J. Fluid Mech. (1990), *vol.* 213, *pp.* 227–249
Printed in Great Britain

A comparative computational and experimental study of chaotic mixing of viscous fluids

By P. D. SWANSON AND J. M. OTTINO
Departments of Chemical Engineering and Polymer Science and Engineering, University of Massachusetts, Amherst, MA 01003, USA

(Received 14 March 1989)

The objective of this work is to develop techniques to predict the results of experiments involving chaotic dispersion of passive tracers in two-dimensional low Reynolds number flows. We present the design of a flow apparatus which allows the unobstructed observation of the entire flow region. Whenever possible we compare the experimental results with those of computations. Conventional tracking of the boundaries of the tracer is inefficient and works well only for low stretches (order 10^2 at most). However, most mixing experiments involve extremely large perturbations from steady flow since this is where the best mixing occurs. The best prediction of widespread mixing and large stretching (order 10^4–10^5) is provided by lineal stretching plots; surprisingly the technique also works for relatively low numbers of periods (as low as 2 or 3). The second best prediction is provided by a combination of low-order unstable manifolds – which indicate where the tracer goes, especially for short times – and the eigendirections of low-order hyperbolic periodic points – which indicate the alignment of striations in the flow. On the other hand, Poincaré sections provide only a gross picture of the spreading: they can be used primarily to detect what regions are inaccessible to the dye. Comparison of computations and experiments invariably reveals that bifurcations within islands have little impact in the mixing process.

1. Introduction

An insight into the mixing mechanisms of viscous fluids can have major consequences in fields as diverse as earth sciences, physiology, chemical engineering, and material sciences. The problem, however, is a complex one and its understanding requires a symbiotic use of analytical, computational, and experimental techniques. There is reason to believe that a unified theory is emerging and mixing in slow flows is beginning to be understood. However, until now there has not been any thorough combined experimental–computational study of chaotic mixing in slow flows.

We are of the opinion that in order to make further progress it is necessary to thoroughly characterize/understand the mixing in a few prototypical systems. In this study we focus on the mixing of similar fluids in the flow field of two non-concentric cylinders (a journal bearing) using an array of analytical, computational and experimental tools. The journal bearing is an ideal system for such studies: it is one of the few flows which is experimentally realizable and which has an analytic solution for the stream function. To date there is only one combined experimental–computational study of chaos in the journal bearing flow (Chaiken *et al.* 1986). Other, strictly computational results have been reported by Chaiken *et al.*

(1987) and Aref & Balachandar (1986). The purpose of these works was to show that it is possible to have chaos in a physically realizable Stokes flow. However, little emphasis was placed on the relationship and possible use of dynamical systems concepts to the study of mixing.

In this study, we reduce background information to a minimum, assuming that the reader is somewhat familiar with the fundamentals of mixing and chaos (for background and applications see Ottino 1989). The objective is to present a rather comprehensive study of possible tools of analysis using a particular flow geometry to demonstrate their application. In fact, in order to emphasize the relationship between various methods of analysis and the experiments we shall not explore a large region of the operating space. Instead we shall concentrate on a few operating conditions. Whenever possible we shall compare the experimental results with those of computations. (Figures, 4, 9, 10, 11 and plates 1 and 3 correspond to the *same* experiments.)

Undoubtedly traditional dynamical systems tools are of relevance to mixing in chaotic flows. However, their importance can be overestimated. There are two significant differences between standard studies of dynamical systems and mixing. The first difference is that traditional studies focus on long-time and asymptotic behaviour (for example, attractors in dissipative systems or Liapunov exponents in conservative or dissipative systems); but in mixing we are primarily interested in short-time behaviour, since we want rapid mixing. The second difference is that in most physical applications chaos is something to be avoided, so studies are done on small perturbations from regular behaviour. However, in mixing we want as much chaos as possible so the perturbations must be large; this restricts the usefulness of analytical tools based on small perturbations from integrability. Finally, the global behaviour of the system is much more important in mixing that in ordinary studies of dynamical systems. It will be apparent from our results that the mixing problem cannot be solved in a complete form within the framework of current analytical tools and that new tools will have to be developed.

2. Journal bearing system

2.1 *Geometry*

The journal bearing or eccentric cylinder geometry is shown in figure 1 (*a*). The flow region is contained in the annulus between two infinitely long cylinders whose central axes are parallel but not coincident, i.e. the cylinders are eccentric. The geometry of the system is completely determined by two dimensionless parameters, the ratio of the radius of the inner cyclinder to the radius of the outer cylinder, $r = r_{\text{in}}/r_{\text{out}}$, and the distance between the centres of the two cylinders, e, normalized by the radius of the outer cylinder $\epsilon = e/r_{\text{out}}$. Note that $0 < r+\epsilon < 1$.

If the fluid between the cylinders is incompressible the problem can be solved through the use of the stream function. If we further assume that the flow is steady and inertial effects are negligible then the stream function can be found by solving the biharmonic equation, $\nabla^4 \Psi = 0$, using the appropriate boundary conditions. Analytic expressions for the stream function can be found in a variety of sources (Jeffery 1922; Duffing 1924; Müller 1942 (*a*, *b*); Wannier 1950; Ballal & Rivlin 1976). We choose the solution presented by Wannier (1950) because it leads to faster numerical simulations. Also, it is given in rectangular coordinates, in contrast to the other solutions which are given in bipolar coordinates. Rectangular coordinates simplify several computational aspects of the problem such as the graphics.

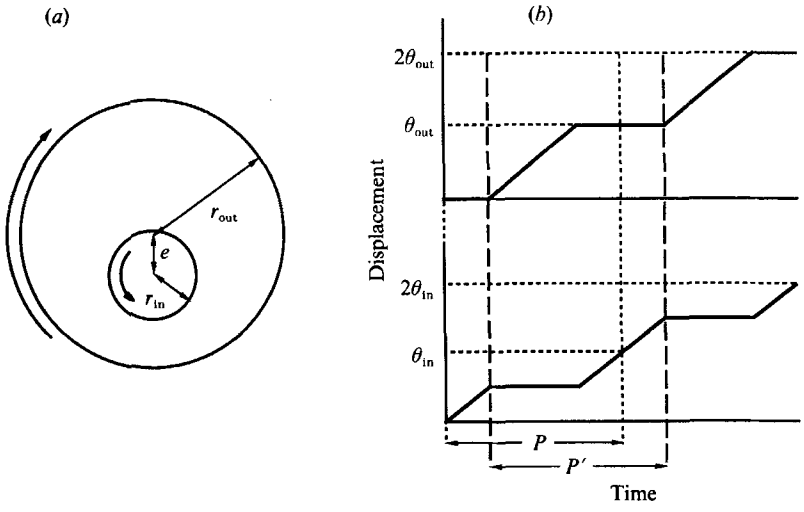

FIGURE 1. Journal bearing geometry. (*a*) The basic geometry is determined by two parameters, the ratio of the radii of the two cylinders, $r = r_{in}/r_{out}$ and the ratio of the eccentricity to the outer cylinder $\epsilon = e/r_{out}$. (*b*) The discontinuous velocity protocol. When a cylinder is in motion, it moves with a constant velocity. This is not a requirement however. If just one cylinder at a time moves the specific velocity profile of the cylinder is unimportant; only the displacement effects the results. P' shows the most obvious method of creating the discontinuous velocity protocol. First the outer cylinder rotates its full displacement followed by the inner cylinder moving for its full displacement. This leads to a non-symmetric mapping. If instead we define the period as P, i.e. move the inner cylinder for half of its total displacement first, then move the outer cylinder, and finally move the inner cylinder for the other half of its displacement, the resulting mapping is symmetric.

Since the problem is linear, the stream function can be written as a sum of terms due to the motion of the inner and outer cylinders:

$$\Psi(x, y) = \Psi_{in}(x, y)\, v_{in} + \Psi_{out}(x, y)\, v_{out},$$

where v_{in} and v_{out} are the angular speeds of the cylinders.

Alternatively, if both cylinders are in motion, which is the base case in this research†, one of the velocities can be divided out, recasting the equations in this form:

$$\frac{\Psi(x, y)}{v_{out}} = \Psi_{in}(x, y)\Omega + \Psi_{out}(x, y),$$

where $\Omega = v_{in}/v_{out}$. The velocity field is derived from the stream function in the usual way:

$$\frac{dx}{dt} = v_x = -\frac{\partial \Psi}{\partial y}; \quad \frac{dy}{dt} = v_y = \frac{\partial \Psi}{\partial x}.$$

Since there is an analytic expression for the velocity field, the equations of motion are a pair of coupled ordinary differential equations. The mixing in the flow is determined by the integrated velocity field (the so-called motion). Dividing these equations by the velocity of the outer cylinder and setting $d\theta = v_{out}dt$ the equations of motion become

$$\frac{dx}{d\theta} = -\frac{\partial(\Psi/v_{out})}{\partial y}; \quad \frac{dy}{d\theta} = \frac{\partial(\Psi/v_{out})}{\partial x}.$$

† In order to generate chaos in a flow the streamlines, not just the velocity, must be time dependent. A time-dependent flow with just one cylinder moving would not result in chaotic motion.

Physically, θ is the angular displacement of the outer cylinder and it is clear from this last equation that the position of a fluid element is dependent on the *displacement* θ of the cylinders and not the actual velocity. This seemingly trivial point will have important consequences when we define our method for parameterizing the perturbation of the velocity field in the next section.

2.2. *Boundary motion*

A two-dimensional Hamiltonian system cannot be chaotic. Therefore, we must add another dimension to the system if there is to be any possibility of chaos. The geometry fixes the number of spatial dimensions, so in order to make the system chaotic it must be made time dependent. Strictly speaking when the system is made time dependent there is no longer an explicit solution for the velocity field. However, we make a pseudo-steady-state assumption such that at any instant in time the velocity field is given by the instantaneous velocities of the inner and outer cylinders. In this case the stream function can be written

$$\Psi(x, y, t) = \Psi_{\text{in}}(x, y)\, v_{\text{in}}(t) + \Psi_{\text{out}}(x, y)\, v_{\text{out}}(t),$$

where $\Psi_{\text{in}}(x, y)$ and $\Psi_{\text{out}}(x, y)$ are the steady-state stream functions due to motion of the inner and outer cylinders respectively (Aref & Balachandar 1986). To further simplify the problem we shall only consider periodic motion of the boundaries.

The addition of an explicit time dependence increases the dimension of the system to three; however, since the time dependence is periodic the problem can be reduced to a two-dimensional map (see Guckenheimer & Holmes 1983). When the motion is made periodic the period becomes an additional parameter in the problem. It may appear that the period should be thought of in terms of time. But as was shown in the previous section, as long as inertial forces are negligible, the relevant parameter is displacement rather than time. Therefore the period of the perturbation will be parameterized by the rotation of the outer cylinder per period, i.e. if T is the period of a particular periodic velocity for $v_{\text{out}}(t)$ and $v_{\text{in}}(t)$ (throughout this work $v_{\text{in}}(t)$ and $v_{\text{out}}(t)$ will always have the same period), then the period of the perturbation is defined as

$$\theta = \frac{1}{2\pi} \int_0^T v_{\text{out}}(t)\, \mathrm{d}t.$$

Another change caused by the time-dependent boundary motion is in the interpretation of Ω. It is now the ratio of the *average* velocities of the inner and outer cylinders or, equivalently, the ratio of the inner-cylinder to outer-cylinder displacement per period.

Obviously, there are an infinite number of periodic motions that we could use but for simplicity we primarily use a discontinuous velocity protocol (see figure 1*b*). The most general definition of a discontinuous velocity protocol is one where only one cylinder is in motion at any particular instant of time. For this protocol the actual wave form of the cylinder velocities is unimportant (provided creeping flow conditions, low Reynolds and Strouhal numbers, are always met). Only their *displacements* affect the results. Note that this is only true for the discontinuous velocity protocol. If we choose a velocity protocol where the cylinders can be in motion at the same instant of time then the results will depend on the actual waveform (some comments regarding this issue will be made in §5). The straightforward discontinuous velocity protocol moves one cylinder for its periodic displacement and then move the other for its periodic displacement. Such a protocol was used by Chaiken *et al.* (1986, 1987) and Aref & Balachandar (1986). However, the

mapping that results from this motion is not symmetric† and makes the analysis unnecessarily complicated. If instead of moving the first cylinder its entire displacement per period we initially move it one half of its displacement per period, then move the other cylinder for its periodic displacement, and finally move the first cylinder for the other half of its displacement per period, the resulting map will be symmetric about the line formed by the centres of the two cylinders. A symmetric map simplifies the implementation of several of the methods of analysis, most notably the search for periodic points.

3. Experimental set-up

The experimental apparatus was designed in order to accomplish two key objectives. The first is to be able to rotate the inner and outer cylinders independently and at variable speeds according to some arbitrary wave form. The second is to be able to view, without obstruction, the entire flow domain. The first is relatively easy to implement; the second requires some forethought to implement owing to constraints created by the first objective. However, with the help of a machine designer, Mr Ed O'Brien, we have fabricated an apparatus which meets both objectives.

The experimental system is shown schematically in figure 2. The key element of the design resides in the outer cylinder, which is a stainless steel tube 15 cm long and of 15 cm inner diameter. Two large bearings fix the cylinder to the frame of the apparatus while a large gear is used to rotate it. The bearings and the gear are all attached to the outer surface of the cylinder so that there is a clear view through the interior of this cylinder. The inner cylinder is solid, also 15 cm long with a 5 cm diameter, and is connected to an arm extending over the top of the outer cylinder. The arm is attached to the frame with bolts in slotted holes such that the inner cylinder can be positioned anywhere within the outer cylinder. The bottom of the outer cylinder is a glass plate thereby allowing the flow domain to be viewed from the bottom of the apparatus as well as the top. If the flow is viewed from the top there will be an obstruction caused by the arm which holds the inner cylinder (see figure 2*b*). Although we have the ability to choose which part of the flow the arm for the inner cylinder will obstruct we do not know *a priori* if significant events will take place in the obstructed region. (Note that a very important point in the flow, a period-1 hyperbolic point, is hidden in the most conventional design.) However, the view from the bottom is completely unobstructed and consequently no information is lost when recording experiments on film.

The cylinders are rotated by identical motors; Bodine NSH-11D5. The inner cylinder is driven with a timing belt whereas the outer cylinder is driven directly through gears. The gearing of the two cylinders is such that their linear velocities have nearly identical ranges. Tachometers are attached to the motors to monitor their speeds and allow feedback control. The displacements of the cylinders are calculated by numerically integrating the speeds. The motors are controlled by PC's LIMITED 286 personal computer with a DSP-16 Data Acquisition Processor board from Ariel Corp.

The cylinder velocities are created in the following manner. First the computer approximates the desired wave form by generating a look-up table of values which

† A periodic motion is symmetric if running the flow forward and the reflecting across the line of symmetry (the line through the centres of the two cylinders in this case) is equivalent to running the flow backward (see Franjione, Leong & Ottino 1989 for a thorough discussion).

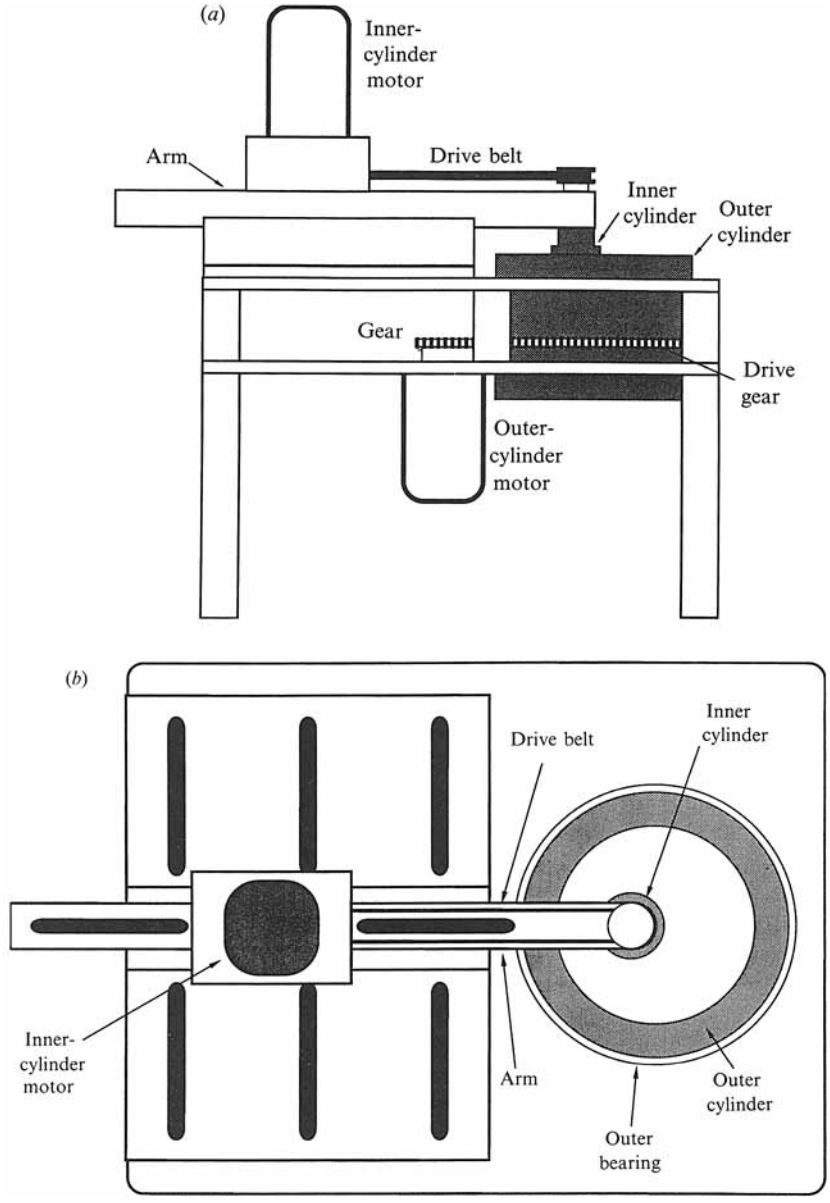

FIGURE 2(*a*, *b*). For caption see facing page.

are to be converted into voltage and output every 0.01 s by the DSP-16. This table is then down loaded to the DSP-16. The DSP-16 has its own microprocessor which has been programed to send the wave form to a D/A converter which in turn outputs a voltage to the controller of the appropriate motor. At the same time the motor speeds are determined from the tachometers and converted to digital values by the DSP-16 board. These values are summed to give the total displacement of each cylinder. The output voltages for the next period are determined from these displacements through a simple proportional control scheme. Since the time lag in the control is so great (a whole period) it is useful only in providing the desired rotation of the cylinders and correcting for drift. However, these are the only two

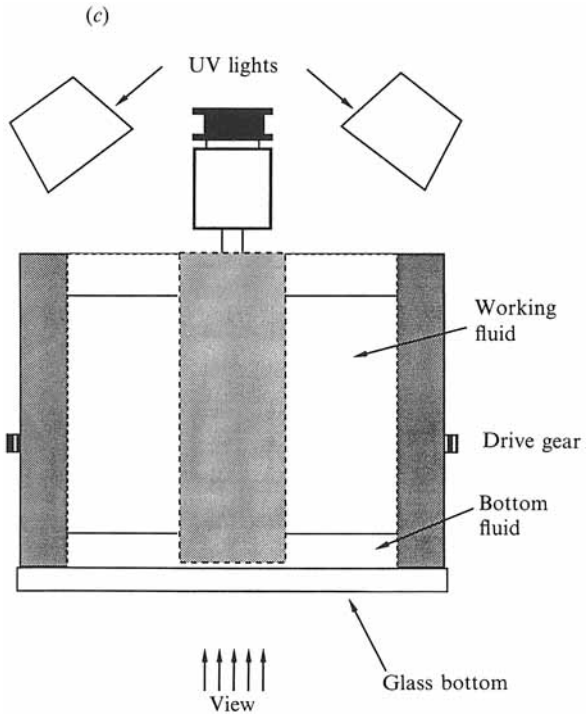

FIGURE 2. Schematics of the experimental apparatus: (*a*) the side view, (*b*) the view from the top, and (*c*) a cutaway view of the flow region. Two large bearings (not shown) connect the outer cylinder to the apparatus. The bottom of the outer cylinder is made of glass which allows an unobstructed view of the flow domain when viewed from the bottom. The inner cylinder is attached to the base of the apparatus such that it can be positioned anywhere within the outer cylinder.

major control problems so this simple proportional control scheme is more than adequate.

The working fluid is glycerin which has a viscosity of approximately 7 P and a density of 1.2 g/cm^3. This results in a Reynolds number of 0.8 and a Strouhal number ranging from 0.2 to 0.026, where the lengthscale is based on the difference in radii of the two cylinders, the velocity scale is based on the sum of the speeds of the cylinders, and the characteristic time is the period of the perturbation. The cylinders are only 15 cm deep so in order to eliminate bottom effects the glycerin is floated on a 2 cm thick layer of 1-iodo-3-methylbutane which has a viscosity of approximately 1 cP and a density of 1.5 g/cm^3. The upper surface of the glycerin is open to the atmosphere. Experiments have verified that the flow in the glycerin is independent of depth. Therefore the bottom effects are confined to the bottom organic layer and the effects of the upper free surface are negligible.

The basic experiments are performed by marking a region of the flow with a fluorescent dye (Cole-Parmer; type J295-05) dissolved in glycerin and observing how this region deforms when subjected to a specified velocity protocol. The properties of the glycerin with and without the fluorescent dye are nearly identical so effects due to viscosity difference and surface tension can be ignored (in other related work we magnify the effects of surface tension to study drop breakup in chaotic flows). Also, the diffusion coefficient of the dye is quite small (10^{-8} cm^2/s) so the spreading of the dye shown in the results is strictly due to convection. To get maximum contrast between the marked and unmarked fluid the flow region is illuminated with long-

wave UV light and all visible sources of light are removed. In addition, a UV filter is used on the camera to eliminate any direct UV light. With this technique only the dye is observable. We call these experiments deformation experiments or sometimes blob deformation experiments. The experimental results are recorded on colour slide film; a picture of the flow region is taken after every period. The film used is Fujichrome ISO 100 color slide film. Prints of the slides were made on CIBACHROME-AII photographic paper.

4. Results

4.1. *Experimental results*

Figure 3 (plate 1) shows the results of deformation experiments where the dye is initially placed arbitrarily somewhere in the 'chaotic' region. (For the purposes of discussion in this section the 'chaotic' region is the region that the dye spreads over in figure 3. The regular regions are the regions void of dye. We have observed good agreement between the regions the dye spreads over in figure 3 and regions of chaos. However, this agreement may not exist in general.) The exact initial location of the drop has no noticeable effect on the results of figure 3 provided it is placed in the chaotic region. All the experiments correspond to $r = \frac{1}{3}$, $\epsilon = 0.3$, and $\Omega = -3.0$, and to the discontinuous velocity profile shown in figure 1(*b*). The only parameter varied is θ, the amount that the cylinders are rotated per period. The values of θ vary from 90° to 720° by increments of 90° in figures 4, 9, 10, 11 and plates 1 and 3.

It is apparent from visual inspection that the area covered by the dye increases with increasing period. This may at first glance appear to violate the area-preserving character of the flow. However, this effect is mainly due to high concentration of the dye. As it stretches it appears to maintain a constant fluorescent intensity even though its concentration is decreasing. There is also an effect due to parallax. Since the dye is a three-dimensional object it must have some depth. Initially the dye is spherical so there is no parallax. However, as the experiment proceeds the dye is stretched into a very narrow ribbon. After a few periods the width of the dye ribbon is (in general) several orders of magnitude less than the depth. Viewing such a ribbon from a finite distance results in an apparent width greater than the actual width. As the dye is stretched the width observed is due to viewing the striations from a position slightly off vertical. This is an unavoidable drawback of the photography. However, these effects can be used to our advantage. This apparent area increase can be measured, thereby providing a quantitative measure of the mixing in the flow (independent measurements of the perimeter and area are proportional for area coverages of less than 50%, Leong & Ottino 1989). We should also explain that the pictures in figure 3 represent an apparent steady state in the mixing process; if the number of periods is increased the large-scale aspects of the figures, folds and islands, are nearly unchanged. (Obviously it is not a true steady state since the dye will eventually diffuse throughout the entire flow domain. Rather, it is a crossover from convection-dominated mixing to diffusion-dominated mixing. In the timescale of our experiments diffusion is unimportant.)

At $\theta = 90°$ the dye spreads over a small region corresponding to the location of the separatrixes for the unperturbed flow. At $\theta = 180°$ the shape of the region which the dye has spread over is identical to that at $\theta = 90°$ but it is now approximately twice as thick. However, when θ is increased to 270° there is a significant increase in the area coverage. The shape of the 'chaotic' region has also been greatly altered. For $\theta = 90°$ and $\theta = 180°$ there was a single large 'regular' region in the middle of the flow

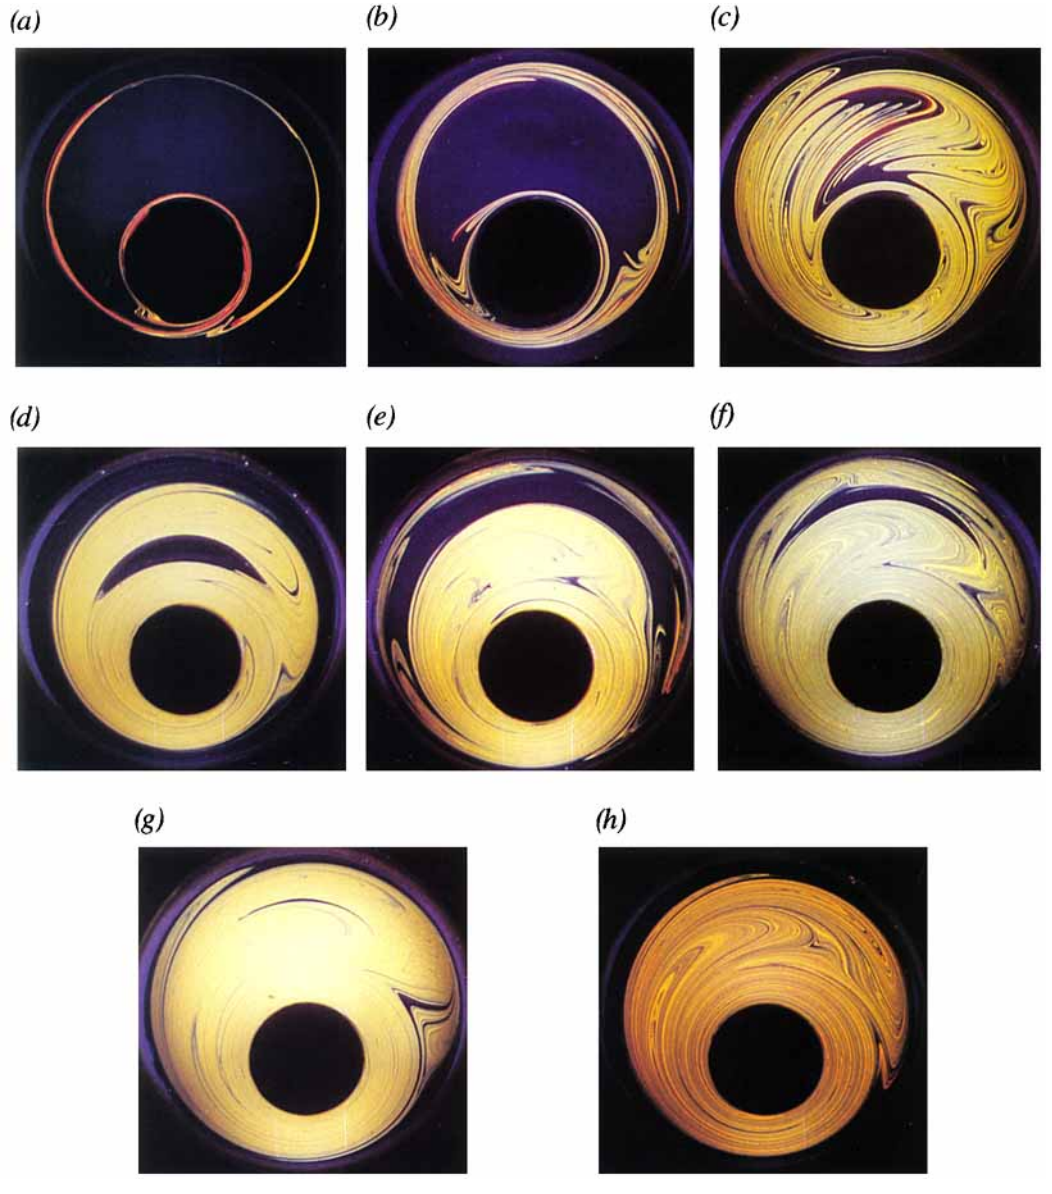

FIGURE 3. Deformation experiments for the discontinuous velocity protocol with values of θ ranging from 90° to 720°. The remaining parameters are held constant at $\varepsilon=0.3$, $r=1/3$, and $\Omega=-3.0$. *(a)* $\theta=90°$; *(b)* $\theta=180°$; *(c)* $\theta=270°$; *(d)* $\theta=360°$; *(e)* $\theta=450°$; *(f)* $\theta=540°$; *(g)* $\theta=630°$; *(h)* $\theta=720°$. *(a-g)* Deformation after 10 periods, *(h)* deformation after 5 periods. In all cases the dye initial condition was a small drop of dye which became stretched by several orders of magnitude. The Poincaré section with the appropriate parameters was used to determine the initial location of the drop.

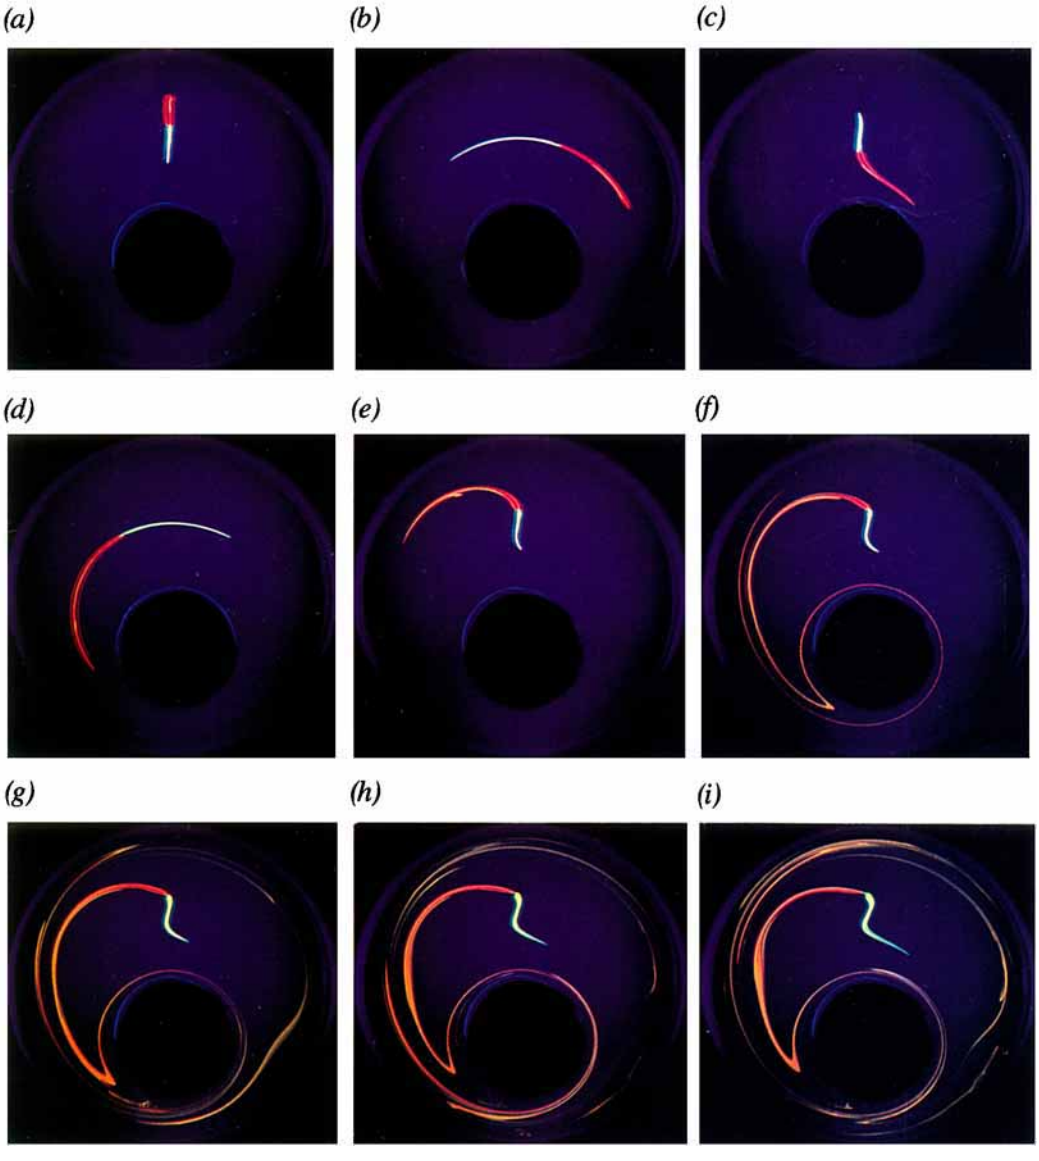

FIGURE 7. Experiments demonstrating the near solid-body rotation in the regular region and the large difference in stretching between the regular and chaotic regions. The parameter values are $\varepsilon=0.3$, $\Omega=-3.0$, $r=1/3$, $\theta=150°$. This matches the Poincaré section shown in figure 6(*e*). The initial placement of a line segment of dye is shown in (*a*). Two colours are used to differentiate the ends of the line segment. The white section passes directly through the middle of the central island while the red segment passes through the boundary of the island and into the chaotic region. (*b-e*) Show the deformation after 1, 2, 3, and 4 periods respectively. After 4 periods the white segment has come back very near to its original position, demonstrating that the rotation in the central island is approximately period 4 and also nearly a solid-body rotation, at least in the interior. The red dye shows that near the boundary of the island there is a steep gradient in the rotation rate unlike the solid-body rotation of the interior of the island. (*f-i*) Show the results after 8, 12, 16, and 20 periods respectively. Even after 20 periods the white dye is deformed only slightly whereas the small amount of the red segment in the chaotic region has stretched by several orders of magnitude.

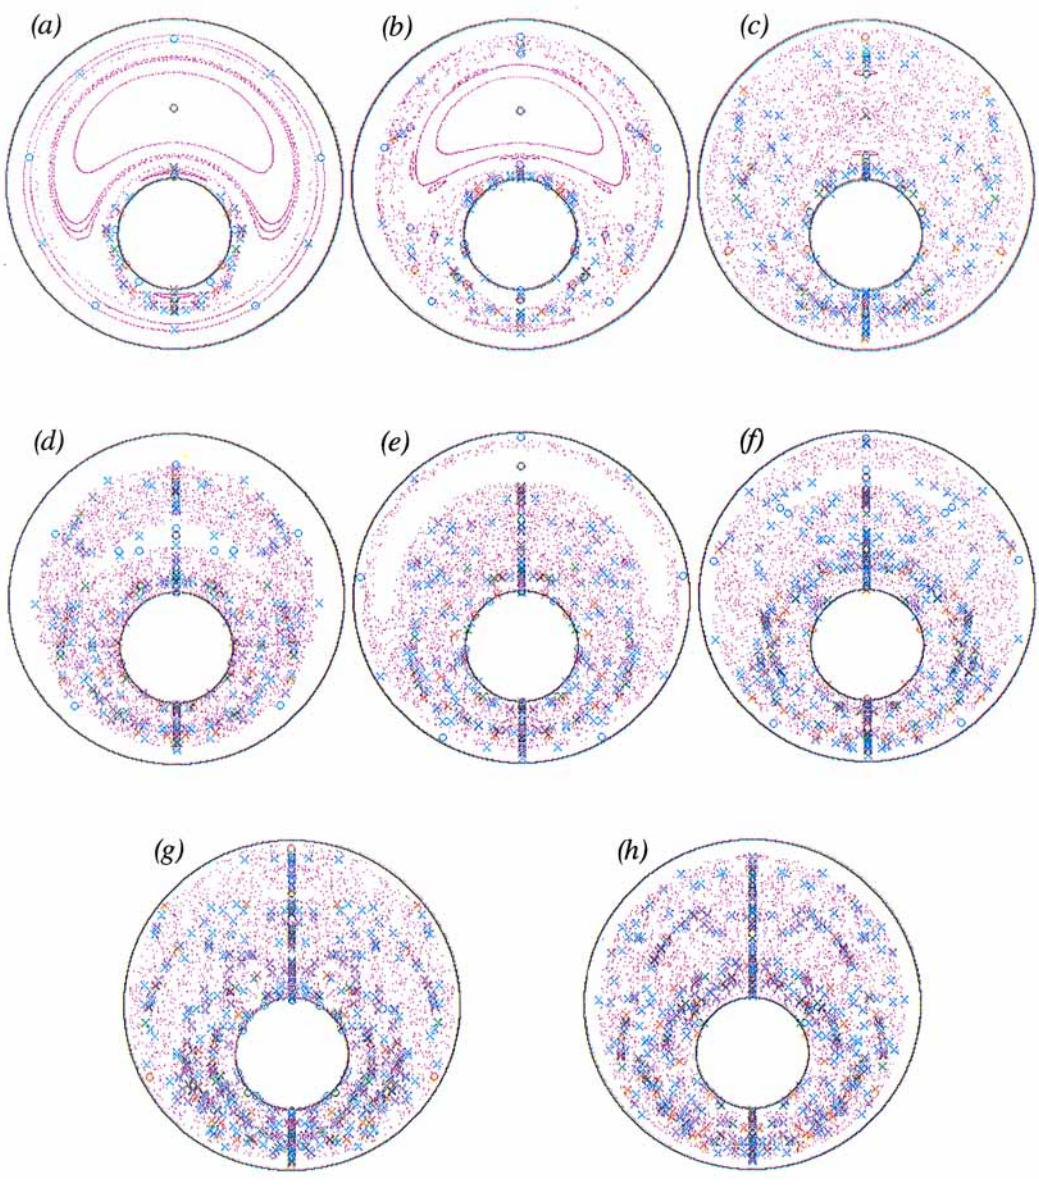

FIGURE 8. Plots of periodic points up to period 5 for values of θ ranging from 90° to 720° as in figure 3. The other parameters are $\varepsilon=0.3$, $r=1/3$, $\Omega=-3.0$. The circles represent elliptic periodic points and the crosses represent hyperbolic points. The points are colour-coded by period; black marks period-1 points, green period-2, red period-3, dark blue period-4, light blue period-5. The Poincaré sections of figure 4 are overlaid to show the regular and chaotic regions of the maps. The figures all show an excellent agreement between the location of periodic points and the features of the Poincaré sections. The number of periodic points increases for all periods as θ is increased but the number of elliptic periodic points seems to go through a maximum. At $\theta=720°$, (*h*), there are no elliptic periodic points of periods 1–5. It appears that all of the elliptic periodic points in the previous maps have become hyperbolic. This might be expected based on the widely accepted idea that as the perturbation is increased islands tend to break down into smaller islands of higher period.

domain. At $\theta = 270°$ this central island has disappeared. Instead there are two smaller islands, one above the other, on the line of symmetry. It is clear that at $\theta = 270°$ there is something fundamentally different from the previous values of θ.

When θ is increased to 360° the picture becomes similar to $\theta = 90°$ and 180° again; the two islands have disappeared and been replaced by one large central island. The total area that the dye spreads over also appears to have decreased slightly from $\theta = 270°$. At $\theta = 450°$ there is still a central island which is devoid of dye; however, it does not appear to have evolved from the $\theta = 360°$ island; it is not centred anywhere near the expected location. Instead it looks as if part of the regular region which was previously next to the outer wall breaks away and is moving into the chaotic region. At the same time the previous central island has disappeared. This new island continues to shrink as θ goes through 540° and 630°. Finally at 720° the only regular region is a thin strip next to the outer wall. These changes are difficult to completely predict by any conventional technique and new methods of analysis need to be developed (see §4.5). The most straightforward simulation technique would be to represent a blob by a collection of points. However, this is out the question; the storage and time necessary to resolve thin striations makes the problem completely impractical (see Franjione & Ottino 1987).

Several other aspects of the problem are worthy of note. The first is the dependence of the behaviour near the outer cylinder on the rotation of the outer cylinder. There must always be a regular region next to the outer cylinder, although it may be quite thin, because at some, perhaps infinitesimal, distance from the wall the rotation of the inner cylinder will have a negligible effect on the flow (the same holds true for the effects of the outer cylinder next to the inner wall). But the thickness of this region varies greatly depending on the rotation of the outer cylinder and does not precisely follow the simple rule that large period creates more chaos and consequently smaller regular regions. Rather, the key parameter seems to be the 'phase' of the outer cylinder (how close it is to an integer number of rotations per period). For the periods of $\theta = 360°$ and $\theta = 720°$ the outer cylinder moves exactly one and two revolutions per period respectively. Notice that the regular region next to the wall is thicker for these periods than for the periods immediately preceding and following them.

Another aspect is the evolution of the striations and folds as the period is increased. We have stated that the macroscopic picture does not change in a continuous fashion, i.e. the region that the dye spreads over does not evolve in a 'linear' fashion with the changing period. However, if we adopt a different viewpoint of the mixing process and focus on the folding and nesting pattern of the dye striations then the pictures do seem to change in a continuous manner; each fold 'folds more' as the period of the perturbation is increased. Each folding pattern can be thought of as local event, affecting its immediate neighbourhood. The folds of a particular pattern will propagate through the flow until they 'collide' with the folds of a different pattern. The global picture is then determined by how folds of different patterns are able to fit together. If the folds fit together well the dye will spread over a large region and the chaotic region is large. If the folds start approaching each other at blunt angles there will be gaps where no dye can penetrate; these gaps are the regular regions. The sizes of the regular islands are determined by the misfit of the folds.

4.2. *Poincaré sections*

A Poincaré section shows the long-time behaviour of trajectories of several different initial conditions. If the initial conditions are chosen well they will indicate the behaviour of all initial conditions. In a typical Poincaré section (e.g. figure 4) we see

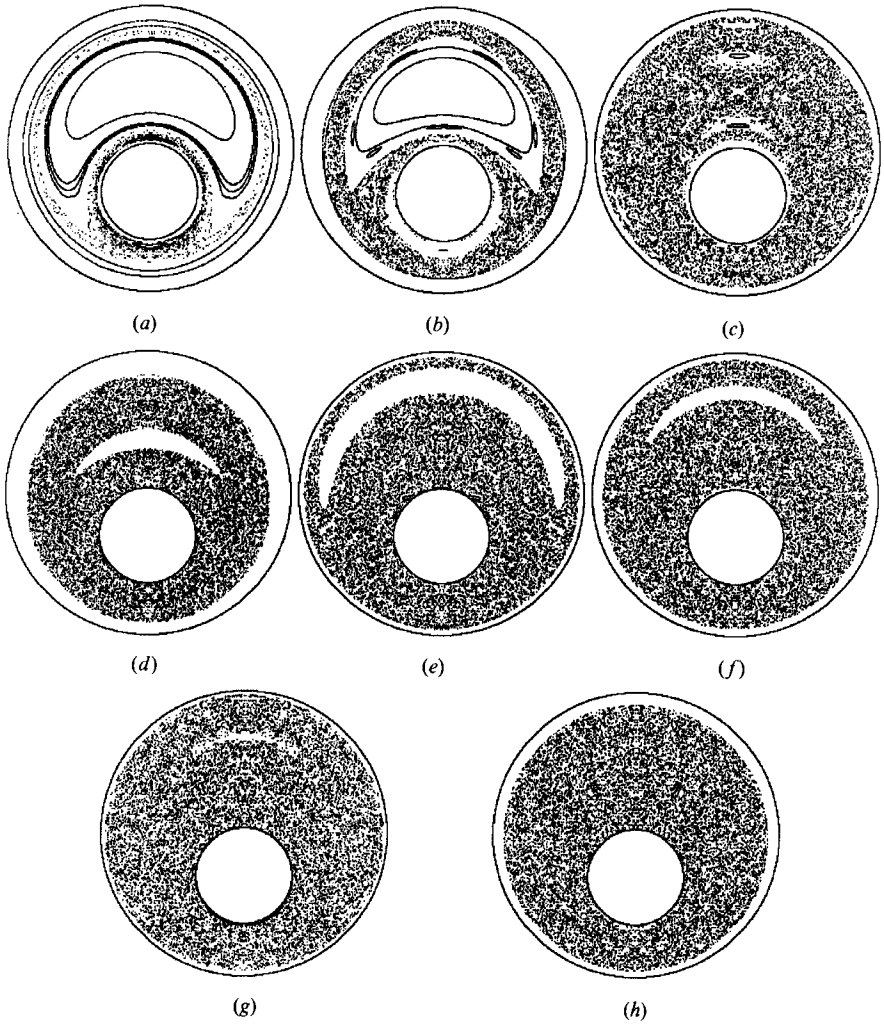

FIGURE 4. Poincaré sections for a discontinuous velocity protocol with values of θ ranging from 90° to 720° in 90° increments (*a–h* respectively). The remaining parameters are $\epsilon = 0.3, r = \frac{1}{3}, \Omega = -3.0$ and they are held constant. Each Poincaré section is made with the same eight initial conditions and the initial conditions are iterated for 1000 periods. Since the mapping is symmetric each iterate is reflected across the line of symmetry to give a total of 2000 points per initial condition or a total of 16000 points. The Poincaré section of $\theta = 90°$, (*a*), is mainly regular. Only a thin strip in the region where the separatrixes have broken apart is chaotic. In (*b*), $\theta = 180°$, the chaotic region has grown but is still confined to a strip (thicker than $\theta = 90°$) around the steady-state separatrixes. The central island is reduced in size and the boundary has developed pointed ends. When θ is increased to 270°, (*c*), the central island disappears and two period-2 islands, much smaller in total area, appear. As a result the chaotic region now covers nearly the entire flow domain. It appears that the central island has bifurcated into the two period-2 islands. However, at $\theta = 360°$, (*d*), the central island returns, the period-2 islands disappear, and the regular region next to the outer wall grows in size. The net result is that the chaotic region is smaller for $\theta = 360°$ than for $\theta = 270°$. At $\theta = 450°$, (*e*), there is still a large regular region but it is unrelated to the central island of the previous plots. Rather it appears to be a section of the regular region next to the outer wall which has broken off and moved into the chaotic region. At the same time the chaotic region extends closer to the outer wall. At $\theta = 540°$ (*f*) the new period-1 island still exists but now it has been reduced in size. The regular region next to the outer wall has increased in size. At $\theta = 630°$ (*g*) the system is completely chaotic except for a small period-1 island which is the same island as in $\theta = 450°$ and 540°. Finally, at $\theta = 720°$, (*h*), the Poincaré section is completely chaotic except for a thin band next to the outer cylinder (compare with plate 1 and figure 11).

that a trajectory has one of two distinct types of behaviour. With the first type of behaviour all of the points of a trajectory lie on a smooth surface. This implies that the trajectory is confined to a lower-dimensional subspace of the phase space (since we have a two-dimensional map the only subspace is a curve). This is the type of behaviour that one expects from unperturbed Hamiltonian systems so it is usually referred to as a 'regular' trajectory. With the second type of behaviour the points of a trajectory are scattered throughout some area of the Poincaré section. There is no apparent structure to the trajectory with this type of behaviour (although this is in part due to some shortcomings of Poincaré sections), and so it is called 'chaotic'.

These definitions of regular and chaotic are not rigorous in a mathematical sense and it is possible for regular regions to look chaotic and chaotic regions to look regular. A regular orbit may look chaotic because its exact trajectory cannot be determined owing to round-off error. On the other hand, a chaotic orbit may appear to be regular due to the chaotic region having a very thin width (for example, similar to figure 4*a*, but with a smaller θ). It is tempting to define chaotic regions and islands of regular behaviour in terms of KAM surfaces (Guckenheimer & Holmes 1983). However, in terms of the experiments this definition is not very practical. Our islands manifest themselves in ~ 10 periods whereas KAM surfaces become visible in computations typically after 10^3–10^4 periods.

4.2.1. *Agreement between Poincaré sections and experiments*

Figure 4 shows the Poincaré sections with values of θ from 90° to 720°. The other parameters are $\epsilon = 0.3, r = \frac{1}{3}, \Omega = -3.0$ and they are held constant. Comparing the Poincaré sections to the experiments, plate 1, we see that for all eight values of θ the region which the dye spreads over is completely contained in the chaotic regions of the Poincaré sections. However, the converse is not true; the dye does not completely cover the chaotic region of the Poincaré section. This is due to the much lower number of periods in the experiments (10 or fewer) compared to the Poincaré sections (1000+ periods). In theory, if one ran an experiment for as many periods as the Poincaré section the dye should spread out over the entire chaotic region. In practice it is not feasible, or in fact useful, to run an experiment for so many periods. The goal is to characterize the mixing ability of a flow. Therefore we are interested in the behaviour after just a few periods (order of magnitude 10). The question we must answer is 'what can the Poincaré section tell us about the short-time behaviour of a system?' A Poincaré section does not give precise information about where a particular set of initial conditions (such as a drop of dye) will go, but it shows regions where initial conditions are forbidden from entering regardless of how many periods the experiment is run. If dye is placed in the chaotic region it will not go into any regular regions. If dye is placed in a regular region it stays within the regular region (experience tells us that a blob placed in a regular region stretches very little) and never enters the chaotic regions or other regular regions.

4.2.2. *Drawbacks*

Hindered communication. Incorrect analysis can result if the Poincaré section data are not properly examined. Generally Poincaré sections are shown with trajectories of many initial conditions on the same plot. Since the plots are usually done in black and white there is no way to differentiate between the trajectories (points) corresponding to different initial conditions. Since regular orbits tend to appear as closed curves, each initial condition in a regular region will have its own distinct orbit and therefore is difficult to misinterpret. However, chaotic orbits show up as a haze

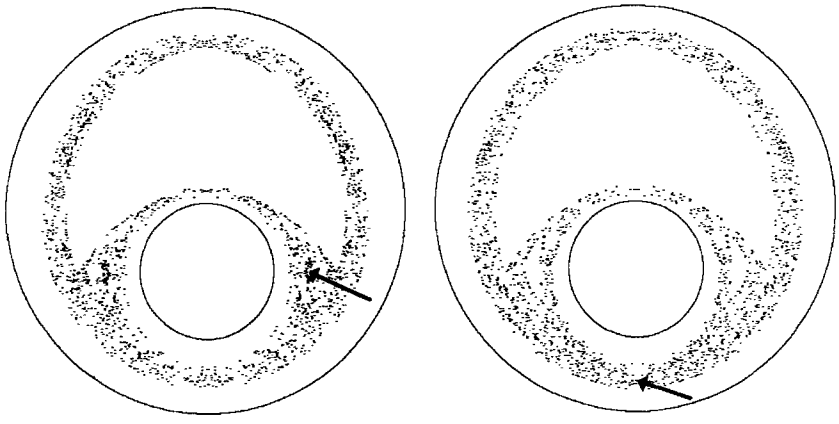

(a) (b)

FIGURE 5. Non-homogeneity in the chaotic region. (a) and (b) show two individual orbits of the Poincaré section shown in figure 4(b), $\theta = 180°$. The initial locations of the trajectories are indicated by the arrows. The chaotic trajectories in the combined Poincaré section appear to homogeneously cover the chaotic region. However, a close inspection of (a) reveals regions of high point concentration, most notably on either side of the inner cylinder. Note that one of these regions is where the point was initially placed. In (b) these regions have no points in them. This is evidence of hindered communication between these small chaotic regions and the large chaotic region.

of points scattered over an area. Parts of a chaotic region can have hindered communication with other parts of the chaotic region. In a Poincaré section made with several different initial conditions on the same plot it is not possible to differentiate between chaotic regions where all initial conditions sample the entire region in a uniform manner and chaotic regions where initial conditions can get trapped in a subregion of the chaotic region for some length of time. The amount of time that a trajectory spends in such a subspace determines the extent to which the communication is 'hindered'.

An example is given in figure 5(a, b) which shows individual trajectories in the chaotic region of figure 4(b), $\theta = 180°$. Careful examination reveals that there are regions of high concentration of points on either side of the inner cylinder in figure 5(a) whereas these same regions are devoid of any points in figure 5(b). When they are combined, as in figure 4(b), the chaotic region appears uniformly covered by the trajectories.

For some applications hindered communication between different parts of a chaotic region may be unimportant. However, for mixing we are interested in short-time behaviour and in this case barriers that only hinder communication on a long timescale may keep regions completely separate for the timescale we are interested in. Therefore it is imperative that the Poincaré section data are properly analysed. This can be accomplished in one of two ways. The first is to look at the Poincaré section of each initial condition separately as is done in figure 5. Any non-uniformity in the chaotic regions will become obvious. A better way of doing this is to give each initial condition its own colour, if one has the luxury of some sort of colour graphics device. Non-uniformities in chaotic regions can be spotted easily in such colour plots. (Colour plots are more than just aesthetically pleasing: the information density can be many times greater than a black and white plot.)

Bifurcation in central island. Conversely, not everything that we see in a Poincaré

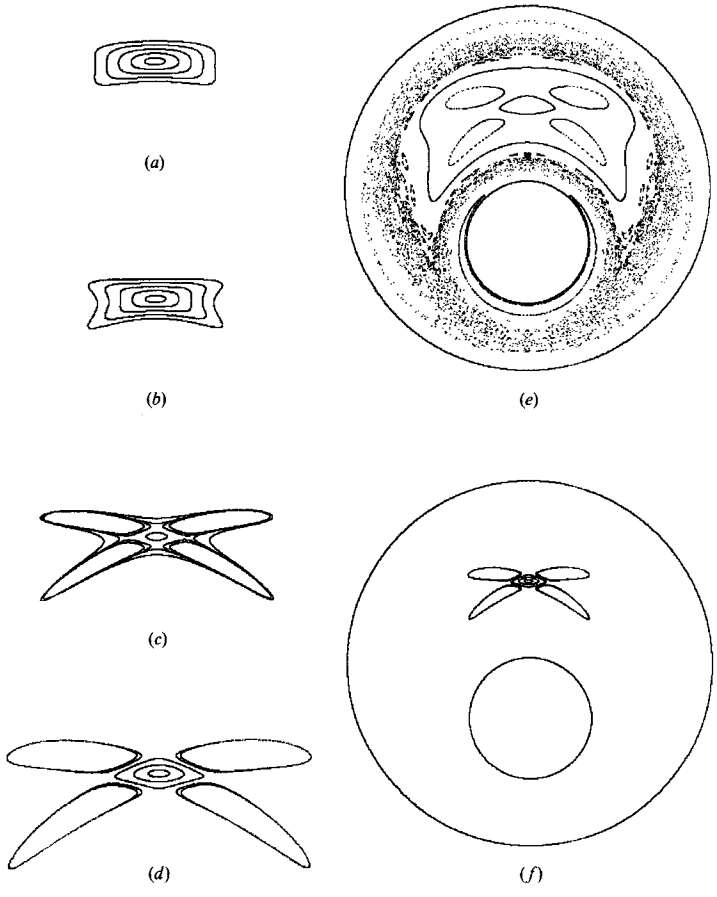

FIGURE 6. The appearance of period-4 orbits in the centre island. For all figures $\epsilon = 0.3$, $r = \frac{1}{3}$, $\Omega = -3.0$. (*a–d*) are centred at the elliptic point of the central island and are magnified by a factor of 2.5 (*a*) $\theta = 149.3°$, (*b*) $\theta = 149.4°$, (*c*) $\theta = 149.5°$, (*d*) $\theta = 149.6°$, (*e*) $\theta = 150.0°$, (*f*) shows figure (*d*) at actual size relative to the cylinders. When $\theta = 149.3°$ the orbits are elliptic in shape. At $\theta = 149.4°$ the outer orbit develops four dimples, an indication that a bifurcation is about to take place. $\theta = 149.5°$ shows some rather dramatic effects of this bifurcation. Period-4 islands have appeared and grown quite large. At $\theta = 149.6°$ the period-4 islands are still larger in size. They grow and move away from the central hyperbolic point very quickly. In fact it is clear that even a change of 0.1° (0.07%) can lead to dramatic changes in the Poincaré section.

section has a useful meaning in terms of mixing. For example, figure 6 shows the bifurcation and appearance of a chain of period-4 islands in the central period-1 island. The size of the period-4 islands grows very quickly with a very small change in θ (from 149° to 150°). This might lead to the belief that small inaccuracies in our control of the experiment can lead to very large changes in the mixing picture. However, figure 7 (plate 2) clearly demonstrates that this is not the case. The figure shows the evolution of a line consisting of two colours, red and white, placed through the centre of the regular region. The colours show the orientation of the line and demonstrate that the line rotates with a period of approximately 4 (it would be difficult to differentiate between period-2 rotation and period-4 if the line were made of just one colour). Further insight can be gained because the tip of the red segment was placed in the chaotic region. This provides a dramatic contrast between behaviour in the regular and chaotic regions. The portion of the line in the central

island tends to move as a solid body, stretching very little even after twenty periods. The portion of the line in the chaotic region behaves in a substantially different manner. It stretches and folds after a few periods, elongating to many times its initial length after twenty periods.

The conclusion is that the period-4 islands which appear in the central island have a long timescale associated with them (it can take more than 1000 periods to make one rotation about some of the orbits of the period-4 islands) and therefore the dramatic changes seen in the Poincaré sections will only be seen experimentally at very high periods. Since we are interested in mixing, long-time behaviour is unimportant and so this behaviour in the central island is unimportant. Further analysis of the regular regions for different parameter conditions support the conclusion that regular regions do not mix (i.e. deform blobs of dye). Regular regions can be thought of as dead zones where any dye initially placed in them will stretch very little if at all.

(Lack of) rate information. If our ultimate goal is to describe mixing in real applications, rate information is of prime importance. However, it is not possible to determine rates from Poincaré sections. This flaw is fundamental to Poincaré sections and so it is necessary to develop other methods which can provide this information. Another area of interest where Poincaré sections provide no information is in describing the shape of the striation patterns, e.g. folds, in a chaotic region. A first step towards both of these objectives is to locate the most important (low-order) periodic points of the flow.

4.3. *Periodic points*

One can think of finding periodic points as a minimization problem. A periodic point of a map is a solution to the equation

$$\boldsymbol{f}^n(\boldsymbol{x}) = \boldsymbol{x}, \tag{1}$$

where $\boldsymbol{f}^n(\boldsymbol{x})$ is the nth iteration of the map. The point $\boldsymbol{x}$ is a period-m point if it is a solution to (1) for $n = m$ but not a solution for any $n < m$. This definition suggests that a straightforward way to find periodic points is to find stationary points of equation (1). There are numerous methods that will find solutions to such an equation. A simple method is to define a function

$$d(\boldsymbol{x}) = \|\boldsymbol{f}^n(\boldsymbol{x}) - \boldsymbol{x}\|^2$$

which has the properties of being positive definite and having global minima at the periodic points. This is a standard idea for locating extrema of a vector field. In the context of a map d has a simple interpretation, it is the distance between a point and its image after n periods. When d is zero the point has been mapped back to itself and therefore must be a periodic point. To find the zeros of d one can follow its gradient down to a minimum. While this method cannot guarantee that the minima found are global minima, the method appears to be quite adequate in this case (every minimum we have found has been a periodic point).

Without any additional information this method leads to a search for periodic points in two dimensions. However, the symmetry of the mapping places severe restrictions on the locations of periodic points. Period-1 periodic points must mainly fall on the line of symmetry. If a period-1 periodic point was not on the line of symmetry there would have to be a matching period-1 periodic point one the other side of the symmetry line and although there is nothing preventing this, it is rarely observed in practice. Therefore, we can find most of the period-1 periodic points just

by searching the line of symmetry. Thus, we have reduced the two-dimensional search to a one-dimensional search. Periodic points of higher period must always occur in chains. For example, a period-5 periodic point must occupy four other positions before it returns to its initial location. The periodic point can only go to these four other positions or uniqueness will be violated. Therefore these four other positions are also periodic points of period 5. Together, this set of periodic points forms a 'chain'. We can simplify the search for periodic points of odd period using arguments similar to those given for the period-1 case. There is always an unpaired point in an odd chain of periodic points and this point will almost certainly fall on the line of symmetry. If this point is located we can map it forward enough times to locate the rest of the points in the chain. Thus we can reduce the search for odd-period periodic points to just a search of the line of symmetry.

Even-period periodic points present a greater difficulty since it is possible for all members of a chain to lie off of the line of symmetry without the existence of a mirror-image chain. However, in many cases points of a chain of even period will lie on the axis of symmetry (there must be an even number of points on the line of symmetry) and in these cases the chain can be found in a manner analogous to the odd-period chains. When none of the points lie on the line of symmetry the search for chains of even-period periodic points can still be simplified but we must use an additional piece of information to do so. Our mapping is derived from a continuous flow and therefore when a point is moving it must follow a continuous path. Since chains of even-period periodic points lie symmetrically across the line of symmetry there must be a path that crosses the line of symmetry at some intermediate time during the period. It turns out that this intermediate time is exactly half-way through the period. Therefore, if we search the line of symmetry at this time we can locate those even-period periodic points that never lie on the axis of symmetry after whole periods. Again, the search for periodic points is reduced from two dimensions to one dimension. This is one of the great benefits of using a mapping that is symmetric.

The character of the periodic point is determined from the eigenvalues of the Jacobian evaluated at the point. Since the map is area preserving, the character of the periodic point can be determined from the trace of the Jacobian. If the trace is greater than two the point is hyperbolic and it will stretch nearby fluid elements. If the trace is less than two the point is elliptic and it will tend to rotate nearby fluid elements.

4.3.1. *Agreement between periodic points and Poincaré sections (and experiments)*

Figure 8 (plate 3) shows plots of the periodic points superimposed upon the associated Poincaré section. The circles represent elliptic periodic points and the crosses hyperbolic periodic points. The agreement between the location of the periodic points and the various types of behaviour in the Poincaré sections is remarkable. Every island surrounds an elliptic periodic point of the appropriate period. On the other hand, there is very little basis on which to judge the agreement between the hyperbolic periodic points and the Poincaré sections. This is due to an inherent flaw of Poincaré sections: they give no information on structure in the chaotic regions (experiments do). However, there are instances where comparisons can be made. For the case $\theta = 180°$ there is a chain of period-5 islands with period-5 elliptic points in the middle on the outer edge of the chaotic region. In between these islands there is a chain of period-5 hyperbolic points. This is precisely the type

of behaviour predicted by the Poincaré–Birkhoff theorem (see for example, Lichtenberg & Lieberman 1983, p. 169) which describes the way in which rational orbits break down into sequences of hyperbolic and elliptic periodic points.

4.3.2. *Agreement between eigendirections of periodic points and striations in experiments*

Several things can be learned directly from the periodic points. The largest eigenvalue (in absolute value) of a hyperbolic periodic point gives an indication of the rate of stretching in the neighbourhood of that point. Indeed, the eigenvalue seems to be a reasonable measure of the rate along the entire unstable manifold of the hyperbolic periodic point. Since the manifold tends to stretch throughout the entire chaotic region the eigenvalue may give a good indication of the mixing rate throughout the chaotic region. Also, the eigenvalues of a periodic point vary with θ proportionally to the variation of the mixing rate with θ. This can be used to predict the difference in mixing rate for two values of θ based on the difference in the eigenvalues of a particular periodic point.

We can get an idea of the alignment of the striations in the chaotic region by plotting the eigendirection associated with the maximum (in absolute value) eigenvalue of hyperbolic periodic points. Plots are shown in figure 9. It is clear when comparing figure 9 to the experiments shown in figure 3 that the eigendirections do give a very good indication of the direction of the striations in the neighbourhood of the periodic points. Unfortunately the periodic points are not distributed uniformly throughout the chaotic region: some areas have a plethora whereas others are completely void. This makes it impossible to speculate on the complete striation structure, so some other method must be used to obtain such a picture.

4.4. *Manifolds*

The manifolds of the periodic points can be of great use in determining the shape of the dye striations in the chaotic region. In order to find the unstable manifold of a particular hyperbolic periodic point we can simply surround the point with a circle and then 'turn on' the flow. The circle will deform but will always surround a segment of the unstable and stable manifolds. As the number of periods increases the length of the unstable manifold enclosed by the circle increases whereas the length of the stable manifold decreases. After a certain number of periods (depending on the eigenvalues of the periodic point and the initial radius of the circle surrounding the periodic point) the circle will appear as a single curve surrounding part of the unstable manifold. Note that manifolds are infinitely long and this method will only show a segment of the unstable manifold.

In order to make this method computational, the circle surrounding the periodic point must be represented as a collection of discrete points. Therefore we do not compute the deformation of circle but rather the deformation of a circle of points. Figure 10 shows manifolds of the period-1 hyperbolic point corresponding to the saddle point of the steady counter-rotating flow. One complication which arises very quickly is that the distance between points becomes so large that the resolution in our representation of the manifold is not acceptable and we can no longer determine where the manifold is. Such a complication is inherent to the method of computation because we are interested in stretching the circle to several orders of magnitude larger than its initial radius. Obviously this causes the distance between points to increase by several orders of magnitude as well. There are at least two possible

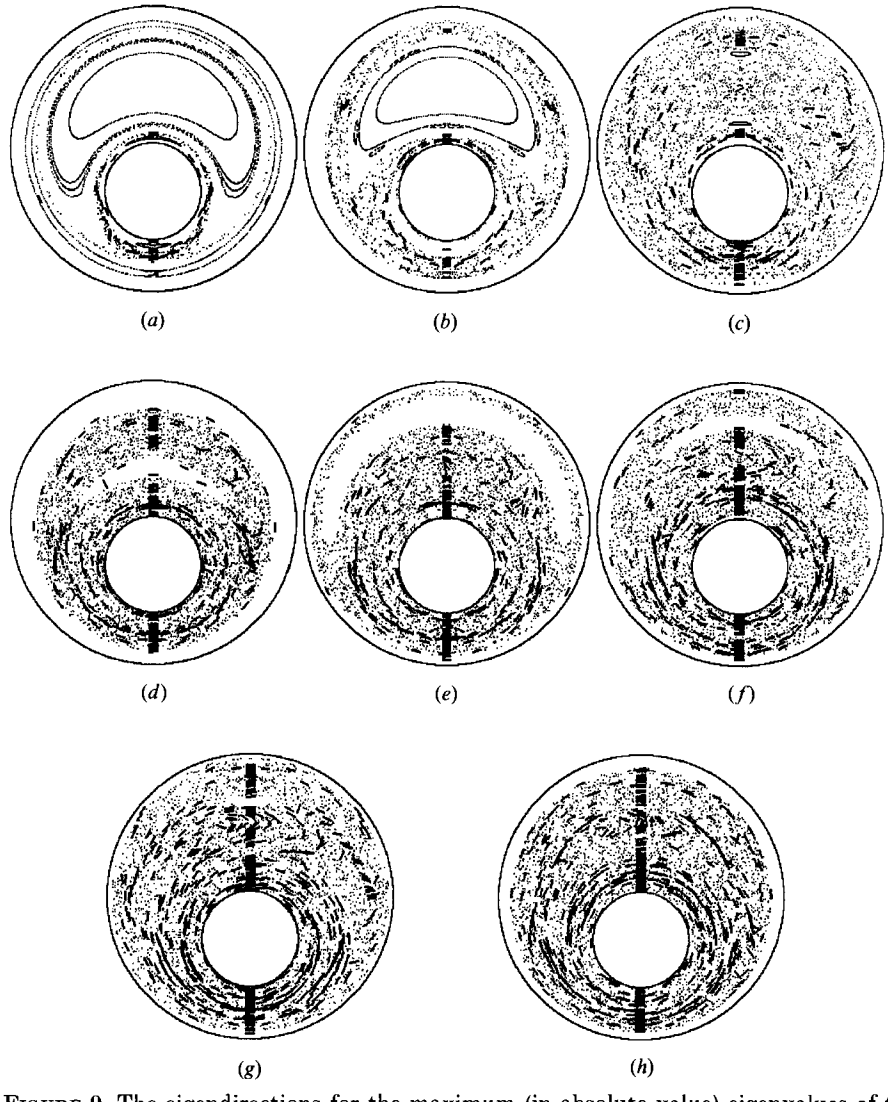

FIGURE 9. The eigendirections for the maximum (in absolute value) eigenvalues of the hyperbolic periodic points up to period 5. In all cases $\epsilon = 0.3$, $r = \frac{1}{3}$, $\Omega = -3.0$.

methods of dealing with this problem. The first is to make the initial distance between points very small. This means that we must use a very large number of points in the initial circle. This is not very efficient because the stretching is non-uniform and it is very difficult to determine *a priori* how much each pair of points will separate. To ensure that no two points spread too far apart we must use an excessive number of points. A much better way to deal with the 'resolute problem' is to add extra points dynamically, as needed. With this method we save time and storage space by avoiding computations with unnecessary points. We have implemented this second method in our manifold algorithm. Even using the second method the number of points necessary to represent the manifold increases geometrically and for the particular examples shown in figure 10 it was excessively time consuming to run the algorithm for more than three or four periods depending

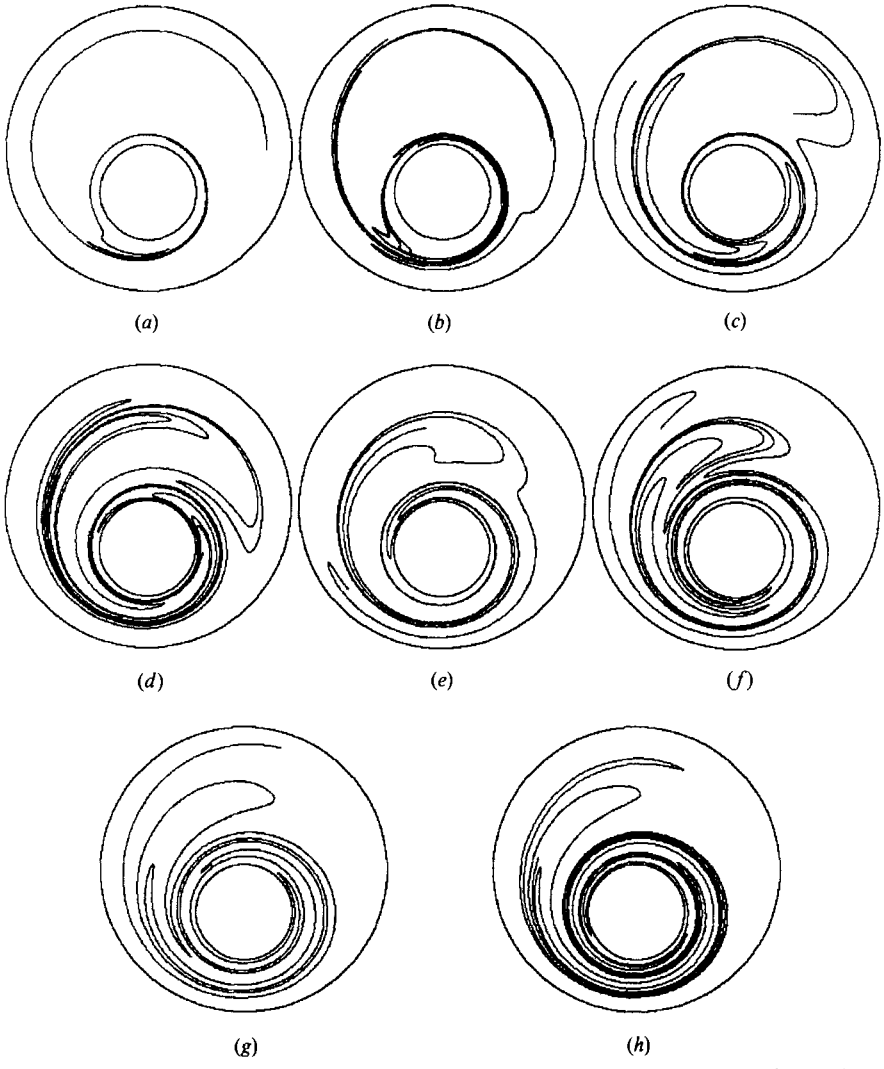

FIGURE 10. Unstable manifolds of the perturbed saddle point of the steady flow for various values of θ ranging from 90° to 720°. The other parameters are $\epsilon = 0.3$, $r = \frac{1}{3}$, $\Omega = -3.0$. Initially, for (a, b), $\theta = 90°$ and $\theta = 180°$, the unstable manifold stays very close to the separatrixes of the steady state case. When $\theta = 270°$, (c), the manifold begins to wander over a large part of the flow domain; it is no longer restricted to a neighbourhood of the steady-state separatrixes. As θ is increased further the manifolds spread over a larger region of the flow domain. (Compare with the striation patterns in plate 1.)

on the value of θ and the initial radius of the circle. In any event, displaying a larger segment of the manifold than we have shown would tend to obscure rather than clarify the results.

A comparison between figure 3 and figure 10 demonstrates the very good agreement between the shape of the striations and the shape of the manifolds. The dye has many more folds than the sections of the manifolds shown here but that is because the dye is allowed to evolve for many more periods (the experiments not only evolve for more periods, 10 compared to 3 or 4, but the initial radius of the drop is some orders of magnitude greater than the radius of the initial circle). Pictures

of the dye after a short number of periods show patterns that are strikingly similar to the manifold plots. The 'shape' that the dye will adopt becomes apparent after a very few periods; subsequent periods merely fill in the picture. Undoubtedly if we were to plot a very long length of the manifold it would look nearly identical to the dye deformation experiments.

One analytical method that exists for studying manifolds is the Melnikov method. The Melnikov method provides a measure of the perpendicular (relative to the unperturbed manifold) distance between the stable and unstable manifolds (for details of the method see Guckenheimer & Holmes 1983; Wiggins 1988). The method is valid only for small perturbations and therefore must be used with extreme care in the analysis of mixing. If the perpendicular distance between the manifolds becomes multivalued the Melnikov method is no longer valid (for example if one of the manifolds has a 'kink' in it). In fact, there is always some neighbourhood of a hyperbolic periodic point within which the Melnikov method can no longer accurately approximate its manifolds. However, Rom-Kedar, Leonard & Wiggins (1989) have shown that useful information about the mixing can be gained if the Melnikov method is valid for as little as one 'lobe' (a lobe is the area enclosed by the segments of the stable and unstable manifolds between neighbouring zeros of the Melnikov function) of the intersection of the stable and unstable manifolds. By viewing plots of the unstable and stable manifolds (not shown) we have determined that the Melnikov method may be useful (i.e. there is at least one lobe without 'kinks' in it) for values of θ up to about 180°. Since the region of most interest (i.e. best mixing) occurs well beyond 180°, the Melnikov method is of little use to our mixing studies. However, suitable extensions of the method might be developed which could be of use in mixing problems.

4.5. *Stretching*

A necessary condition for mixing to occur is that fluid elements must be stretched. In the experiments shown in figure 3 the initial drops have been stretched to a length several orders of magnitude greater than their initial radius. A natural question to ask in this context is 'where does most of the stretching occur?' The results presented so far would suggest that the stretching is highly non-uniform throughout the flow domain.

To determine the stretching consider an infinitesimal fluid element (a point in the flow domain). The stretching of an infinitesimal vector of arbitrary initial orientation located at the element (at this point) can be computed if one knows the deformation tensor. The deformed infinitesimal vector at some later time t is given by

$$d\boldsymbol{x}(t) = \boldsymbol{F}(t)\cdot d\boldsymbol{X},$$

where $d\boldsymbol{X}$ is the initial vector, $d\boldsymbol{x}(t)$ is the deformed vector, and $\boldsymbol{F}$ is the deformation tensor. To find $\boldsymbol{F}$ we solve the following set of differential equations:

$$\dot{\boldsymbol{F}} = (\nabla \boldsymbol{v})^{\mathrm{T}}\cdot \boldsymbol{F}, \quad \dot{\boldsymbol{x}} = \boldsymbol{v};$$

$$\boldsymbol{F}_{t=0} = \boldsymbol{I}, \quad \boldsymbol{x}_{t=0} = \boldsymbol{X},$$

where $\boldsymbol{x}$ is the location of the fluid element initially located at $\boldsymbol{X}$, $\boldsymbol{v}$ is the velocity, and $\boldsymbol{I}$ is the identity matrix. One detail should be emphasized: the fluid element is convected by the flow so the gradient of the velocity field is time dependent. The equations of motion for the fluid element must simultaneously be solved to compute $(\nabla \boldsymbol{v})^{\mathrm{T}}$. The length stretch is defined as the ratio of the magnitudes of $d\boldsymbol{x}$ and $d\boldsymbol{X}$:

$$\lambda = \|d\boldsymbol{x}\|/\|d\boldsymbol{X}\|.$$

Ideally we would like to have a closed-form solution for $\boldsymbol{F}$, and hence an expression for λ, as a function of position and time. In practice, we are unable to find such a solution and so we must compute $\boldsymbol{F}$ numerically. This means that we can only determine λ at discrete points in the flow domain. In order to determine the stretching field we discretize or pixelate the flow and approximate the stretching in a pixel as the stretching of a single point within the pixel. Note that this assumes that the stretching is relatively homogeneous within a pixel. Such an assumption is adequate for the type of results shown here. A value of stretching is arbitrarily chosen above which the stretching is considered 'good' and any pixel with stretching larger than this value is coloured white. This value is called the 'cutoff value'. Note that although our choice is arbitrary, the results are relatively insensitive to it over a very large range of values (generally at least an order of magnitude). Pixels whose values of stretching are less than the cutoff value are coloured black.

One further detail must be resolved. We must choose whether to plot the stretching as a function of initial location or final location, i.e. we can plot how much a point initially located in a particular pixel will stretch or how much a point currently located in a particular pixel has stretched. The results shown here show the stretching as a function of final location; however, as a consequence of the symmetry of the system it can be shown that plots according to initial location and final location are symmetric with respect to each other. Note that if we plot the results as a function of final position there is no information on the initial position of the points. Likewise, a plot of the results as a function of initial position would give no information about the final position.

Figure 11 shows the results of the stretching computations plotted according to final position. It is significant that there is no information in the plots indicating where the points originated. The plots only show how much each point has stretched after some number of periods and then only up to some cutoff value. There may be points which stretch several orders of magnitude more than the cutoff value but this information is hidden by the plotting method. In the light of these provisos let us compare the stretching results to the Poincaré sections. For all cases the white regions (where the amount of stretching is large) lie within the chaotic regions as marked by the Poincaré sections. This is exactly what we expect.

When the stretching plots are compared with the experiments (figure 3) something quite unexpected becomes evident; the regions of high stretch in the stretching plots match, to within their limited resolution, the regions that dye spreads over in the experiments. This correlation is remarkable because the stretching computations and the experiments show completely different types of results. The experiments show the evolution of a continuous set of initial conditions (the initial drop of dye) whereas the stretching computations show the degree of stretching in the flow domain but contain no information about initial location. Further studies have shown that the agreement between experiments and stretching computations is valid even for a very short number of periods, two or so, where the dye covers only a small fraction of the chaotic region. There must be some underlying reason behind this remarkable match of such dissimilar processes.

A possible explanation for the agreement between the stretching plots and the experiments occurs when the stretching plots are compared to the manifold portraits (figure 10). The stretching plots seem to have shapes like those of the manifolds. However, we have already stated that the highest stretching in the flow occurs on these manifolds (the manifolds of period-1 hyperbolic points) so it is not surprising that the stretching plots should show maximum stretching near the

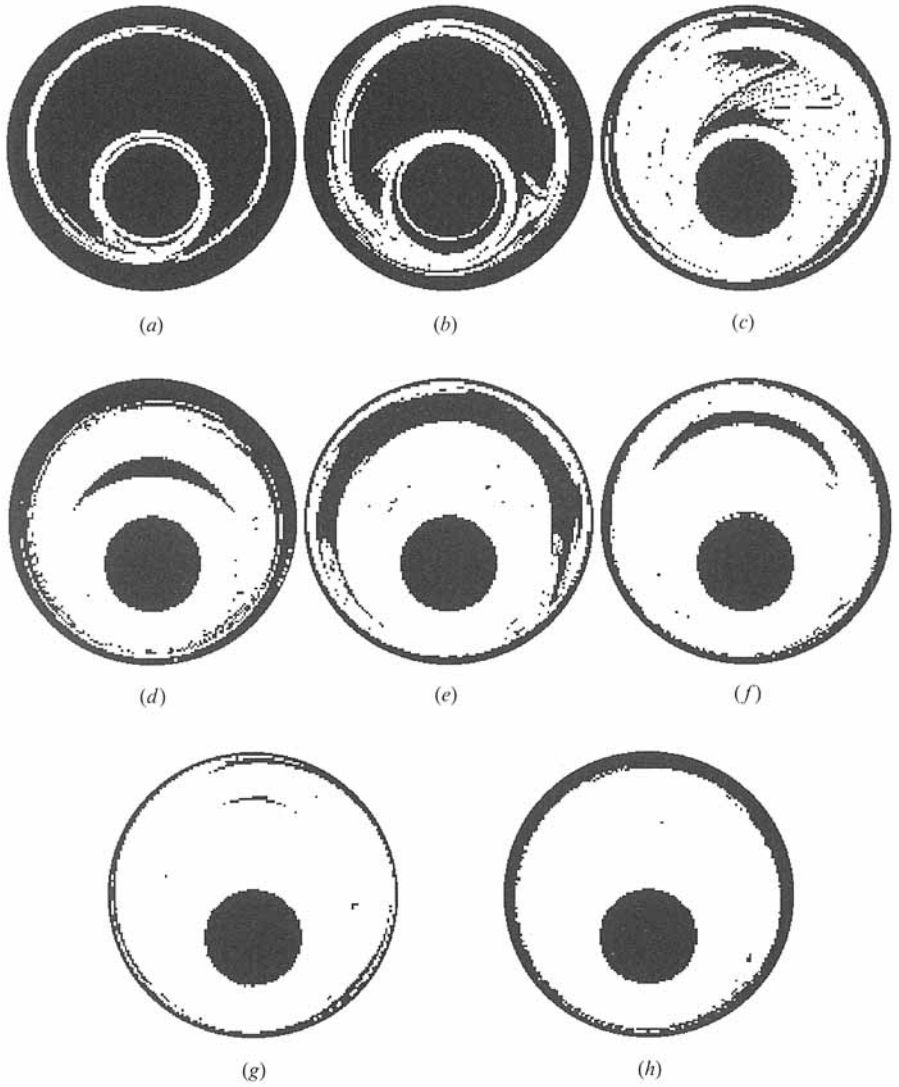

FIGURE 11. Stretching plots for the discontinuous velocity protocol for values of θ ranging from 90° to 720°, with the other parameters held constant at $\epsilon = 0.3$, $r = \frac{1}{3}$, $\Omega = -3.0$. To create these plots the flow domain is subdivided into squares of side $\frac{1}{64}$ of the outer cylinder radius. Two initial conditions are placed into each box and mapped backward for 10 iterations. At the same time the deformation tensor associated with each point is computed. From the deformation tensor it is possible to calculate the stretch of an infinitesimal vector associated with the point of any initial orientation. In these computations the stretching from 100 initial orientations are averaged to determine an average value of stretching associated with each initial condition. Then the two initial conditions of each square are averaged to obtain an average value of stretching for each square. The squares are colour-coded according to the amount of stretching they undergo. Any region where the stretching is greater than the cutoff value is white while any region where the stretching is less than the cutoff value is black. The cutoff value is chosen in an *ad hoc* manner but the results are relatively insensitive to it over a large range of cutoff values (compare with plate 1).

manifolds. It was also noted that the shape of the manifolds match the shape of the striations in the experiments. So the underlying reason for the agreement between the stretching plots and the experiments may be a result of the manifold structure.

5. Conclusions

The overall goal of this research is to find computational methods to predict what sort of mixing, as shown by dye experiments, a given flow will exhibit. It is clear that no one particular computational method can provide a complete prediction. If we just want to get a gross picture of the spreading of the dye, then the Poincaré section is adequate. The chaotic regions of the Poincaré sections match the regions that the dye spreads over in experiments. However, the match is not perfect and there is no information about the rate of mixing in a Poincaré section. Further, some phenomena which look quite spectacular in the Poincaré sections do not show up in any significant way in the experiments (i.e. the Poincaré sections can be misleading if viewed on their own). The plot of the eigendirections associated with the hyperbolic periodic points shows a strong correlation with the alignment of the striations in the mixing experiments. However, since the periodic points are not distributed uniformly throughout the flow, they cannot form a complete picture of the striations. A more complete picture is given by the unstable manifolds of the hyperbolic periodic points; the manifolds of period-1 manifolds give a template for the folding of the dye throughout the chaotic region but the method can only give a template, not the full striation pattern.

The best indicator of dye spreading is given by stretching plots. However, since we do not have an adequate explanation for the correlation between the stretching plots and the experiments, stretching plots should be used with caution. Also one should keep in mind that the appearance of the stretching plots depends, to a small degree, on the specific cutoff value used (a good rule of thumb is to make the cutoff value equal to the sum of the total displacements of the cylinders for the experiment). In spite of these imperfections the stretching plots seem to match the experiments extraordinary well. In fact, if we increase the resolution of the stretching plots the match with the experiments is even better. This match does not come without a cost however; the computational time for these stretching plots is at least an order of magnitude greater than any other method and becomes even greater if more resolution is desired.

Note that all of these methods deal with the chaotic region. From the viewpoint of mixing it is a reasonable approximation to consider fluid elements in the regular regions to not stretch or mix at all; the regular regions are best thought of as 'dead' zones. Note also that the results presented here were produced using a discontinuous velocity protocol. We chose this protocol because it was very simple to produce both experimentally and computationally. However, we have done computations based on several other similar velocity protocols (similar in the sense that they create a symmetric mapping) and have found the results to be remarkably similar to the discontinuous velocity protocol. In fact, for many sets of parameter values it requires very close inspection to differentiate between results from different velocity protocols.

Finally, the time it takes to run a particular numerical experiment may be a very important factor when attempting to decide which of the methods to use. The Poincaré sections shown here took from 5 to 50 min to complete on our Sun 3-160 (with a floating point accelerator). It took approximately the same time to find the

periodic points up to period 5. The time to find periodic points of a given period goes up exponentially with the period. It took from 20 to 50 min to find a length of manifold approximately 80 times the radius of the outer cylinder. The stretching plots took far longer to compute, ranging in time from seven hours to two days. Note that our Sun microcomputer works at about half the speed of a VAX 8600 on this particular problem.

The authors take pleasure in acknowledging the financial support of the Department of Energy (DE-FG02-ER 13333).

REFERENCES

AREF, H. & BALACHANDAR, S. 1986 Chaotic advection in a Stokes flow. *Phys. Fluids* **29**, 3515–3521.

BALLAL, B. Y. & RIVLIN, R. S. 1976 Flow of a Newtonian fluid between eccentric rotating cylinders: inertial effects. *Archiv. Rat. Mech. Anal.* **62**, 237–294.

CHAIKEN, J., CHEVRAY, R., TABOR, M. & TAN, Q. M. 1986 Experimental study of Lagrangian turbulence in a Stokes flow. *Proc. R. Soc. Lond.* A **408**, 165–174.

CHAIKEN, J., CHU, C. K., TABOR, M. & TAN, Q. M. 1987 Lagrangian turbulence and spatial complexity in a Stokes-flow. *Phys. Fluids* **30**, 687–694.

CHIEN, W.-L., RISING, H. & OTTINO, J. M. 1986 Laminar mixing and chaotic mixing in several cavity flows. *J. Fluid Mech.* **170**, 355–377.

DUFFING, G. 1924 Beitrag zur Theorie der Flüssigkeitsbewegung zwischen Zapfen und Lager. *Z. Angew. Math. Mech.* **4**, 297–314.

FRANJIONE, J. G., LEONG, C. W. & OTTINO, J. M. 1989 Symmetries within chaos: a route to effective mixing. *Phys. Fluids* A, **1**, 1772–1783.

FRANJIONE, J. G. & OTTINO, J. M. 1987 Feasibility of numerical tracking of material lines in chaotic flows. *Phys. Fluids* **30**, 3641–3643.

GUCKENHEIMER, J. & HOLMES, P. 1983 *Nonlinear Oscillations, Dynamical Systems, and Bifurcation of Vector Fields*. Springer.

JEFFERY, G. B. 1922 The rotation of two cylinders in a viscous fluid. *Proc. R. Soc. Lond.* A **101**, 169–174.

LEONG, C.-W. & OTTINO, J. M. 1989 Experiments on mixing due to chaotic advection in a cavity. *J. Fluid Mech.* **209**, 463–499.

LICHTENBERG, A. J. & LIEBERMAN, M. A. 1983 *Regular and Stochastic Motion*. Springer.

MÜLLER, W. 1942*a* Über die Drehung zweider Zylinder in einer Zähen Flüssigkeit und die Theorie der Kräfte am Rotaionsviskosimeter mit exzentrischen Zylindern. *Ann Phys.* **41**, 335–354.

MÜLLER, W. 1942*b* Beitrag zur Theorie der Langsamen Strömung Zweier exzentrischer Kreizylinder in der Zähen Flüssigkeit. *Z. Angew. Math. Mech.* **22**, 177–189.

OTTINO, J. M. 1989 *The Kinematics of Mixing: Stretching, Chaos, and Transport*. Cambridge University Press.

ROM-KEDAR, V., LEONARD, A. & WIGGINS, S. 1989 An analytical study of transport, mixing and chaos in an unsteady vortical flow. California Institute of Technology, preprint.

WANNIER, G. H. 1950 A contribution to the hydrodynamics of lubrication. *Q. Appl. Maths* **8**, 1–32.

WIGGINS, S. 1988 *Global Bifurcations and Chaos*. Springer.