J. Fluid Mech. (1991), vol. 233, pp. 83–118
Printed in Great Britain

Short-time Lyapunov exponent analysis and the transition to chaos in Taylor–Couette flow

By JOHN A. VASTANO[1] AND ROBERT D. MOSER[2]
[1] Center for Turbulence Research, Stanford University, Stanford, CA 94305, USA
[2] NASA-Ames Research Center, Moffett Field, CA 94035, USA

(Received 17 May 1990 and in revised form 29 April 1991)

Short-time Lyapunov exponent analysis is a new approach to the study of the stability properties of unsteady flows. At any instant in time the Lyapunov perturbations are the set of infinitesimal perturbations to a system state that will grow the fastest in the long term. Knowledge of these perturbations can allow one to determine the instability mechanisms producing chaos in the system. This new method should prove useful in a wide variety of chaotic flows. Here it is used to elucidate the physical mechanism driving weakly chaotic Taylor–Couette flow.

Three-dimensional, direct numerical simulations of axially periodic Taylor–Couette flow are used to study the transition from quasi-periodicity to chaos. A partial Lyapunov exponent spectrum for the flow is computed by simultaneously advancing the full solution and a set of perturbations. The axial wavelength and the particular quasi-periodic state are chosen to correspond to the most complete experimental studies of this transition. The computational results are consistent with available experimental data, both for the flow characteristics in the quasi-periodic regime and for the Reynolds number at which transition to chaos is observed.

The dimension of the chaotic attractor near onset is estimated from the Lyapunov exponent spectrum using the Kaplan–Yorke conjecture. This dimension estimate supports the experimental observation of low-dimensional chaos, but the dimension increases more rapidly away from the transition than is observed in experiments. Reasons for this disparity are given. Short-time Lyapunov exponent analysis is used to show that the chaotic state studied here is caused by a Kelvin–Helmholtz-type instability of the outflow boundary jet of the Taylor vortices.

1. Introduction

At any given time many physical processes may be active in a turbulent flow. To correctly analyse such a flow, it is necessary to discover which of these processes are of primary importance in producing and maintaining the turbulence. This daunting task has previously been attempted by making *a priori* assumptions about the nature of the turbulence. It is then a secondary task of the researcher to show that these assumptions are justified *a posteriori*. In this paper we report on the first test of short-time Lyapunov exponent analysis, a method we have developed for determining the important physical instabilities in complex fluid flows. This new technique makes use of the relatively recent insights into the nature of turbulence provided by dynamical systems theory.

Temporal Fourier analysis of turbulent flows results in broadband power spectra, which until recently was taken to imply that a continuum of temporal frequencies of motion are excited in turbulent flow. It is now generally accepted that this is not the case. Thus turbulence is neither multiply periodic nor stochastic, but *is* aperiodic. The phase-space attractors for this aperiodic motion, now termed chaos, have a complicated (fractal) structure which led Ruelle & Takens (1971, see also Newhouse, Ruelle & Takens 1978) to call them 'strange attractors'. Fully developed turbulence in a finite domain thus involves a large but finite number of degrees of freedom.

The first confirmation of these mathematical ideas in a fluid dynamic system was in experiments on Taylor–Couette flow (Gollub & Swinney 1975; Fenstermacher, Swinney & Gollub 1979). These experiments detailed a direct transition from two-frequency quasi-periodicity to aperiodic motion with a broadband power spectrum. Brandstater and co-authors (Brandstater *et al.* 1983; Brandstater & Swinney 1987) analysed experimental time series data at Reynolds numbers just past the onset of aperiodicity and found that to within experimental uncertainty the system exhibited low-dimensional chaos. The experimentally determined dimension of the chaotic attractors increased slowly with Reynolds number. Since these results appeared, there has been an explosion of theoretical, experimental, and computational research on chaotic systems. It is therefore ironic that the first example of chaotic transition in hydrodynamics, Taylor–Couette flow, has remained relatively poorly understood.

The short-time Lyapunov exponent analysis technique described in this paper is used to determine the mechanism underlying the transition to chaos in Taylor–Couette flow as described by Brandstater *et al.* (1983). In addition to the general interest in Taylor–Couette flow as an important problem in the study of transition to turbulence, this flow has two advantages for testing the short-time Lyapunov exponent analysis. First, the transition to chaos is gradual; that is, immediately beyond the transition to chaos the system is only weakly chaotic, making the analysis simpler. Secondly, there is a wealth of experimental data on this flow, and in particular the transition to chaos studied here has been well documented by Brandstater *et al.*

To present the results of this study we must draw on two large bodies of research, the first on Taylor–Couette flow and the second on dynamical systems theory. The succeeding section contains the necessary reviews of these topics; in addition, §2.2 provides the mathematical basis for the short-time Lyapunov exponent analysis. In §3 the results of the simulations are presented and compared with the experiments. An analysis of the transition to chaos using short-time Lyapunov exponent analysis is presented in §4, and a summary of this research and concluding remarks are provided in §5.

2. Background

2.1. *Taylor–Couette flow*

Flow between concentric rotating cylinders, Taylor–Couette flow, has been studied by experimentalists and theorists for over 100 years. A brief overview of the experimental and theoretical results relevant to the current study are given below. The review by DiPrima & Swinney (1981) and the references therein should be consulted for further details. In the cases considered here, the outer cylinder is fixed and the inner cylinder rotation frequency is Ω. Consider the (r, θ, z) cylindrical coordinate system with the axis of the cylinders on the z-axis. Let the inner cylinder radius be r_i, the outer cylinder radius be r_o, and the gap be $\delta = r_o - r_i$. The dimensionless parameters that characterize the flow with the outer cylinder fixed are the radius ratio $\eta = r_i/r_o$, the aspect ratio $\Gamma = L/\delta$, and the Reynolds number $R =$

$2\pi r_i \Omega\delta/\nu$. Here L is the working axial height, and ν is the kinematic viscosity. For ease of reference, we define a reduced radius $\hat{r} = (r-r_i)/\delta$, and a reference velocity $U = 2\pi r_i \Omega$ (velocity of the inner cylinder).

For moderately large radius ratios ($\eta \approx 7/8$) and aspect ratios ($\Gamma \geqslant 20$), a standard set of laminar flow states has been observed. Couette flow, observed for R near zero, is a steady state with no axial or aximuthal variations. At a critical Reynolds number, R_c, Taylor vortices appear. This steady state has a well-defined axial periodicity. At higher Reynolds numbers, first one and then two travelling waves appear on the Taylor vortices. These flows are, respectively, temporally periodic (wavy Taylor vortex flow) and quasi-periodic (modulated wavy vortex flow). Finally, the quasi-periodic flow gives way to a chaotic flow.

Each successive transition introduces new flow features that are superposed on the previous flow. Couette flow has only an azimuthal component of velocity. Taylor vortex flow retains the azimuthal symmetry, but horizontal toroidal vortices are superposed on the basic Couette flow, with neighbouring vortices of opposite circulation. The borders between the vortices are alternating radial outflow and inflow jets. While the jets are identified by their radial velocity components, the strongest velocity component in both jets is azimuthal, and the jets carry a large fraction of the angular momentum flux in the flow. The periodic and quasi-periodic states result from azimuthal travelling waves on the boundary jets. In wavy Taylor vortex flow the travelling wave has an integer azimuthal wavenumber m_1 and constant azimuthal rotation frequency $\omega_1/m_1\Omega$. Modulated wavy vortex flow has two travelling waves, and the second wave also has a well-defined wavenumber m_2 and rotation frequency $\omega_2/m_2\Omega$. Since the flow is quasi-periodic, ω_2/ω_1 is irrational.

The most detailed experimental investigation of the transition from quasi-periodicity to chaos was performed by Brandstater and coworkers (Brandstater 1984; Brandstater *et al.* 1983, Brandstater & Swinney 1987). The physical parameters of the experimental system were $\eta = 0.875$ and $\Gamma = 20$. The flow state chosen for study had eight pairs of Taylor vortices, giving an average axial wavelength of 2.5δ. For this case, the transition to Taylor vortex flow occurred at $R_c = 118.4$ and modulated wavy vortices with wavenumbers $m_1 = 4$ and $m_2 = 4$ appeared at $R/R_c \approx 10.0$. A transition to aperiodic flow was observed at $R/R_c = 11.7$. Extensive tests of the aperiodic state at and past transition showed that the flow exhibited low-dimensional chaos. Henceforth, any discussion of experimental results refers to the experiments of Brandstater & Swinney (1987) unless otherwise noted.

It should be noted that flow states other than those discussed above are possible and indeed common. Even for fixed values of aspect ratio and axial and azimuthal wavelength, a variety of bifurcation sequences are possible, depending on the history of the flow. Coughlin (1990) investigated bifurcations from quasi-periodicity to three-frequency motion as well as a period-doubling transition to a chaotic state different from that observed by Brandstater & Swinney (1987).

2.2. *Lyapunov exponent theory*

The spectrum of Lyapunov exponents is a fundamental description of the phase-space evolution of a dynamical system. In recent years Lyapunov exponents have been used to characterize the evolution of many low-dimensional chaotic systems. However, the numerical expense of computing accurate Lyapunov exponents for high-dimensional systems (e.g. our simulation of Taylor–Couette flow) has with a few exceptions (e.g. Keefe, Moin & Kim 1990) precluded their use in such systems. In this paper we show that an analysis of the short-time contributions to the long-time average exponents and the associated perturbations can provide useful information about the physical mechanisms at work in a flow. This analysis does not suffer from the computational difficulties associated with determining the Lyapunov exponents.

2.2.1. *Lyapunov exponents*

The review article of Eckmann & Ruelle (1985) gives a good introduction to the ideas of dynamical systems theory and in particular to the theory of the Lyapunov exponent spectrum, and should be consulted for details. For simplicity, the description here will be couched in the language of ordinary differential equations. Recent theoretical results (e.g. Constantin, Foias & Temam 1988) indicate that the analysis carries over directly to solutions of the Navier–Stokes equations, at least in physically relevant cases.

Consider a solution vector $\boldsymbol{x}(t)$ of a forced dissipative system. If the solution is bounded, then it will almost certainly relax to some subset of the phase space called an attractor. Simple attractors include fixed points (steady states), periodic orbits (one-frequency motion), and n-tori (n-frequency motion). Strange or chaotic attractors are the type of interest here, and they are more difficult to define or describe (see the review by Eckmann & Ruelle 1985). The primary characteristic of chaos is that the system is unstable on the attractor; that is, the system is confined to the attractor and perturbations off the attractor decay, but infinitesimal perturbations in one or more directions on the attractor grow at an exponential rate. Note that chaotic attractors for ideal mathematical systems are fractals, and thus have structure at all lengthscales in phase space; however, in real physical systems this structure is not accessible because at some level it is overcome by noise (e.g. from instruments, round-off, external sources, etc.).

Lyapunov exponents are generalized quantities analogous to the eigenvalues of the linearized stability equations at a steady state. Let the time evolution of a system be governed by

$$\dot{\boldsymbol{x}}(t) = \boldsymbol{F}(\boldsymbol{x}),$$

with an initial condition $x(0)$ on the attractor. An infinitesimal perturbation $\delta\boldsymbol{x}$ to the system state will evolve according to

$$\delta\dot{\boldsymbol{x}}(t) = J(\boldsymbol{x}(t))\,\delta\boldsymbol{x}(t),$$

where $J = \mathrm{d}F/\mathrm{d}\boldsymbol{x}$ is the Jacobian. Given the base flow $\boldsymbol{x}(t)$ and an initial perturbation $\delta\boldsymbol{x}(0)$, the future state of the perturbation can be formally expressed as

$$\delta\boldsymbol{x}(t) = M(t, 0)\,\delta\boldsymbol{x}(0),$$

where

$$M(t, 0) = \exp \int_0^t J(\boldsymbol{x}(s))\,\mathrm{d}s.$$

If the initial norm of the perturbation is unity, then at time t the average exponential growth rate is

$$\overline{\lambda}(t) = \left(\frac{1}{t}\right)\log\left(\|\delta\boldsymbol{x}(t)\|\right),$$

and the long-time average exponential growth rate is

$$\lambda = \lim_{t\to\infty}\overline{\lambda}(t).$$

The long-time average λ is independent of the initial flow state $x(0)$, but *is* dependent on the choice of initial perturbation $\delta x(0)$. The possible values for λ comprise the Lyapunov exponent spectrum of the attractor, ordered from largest to smallest, so $\lambda_1 > \lambda_2 > \lambda_3 \ldots$. An attractor is defined to be chaotic when λ_1 is positive, indicating that there is at least one direction of exponential growth. However, for dissipative systems the sum of the Lyapunov exponents is negative, so there must also be negative exponents. A zero Lyapunov exponent corresponds to non-exponential growth. In particular, a perturbation exactly tangent to the trajectory is associated with a zero exponent. The spectrum of Lyapunov exponents can be used to estimate the fractal dimension of the attractor in phase space using the Kaplan–Yorke conjecture (Fredrickson *et al.* 1983). Assuming unit multiplicity (see below), the sum of the first k exponents is non-negative and the sum of the first $(k+1)$ exponents is negative, then

$$D_{\mathrm{KY}} \equiv k + \frac{\sum_{j=1}^{k} \lambda_j}{|\lambda_{k+1}|}.$$

2.2.2. *Lyapunov perturbations*

The short-time Lyapunov exponent analysis presented here relies on the properties of the perturbations $\delta x(t)$ associated with the Lyapunov exponents. The space of all possible (infinitesimal) perturbations at $x(t)$ will be called the tangent space at $x(t)$. To simplify the discussion, it is again assumed that the phase space and the tangent spaces are $\Re^N$, for N finite.

The multiplicative ergodic theorem, first proven by Oseledec (1968), states that the tangent space at any time t (in particular, $t = 0$) may be decomposed into a nested set of subspaces E_i, such that $\Re^N \equiv E_1 \supset E_2 \supset \ldots$, and

$$\lim_{t\to\infty} \left(\frac{1}{t}\right) \log(\|M(t,0)\,\delta x\|) = \lambda_i \quad \text{for} \quad \delta x \in E_i \text{ and } \delta x \notin E_{i+1}.$$

Thus any initial perturbation that is not contained in E_2 will grow at the average rate λ_1, while perturbations in the restricted subspace E_2 and not in E_3 will grow as λ_2. This can be carried iteratively to any number of exponents.

Now consider the Lyapunov exponents for the time-reversed system. An attractor for the time-forward system will become unstable when time is reversed, but it remains an invariant set, so the multiplicative ergodic theorem applies. Lyapunov exponents for the time-reversed system are exactly the same as for the time-forward case, but with opposite sign. The tangent space at $t = 0$ can then be divided into a different set of subspaces $\bar{E}_i$, with $\bar{E}_1 \subset \bar{E}_2 \ldots \subset \bar{E}_N \equiv \Re^N$, and

$$\lim_{t\to-\infty} \left(\frac{1}{|t|}\right) \log(\|M(t,0)\,\delta x\|) = -\lambda_i \quad \text{for} \quad \delta x \in \bar{E}_i \text{ and } \delta x \notin \bar{E}_{i-1}.$$

Thus for the time-reversed system the least-constrained subspace corresponds to λ_N, and almost any perturbation will grow at the average exponential rate $-\lambda_N$.

These two sets of subspaces can be used to form a set of disjoint subspaces $F_i \equiv E_i \cap \bar{E}_i$ (Eckmann & Ruelle 1985). If $\delta x \in F_i$, then the long-time average exponential growth rate of δx will be $\pm\lambda_i$, forward and backward in time. The dimension of F_i is the multiplicity of the exponent λ_i and is independent of x. There are three important properties of the F_i. First, basis vectors in each F_i can be chosen to comprise a basis set for the tangent space. Second, the F_i are in general *not* orthogonal. Third, the decomposition of tangent space into the F_i depends on the phase-space location x. In particular, if at $x(0)$ we have $\delta x \in F_i(x(0))$, then

$$M(t,0)\,\delta x \in F_i(x(t)).$$

The short-time contribution $\lambda_i(t)$ to λ_i (with unit multiplicity) is defined as the instantaneous exponential growth rate of $\delta x \in F_i$. That is

$$\lambda_i(t) = \frac{1}{\|\delta x(t)\|}\frac{\mathrm{d}\|\delta x(t)\|}{\mathrm{d}t} \quad \text{for} \quad \delta x(t) \in F_i(x(t)).$$

Now, if we knew the $F_i(x(t))$ and the $\lambda(t)$, we would know what perturbations on the attractor were unstable in the long term (i.e. the space spanned by the F_i associated with positive exponents or the 'unstable subspace'), and when (or where in phase space), the instabilities are most active. By analysing these data it may be possible to determine the cause of the instabilities underlying the chaos. For example, by examining a perturbation in F_1 at a time when $\lambda_1(t)$ is a large local maximum and relating the perturbation to features in the base flow $x(t)$, the features of the base flow responsible for the most dangerous instability may be determined. Note that these perturbations are in general different from the eigenfunctions of the local Jacobian at time t. The instantaneous growth of the local eigenfunctions can be much larger than that of the perturbations in the subspace F_1. However, in the long run the growth of perturbations in F_1 will dominate. The F_i provide more useful information regarding the stability of the system than the local eigenfunctions because they depend on the history of the flow rather than its form at a single point.

2.2.3. *Algorithmic considerations*

The properties of the F_i can be exploited to compute the entire Lyapunov exponent spectrum (of a finite-dimensional system). A number of basically similar numerical approaches have been suggested (Benettin *et al.* 1980; Wolf *et al.* 1985; Eckmann & Ruelle 1985; Greene & Kim 1987; Goldhirsch, Sulem & Orszag 1987). All of the methods evolve a set of perturbations and obtain information about the growth rates associated with the F_i. The differences between the approaches are strictly numerical and are negligible when the methods are applied correctly. Following Benettin *et al.*, an initially orthogonal set of N perturbations is evolved by integrating the equations of motion for both the base flow $x(t)$ and the perturbations. Each perturbation can be decomposed onto the set of basis functions chosen from the F_i. As time proceeds, the components grow or shrink exponentially, according to $\lambda_i(t)$, and since λ_1 dominates, all of the perturbations ultimately tend towards F_1. Before the set of perturbations become numerically indistinguishable, a Gram–Schmidt reorthogonalization is performed. The first perturbation in the set is simply renormalized, so it will evolve to F_1, and thereafter the normalization factor is the growth undergone by perturbations in F_1 since the last renormalization; thus, the logarithm of the factor is an approximation to the short-time contribution to λ_1,

$$\lambda_1(t) \approx \left(\frac{1}{\Delta t}\right)\log\left(\frac{\|\delta x(t+\Delta t)\|}{\|\delta x(t)\|}\right).$$

The second perturbation in the set is then made orthogonal to the first, which has the effect of subtracting out the growth of the component in F_1. The fastest growing component of the remainder is thus in F_1 and F_2 (i.e. $F_1 \otimes F_2 = E_2$), and will be orthogonal to F_1. Area elements in this subspace grow at the rate given by the sum of the first two Lyapunov exponents. The normalization factor for the orthogonalized second perturbation yields the short-time contribution $\lambda_2(t)$. Proceeding iteratively, the normalization of the third perturbation will give the short-time contribution $\lambda_3(t)$, and in principle all N Lyapunov exponents can be computed. The set of perturbations obtained in this manner is identical to the set obtained by orthogonalizing the F_i, starting with F_1 and are called the Lyapunov perturbations (L_i).

The unavoidable problem with computing Lyapunov exponents, especially for model systems with many degrees of freedom, is that the convergence of the running average exponents to their asymptotic values is slow. Goldhirsch *et al.* (1987) conjectured that the average converges as $(1/t)$, but the convergence may be even slower (see §3.3). This difficulty is bypassed when the short-time contributions to the Lyapunov exponents are studied, because the convergence of the initial perturbations to the appropriate subspaces is exponentially fast. Therefore, after a brief initial relaxation the Lyapunov perturbations and the short-time contributions are available. However, as discussed above, L_i is not in general an element of F_i (though L_1 is in F_1) because the F_i are not orthogonal. The important property of the F_i for the short-time Lyapunov exponent analysis is that they span the unstable subspaces. As pointed out by Greene & Kim (1987), the orthogonalized Lyapunov perturbations span the same spaces and will therefore serve equally well in the short-time analysis. It is these Lyapunov perturbations which are examined in §4.

3. Numerical simulations

3.1. *Numerical method*

To simulate Taylor–Couette flow we solved the three-dimensional incompressible Navier–Stokes equations in an axially periodic domain between concentric cylinders. A spectral method using expansion functions that satisfy the boundary conditions and continuity equation had previously been developed for this flow (Moser, Moin & Leonard 1983) and used extensively for computations of a curved turbulent channel (Moser & Moin 1987). The linear part of the equations is time-advanced using the Crank–Nicolson (implicit) scheme, while the nonlinear term is time-advanced using a compact storage, third-order Runge–Kutta (explicit) scheme devised by A. Wray (private communication). This code was modified to allow the computation of N Lyapunov exponents by including N perturbation fields which evolve according to the linearized equations.

Lyapunov exponents are invariant to smooth transformations of the phase-space coordinates. We are thus free to choose a computationally convenient inner product for the orthogonalization procedure discussed in §2.2.3. The orthogonalization procedure is performed at a point in the computation at which the perturbation fields are represented as a finite Fourier series in θ and z with wavenumbers k_θ and k_z for each of $N_r + 1$ radial locations $\hat{r}_j = \tfrac{1}{2}(\cos(\pi j/N_r) + 1)$. Let $\delta\hat{u}_i(k_\theta, k_z, j)$ and $\delta\hat{v}_i(k_\theta, k_z, j)$ be the Fourier transforms of the ith velocity components of the two perturbations δu and δv. The inner product is then defined as

$$(\delta u, \delta v) = \sum_{i=1}^{3} \sum_{j=0}^{N_r} \sum_{k_\theta \geqslant 0} \sum_{k_z} \mathrm{Re}\,\{\delta\hat{u}_i(k_\theta, k_z, j)\, \delta\hat{v}_i^{*}(k_\theta, k_z, j)\}.$$

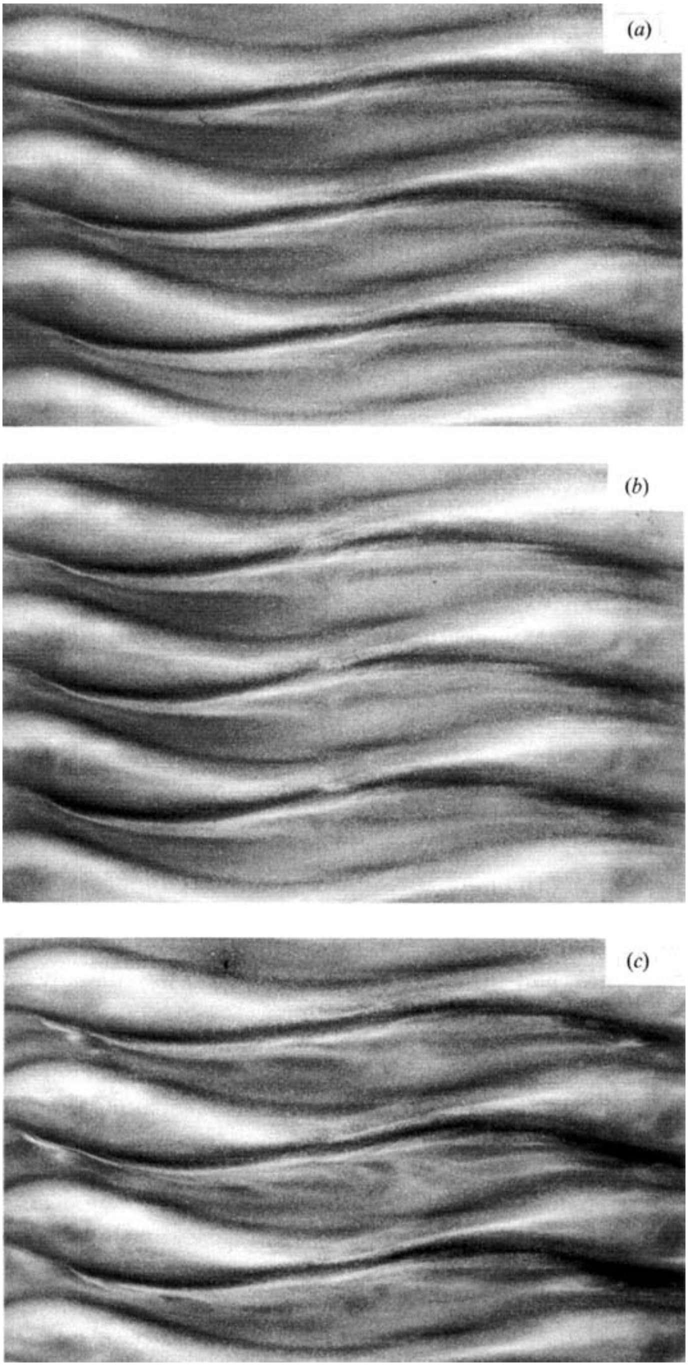

FIGURE 1. Photographs of the experimental flow near the transition from quasi-periodicity to chaos. The values of R/R_c are (*a*) 10.6, (*b*) 11.2, and (*c*) 11.9. The flows are visualized with a suspension of small platelets, and the dark bands indicate inflow and outflow jets. Photographs courtesy of Anke Brandstater and Harry Swinney (see Brandstater & Swinney 1987).

Note that this is not a consistent approximation of the L^2 inner product of the two perturbations, but is closely related to the Chebychev norm in $\hat{r}$.

Because of the azimuthal and axial periodicity of the model system, perturbations corresponding to axial or azimuthal rotations will neither grow nor decay, and will thus yield zero Lyapunov exponents. There is no need to include them in the computations. To eliminate the axial and azimuthal shifts, all perturbations are made orthogonal to them in a pre-conditioning step of the orthogonalization procedure discussed in §2.2.3. This reduces the number of perturbations which must be carried in the computations by two. For the case of periodic and quasi-periodic Taylor–Couette flow, the azimuthal shift is a perturbation on the attractor since these states consist of azimuthal travelling waves. Thus, by eliminating the azimuthal shift in these cases, we eliminate a zero exponent of the flow. The periodic flow will have no zero exponents and the quasi-periodic flow will have one zero exponent (instead of one and two respectively). In all cases the axial shift, and in the chaotic case the azimuthal shift as well, is off the attractor; these shifts correspond to Lyapunov exponents which are in some sense 'parasitic', and can be ignored.

3.2. *Validation of computations*

The numerical code used here was validated by Moser and co-workers (Moser *et al.* 1983; Moser & Moin 1984) by comparisons with numerous experimental results (including King *et al.* 1984) as well as the code of Marcus (Marcus 1984*a*, *b*). For example, the travelling wave frequencies found with the code used here agree to within one part in 10^3 with the experiments of King *et al.* and to within one part in 10^4 with the results of Marcus (Moser & Moin 1984).

The computational effort necessary to estimate N Lyapunov exponents is at least $N+1$ times that required to compute the base flow. The simulations must therefore be performed with as few numerical degrees of freedom as possible. The standard resolution used here was 16 Chebychev modes radially and 32 Fourier modes both azimuthally and axially. The nonlinear terms were computed on a physical grid and dealiased using the 3/2 rule.

To verify that this resolution was sufficient to accurately represent the flow, several states at Reynolds numbers below the transition to chaos was computed and compared with experimental results, as well as to the computations of Coughlin (1990). In addition, higher-resolution simulations were performed to confirm the adequacy of this resolution. To match the experimental conditions (§2.1), the axial wavelength was fixed at 2.5δ and the azimuthal period was set to $\frac{1}{2}\pi$. This allows the simulation of one pair of Taylor vortices and one period of the travelling waves in the $m_1 = m_2 = 4$ quasi-periodic state. Any axial or azimuthal symmetry-breaking bifurcations in the experiment are therefore suppressed in the simulations.

In the experiments, the quasi-periodic state was observed from $R/R_c = 10$ until a transition to aperiodicity at about $R/R_c = 11.7$. Figure 1 shows photographs of the experimental flow at three Reynolds numbers. To mimic the experiments, the simulation was initialized at $R/R_c = 10.6$ with random noise and evolved until the asymptotic quasi-periodic flow developed. Using this case as an initial condition, the Reynolds number was increased or decreased in small steps to determine the range of stability of the quasi-periodicity. An illustration of the flow at $R/R_c = 10.6$ is shown in figure 2, which should be compared to figure 1(*a*). The large-scale features of the flow, namely the Taylor vortices and the travelling waves on the vortex boundaries, are readily apparent. Note the high-speed and low-speed azimuthal jets

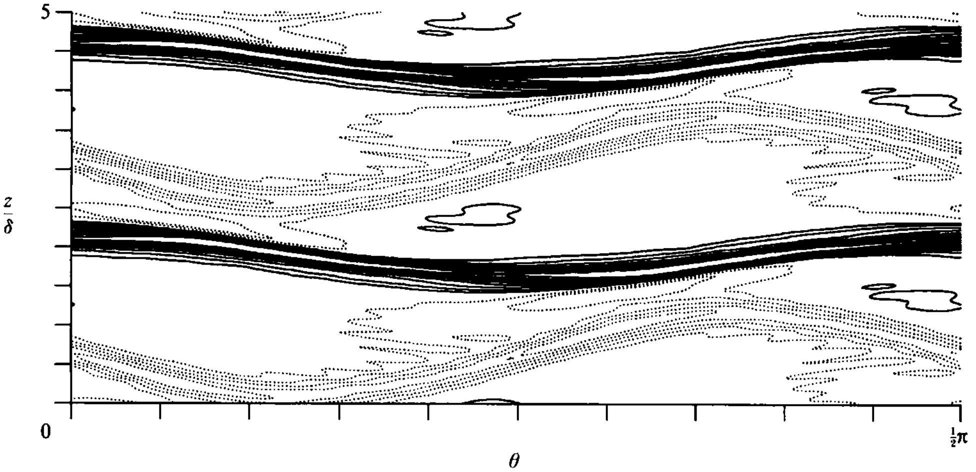

FIGURE 2. Azimuthal velocity contours in the r-plane at $\hat{r} = 0.31$ for the quasi-periodic flow in the simulation at $R/R_c = 10.6$. Solid contours indicate velocities greater than the plane average, dotted contours are less than the plane average. The contour increment is $0.03U$.

on the vortex boundaries in figure 2. These are the outflow and inflow boundary jets respectively (see §2.1), which will be of great interest in §4.

In both the experiments and the simulations the first travelling wave frequency was $0.33\, m_1 \Omega$. The ratio of the two travelling wave frequencies (ω_1/ω_2) in the experiment and computations agrees to 2% over the Reynolds-number range for which they are both quasi-periodic. Also, this ratio changed by less than 1% when the resolution was doubled. As the Reynolds number was decreased in the simulation, the quasi-periodic state remained stable to $R/R_c \approx 9.4$. Below this Reynolds number, a third frequency (ω_3) of motion appeared. Coughlin (1990) has studied this case in detail with a different spectral code at higher resolution, and confirms the appearance of ω_3. The good agreement of our simulation both with experiments and with simulations at higher resolutions gives us confidence that our resolution is sufficient for the weakly chaotic state to be studied here.

The Lyapunov exponent computations were also validated in tests on the quasi-periodic flow, for which λ_1 and λ_2 should be zero, corresponding to the two independent frequencies of motion. Recall that the code for computing Lyapunov exponents (see §3.1) removes perturbations corresponding to azimuthal and axial rotations. This should remove one of the two zero exponents for the quasi-periodic flow. That is in fact what occurs. To check that the code was performing correctly, the perturbations corresponding to azimuthal rotations were left in, and a second zero exponent appeared. Finally, a periodic (one travelling wave) state for a different axial wavelength and azimuthal wavenumber regime was computed, and the removal of perturbations corresponding to azimuthal rotations resulted in the removal of the (single) zero Lyapunov exponent.

3.3. *Transition to aperiodicity* (*chaos*)

In this section a transition to a chaotic flow from quasi-periodicity will be illustrated using the standard tools of nonlinear dynamics. This transition was approached both from below and from above in Reynolds number, and no hysteresis was observed, indicating that the bifurcation to chaos is supercritical for the parameter values studied here. The asymptotic flow at $R/R_c = 10.98$ is quasi-periodic. At $R/R_c = 11.15$, a quasi-periodic initial condition taken from the asymptotic flow at $R/R_c = 10.98$ did not relax to quasi-periodicity even for an evolution of several hundred cylinder periods. At a slightly higher Reynolds number, the aperiodic part of the velocity variation is larger. The transition to aperiodicity thus occurs in the Reynolds-number range $10.98 < R/R_c < 11.15$.

The computer resources necessary to accurately compute a partial Lyapunov exponent spectrum are too great to study many different Reynolds-number cases. In this study, two cases were chosen: a quasi-periodic state at $R/R_c = 9.8$, and an aperiodic state at $R/R_c = 11.3$. The experimental flows were studied using a time series of a single observable, the radial velocity at an arbitrarily chosen point on the midplane of the gap between the cylinders. As a first step in our analysis we repeat the experimental analysis using the same observable. The velocity time series from the simulations and the corresponding power spectra derived from them are shown in figure 3. The large-scale variations of velocity in the chaotic time series are essentially similar to those in the quasi-periodic time series, except that there are many more small-scale fluctuations in the chaotic case. The power spectrum of the chaotic system displays a broadband noise component at a level of approximately 10^{-9} at low frequencies, whereas the quasi-periodic case has a noise level at 10^{-14} at low frequencies. The rise in the background noise level at low frequencies in the quasi-periodic case is just due to the limited length of the time series ($220/\Omega$) used to compute the spectrum (a Hanning-window fast Fourier transform was used). The sharp peaks in the spectrum are not fully resolved, resulting in the apparent noise. The chaotic spectrum was computed with a similar length time sample ($220/\Omega$) and would thus have a similar noise level, but the noise due to the chaos is several orders of magnitude higher.

Time-delay reconstructions (Packard *et al.* 1980; Takens 1981) of the attractors from the radial velocity time series are shown in figure 4. Given a scalar time series $\{x(t)\}$, an n-dimensional vector time series is formed by successive time delays: $\boldsymbol{x}(t) = (x(t), x(t+\tau), \dots, x(t+(n-1)\tau))$, where τ is the time delay. The choice of τ for the reconstructions shown here was made using the information-theoretic criteria developed by Fraser & Swinney (1986). The reconstructed chaotic attractor in figure 4(*b*) is more disordered than the quasi-periodic attractor in figure 4(*a*). These reconstructions should be compared to figure 6 in Brandstater & Swinney (1987).

The evidence presented above that our two cases are quasi-periodic and chaotic is qualitative. Computations of the Lyapunov exponent spectra for these cases yield quantitative measures of the attractors. The Lyapunov exponent estimates were obtained using the following procedure: the base flow was evolved until an asymptotic regime was reached, and then several perturbation fields were added to begin the Lyapunov exponent computation. The base flow and the perturbations were then evolved until the perturbations had relaxed to *their* asymptotic evolution, and at that point the long-time running averages of the exponents were begun.

For the quasi-periodic flow, which will be discussed later, the asymptotic regime of the base flow was reached when the azimuthal wave speed relaxed to periodic oscillations. This required approximately 40 cylinder revolutions when the base flow was initialized with energy in the appropriate azimuthal wavenumbers. For the chaotic flow, there is no reliable way to determine if an asymptotic regime has been reached, but an evolution of 50 cylinder periods was sufficient to produce reconstructed attractors and power spectra that did not vary significantly at later times.

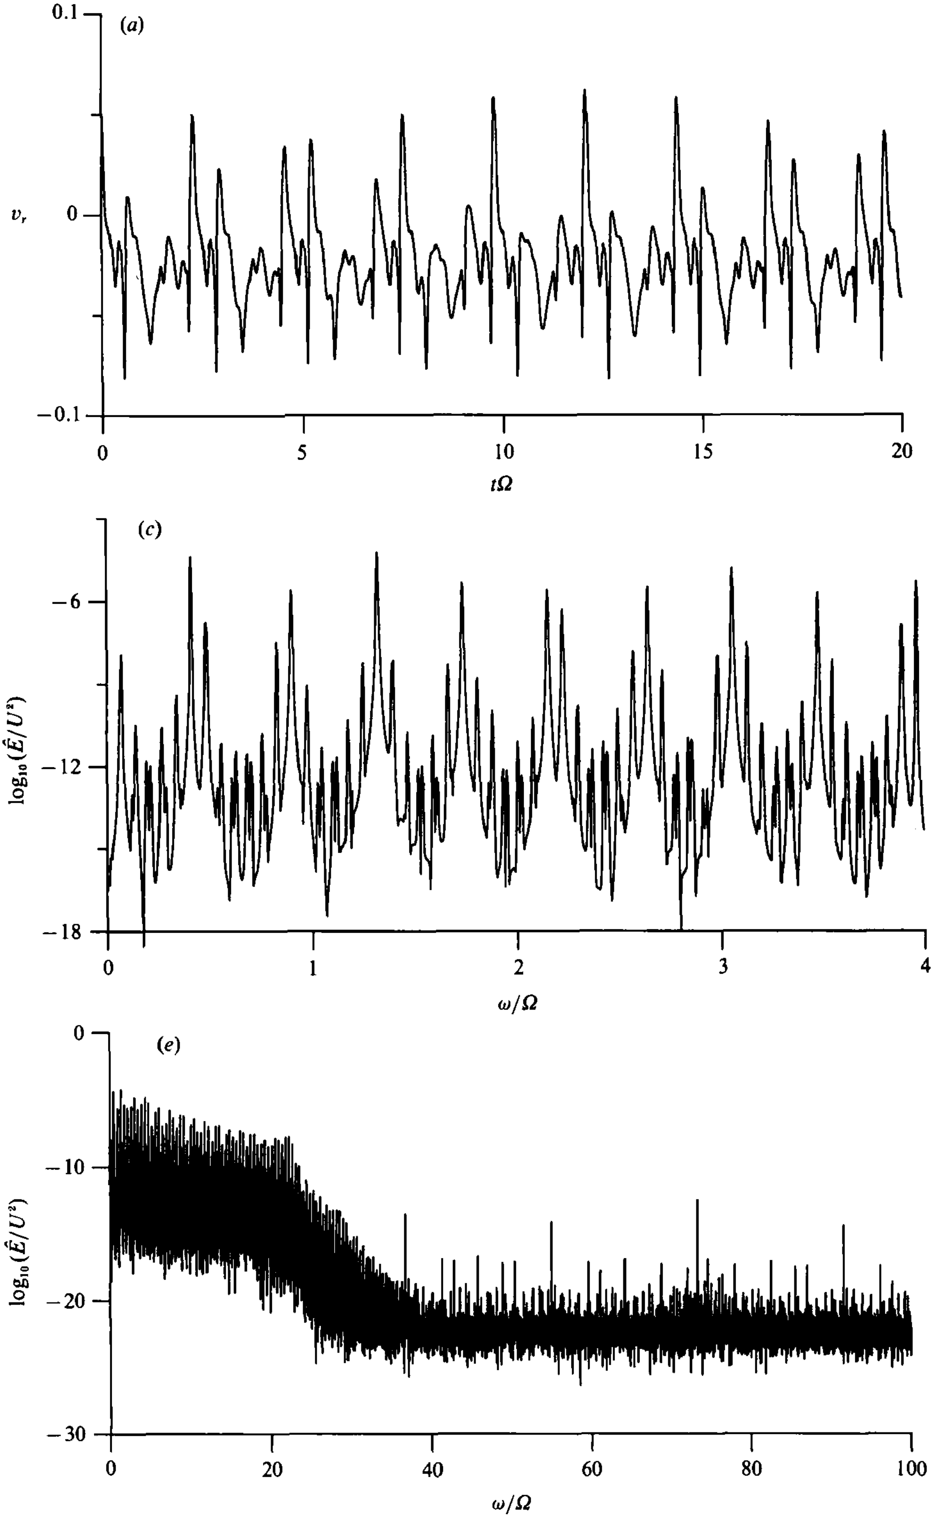

FIGURE 3(a, c, e). For caption see facing page.

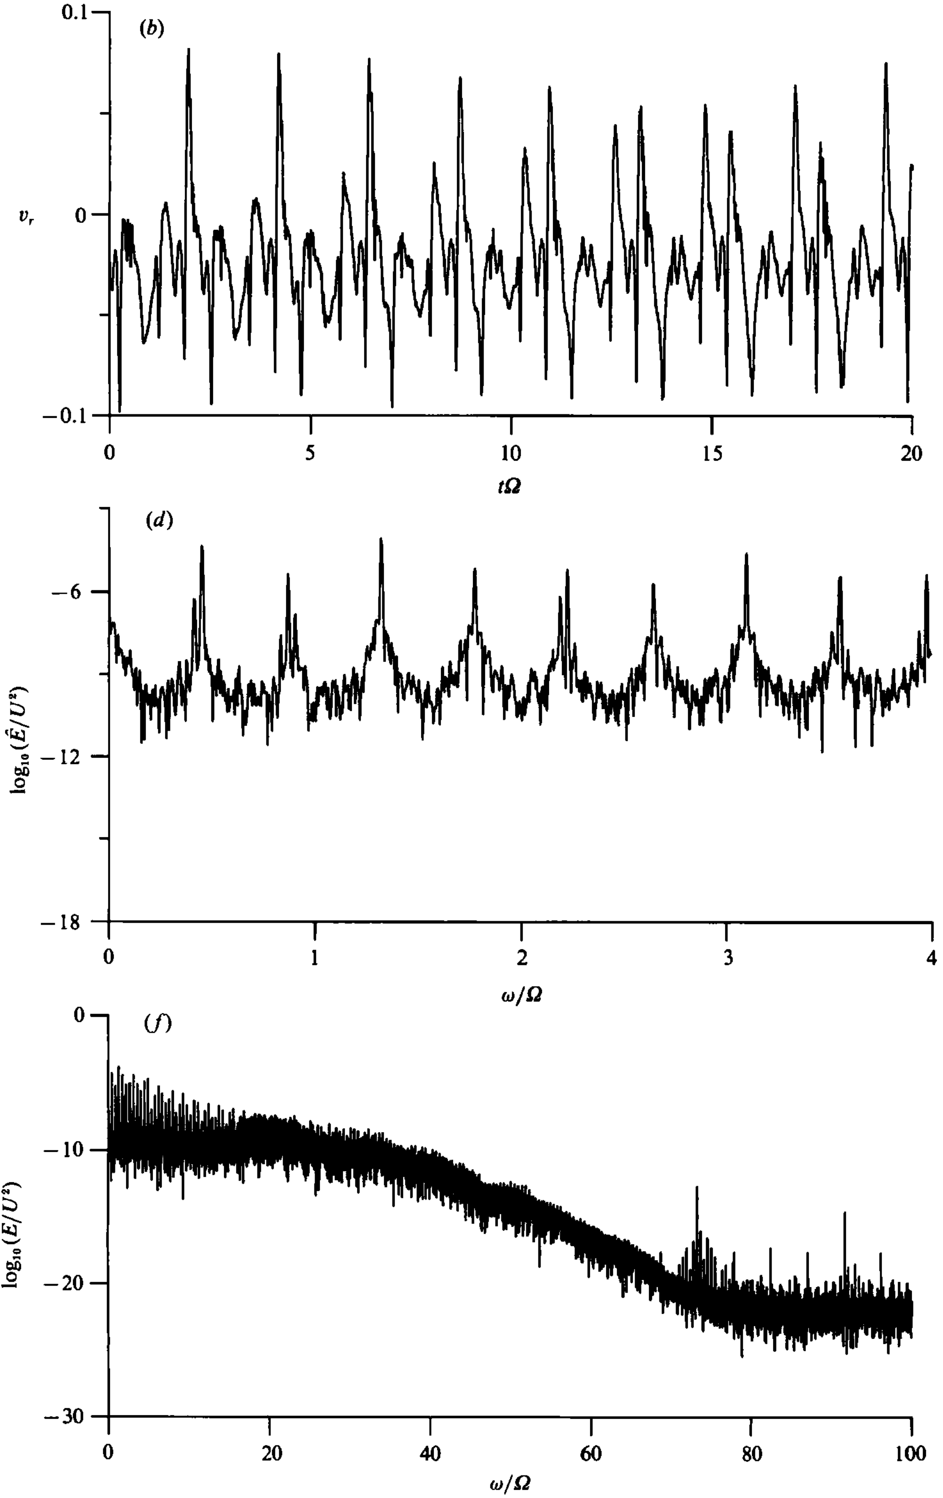

FIGURE 3. Radial velocity time series $v_r(t)$ at midgap for the computed (*a*) quasi-periodic flow ($R/R_c = 9.8$) and (*b*) chaotic flow ($R/R_c = 11.3$), with graphs of the associated power spectra in (*c*, *e*) and (*d*, *f*), respectively.

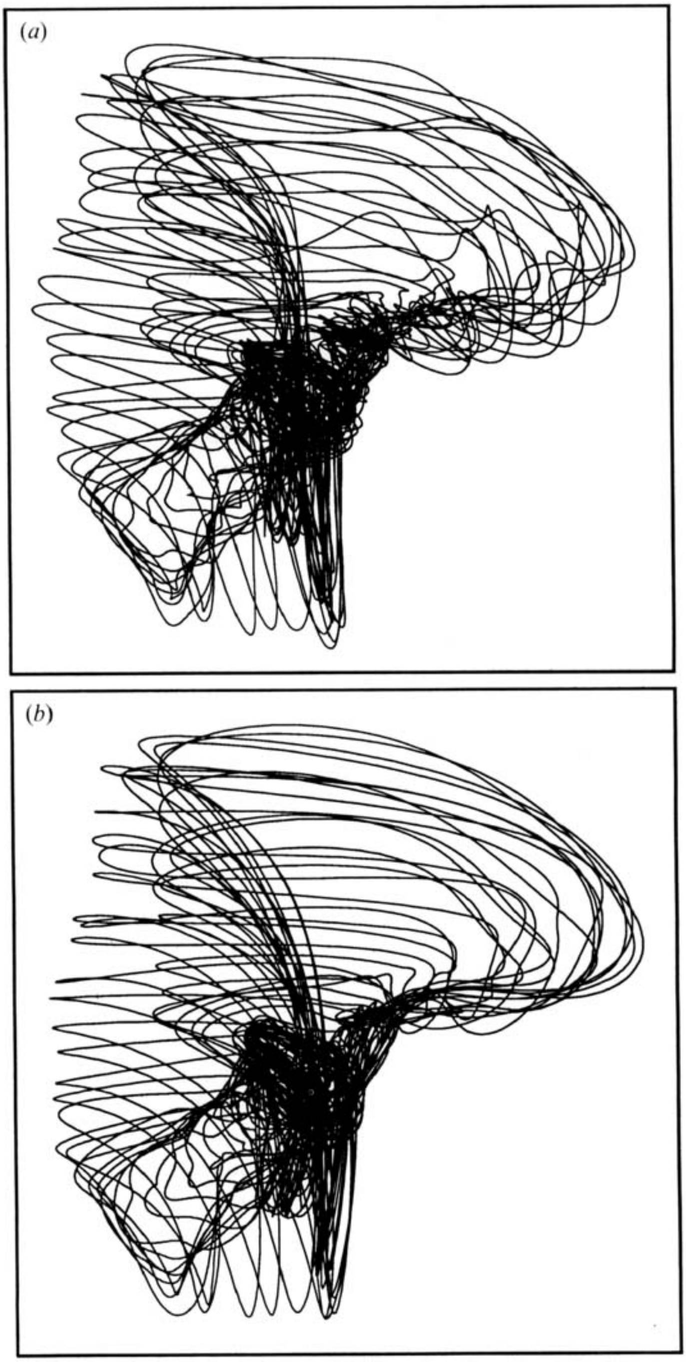

FIGURE 4. Time-delay-coordinate reconstructions of (*a*) the quasi-periodic attractor at $R/R_c = 9.8$ and (*b*) the chaotic attractor at $R/R_c = 11.3$. For both cases a time series of the radial velocity near the midgap was used with a time delay of $0.09/\Omega$.

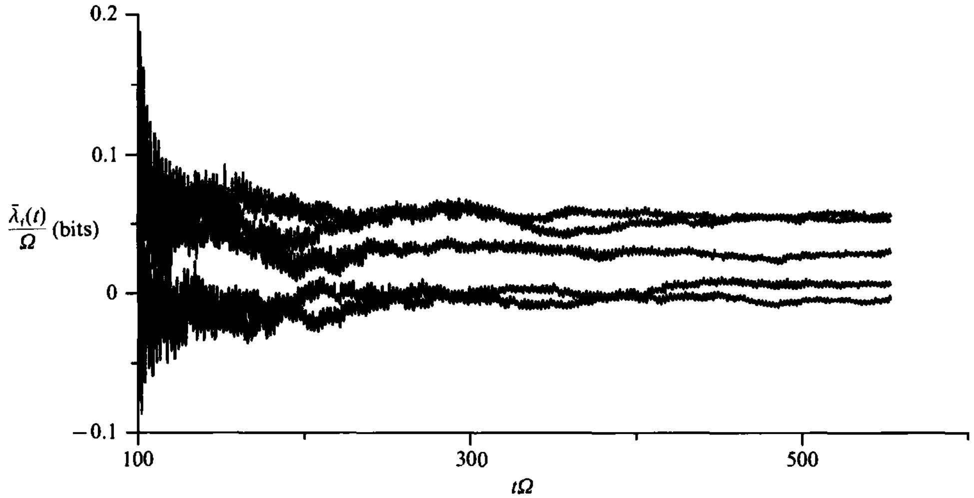

FIGURE 5. Convergence of the first five Lyapunov exponent estimates $\bar{\lambda}_i(t)$ towards their asymptotic values λ_i for the chaotic flow at $R/R_c = 11.3$.

For the chaotic flow, an initial set of five perturbations was used to obtain estimates of the five largest Lyapunov exponents. When it became clear that more exponents would be required to obtain a dimension estimate using the Kaplan–Yorke conjecture, an additional nine perturbations were added to the computation, yielding (less well-converged) estimates of the next nine Lyapunov exponents. Recall that the inclusion of additional perturbations does not affect the original set of perturbations in any way.

The results of the Lyapunov exponent computations are summarized in table 1. The evolution time over which the running averages were performed is indicated for each set of exponents. Note that the quasi-periodic attractor has the requisite double zero Lyapunov exponent, and no positive exponents. There are no more than five positive Lyapunov exponents for the chaotic attractor (recall that there must be at least one zero exponent). The Kaplan–Yorke dimension of this attractor is estimated to be 8.2 ± 0.9.

The convergence of the running average of the first five exponents is shown in figure 5. Note that the exponents converge extremely slowly. In figure 6 the convergence of λ_1 with inverse time is shown. If the conjectured $1/t$ convergence of Goldhirsch *et al.* (1987) holds, then it does so only on computationally inaccessible timescales. Even at the minimal resolution used here, the evolution of the base trajectory and five perturbations for one cylinder period required over 2000 s on a Cray-YMP.

Because the long-time averages converge slowly, it is difficult to quantify the uncertainty associated with an exponent average at finite time. For the quasi-periodic case in table 1, the exponent contributions relax to a periodic form (see §4.3) and the long-time averages can therefore be found with excellent accuracy. The uncertainty estimates for the chaotic case in table 1 were derived by assuming that the average of the contributions over one period of the waveform flex (see §4.1) is uncorrelated with the average over other periods. The correlation of the average over successive periods was found to be less than 0.2 for the exponents we calculated. With this assumption the variance of a Lyapunov exponent averaged over N_{per} periods is just $1/N_{\text{per}}$ times the variance of the average over one period. This was used

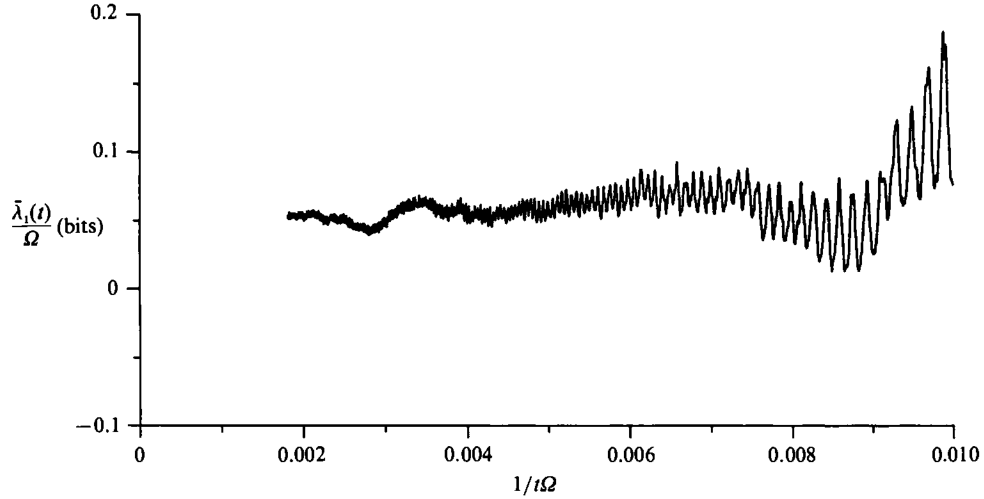

FIGURE 6. Convergence of the $\bar{\lambda}_1(t)$ towards λ_1 for the chaotic flow at $R/R_c = 11.3$, plotted versus inverse time.

λ/Ω (bits)	$R/R_c = 9.8^a$	$R/R_c = 11.3^b$
λ_1	0^c	0.052 ± 0.011
λ_2	0.000 ± 0.001	0.055 ± 0.010
λ_3	-0.169 ± 0.001	0.029 ± 0.010
λ_4		0.006 ± 0.010
λ_5		-0.004 ± 0.010
λ_6		-0.006 ± 0.033
λ_7		-0.041 ± 0.031
λ_8		-0.096 ± 0.025
λ_9		-0.050 ± 0.030
λ_{10}		-0.123 ± 0.037
λ_{11}		-0.128 ± 0.028
λ_{12}		-0.137 ± 0.033
λ_{13}		-0.186 ± 0.030
λ_{14}		-0.183 ± 0.029

[a] Averaging time was $295/\Omega$.
[b] Averaging time was $572/\Omega$ for $\lambda_1 - \lambda_5$ and $58/\Omega$ for $\lambda_6 - \lambda_{14}$.
[c] By definition (not computed, see text).

TABLE 1. Estimated Lyapunov spectrum for the two Reynolds numbers studied. The units are bits, where bits indicates the logarithm base two. See text for an explanation of the uncertainty estimates.

to determine the variance of each exponent estimate in table 1. The deviation is reported as the uncertainty in the table. Note that the assumptions underlying this analysis imply that the exponents converge like $1/t^{\frac{1}{2}}$ rather than the $1/t$ convergence conjectured by Goldhirsch *et al.* (1987).

3.4. *Comparison to experiments*

Brandstater & Swinney (1987) studied a large number of quasi-periodic and chaotic states over a wider range of Reynolds numbers than has been explored in our simulations. Their principal result was that the transition to chaos is gradual; that is, the fractal dimension of the chaotic attractor increases slowly with Reynolds number. Our simulation does not contradict this observation. However, chaos was detected in the simulation at a Reynolds number 5% lower than the observed transition in the experiment, and the Lyapunov exponent calculation showed a significantly larger value for the fractal dimension near the onset of chaos. These differences are not surprising. Brandstater & Swinney comment that small-scale fluctuations were observed in visualizations of the quasi-periodic states prior to the Reynolds number at which the dimension estimate began to rise above two; this is visible in the photograph of their flow at $R/R_c = 11.2$ (see figure 1). These fluctuations did not affect the dimension estimates because they were below the noise level in the experimental measurements.

The power spectra from the experiment at $R/R_c = 11.4$ (see figure 3 of Brandstater & Swinney 1987) indicate no more than five decades of signal-to-noise separation between the discrete frequency peaks and the broadband noise level. In contrast, the power spectrum presented in figure 3 for the chaotic attractor found in the simulation at $R/R_c = 11.3$ has at least six decades of separation. If chaotic broadband noise at this level were present in the experiments, it would not have been detected. Also, if a dimension-estimating algorithm is correctly applied to the data, it will only measure the fractal-like scaling of the attractor above the experimental noise floor. Thus any experimental (including numerical experiments) determination of dimension is only a statement about the attractor's structure above some phase-space lengthscale. The correct interpretation of Brandstater & Swinney's results, as they carefully state, is that chaos appears above the noise in their system at $R/R_c \approx 11.7$.

Experimental time series data are unavoidably contaminated with (at least) measurement noise. One standard technique for removing noise is to low-pass filter the time series. The filter cutoff used in the experiments of Brandstater & Swinney was about 20 times the inner cylinder frequency. In the simulations, the computational noise floor is extremely low, and small-scale, relatively high-frequency fluctuations are observable, e.g. the time series in figure 3. To demonstrate the effect of low-pass filtering and limited resolution the attractor is reconstructed in figure 7 from the simulation time series after it was low-pass filtered with cutoffs at 20 and 15 times the cylinder frequency. These should be compared to the original reconstruction in figure 4. Clearly, some dynamics on the attractor are removed by the low-pass filtering. If this also occurred when the experimental data were filtered, some part of the true chaotic dynamics would be deleted along with the high-frequency noise. Thus, beyond the apparent transition to chaos, the experimental dimension estimates are likely to underestimate the attractor dimension.

There are also two aspects of the simulations which should be considered in any comparison with the experiments. First, axial periodicity is assumed in the computational model, while in the experiments the finite length of the cylinder provided end effects and axial wavelength modulations. Second, the computations are done with finite though adequate spatial and temporal resolution. Care was taken to minimize these limitations; however, they cannot be eliminated as a source of discrepancies with experiments.

It is important to note that a four-fold azimuthal symmetry was imposed in the computation model. A transition to chaos that broke this symmetry would therefore not be allowed, and could not be observed in the simulations. If the chaotic transition in the experimental system required breaking the azimuthal symmetry of the quasi-periodic state, the transitions observed in the simulation and the experiment would

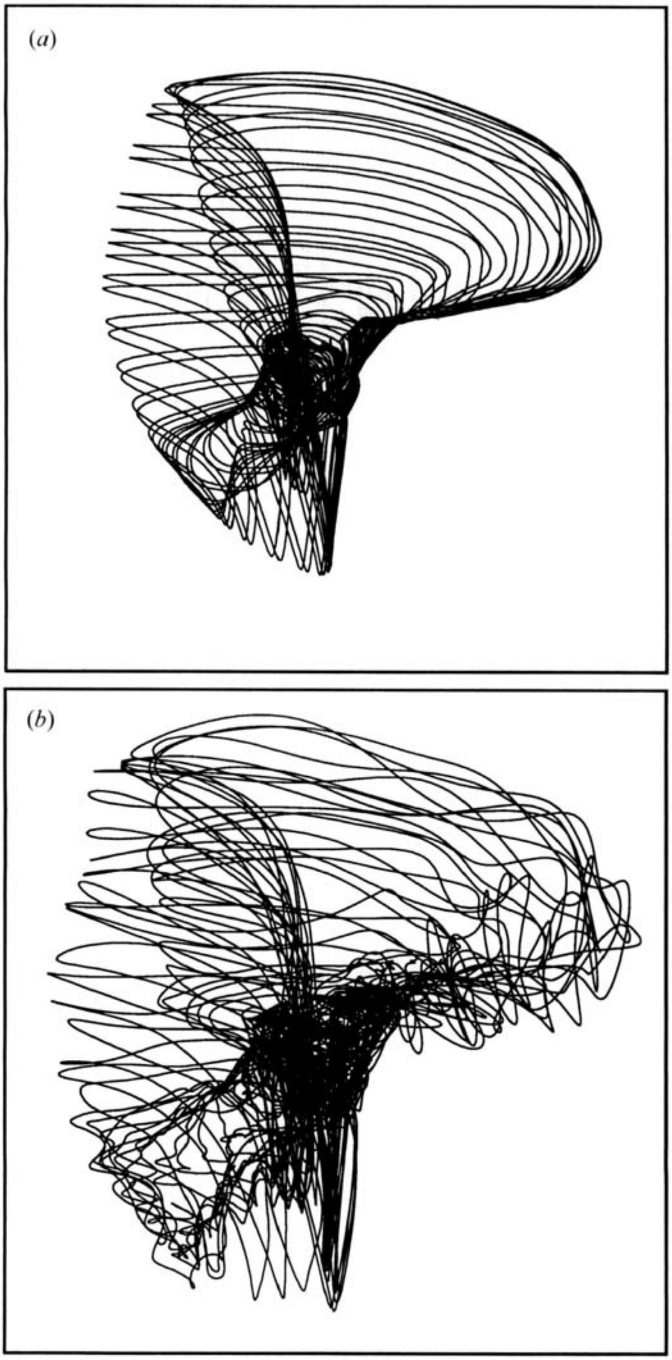

FIGURE 7. Time-delay-coordinate reconstructions of the chaotic attractor at $R/R_c = 11.3$ using low-pass filtered data. In (*a*) the filter cutoff frequency is 20Ω; in (*b*) the cutoff frequency is 15Ω. The time delay is $0.09/\Omega$.

be vastly different. This is not the case. The close correspondence in the parameter values and form of the observed transitions indicates that the azimuthal symmetry need *not* be broken at the transition to chaos in Taylor–Couette flow.

In summary, the current simulation results, which indicate a transition to chaos slightly below $R/R_c = 11.15$ are consistent with the experimental results of Brandstater & Swinney (1987). The apparent disagreement between simulations and experiments in the transition location and the estimated dimensions may be caused by the limited precision of the experimental measurements. The assumption of axial and azimuthal periodicity in the simulations may also preclude an exact correspondence with the experiments.

4. Short-time Lyapunov exponent analysis

The estimation of even a small number of Lyapunov exponents for an attractor of a large-scale fluid dynamics model consumes enormous computer resources. If the result of this expenditure is only a handful of numbers, and with those an estimate of the attractor dimension, then it is not clear that the expense is worthwhile. Fortunately, as was discussed in §2, the short-time contributions $\lambda_i(t)$ to the long-time average exponents as well as the perturbation fields themselves are local *phase-space* properties on the attractor and can be used to gain information about the mechanisms at work in the fluid flow. Any disturbance with a component in the subspace spanned by the Lyapunov perturbations corresponding to the positive Lyapunov exponents will exhibit long-term exponential growth. The task is then to discern what mechanisms are responsible for the exponential growth by examining the perturbations which grow. In general this would be difficult; however, if the growing perturbations have a localized structure in *physical* space the analysis could be greatly simplified.

In this section we will apply short-time Lyapunov exponent analysis to our simulations of Taylor–Couette flow. The Lyapunov perturbation fields will be shown to be highly structured spatially, and these structures will be used to discover the physical mechanism underlying the chaos. This mechanism is a Kelvin–Helmholtz instability originating in the high-shear region of the outflow jet. The short-time Lyapunov exponent contributions $\lambda_i(t)$ vary enormously, but a simplified one-dimensional stability analysis of the outflow jet will show that the jet instability is always present, with a relatively small difference in stability as the base flow evolves. This apparent paradox is resolved by noting that the $\lambda_i(t)$ and the stability eigenvalues are different quantities. The $\lambda_i(t)$ are volume-averaged growth rates, while the stability eigenvalues are spatially local. The actual value of $\lambda_i(t)$ will be shown to depend on the magnitude of the perturbation energy in the unstable jet and alignment of the perturbation with the eigenmode of the instability.

4.1. *The chaos-producing mechanism*

Figure 8 shows the contributions $\lambda_i(t)$ to λ_1 and λ_5 over a short time interval for the chaotic attractor at $R/R_c = 11.3$. The long-time average λ_1 is 0.052Ω bits (where bits refers to the logarithm base two) but $\lambda_1(t)$ can be almost two orders of magnitude larger than that, and of either sign. Recall that these are exponential growth rates. Although λ_1 is positive, the perturbations corresponding to λ_1 will decay exponentially for short time intervals. At other times, it will grow at exponential rates much larger than average. This wide variation in short-time growth rate is seen for all Lyapunov exponents computed in this problem. The long-time average λ_5 is

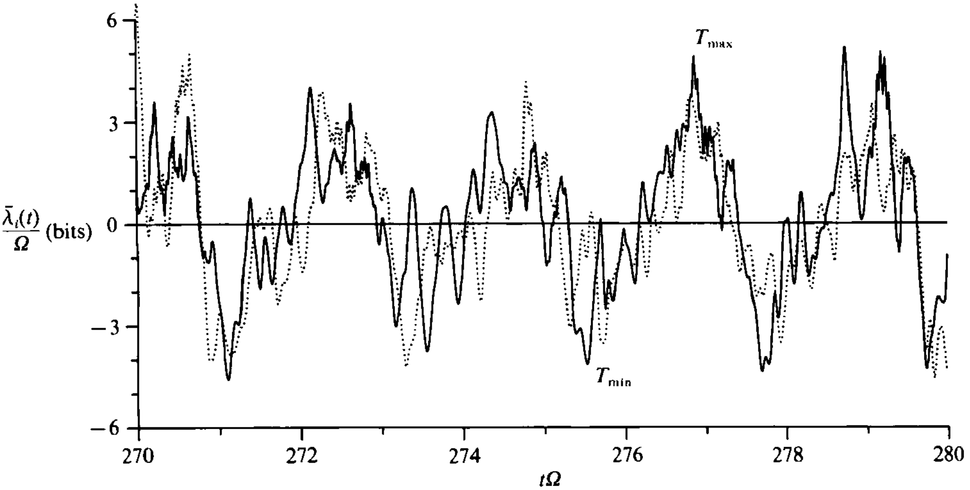

FIGURE 8. Short-time contributions $\lambda_i(t)$ to λ_1 (solid line) and λ_5 (dotted line) for the chaotic flow at $R/R_c = 11.3$. The extrema in $\lambda_1(t)$ that will be studied in detail are labeled $T_{\min}$ and $T_{\max}$. Note that $\lambda_1 = 0.052\Omega$ bits and $\lambda_5 = -0.004\Omega$ bits. The running averages $\bar{\lambda}_i(t)$ obtained from $\lambda_i(t)$ are shown in figure 6.

nearly zero, but $\lambda_5(t)$ can be large and of either sign. The large fluctuations of $\lambda_5(t)$ correspond in time with those of $\lambda_1(t)$. In fact, short-time contributions to the first to fourteenth exponents all show approximately the same large-scale variation (with different detailed fluctuations). This suggests that the same physical process is driving the instantaneous instability of all these perturbations (see below).

In figure 9, the streamwise (azimuthal) velocity of the base flow is illustrated at times $T_{\min}$, corresponding to the large negative peak in $\lambda_1(t)$ in figure 8, and $T_{\max}$, corresponding to the succeeding large positive peak. The analysis presented below will concentrate on these two times. Other times have been examined to ensure that the conclusions drawn from this analysis are valid. The two times $T_{\min}$ and $T_{\max}$ are seen to correspond to the relaxation and flection, respectively, of the underlying waveform. This 'flex' is the interference of the two $m = 4$ azimuthal travelling waves, and its period is the difference between the travelling wave frequencies ω_1 and ω_2.

Examination of the perturbation field corresponding to λ_1 at $T_{\min}$ and $T_{\max}$ provides further evidence about the instability. In figure 10, the azimuthal velocity of the perturbation is shown for the same r-plane as in figure 9. At both times, and all other times examined, the perturbation is localized on the outflow or high-speed boundary jet and along the outer wall (see §4.2), with some weak activity on the inflow or low-speed jet. The structure of the perturbation on the outflow jet is reminiscent of the most unstable eigenfunction for a Kelvin–Helmholtz instability of a plane jet: the structure is dominated by regions of streamwise velocity that alternate in sign along the outflow jet, with opposite sign across the jet. In addition, the wavelength of the Lyapunov perturbation is close to the wavelength of the maximally unstable Kelvin–Helmholtz eigenfunction for a jet of this thickness (see below). In the remainder of this section we will present evidence supporting the assertion that a Kelvin–Helmholtz-type instability drives the transition to chaos in Taylor–Couette flow.

An examination of the base flow reveals the presence of nearly two-dimensional

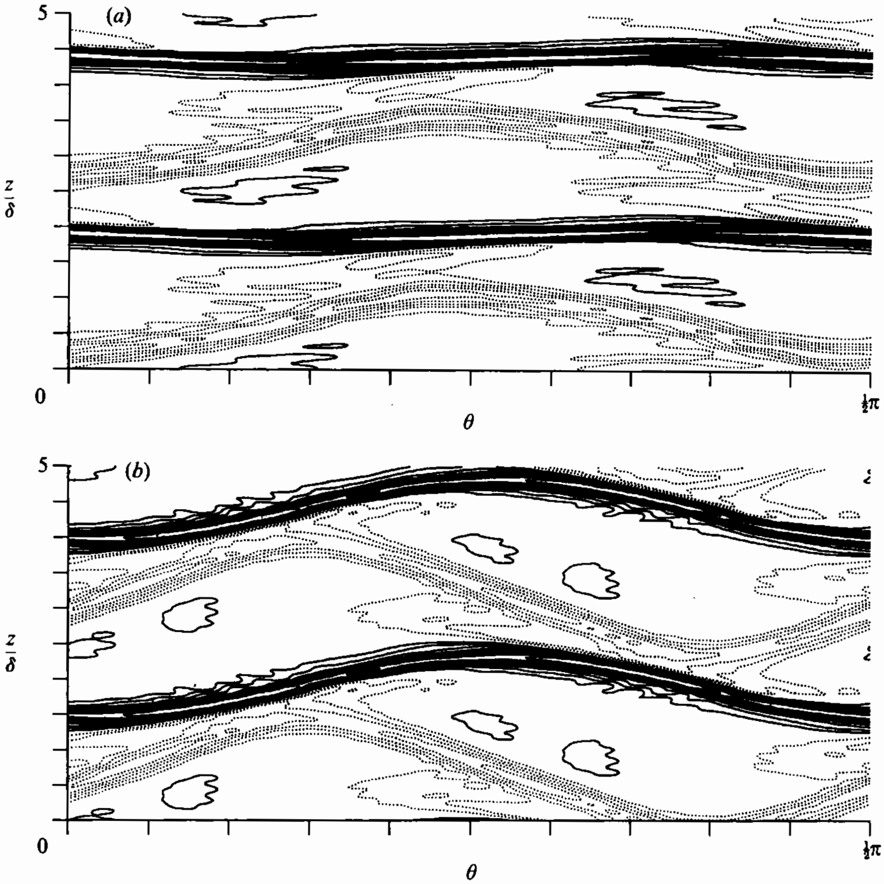

FIGURE 9. Azimuthal velocity contours at $\hat{r} = 0.31$ for the base flow at times (*a*) $T_{\min}$ and (*b*) $T_{\max}$. Solid contours indicate velocities greater than the plane average, dotted contours are less than the plane average. The contour increment is $0.03U$.

shear layers in the outflow boundary jet. Recall that both jets are dominated by the azimuthal component of velocity, and therefore the radial component of vorticity is dominant. The radial vorticity field of the base flow at times $T_{\min}$ and $T_{\max}$ is shown in figures 11 and 12. The r-plane used in figure 11 was chosen because the maximum of ω_r occurs at or near this radius at all times (figures 9 and 10 are in this same plane); also, the θ-planes shown in figure 12 are those for which ω_r is maximum. At both times ω_r is concentrated in thin, relatively flat sheets that are approximately aligned in the radial and azimuthal directions, though they are slightly inclined to the r- and θ-axes at time $T_{\max}$. These sheets correspond to the shear regions of the outflow boundary jet (one on each side of the jet). They extend across most of the gap width and through at least half the domain in θ. The extent of the sheets is many times their thickness; therefore, the radial vorticity is a good approximation of the shear in the jet. At time $T_{\max}$, the shear (vorticity) is stronger on the side of the outflow jet that is close to the inflow jet; for an asymmetric jet such as this, the expected shear instability would be stronger on the high-shear side. This asymmetry favouring the high-shear side can be seen in the perturbation field at time $T_{\max}$ in figure 10(*b*). In contrast, at $T_{\min}$ the outflow is straight, the shear is more nearly equal on the two sides

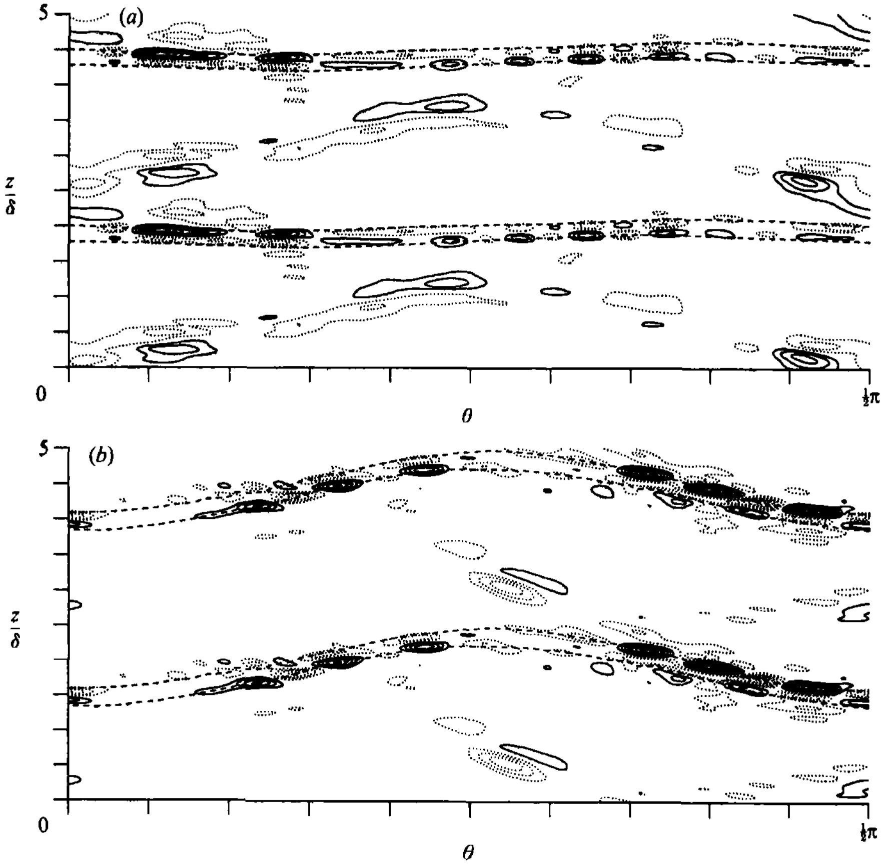

FIGURE 10. Azimuthal velocity contours at $\hat{r} = 0.31$ for the first Lyapunov perturbation field at times (a) $T_{\min}$ and (b) $T_{\max}$. Positive velocity contours are solid; negative velocity contours are dotted. The contour increment is 0.3. The dashed lines indicate the approximate boundaries of the outflow jet.

of the jet, and the perturbation strength at this time in figure 10(*a*) is more nearly symmetric across the jet. At both times, the azimuthal location of the region of maximum shear (near where the jets are closest) is where the perturbation field strength is largest. Also, as will be shown below, the perturbation gains energy only in the shear regions of the jet. All of this evidence reinforces the hypothesis that the exponentially growing perturbations associated with the chaos in this system are due to a shear instability of the outflow boundary jet.

The shear layer is essentially flat and uniform radially. If the Lyapunov perturbation field corresponds to an instability of the shear layer, it should then be largely two-dimensional (corresponding to the most unstable disturbances for a plane free shear layer), with little or no structure in the radial direction. Contours of azimuthal velocity for the perturbation in a plane roughly coincident with the strongest thin shear layer at time $T_{\max}$ are shown in figure 13. The only significant variation of the perturbations in the 'radial' direction (except for very near the wall) is the tilting of the disturbance contours away from vertical. This tilting is not surprising because the jet has a radial velocity component, and the most unstable disturbance in a jet has no variation perpendicular to the jet velocity and the shear

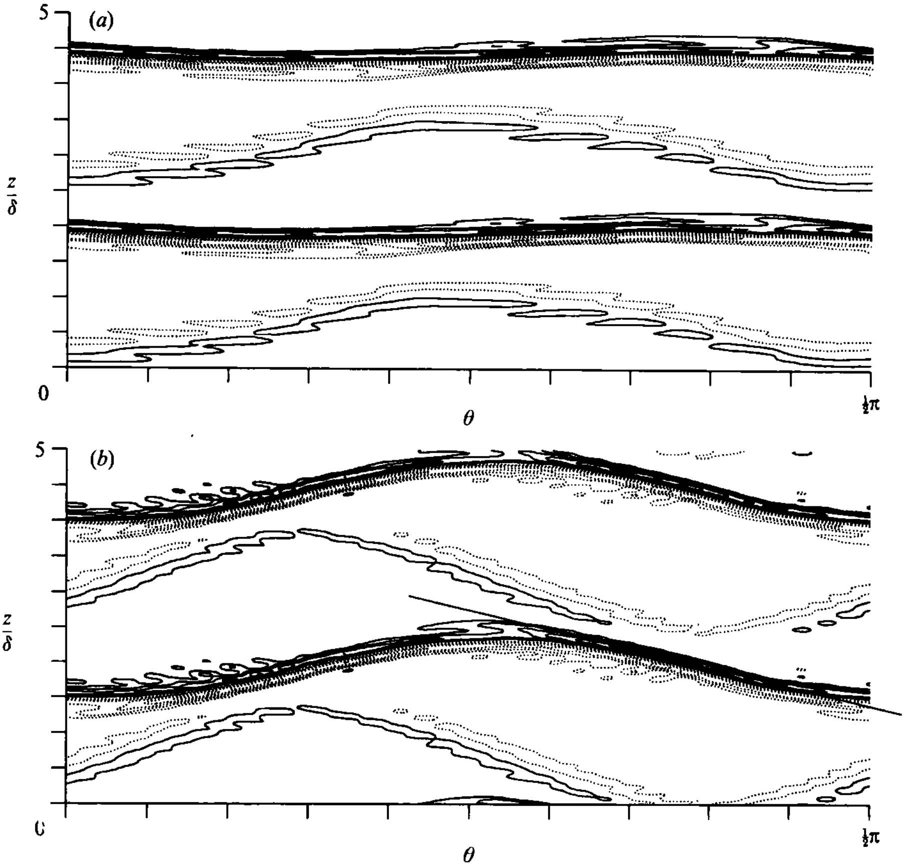

FIGURE 11. Contours of the radial vorticity ω_r of the base flow at $\hat{r} = 0.31$ for times (*a*) $T_{\min}$ and (*b*) $T_{\max}$. Positive vorticity contours are solid; negative vorticity contours are dotted. The contour increment is $0.6U/\delta$. In (*b*), the diagonal line is the intersection of the plane shown in figure 13 with this plane.

direction. The vector superposed on the contours in figure 13 is the in-plane component of the base flow velocity at that point (midgap point). The perturbation is approximately aligned as expected for a jet instability.

The localization of the instability to the outflow jet is also supported by the experiments of Gorman & Swinney (1982). They noticed that as the transition to chaos was approached, 'fuzz' appeared on the outflow jet. As we have shown, the chaotic part of the flow commences as small-spatial-scale disturbances on the outflow jet – this is probably what Gorman & Swinney observed. In another parameter regime, Coughlin (1990) showed that instabilities of the outflow jet also underlie several bifurcations to ordered states.

4.2. *Temporal evolution of the first Lyapunov perturbation*

So far, our examination of the base flow and perturbation fields has provided qualitative evidence linking the chaotic instability and a Kelvin–Helmholtz-type instability of the outflow boundary jet. We now turn to a quantitative analysis to obtain a more complete description of the evolution of the perturbation. In particular, it is essential to understand the difference in the short-time Lyapunov

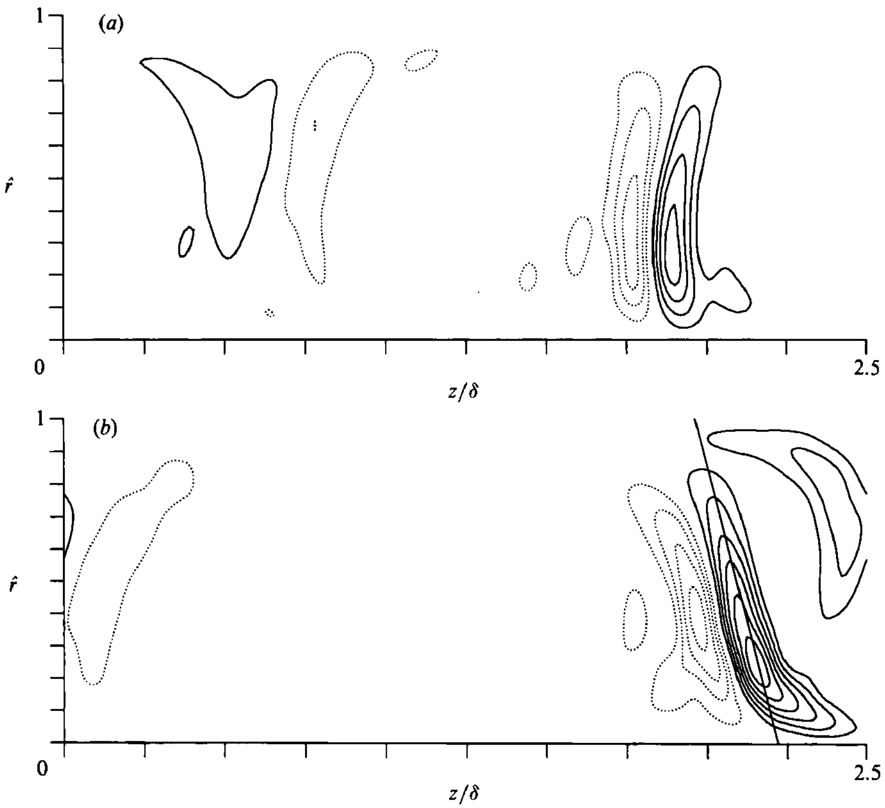

FIGURE 12. Contours of ω_r for the base flow at times (a) $T_{\min}$ and (b) $T_{\max}$, in the (r, z)-plane of maximum ω_r. Negative contours are dotted. The contour increment is $0.6U/\delta$. In (b), the diagonal line is the intersection of the plane shown in figure 13 with this plane.

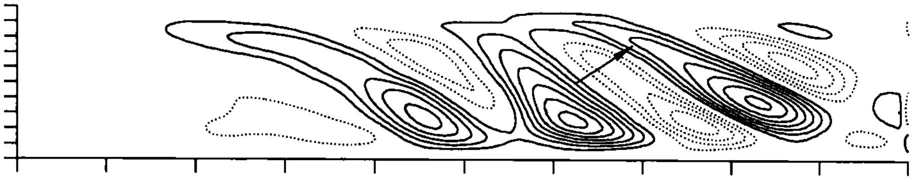

FIGURE 13. Azimuthal velocity contours for the first Lyapunov perturbation field in a plane approximately tangent to the shear layer at the point of maximum ω_r at time $T_{\max}$. Positive velocity contours are solid, negative velocity contours are dotted. The contour increment is 0.4. The arrow is the in-plane component of the base flow velocity vector at that point. This plane intersects the planes shown in figures 11(b) and 12(b) at the lines shown in those figures. The vertical axis extends from the inner wall to the outer wall.

exponent contributions $\lambda_1(t)$ between times $T_{\min}$ and $T_{\max}$. This is of interest because there is no reason from stability considerations to expect such a large variation in the short-time contribution. In the course of this discussion the difference between the short-time Lyapunov exponent analysis and a linear stability analysis about an unsteady flow at an instant of time will be made clear.

To understand the difference in short-time Lyapunov contributions, we first consider the convective derivative of the first perturbation energy density ($\mathrm{D}E/\mathrm{D}t$ where $E = \frac{1}{2}\delta\boldsymbol{u}\cdot\delta\boldsymbol{u}$), using the base-flow convection velocity. The integral in θ and z

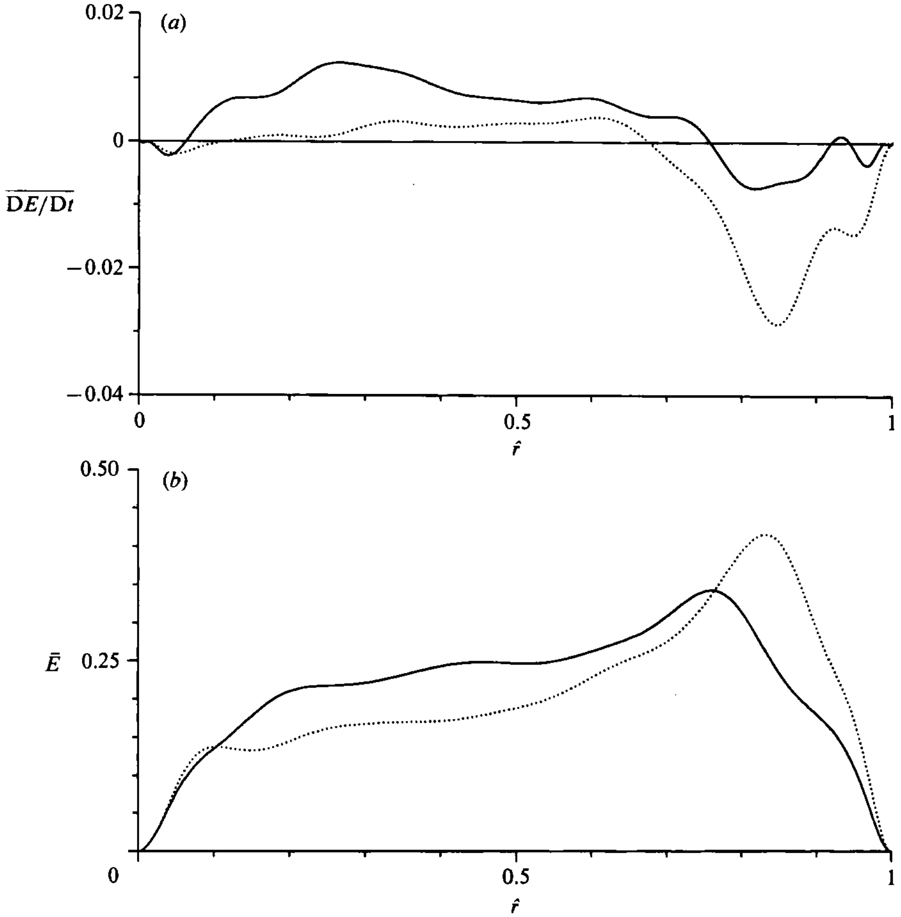

FIGURE 14. (a) The (θ, z)-integral $\overline{DE/Dt}$ of the convective derivative of energy density and (b) the (θ, z)-integral of energy density $(\bar{E})$ as a function of r for the first Lyapunov perturbation field, at times T_{max} (solid curve) and T_{min} (dotted curve).

of the convective derivative of the energy density at T_{min} and T_{max} is shown in figure 14; the integrated perturbation energy densities are also shown. At both times the perturbation energy is growing at the r-locations where the shear layers are present (see figure 12), while the energy is decaying near the outer wall. The growth rate near $\hat{r} = 0.3$ (where ω_r is maximum) is higher for the perturbation at T_{max}, while the decay rate near the outer wall is larger at T_{min}. In addition, the integrated energy in the perturbation, shown in figure 14(b), is relatively higher in the growth region at time T_{max} and relatively higher in the decay region at T_{min}. This analysis shows that the perturbation energy is generated away from the walls in the region where the shear layers are strongest, and dissipates near the outer wall. Recall that the high shear and the perturbation are associated with the radial *outflow* boundary jet. The radial velocity associated with the jet convects the energy in the perturbation towards the outer wall. As it leaves the high-shear regions of the outflow jet and enters the viscously dominated near-wall region, the energy dissipates. Note also that at both times the perturbation energy is maximum near the outer wall, not in the high-shear region. Apparently, perturbation energy collects near the outer wall as it is being dissipated there.

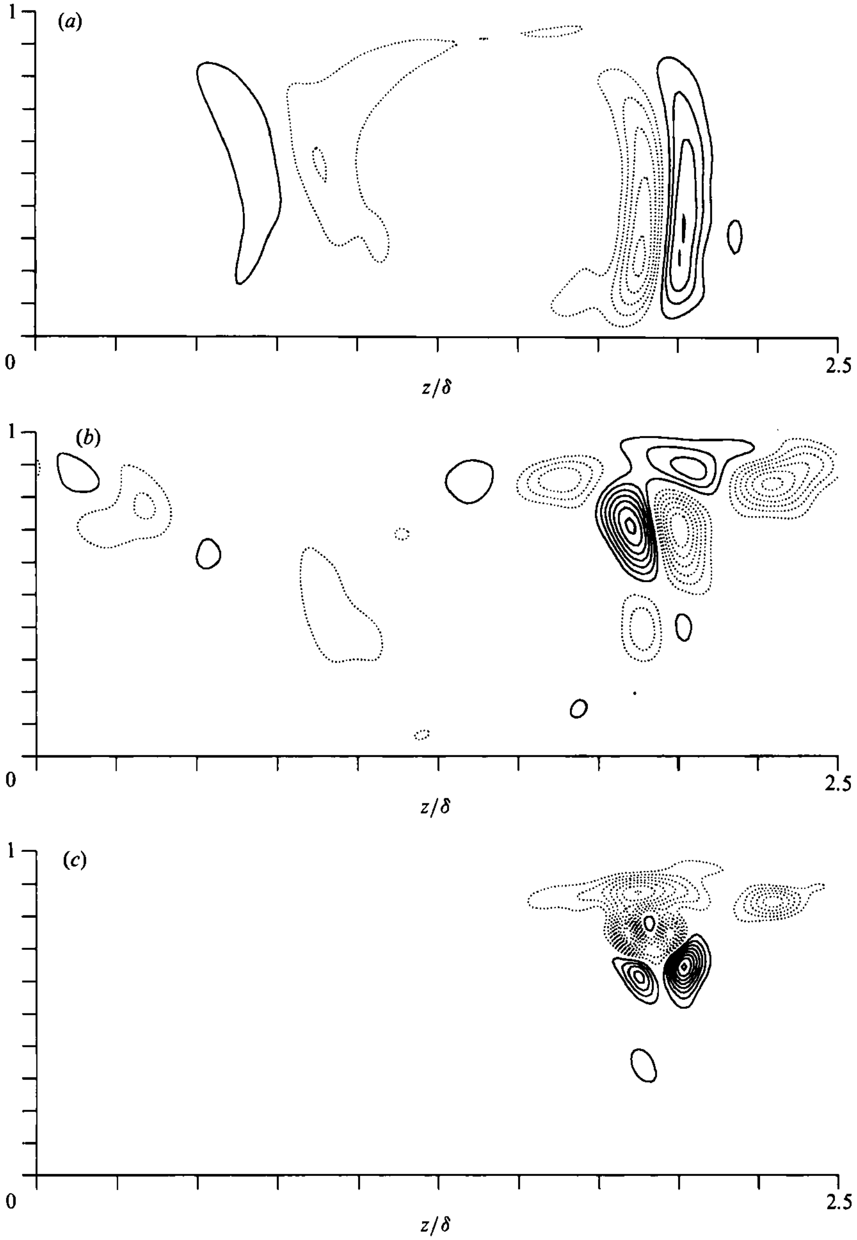

FIGURE 15. Contours of (*a*) the base-flow radial vorticity (*b*), the perturbation azimuthal velocity, and (*c*) the convective derivative of the perturbation energy density at time $T_{\min}$. Dotted contours are negative. Contour increments are (*a*) $0.6U/\delta$, (*b*) 0.3, and (*c*) 0.2.

The integrated energy gives a general description of the perturbation evolution. A more detailed picture of the evolution can be obtained by examining the spatial structure of the convective derivative of the energy density. Contours of the base-flow radial vorticity, the perturbation azimuthal velocity and the convective derivative of the perturbation energy density in the (r, z)-plane passing through the point of maximum energy growth are shown in figures 15 and 16 for times $T_{\min}$ and

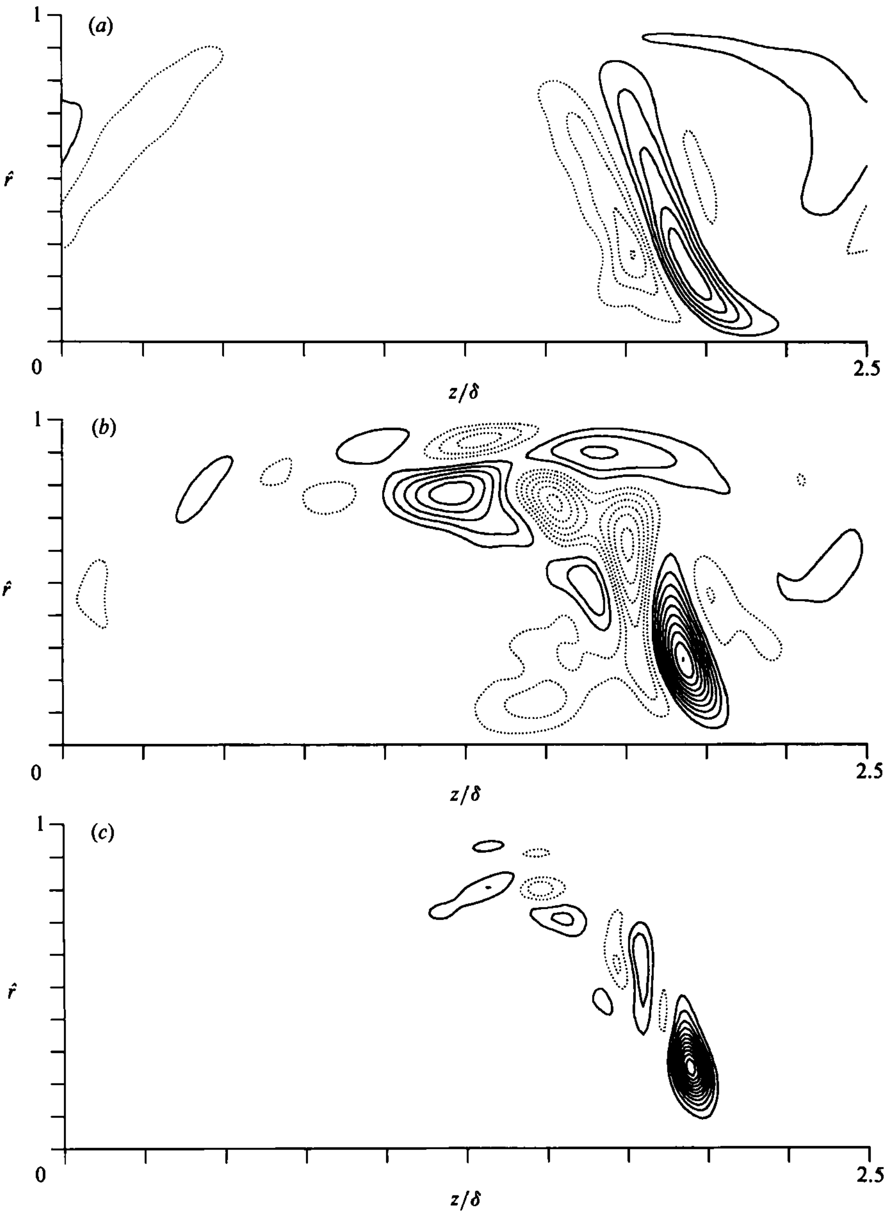

FIGURE 16. Contours of (*a*) the base-flow radial vorticity, (*b*) the perturbation azimuthal velocity, and (*c*) the convective derivative of the perturbation energy density at time $T_{\max}$. Dotted contours are negative. Contour increments are (*a*) $0.6U/\delta$, (*b*) 0.3, and (*c*) 0.4.

$T_{\max}$ respectively. At time $T_{\min}$ this plane also happens to contain the point of maximum decay. These maximum-growth points are not far from the points of maximum radial vorticity (displaced by 0.95δ along the jet) at both times; however, at time $T_{\min}$ the maximum-growth point is also displaced outward in r from the maximum-vorticity point (by 0.18δ). The contour plots show that the high-growth regions are indeed centred on the shear layers. The maximum perturbation energy as indicated by the azimuthal velocity (which dominates) occurs in the shear region at

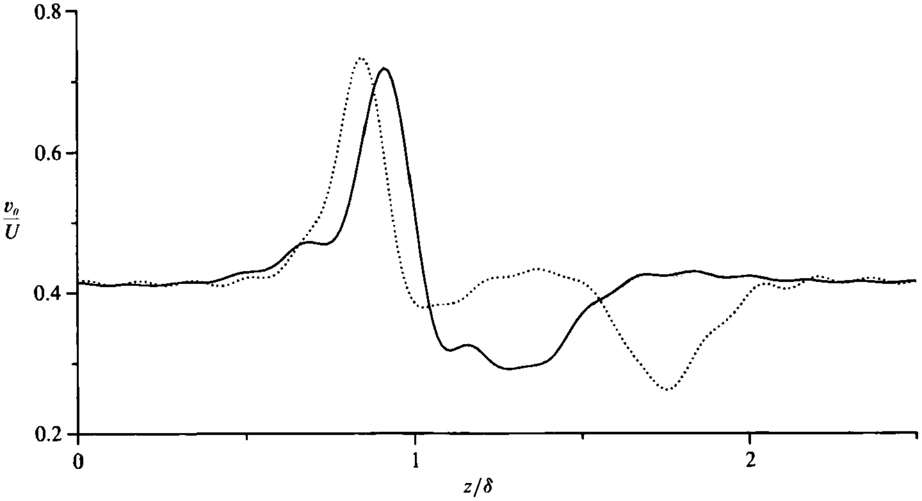

FIGURE 17. Azimuthal velocity profiles as a function of z for the base flow at the (r, θ)-locations of maximum radial vorticity for times $T_{\max}$ (solid curve) and $T_{\min}$ (dotted curve).

time $T_{\max}$, though there is substantial perturbation energy spread out along the outer wall as well. At time $T_{\min}$, most of the energy is spread out along the outer wall. As the energy in the perturbation is carried to the outer wall by the outflow jet, it is also carried along the wall by the Taylor–Couette roll cells. Thus, the near-wall energy and decay is spread out over a large area. In fact the maximum pointwise energy decay at times $T_{\min}$ and $T_{\max}$ differs by only 25% (1.8 and 1.4, respectively), suggesting that the large integrated decay shown in figure 14 is just due to the larger volume over which decay is occurring. In contrast, the maximum pointwise growth at $T_{\min}$ and $T_{\max}$ (1.6 and 4.2, respectively) differ by a factor of 2.5. This is consistent with the observation that away from the outer wall, the energy and growth regions are concentrated in the outflow jet shear layers, so that the difference in integrated growth between times $T_{\min}$ and $T_{\max}$ is due to stronger pointwise growth at time $T_{\max}$. However, the logarithmic convective derivatives of the energy density $(1/E)(\mathrm{D}E/\mathrm{D}t)$ at the maximum-growth points differ by just 32% (25Ω and 33Ω respectively, in bits). This implies that the difference in local growth rates in the shear layer is largely due to the lower perturbation energy density in the shear layer at time $T_{\min}$.

To determine if the differences in stability of the outflow jet between times $T_{\min}$ and $T_{\max}$ could be responsible for the 32% difference in pointwise growth rates, a one-dimensional stability analysis of the jets was performed. Azimuthal velocity profiles were taken from the base flow along an axial line passing through the point of maximum shear at times $T_{\min}$ and $T_{\max}$ and used for Orr–Sommerfeld calculations. The point at which the magnitude of the radial vorticity (and therefore the shear) is maximum at time $T_{\max}$ is located at $\hat{r} = 0.31$, while at $T_{\min}$ it is located at $\hat{r} = 0.25$. The maximum radial vorticity (shear) is 15% larger at $T_{\max}$ than at $T_{\min}$, apparently because the inflow and outflow jets are closer at $T_{\max}$. The azimuthal velocity as a function of z of the $(r, \theta$-locations of the radial vorticity maximum is shown in figure 17 for both times. The inflow and outflow jets are apparent. Also, at $T_{\max}$ the two jet shear layers have merged, resulting in a stronger shear layer with a larger velocity jump. The one-dimensional stability computations show that the jets at both times

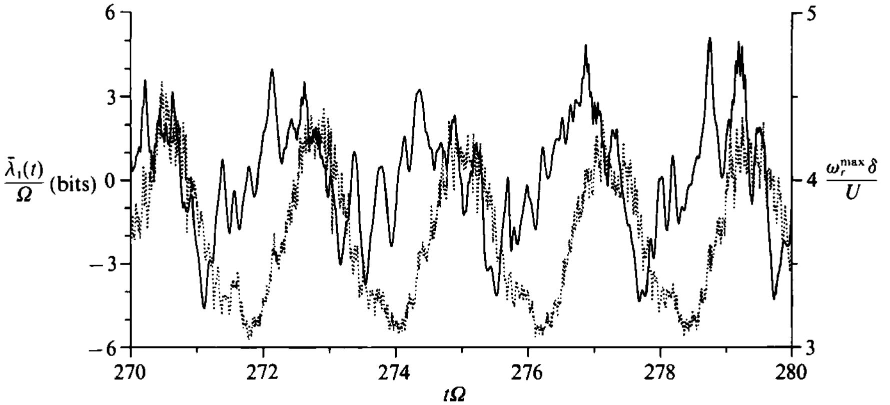

FIGURE 18. The short-time contribution $\bar{\lambda}_1(t)$ (solid curve) and the maximum radial vorticity in the base flow $\omega_r^{\max}$ (dotted curve) for the chaotic flow at $R/R_c = 11.3$.

are unstable with amplification rates of 32.3Ω and 33.3Ω at times $T_{\min}$ and $T_{\max}$ respectively, suggesting that the jet is only slightly more unstable at time $T_{\max}$. While these amplification rates are an order of magnitude larger than the maximum value of $\bar{\lambda}_1(t)$, they are in good agreement with the maximum logarithmic convective derivative of the perturbation energy density reported above. Also, the most unstable wavelength in the one-dimensional computation is 0.77δ, in good agreement with the perturbation wavelength measured from figure 10 after accounting for the inclination of the perturbation (about 0.8δ). This is strong evidence that the growth of the Lyapunov perturbation is dominated by the Kelvin–Helmholtz instability as modelled by the one-dimensional stability analysis.

The one-dimensional stability analysis is crude, ignoring as it does all three-dimensional effects. However, if the essence of the chaos-producing mechanism is the instability of the nearly two-dimensional shear regions, as argued in §4.1, then this analysis should at least capture the relative stability of the jet at the two times. This difference in stability is not sufficient to explain the difference in the observed pointwise growth rates at the two times. The explanation is that the perturbation is more closely aligned with the unstable jet eigenmode at $T_{\max}$ than at $T_{\min}$. Not only is there more perturbation energy in the shear region of the outflow jet at $T_{\max}$, but the perturbation energy that is there is aligned to produce a (32%) larger exponential growth rate than at $T_{\min}$.

Our analysis has shown that the large variations in $\bar{\lambda}_1(t)$ are *not* due to a large variation in the stability of the base flow. Rather, the variations are caused by the distribution and alignment of the energy in the first Lyapunov perturbation relative to the base flow. At any time there are perturbations with instantaneous growth rates as large as $\bar{\lambda}_1(t)$ at $T_{\max}$ or larger. We must therefore ask why the first Lyapunov perturbation does not take such a form at all times. The answer is that a rapidly growing perturbation at some time (e.g. $T_{\min}$) does *not* evolve continuously into a rapidly growing perturbation at some later time (e.g. $T_{\max}$). This is not surprising since the base flow is continuously evolving and in particular the outflow jet is moving as the base flow 'flexes'. If the shear regions of the jet are moving, it would be expected that a perturbation would become 'misaligned' as it evolved. Recall that the first Lyapunov perturbation is the one that has undergone and will undergo the fastest long-term exponential growth. It would therefore be expected that $\bar{\lambda}_1(t)$

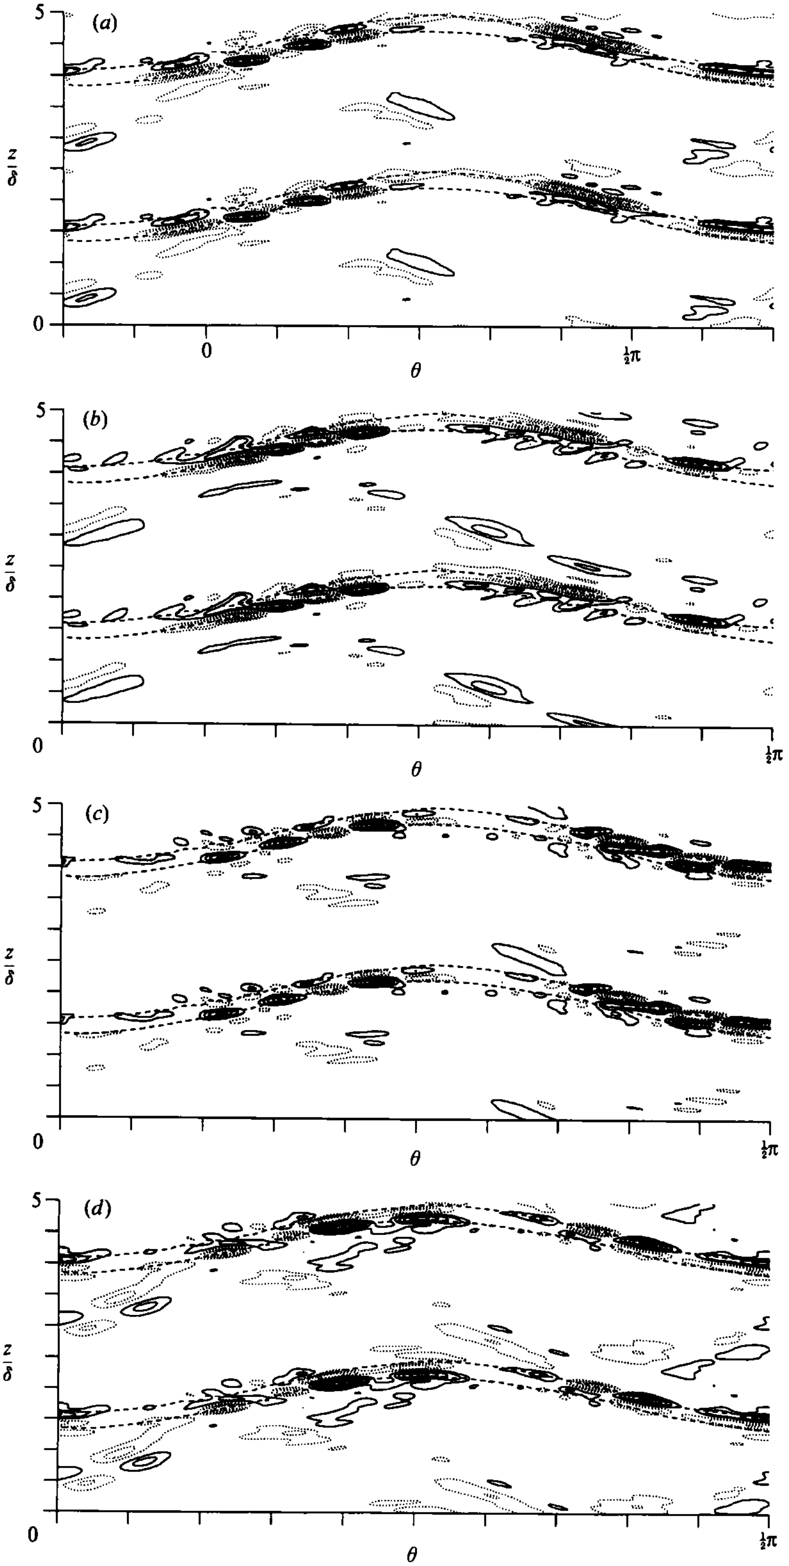

FIGURE 19. For caption see facing page.

would be maximum near the time at which the shear in the jet shear layers (ω_r) is maximum since, other things being equal, this would produce the largest long-term growth. This is indeed the case, as can be seen in figure 18, where the $\lambda_1(t)$ is plotted alone with the maximum in ω_r. The fact that the maximum short-time contribution leads the peak in ω_r by approximately half a cylinder period is apparently caused by dissipation near the outer cylinder increasing faster than the local growth rates in the shear regions. At time $277.1/\Omega$, when ω_r is maximum, the maximum pointwise growth rate is 30% greater than at $T_{\max}$, but the decay rate near the outer wall is also larger.

We now have a clear picture of the evolution of the first perturbation from the maximum-growth time to the maximum-decay time and back to maximum growth. At $T_{\max}$, perturbation energy is being generated by the Kelvin–Helmholtz instability of the outflow boundary jet. The perturbation energy is convected out to the wall until at maximum decay the energy is concentrated at the outer wall, where the decay is occurring. By the time the maximum growth time is reached again, much of the perturbation energy at the outer wall has decayed, and new perturbation energy is being generated in the shear layer. The local decay near the outer wall is not significantly different at $T_{\max}$ and $T_{\min}$, but there is a large difference in the local growth rates in the shear layer.

4.3. *Other Lyapunov perturbations*

To this point we have examined only the perturbation associated with λ_1; however, from table 1 there are from three to five positive Lyapunov exponents, and we must examine the perturbation associated with each of them. The streamwise velocities of the second to fifth perturbations are shown at time $T_{\max}$ in figure 19. Note that they are similar to the first perturbation in that activity is concentrated along the outflow boundary jet and is of similar detailed structure. The major differences are the lengthscale along the outflow jet and the distribution of disturbances along the jet. Also, the higher (slower growing) perturbations contain more energy on the inflow jet. The entire space of long-term unstable disturbances is thus associated with the Kelvin–Helmholtz instability of the outflow (and to a lesser extent the inflow) boundary jet. This being the case, any small-scale structure visible in the chaotic base flow should be concentrated at the outflow boundary jet. This has been observed experimentally (see figure 1 and Gorman & Swinney 1982) and is apparent at $R/R_c = 12.0$ (figure 20) in the current computations. It is difficult to see any of the small-scale structure of the flow in the contour plots at $R/R_c = 11.3$ because of the low magnitude of the fluctuations.

4.4. *Lyapunov analysis of quasi-periodic flow*

The short-time Lyapunov exponent contributions for the first non-zero Lyapunov exponent (λ_3) of the quasi-periodic attractor at $R/R_c = 9.8$, are shown in figure 21. Recall that λ_3 is negative, corresponding to a perturbation that is exponentially damped in the long-term. The contributions are (nearly) periodic in time, since they reflect the average variation in the perturbations over the entire physical domain and are thus effectively in the rotating reference frame. The streamwise velocity of the base flow and the perturbation corresponding to λ_3 at a maximum in the short-

FIGURE 19. Azimuthal velocity contours at $\hat{r} = 0.31\delta$ and time $T_{\max}$ for (*a*) the second, (*b*) third, (*c*) fourth and (*d*) fifth Lyapunov perturbation fields for the chaotic attractor at $R/R_c = 11.3$. Dotted contours are negative. Contour increments are $0.3U$. The dashed lines indicate the approximate boundaries of the outflow jet.

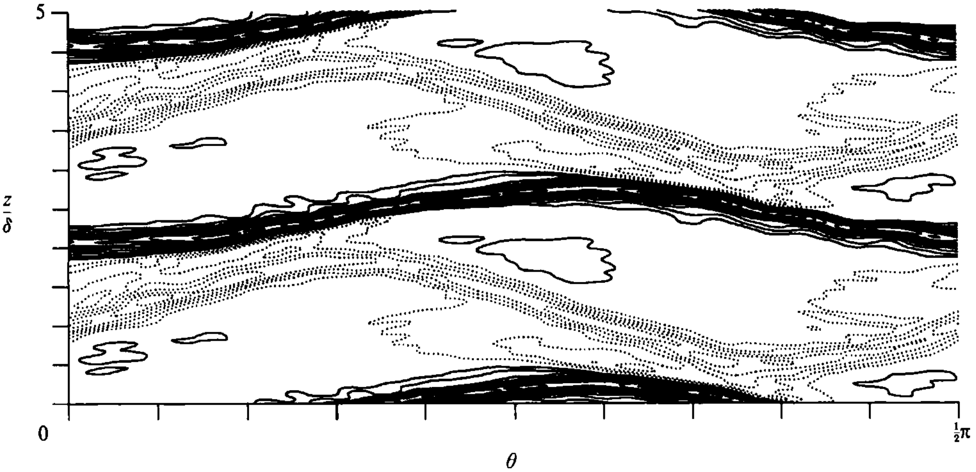

FIGURE 20. Azimuthal velocity contours at $\hat{r} = 0.31$ for the chaotic flow at $R/R_c = 12.0$. High-velocity contours (greater than the plane average of $0.42U$) are solid, low-velocity contours are dotted. The contour increment is $0.03U$.

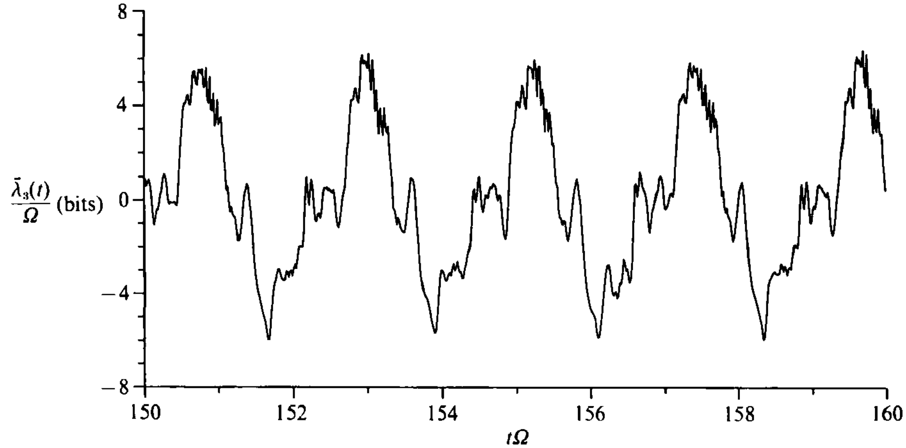

FIGURE 21. The short-time contributions $\lambda_3(\tau)$ for the quasi-periodic attractor in the simulation at $R/R_c = 9.8$.

time contributions are shown in figure 22. Again, the localization of the perturbations to the outflow boundary jet and the detailed structure of the perturbations is similar to that for λ_1 in the chaotic case. The Kelvin–Helmholtz instability of the outflow boundary jet is active in this case as well. The major difference between this quasi-periodic case and the chaotic case is that in the quasi-periodic flow, the intervals of exponential decay outweigh the intervals of exponential growth, leading to an overall decay. As the Reynolds number is increased toward the transition to chaos, the balance between the stable times and the unstable times on the attractor shifts (gradually) towards the unstable, since the Reynolds number of the inflow and outflow jets is increasing. The expected gradual dependence of the Lyapunov exponents on the Reynolds number suggests that the attractor dimension (as measured by the Kaplan–Yorke formula) will increase continuously from two at the

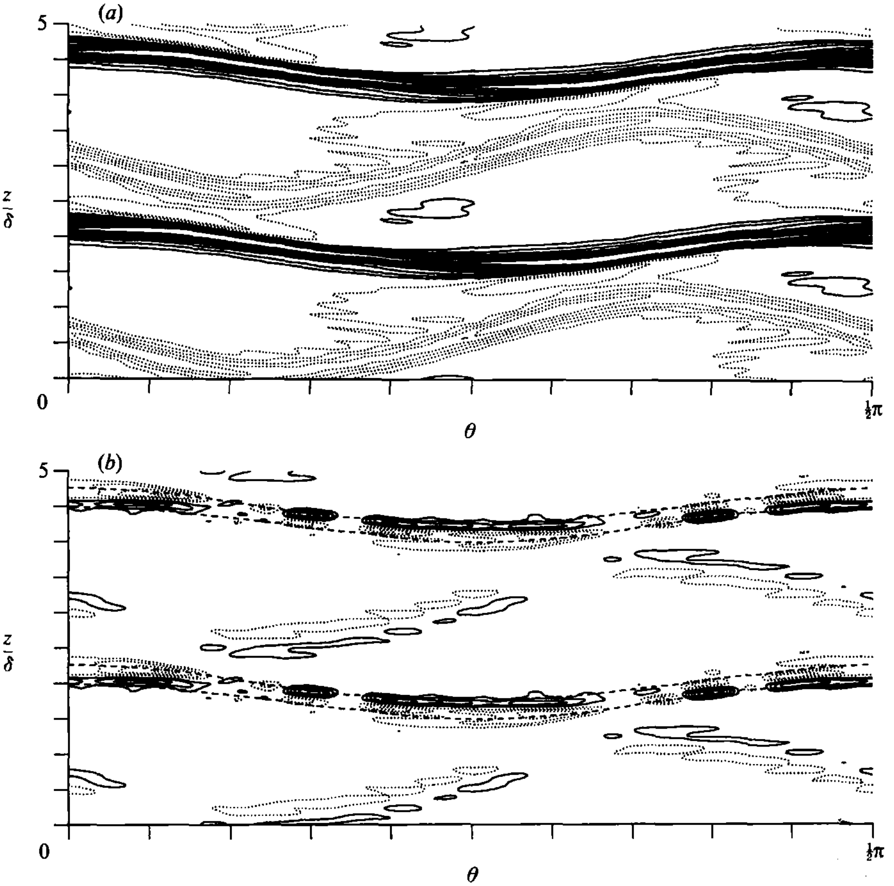

FIGURE 22. The aximuthal velocity contours for (*a*) the base flow and (*b*) the third Lyapunov perturbation at $\hat{r} = 0.31$ for a time corresponding to maximum $\lambda_3(t)$. High-velocity contours in (*a*) are solid, low-velocity contours are dotted. Positive velocity contours in (*b*) are solid, negative velocity contours are dotted. The contour increment is (*a*) $0.03U$, (*b*) 0.3.

transition, unless the first exponent has a multiplicity higher than one. It is difficult to rule out a multiplicity of two for the first exponent, since in figure 5 the first two running averages have not separated. However, there is no reason to expect an exact multiplicity.

5. Conclusions

The results presented here may have broad relevance for the analysis of turbulence and transition. We will first summarize the results for the specific case of Taylor–Couette flow, and then discuss the more general implications of our technique.

5.1. *Chaos in Taylor–Couette flow*

A transition from (4, 4) quasi-periodic Taylor vortices to chaotic Taylor–Couette flow was successfully simulated. It was found that there were quantitative differences in transition Reynolds number and attractor dimension between the current computations and the experiments of Brandstater and co-workers (Brandstater *et al.*

1983; Brandstater & Swinney 1987). The most likely cause of these differences is the relatively high noise level present in the experimental measurements compared to the current computations, and the low-pass filtering that was used in the experiments to eliminate noise. However, the current computations are in qualitative agreement with the experiments. In particular the results presented here support the conclusion that the transition to chaos is gradual; that is, the dimension of the attractor appears to increase continuously from two as the Reynolds number is increased. It was also noted that this transition to chaos occurred without breaking the four-fold symmetry of the quasi-periodic flow.

The results of this study provide strong evidence that the chaos-producing mechanism in this flow is a Kelvin–Helmholtz instability of the outflow jet between the counter-rotating Taylor vortices. The Lyapunov exponent perturbation fields are seen to be concentrated on the outflow jet. Perturbation energy is created in the high-shear regions of the jet toward the inner cylinder and is convected outward. At the outer wall, the perturbations spread out and are dissipated. The instantaneous growth rates of the integrated perturbation energy (the short-time Lyapunov exponent contributions) are marked by a large-scale variation which is two orders of magnitude larger than the long-time-average Lyapunov exponents. This variation is associated with the 'flexing' of the modulated wavy Taylor vortices which are still present in the chaotic flow. Maximum growth occurs near the time of maximum flex, which is also the time of maximum shear in the outflow boundary jet. Eigenvalues from a one-dimensional stability analysis of the outflow jet are in good agreement with the maximum pointwise exponential growth of the perturbation energy in the jet. Both of these local measures of the jet instability show only minor variations in time. The contrast between the variation of the volume-integrated growth rates and the lack of variation in the local stability of the jet is explained by the distribution of perturbation energy at different times. At some times the perturbation energy is concentrated in the outflow jet, producing large integrated growth rates while at other times little energy is in the jet, allowing dissipation near the outer wall to dominate. This is a demonstration of the difference between short-time Lyapunov exponent analysis and instantaneous linear stability analysis of unsteady flows.

It was also noted that the same perturbation characteristics were present in the quasi-periodic flow (and probably at even lower Reynolds numbers). The major difference between the chaotic and quasi-periodic flows is that in the chaotic case the intervals during which perturbations grow slightly outweigh the intervals during which the perturbations decay, leading to long-term growth rather than long-term decay.

5.2. *Lyapunov exponent analysis*

The simulations described here, coupled with a new way of using Lyapunov exponent calculations, provided insight into the chaos-producing mechanism in Taylor–Couette flow. These new results point out the utility of Lyapunov exponent calculations for understanding complex physical phenomena. A large-scale simulation can provide an enormously detailed account of the time evolution of a fluid flow. The difficulty in analysing these data, especially in the case of turbulent flows, is knowing which features of the flow are of dynamical importance. Even at the transition to weak turbulence exhibited by Taylor–Couette flow, the available experimental and numerical data offer a daunting challenge. The short-time contributions to the Lyapunov exponents provided information as to what times were important to analyse. In Taylor–Couette flow, the perturbations corresponding to the Lyapunov exponents were spatially localized, thus providing an indication of which flow features were important. However, it is likely that in some systems the perturbations will not be localized in space. In such cases they may be localized in some transformed space (e.g. Fourier space). It should also be noted that short-time Lyapunov exponent analysis can be applied to any closed flow, but that more work is required to develop similar techniques applicable to open flows.

As we have seen the short-time Lyapunov exponent analysis yields more information than the exponents and dimension usually obtained from Lyapunov exponent calculations. This approach is made even more attractive by its avoidance of the principal difficulty in Lyapunov exponent calculations, namely the slow convergence of the long-time-average exponents. The promise of short-time Lyapunov exponent analysis is its ability to indicatc instances when and regions where interesting phenomena occur. For example, such an analysis could possibly indicate what flow structures are responsible for the randomness of a turbulent flow.

Short-time Lyapunov exponent analysis was successful in the weakly chaotic Taylor–Couette flow. In a more complicated flow, the analysis may be more difficult. In particular, a more chaotic flow may exhibit a large number of instability mechanisms. In that case a large number of Lyapunov perturbations may be needed to capture all of the mechanisms. This could prove to be too expensive in terms of computer resources to be practical. Another possible difficulty arises when the instability mechanism is not localized in space or time. In this case the perturbation fields will not be localized either, and it may be difficult to extract the mechanism via this analysis. In both of these cases, however, valuable information about the flow can be obtained. In the first case the most unstable mechanisms will be captured, while in the case of the non-localized mechanism this information itself is valuable. In the opinion of the authors, the short-time Lyapunov exponent analysis is potentially of such value that at least one Lyapunov perturbation should be computed as a matter of course in simulations of systems that are expected to be chaotic. Computing a single Lyapunov perturbation requires only a factor of two increase in computational effort, and can greatly increase the value of the simulation.

We wish to thank Drs L. Keefe and K. Coughlin, and Professor P. Marcus for helpful discussions of this work. We also thank Drs J. Kim, K. Sharif, and T. Pulliam and Professor H. Swinney for comments on a draft of this paper. We are especially grateful to Professor Swinney for providing the photographs in figure 1.

REFERENCES

Benettin, G., Galgani, L., Giorgilli, A. & Strelcyn, J.-M. 1980 Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. *Meccanica* **15**, 9–20.

Brandstater, A. 1984 Low dimensional chaos in a hydrodynamic system. Dissertation. Christian-Albrechts-Universität, Kiel.

Brandstater, A., Swift, J., Swinney, H. L., Wolf, A., Farmer, J. D., Jen, E. & Crutchfield, J. P. 1983 Low-dimensional chaos in a hydrodynamic system. *Phys. Rev. Lett.* **51**, 1442–1445.

Brandstater, A. & Swinney, H. L. 1987 Strange attractors in weakly turbulent Couette–Taylor flow. *Phys. Rev.* A **35**, 2207–2220.

Constantin, P., Foias, C. & Temam, R. 1988 On the dimension of the attractors in two-dimensional turbulence. *Physica* **30D**, 284–296.

Coughlin, K. 1990 Quasiperiodic Taylor–Couette flow. Dissertation, Harvard University.

DiPrima, R. C. & Swinney, H. L. 1981 Instabilities and transition in flow between concentric rotating cylinders. In *Hydrodynamic Instabilities and the Transition to Turbulence* (ed. H. L. Swinney & J. P. Gollub), pp. 139–180. Springer.

ECKMANN, J.-P. & RUELLE, D. 1985 Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.* **57**, 617–655.

FENSTERMACHER, P. R., SWINNEY, H. L. & GOLLUB, J. P. 1979 Dynamical instabilities and the transition to chaotic Taylor vortex flow. *J. Fluid Mech.* **94**, 103–128.

FRASER, A. M. & SWINNEY, H. L. 1986 Independent coordinates for strange attractors from mutual information. *Phys. Rev.* A**33**, 1134–1140.

FREDERICKSON, P., KAPLAN, J. L., YORKE, E. D. & YORKE, J. A. 1983 The Lyapunov dimension of strange attractors. *J. Diff Equat.* **49**, 185–196.

GOLDHIRSCH, I., SULEM, P.-L. & ORSZAG, S. A. 1987 Stability and Lyapunov stability of dynamical systems: a differential approach and a numerical method. *Physica* **27D**, 311–337.

GOLLUB, J. P. & SWINNEY, H. L. 1975 Onset of turbulence in a rotating fluid. *Phys. Rev. Lett.* **41**, 927–930.

GORMAN, M. & SWINNEY, H. L. 1982 Spatial and temporal characteristics of modulated waves in the circular Couette system. *J. Fluid Mech.* **117**, 123–142.

GREENE, J. M. & KIM, J.-S. 1987 The calculation of Lyapunov spectra. *Physica.* **24D**, 213–225.

KEEFE, L., MOIN, P. & KIM, J. 1990 Application of chaos theory to shear turbulence. In *The Ubiquity of Chaos* (ed. S. Krasner), pp. 56–63. American Association for the Advancement of Science.

KING, G. P., LI, Y., LEE, W., SWINNEY, H. L. & MARCUS, P. 1984 Wave speeds in wavy Taylor-vortex flow. *J. Fluid Mech.* **141**, 365–390.

LEDRAPPIER, F. 1981 Some relations between dimension and Lyapunov exponents. *Commun. Math. Phys.* **81**, 229–238.

MARCUS, P. 1984*a* Simulation of Taylor–Couette flow. Part 1. Numerical methods and comparison with experiment. *J. Fluid Mech.* **146**, 45–64.

MARCUS, P. 1984*b* Simulation of Taylor–Couette Flow. Part 2. Numerical results for wavy-vortex flow with one travelling wave. *J. Fluid Mech.* **146**, 65–113.

MOSER, R. D. & MOIN, P. 1984 Direct numerical simulation of curved turbulent channel flow. *NASA TM* 85974.

MOSER, R. D. & MOIN, P. 1987 The effects of curvature in well-bounded turbulent flows. *J. Fluid Mech.* **175**, 479–510.

MOSER, R. D., MOIN, P. & LEONARD, A. 1983 A spectral numerical method for the Navier–Stokes equations with applications to Taylor–Couette flow. *J. Comput. Phys.* **52**, 524–544.

NEWHOUSE, S., RUELLE, D. & TAKENS, F. 1978 Occurrence of strange Axiom A attractors near quasi-periodic flows on $T^m, m \geqslant 3$. *Commun. Math. Phys.* **64**, 35–40.

OSELEDEC, V. I. 1968 A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems. *Trans. Moscow Math. Soc.* **19**, 197.

PACKARD, N. H., CRUTCHFIELD, J. P., FARMER, J. D. & SHAW, R. S. 1980 Geometry from a time series. *Phys. Rev. Lett.* **45**, 712–716.

RUELLE, D. & TAKENS, F. 1971 On the nature of turbulence. *Commun. Math. Phys.* **20**, 167–192.

TAKENS, F. 1981 Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980*, Lecture Notes in Mathematics, vol. 898, pp. 366–381. Springer.

TAYLOR, G. I. 1923 Stability of a viscous liquid contained between two rotating cylinders. *Phil. Trans. R. Soc. Lond.* A**223**, 289–343.

WOLF, A., SWIFT, J. B., SWINNEY, H. L. & VASTANO, J. A. 1985 Determining Lyapunov exponents from a time series. *Physica* **16D**, 285–317.