmetry reduction. Preprint, arXiv: [1203.3701](https://arxiv.org/abs/1203.3701).
- ZOLDI, S. & GREENSIDE, H. S. 1998 Spatially localised unstable periodic orbits of a high-dimensional chaotic system. *Phys. Rev. E* **57**, R2511.