Journal of Engineering Mathematics **48:** 129–155, 2004.

Mixing measures for a two-dimensional chaotic Stokes flow

M.D. FINN, S.M. COX and H.M. BYRNE
School of Mathematical Sciences, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom

Received 2 June 2002; accepted in revised form 2 September 2003

Abstract. The effectiveness of a large number of protocols for mixing in a two-dimensional chaotic Stokes flow, according to a variety of measures, is investigated. The degree to which the various mixing measures are correlated is computed, and while no single protocol simultaneously optimises all measures, it is found that a small subset of the protocols perform well against most measures. However, it is difficult to elicit general rules for selecting effective protocols: for example, superficially similar protocols are found to exhibit considerably different mixing capabilities. The results presented here suggest that the selection of effective protocols by 'sieving' (*i.e.*, by successively eliminating candidate protocols that fail increasingly discerning mixing measures) may be ineffective in practice.

1. Introduction

High-quality mixing is immensely important in many areas, ranging from food production to paint manufacture and brewing. On an industrial scale, the stirring necessary to achieve good mixing can be enormously expensive, and this justifies the many recent mathematical, numerical and experimental investigations into the quantification and optimisation of the stirring/mixing process (see, *e.g.*, [1, 2]). Of course, there are many practical and economic constraints that may be placed on the mixer (*e.g.*, speed of operation, expense and energy usage), and many different criteria by which the resulting mixtures may be judged (*e.g.*, uniformity of mixing, and yield of a chemical reaction), depending on the application.

In this paper we consider an idealised two-dimensional mixing device [3, 4] that operates in the Stokes-flow (slow flow) regime, in which the flow is chaotic, but laminar; the device approximates the action of mixing a tank of fluid with a long stick of circular cross-section. Laminar mixing is desirable in many applications, such as polymer production and dentistry, where turbulent mixing is either infeasible due to the high viscosity of the working fluid, or damaging to the fluids being mixed. Further, it is now well established that even simple mixing devices, capable of generating Stokes flows with chaotic Lagrangian particle paths, can mix viscous fluids effectively [5, 6]. We examine, in detail, the effectiveness of a large number of 'stirring protocols', which correspond to varying both the path of the stick through the tank and the stick diameter, and do so according to a number of different measures. The central question that we aim to answer is whether one protocol or a small subset of protocols can be identified as somehow mixing 'best' across a range of mixing measures. In order to make our task feasible, we have required that all mixing measures be readily implementable in automatic fashion in a computer code.

We are concerned here with computational mixing measures rather than formal, mathematical ones (see, *e.g.*, [7]). A variety of such tools exist for systematically measuring mixing quality, almost all of which are based on the Lagrangian approach of tracking passive, non-

diffusive, inertia-free tracer particles in the flow. Broadly speaking, they may be classified as dynamical-systems, statistical or physical measures; while such a classification proves useful for some purposes, it is of course a largely arbitrary one. The first class of measures that we study relies on the observation that, since the mixing device operates in time-periodic fashion, it is convenient to work with the mapping that takes tracer particles to their new locations after one period of operation of the mixing device. The dynamical-systems techniques that we use describe certain features of this mapping, and include: iterated mappings [8–11]; periodic-point, manifold and symmetry analysis [11–17], and Liapunov-exponent calculation [10, 18, 19]. The second set of measures that we use are concerned with the end-product of a sequence of mappings, rather than with individual mappings themselves. These statistical techniques include: Kolmogorov [20] and mixing-variance [9, 21, 22] tests of the uniformity of the mixing; diffusive return percentage tests [23], indicating the degree to which the mixing can be undone, and intermaterial area density [24–26], related to the production of thin striations. The third, and final, class of measures that we consider are more physically-based, and include: the energy input required to achieve a given state of mixing [15, 27, 28], the growth rate of an interface across which mixing can take place [29], and the residence time of tracer in the mixer [11]. (One significant measure that lies beyond the scope of the present work is the yield of a chemical reaction taking place in the flow, which for a multi-stage reaction is known to be significantly affected by the details of the mixing [1, 30–35].)

In previous evaluations of the mixing effectiveness of Stokes flow devices, a common approach has been to apply a single measure to a given set of stirring protocols, in order to find the optimum protocol. A more refined strategy is to apply a ‘sieve’ [36], whereby a succession of ever more sophisticated (and computationally expensive) measures are used, at each stage refining the choice of protocol, until an optimum protocol or set of protocols is established. There is no suggestion in [36] that any single optimum protocol need exist; nevertheless, the technique of sieving is conceptually appealing. Central to its success is the notion that different mixing measures broadly agree on their ranking of mixing quality. This latter issue, while crucial to the ‘sieve’ approach, seems to have received little direct attention in the literature; it is, however, a significant concern here. We investigate the extent to which different mixing measures agree, and in particular whether such agreements can lead to a saving in computational time, with a cheap measure serving as a proxy for a more expensive one. We also highlight the discrepancies that may exist between mixing measures, since these indicate the difficulties inherent in simultaneous optimisation of different mixing properties, and hence in implementation of ‘sieving’.

Since this is, to our knowledge, the first systematic study to investigate the correlations between a range of different mixing measures, we focus on a flow where the exact velocity field is known. This provides some significant computational advantages over a flow for which the velocity field itself must be computed at each time step, in terms of both accuracy and CPU time. It also allows us to evaluate thoroughly a large number of stirring protocols. The extent to which the mixing measures considered here could be feasibly calculated when the velocity field is not known exactly is clearly an important practical question, but one whose resolution is left to future studies.

The mixing device used here is the ‘translating, rotating mixer’ (TRM) [3], whose velocity field is expressible in closed form involving finitely many terms and which is (nearly) unique among Stokes flow mixing devices in having the desirable feature of a geometry that changes with time (a notable other such flow being that in a baffled cavity with moving baffles [37], although the velocity field in that case is known only numerically). The TRM is introduced

Table 1. The mixing measures used in this paper, with the abbreviations used in discussions.

Dynamical-systems measures	Iterated mapping	IM
	Liapunov exponent	LE
Statistical measures	Diffusive deviation	DD
	Intermaterial density	ID
	Kolmogorov statistic	KS
	Lamellar widths	LW
	Mixing variance	MV
Physical measures	Energy usage	EU
	Interface stretch	IS
	Residence time	RT

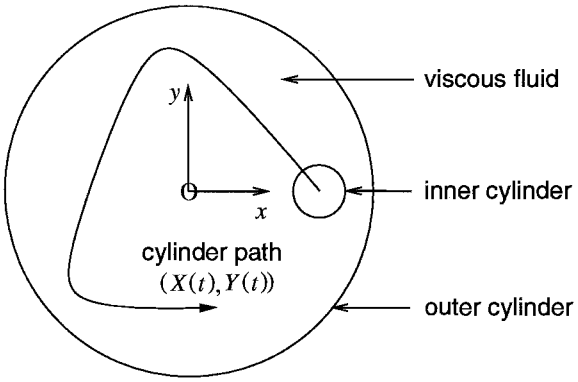

Figure 1. The translating, rotating mixer (TRM) [3]. The coordinate system is chosen such that the origin $(x, y) = (0, 0)$ lies on the outer cylinder axis.

in Section 2, and a set of candidate stirring protocols is constructed by independently varying three of the mixer parameters.

There follows in Section 3 a description of each of the mixing measures that we use. However, it should be emphasised that this paper is not intended as a review of mixing measures, and so these descriptions are necessarily selective; we present sufficient detail to justify the measures and to explain how they are computed in our particular example. Of course, we recognise that the various measures each involve some arbitrariness in their implementation. In order to provide a reference point for later discussion and to preclude confusion due to the abundance of different measures, we summarise the measures used in Table 1. Calculations of the various mixing measures are described in Section 4. Finally, in Section 5, we draw some, necessarily tentative, conclusions for 'optimal' selection of the stirring protocol.

2. Stirring protocols in the translating, rotating mixer

The translating, rotating mixer (TRM) is illustrated in Figure 1. This two-dimensional mixer consists of two parallel circular cylinders, one inside the other. The annular region between the two cylinders is filled with a viscous fluid, and mixing is achieved by translating the inner cylinder through the fluid domain, so mimicking the intuitive method of stirring a tank of fluid using a stick (the outer cylinder is the 'tank', the inner cylinder a 'stick'). Although in principle more complex flow regimes can be generated by additionally allowing the inner and outer cylinders to rotate about their respective axes [3], this additional feature is not considered here, in the interests of limiting the number of parameters to be investigated.

The TRM is a generalisation of the well-studied eccentric annular mixer (EAM), in which the inner cylinder does not translate, but fluid motion is driven only by rotation of the inner and/or outer cylinders about their respective axes [8, 23, 27, 38]. The exact velocity field for the TRM in the Stokes-flow regime has been derived elsewhere [3]; it is readily implemented for numerical simulations, and receives no further discussion here. The availability of an exact expression for the velocity field significantly improves our ability accurately to simulate the mixing, at modest computational expense. The more difficult problem that arises when the velocity field is not known *a priori* does not lend itself to such a detailed study as presented here, with current computing capabilities (but see [4]). A further benefit arising from our knowledge of the exact velocity field is that the power input required to drive the TRM can also be determined exactly [3]; such information is clearly of great importance in an industrial setting, where it provides a constraint on the mixing, and it receives detailed discussion later.

2.1. The 252 stirring protocols

A stirring protocol is represented by the closed curve $(X(t), Y(t))$ along which the axis of the inner cylinder moves, the motion repeating with period T. The protocols that we consider correspond to cycloids of the form

$$(X, Y) = (r_1 \cos(2\pi t') + (0{\cdot}8 - r_1)\cos(2\pi n t'), r_1 \sin(2\pi t') + (0{\cdot}8 - r_1)\sin(2\pi n t'))a_{\text{out}}, \quad (1)$$

where $t' = t/T$, a_{out} is the radius of the outer cylinder, and the parameters r_1 and n can be varied to generate different paths for the inner cylinder. In addition, we consider three different values for the radius ratio $a_{\text{in}}/a_{\text{out}}$, where a_{in} is the radius of the inner cylinder:

$$a_{\text{in}}/a_{\text{out}}: \quad 0{\cdot}05, 0{\cdot}10, 0{\cdot}15. \quad (2)$$

Figure 2 shows the 84 inner-cylinder paths considered in this paper that were obtained by choosing r_1 and n from the following lists:

$$r_1: \quad 0{\cdot}1, 0{\cdot}2, 0{\cdot}3, 0{\cdot}4, 0{\cdot}5, 0{\cdot}6, 0{\cdot}7, \quad (3)$$

$$n: \quad -7, -6, -5, -4, -3, -2, 2, 3, 4, 5, 6, 7. \quad (4)$$

We note here the significance of the parameter n, in that each cycloid has $|n-1|$-th order rotational symmetry. By taking all possible combinations from (2), (3) and (4), we thus generate 252 distinct protocols.

For brevity we use the notation $\mathcal{P}(i, j, k)$ to index the protocols, where $i = 1, 2, 3$ correspond, respectively, to the values of $a_{\text{in}}/a_{\text{out}}$ given in (2), $j = 1, \ldots, 7$ correspond, respectively, to the values of r_1 given in (3) and $k = 1, \ldots, 12$ correspond, respectively, to the values of n given in (4). For example, the protocol with $a_{\text{in}} = 0{\cdot}15a_{\text{out}}$, $r_1 = 0{\cdot}2$ and $n = 4$ is $\mathcal{P}(3, 2, 9)$.

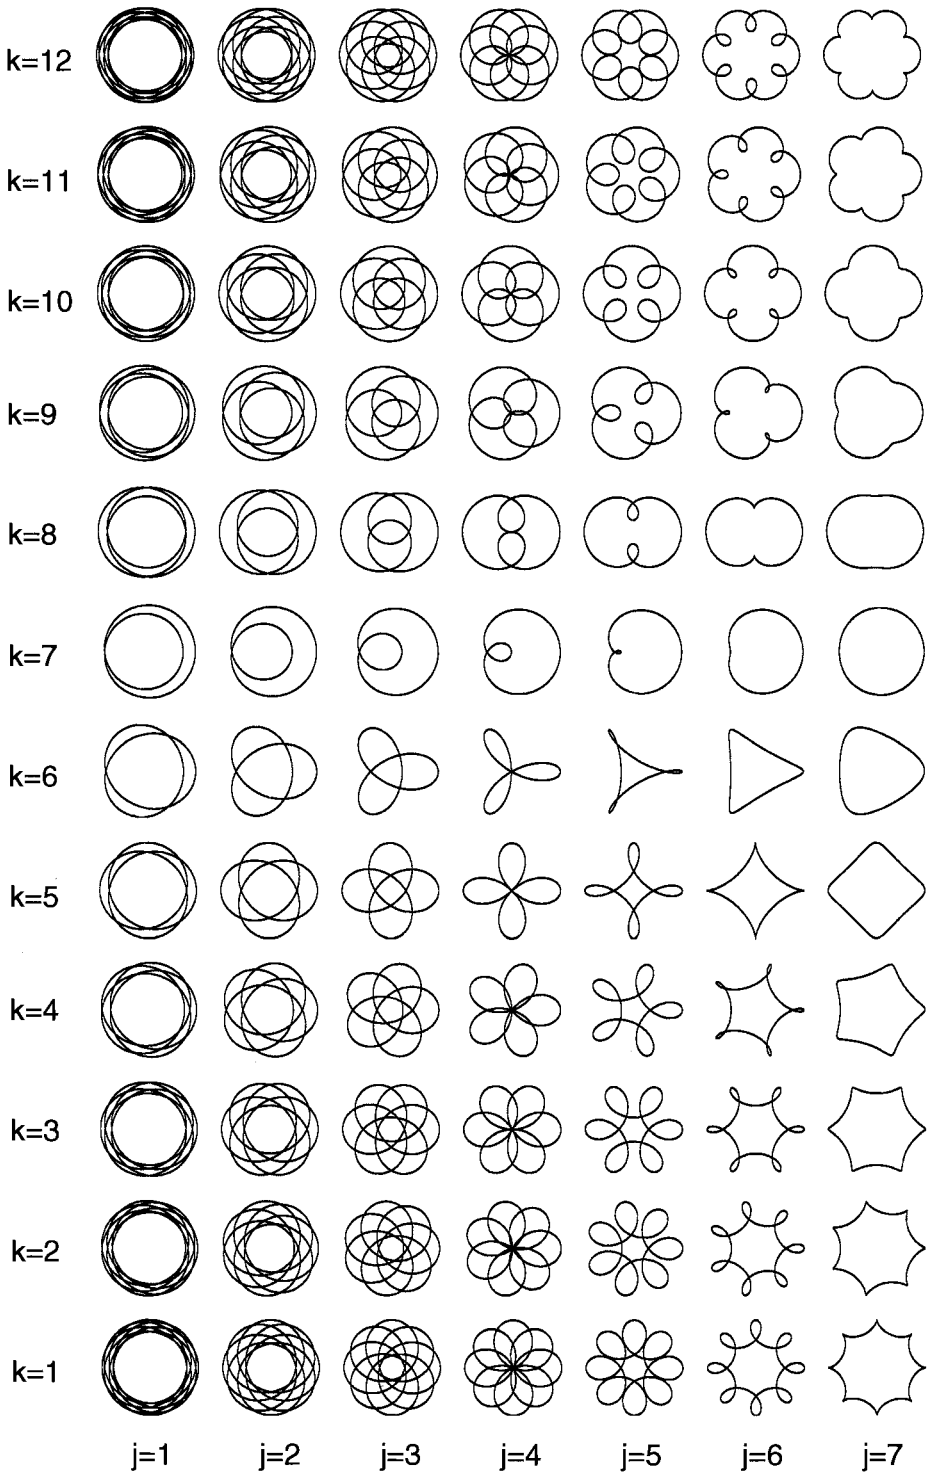

Figure 2. Motion of the inner cylinder axis for the stirring protocols $\mathcal{P}(*, j, k)$. The values of j and k are shown for each protocol. (The value of i is irrelevant for this diagram since it reflects the radius of the inner cylinder, not shown here, and does not influence its path.)

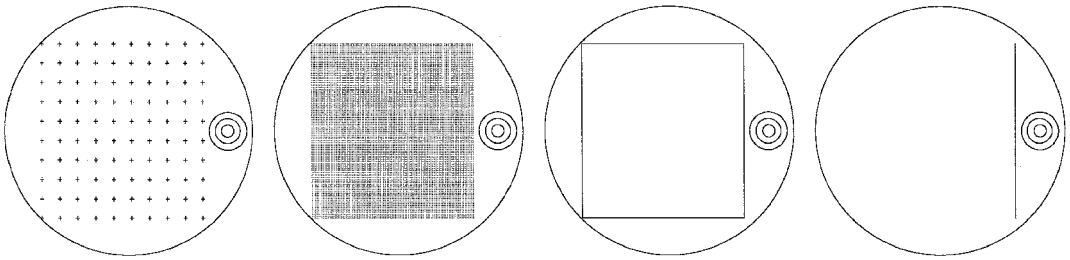

Figure 3. The reference initial conditions used to compute the various mixing measures. The three concentric circles indicate the three inner cylinder radii used in this investigation. From left to right are shown the 10×10 coarse grid, the 100×100 fine grid, the rectangular blob and the linear 'blob'.

3. Mixing measures

In this section we examine some standard mixing measures and describe the ten quantities that will be used to evaluate the mixing achieved by the protocols introduced in Section 2. With two exceptions, the measures are based on tracking the paths of passive, non-diffusive, inertia-less particles, whose positions $(x(t), y(t))$ satisfy

$$\dot{x} = u(x, y, t), \quad \dot{y} = v(x, y, t), \tag{5}$$

where (u, v) is the fluid velocity. (In the two remaining cases, the particles also diffuse, so that their advection according to (5) is supplemented by a random component of motion, described below in Section 3.4.) Depending on the mixing measure, the particles may initially be placed on a regular grid or form the interface to a blob: the four initial conditions that we use are a coarse grid, a fine grid, the interface to a rectangular blob and a linear (and hence infinitely thin) 'blob' (see Figure 3). The coarse grid is a regular 10×10 array of particles, occupying the right rectangle whose lower-left hand corner is at $(-0{\cdot}7a_{\text{out}}, -0{\cdot}7a_{\text{out}})$ and whose upper-right hand corner is at $(0{\cdot}6a_{\text{out}}, 0{\cdot}7a_{\text{out}})$, where the xy-origin is on the axis of the outer cylinder; the fine grid is a regular 100×100 array of particles, occupying the same region. As we explain below, the grid that is adopted depends largely on the computational demands of the particular mixing measure. The rectangular blob interface is defined by placing 100 equally spaced particles along each side of the rectangle used for the coarse and fine grids. Finally, an infinitely thin linear 'blob' is defined by placing 10000 equally spaced particles along the right-hand edge of the same rectangle. The sizes and initial locations of the grids and blobs have been chosen so that they are always caught in the path of the inner cylinder. Clearly other choices could have been made for the precise initial locations of the particles. However, we believe that, while the details might change, the broad thrust of our conclusions would not be altered were different choices to be made.

We now describe each of the mixing measures in turn.

3.1. Dynamical-systems measures

It is convenient, when evaluating the various stirring protocols, to consider instead of the flow (5) the period-one mapping $\mathcal{M}(\mathcal{P}) : (x_0, y_0) \mapsto (x_1, y_1)$, which maps the position (x_0, y_0) of a particle at time $t = 0$ to its position (x_1, y_1) one period later, at time $t = T$. Although the velocity field is known exactly, this mapping requires numerical computation, and is not available in closed form. In view of the periodic nature of the TRM motion, $\mathcal{M}(\mathcal{P})$ also maps

the position of a particle at $t = mT$ to its position at $t = (m + 1)T$, for any integer m. The iterated mapping and Liapunov exponent measures that we describe below are defined in terms of $\mathcal{M}(\mathcal{P})$.

3.1.1. *Iterated mapping measure of chaoticity (IM)*

The simplest and most common dynamical systems tool for evaluating mixing quality is the iterated mapping plot, or Poincaré section [8–11]. This is constructed by taking a small number of tracer particles (here the coarse grid, containing 100 particles) and plotting their positions under repeated application of the mapping $\mathcal{M}$, up to $\mathcal{M}^m$, for some large integer m (here 1023). Regions of poor mixing are revealed as islands of integrable quasi-periodic motion, which surround the elliptic periodic points of $\mathcal{M}$ [27]. Stretching of fluid elements is slow (algebraic in time) in these islands, and no fluid exchange can take place across island boundaries, in the absence of diffusion. By contrast, effective mixing takes place in the chaotic sea, with exponential-in-time separation of initially neighbouring particles. Boundaries between islands and chaotic sea may have an intricate fractal structure, containing infinitely many smaller islands on all spatial scales. This structural complexity gives rise to computational difficulties in distinguishing between genuinely chaotic and merely quasi-periodic particle paths involving long periods. However, at least informally, iterated mappings provide a good visual identification of the degree to which the stirring domain is chaotic.

To systematise the evaluation of the degree of chaoticity in the iterated mapping, we proceed as follows. For a given protocol $\mathcal{P}$ we track the motion of the coarse grid of particles, recording the particle positions under the application of $\mathcal{M}^m$ for $m = 0, \ldots, 1023$ (where $m = 0$ corresponds to the initial locations of the particles). This procedure generates 100 times series, each of the form $\{(x_0, y_0), (x_1, y_1), \ldots, (x_{1023}, y_{1023})\}$, whose periodicity, quasi-periodicity or chaoticity can be determined by taking the discrete Fourier transform (DFT) of the time series of the x- and y-components (see Figure 4). If the DFT appears noisy (*i.e.*, if the peak amplitude in the DFT is comparable in size to the mean amplitude) then we declare the corresponding particle path to be chaotic. If, on the other hand, an amplitude peak exists in the DFT that is significantly greater than the mean (we use a factor of 20 as our threshold ratio) then the corresponding particle path is declared quasi-periodic (exceptionally it may be periodic). Of course, points may be mis-diagnosed by this procedure; in particular, quasi-periodic paths involving very long periods may be mistakenly declared chaotic, but for practical purposes we do not concern ourselves with such fine distinctions.

For a given protocol, we define the iterated mapping mixing measure IM to be the proportion of particles on the coarse grid that give rise to chaotic trajectories, as determined above. This measure provides an estimate of the percentage area of the TRM domain that experiences chaotic advection; high values correspond to good mixing.

3.1.2. *Liapunov exponent (LE)*

In the chaotic sea, neighbouring particles separate exponentially in time, and a corresponding stretching-rate exponent, or Liapunov exponent, can be defined as follows. We first imagine placing an infinitesimal material line element $\mathrm{d}\boldsymbol{x}$ at a point $\boldsymbol{x}$, extending from $\boldsymbol{x}$ to $\boldsymbol{x} + \mathrm{d}\boldsymbol{x}$; then the Liapunov exponent λ at $\boldsymbol{x}$ is

$$\lambda = \lim_{n\to\infty} \limsup_{|\mathrm{d}\boldsymbol{x}|\to 0} \frac{1}{nT} \log \frac{|\mathcal{M}^n(\boldsymbol{x} + \mathrm{d}\boldsymbol{x}) - \mathcal{M}^n(\boldsymbol{x})|}{|\mathrm{d}\boldsymbol{x}|}. \tag{6}$$

Roughly speaking, an infinitesimal material line element of initial length l_0 in a chaotic region will have length $l_n \sim l_0 \exp(\lambda nT)$ at a time nT later. If the flow domain is globally chaotic then

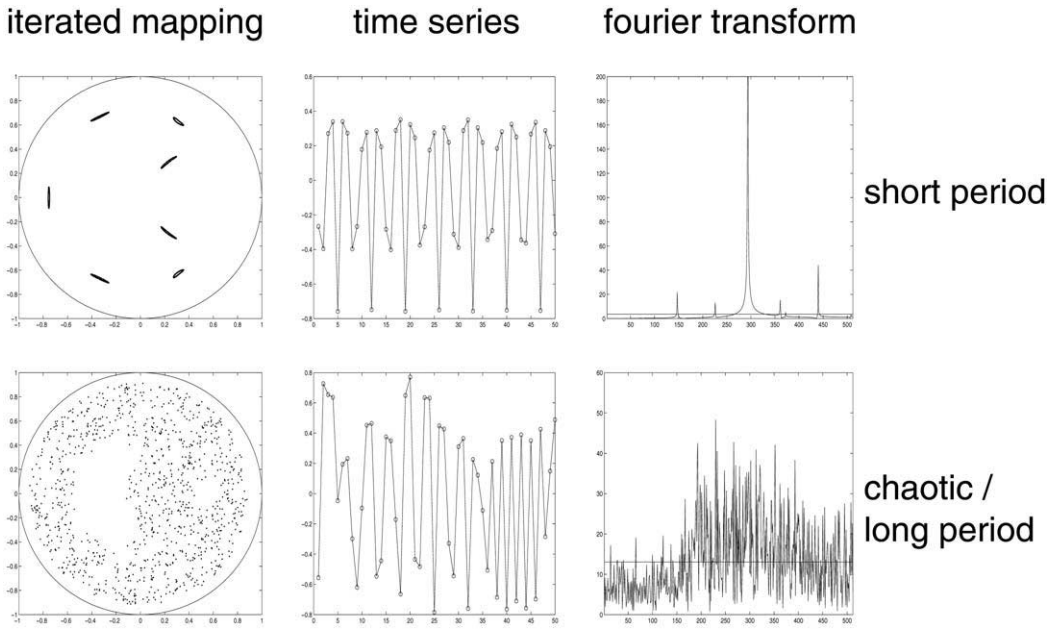

Figure 4. Detecting periodicity using a DFT. In the upper sequence of plots, a quasiperiodic trajectory gives rise to a DFT with well defined, sharp peaks. In the lower sequence, by contrast, an apparently chaotic trajectory yields a full spectrum in the DFT.

ergodicity implies that λ is independent of the starting position $\boldsymbol{x}$ so that a single Liapunov exponent can be used to characterise the protocol. However, in general, we expect a spatial distribution of values of the Liapunov exponent [39].

For the purpose of evaluating our stirring protocols, we compute the *finite-time* Liapunov exponent, obtained from (6) by replacing ∞ with a moderate number N of mappings (here $N \sim 50$), and using a small, but finite, line element (here $|\mathrm{d}\boldsymbol{x}| \sim 10^{-5} a_{\mathrm{out}}$), rather than an infinitesimal one (although in principle the evolution of the latter could be determined by considering the gradient flow). We compute the exponent at all points on the coarse grid and take the mean exponent to be our mixing measure LE; high values of LE correspond to good mixing. Of course, material lines lying within a periodic island experience sub-exponential stretching, and have correspondingly small finite-time Liapunov exponents; this tends to punish those stirring protocols that give rise to large areas of (poorly mixed) islands. Finally, we note that there are other, more sophisticated, related mixing measures, *e.g.*, based on the *distribution* of stretching values, or on Ottino's mixing efficiency [27], which compares local elemental stretching rates against a theoretical upper bound for stretching, but these are not considered here.

3.2. Statistical measures

In contrast to the dynamical systems tools described in Section 3.1, statistical mixing measures are based on the final outcome of the mixing rather than on the details of the evolution of tracer particles in the flow.

3.2.1. *Mixing variance (MV)*
A simple and intuitive statistical mixing measure is the mixing variance. To motivate this measure, we note that perfect mixing would distribute material uniformly over the entire fluid domain, so that in any partition of the domain into regions of equal area, each region would contain an identical amount of material, and thus the variance of the tracer distribution would be precisely zero. Any other distribution of the same material would lead to a non-zero variance, and so the variance of the tracer distribution gives an indication of the mixing quality.

Specifically, we define the mixing variance as follows. For a given protocol $\mathcal{P}$ we take the linear blob as an initial condition and evolve it for ten periods of the stirring protocol, i.e., we apply the mapping $\mathcal{M}^{10}$. Using the rectangular blob to provide the corner points, we then construct a 9×9 array of equal-sized rectangles, and compute the number of particles ultimately found in each of these rectangles. The sample variance of these numbers is calculated, giving the mixing variance measure MV; low values of MV are indicative of good mixing.

3.2.2. *Kolmogorov statistic (KS)*
The Kolmogorov test is an alternative measure of the difference between an actual distribution of stirred tracer particles and a desired (usually uniform) distribution, as, for example, applied by Ling [20] to viscous flow in a rectangular domain stirred by periodic sliding of two of the domain walls. Ling considered rectangles of different sizes, with one corner fixed at some given point in the fluid domain, and recorded the number of particles found in each rectangle. The maximum deviation in this number from that expected under ideal, uniform mixing provided the mixing measure. (Under ideal, uniform mixing, any rectangle should ultimately contain an amount of mixed material in proportion to its area.)

A similar Kolmogorov statistic may be defined for the TRM. As with the MV measure introduced above, we take as initial tracer locations the linear blob shown in Figure 3, and apply $\mathcal{M}^{10}$. We then consider a large number (say, 100) of concentric circles, centred at the TRM origin. For perfectly uniform mixing, a circle of radius r should contain $N_{\text{uniform}} = r^2 N_{\text{total}}/a_{\text{out}}^2$ particles, where $N_{\text{total}} = 10000$ is the total number of particles used for the simulation. (We ignore the presence of the inner cylinder for the sake of a simple calculation.) If the actual number of particles found in a given circle is N_{actual} then the scaled deviation from uniformity is

$$\frac{|N_{\text{uniform}} - N_{\text{actual}}|}{N_{\text{total}}}. \tag{7}$$

The maximum value of this quantity over the set of circles is recorded, and provides the Kolmogorov statistic KS. As for the MV measure, small values of KS correspond to more uniform (*i.e.*, better) mixing.

3.2.3. *Intermaterial interface density (ID)*
When an interface evolves in the chaotic region of the TRM flow, it grows exponentially in length and rapidly becomes highly contorted, by the well known stretching and folding mechanisms of chaotic flows. The way in which the stretched interface is distributed in space is highly non-uniform [26], and the density with which the interface occupies the flow domain varies over many orders of magnitude. Correspondingly, it has been observed [15, 25, 26] that the distribution of striation widths varies over a similarly wide range. This correspondence is natural since the stretching of the interface in one direction is accompanied by a compression

of the striations in another, due to incompressibility of the fluid. (In practice, the production of thin striations leads through diffusion to effective intermingling of the species involved, and hence to good mixing, although, of course, tracer particles do not diffuse in these simulations.) Therefore the uniformity in distribution of a material interface provides a measure of mixing quality. (In keeping with prior nomenclature [26], we refer to the 'intermaterial interface density', rather than explicitly name the measure described below according to its association with the *uniformity* of the interface distribution.)

We define the ID measure as follows. For each protocol, we evolve the linear blob shown in Figure 3 for two periods. We take the same 9×9 array of rectangles as for the MV and KS measures, and compute the length of interface contained within each rectangle after the mixing has taken place. Since chaotic advection causes exponential stretching of the interface, we must be careful to ensure that the interface is at all times adequately resolved, even during just two periods of the flow. To do so, we use the dynamic particle insertion algorithm of Krasnopolskaya *et al.* [15] accurately to resolve the interface into a smooth curve as it is stretched. Our method of computing the length of interface contained within a given rectangle is then as follows: since the interface is well resolved and the dimensions of each rectangle greatly exceed the resolution threshold of the interface, we merely count the number of particles used to represent the interface fragment within each rectangle, after application of the particle insertion algorithm. This gives a reasonably accurate and easy-to-compute approximation to the length of interface contained within the rectangle. The intermaterial interface density ID is then the variance of the distribution of lengths, scaled by the squared mean length of interface in each rectangle. Thus the measure ID is concerned with the uniformity of interface distribution, rather than with its absolute length (for that, see the interface stretch measure below).

3.3. PHYSICAL MEASURES

We now describe some mixing measures which may be categorised as physically based. Indeed, the first of these measures is not concerned with the results of the mixing process, only with the energy input to the device. The remaining two measures described in this section are, however, more directly concerned with mixing quality.

3.3.1. *Energy usage (EU)*

A crucial consideration in large-scale industrial mixers is the energy expended to achieve a certain degree of mixing. Many authors [15, 27] have examined the energetic requirements of laminar mixing devices, sometimes using crude substitutes for energy, such as the distance moved by the boundaries that drive the motion. For the TRM, however, as for some other Stokes-flow mixing devices, an exact expression is available for the power input required to drive the device [3], and so we can compute its energy requirements directly.

If we wish to minimise the energy requirements of the mixer then we need to examine how best to vary the speed of the inner cylinder as it traces out its prescribed path $\mathcal{P}$ in a given time T. The total energy required to drive the mixer can be written as

$$\mathcal{E} = \mu a_{\text{out}}^3 T^{-2} \int_0^T P(s, \dot{s})\, \mathrm{d}t, \tag{8}$$

where μ is the kinematic viscosity of the fluid and P is a dimensionless function related to the power input [3]. Here $s(t)$ denotes arc length along the path $(x, y) = (X(t), Y(t))$ indicated in (1). We wish to choose $s(t)$ to minimise the energy integral (8), subject to the constraints

$$s(0) = 0, \quad s(T) = \Lambda, \tag{9}$$

where Λ is the inner-cylinder path length

$$\Lambda \equiv \oint_{\mathcal{P}} \mathrm{d}s.$$

Since the integrand in (8) is autonomous, it follows that $\mathcal{E}$ is minimised when the Beltrami identity

$$P(s, \dot{s}) - \dot{s}\frac{\partial P(s, \dot{s})}{\partial \dot{s}} = -\kappa \tag{10}$$

is satisfied for some constant κ. For Stokes-flow, the power function can be written in the form $P(s, \dot{s}) = P(s)\dot{s}^2$, where $P(s)$ is a known function [3], and thus (10) reduces to

$$P(s)\dot{s}^2 = \kappa, \tag{11}$$

so that the motion must be carried out in such a way as to maintain a constant power input to the device. In order to satisfy the constraints in (9), κ must take the value

$$\kappa = \frac{1}{T^2}\left(\int_0^T \sqrt{P(s, \dot{s})}\,\mathrm{d}t\right)^2 = \frac{1}{T^2}\left(\int_0^\Lambda \sqrt{P(s)}\,\mathrm{d}s\right)^2,$$

and hence the speed of the inner cylinder for minimum-energy motion is

$$\dot{s} = \frac{\displaystyle\int_0^\Lambda \sqrt{P(s)}\,\mathrm{d}s}{\sqrt{P(s)}\,T}.$$

Finally, from (8) the minimum energy usage measure EU is given by

$$\mathrm{EU} = \min \mathcal{E} = \mu a_{\mathrm{out}}^3 \kappa T^{-1}. \tag{12}$$

While EU is not itself directly a measure of the mixing quality, it is useful to correlate the energy input with genuine mixing measures, to provide an informal measure of mixing efficiency.

3.3.2. *Interface stretch (IS)*

An obvious physical manifestation of effective mixing is a rapid growth in the length of a finite material interface placed in the flow. The interface stretch exponent λ_{is} quantifies the exponential rate of growth of such an interface as it is stretched by the flow [15, 29]: an interface of initial length l_0 has length $l_n \sim l_0 \exp(\lambda_{\mathrm{is}} nT)$ at time nT. This quantity is related conceptually to the Liapunov exponent, although the latter provides a measure of the *local* stretch, whereas the interface stretch exponent gives a global stretching measure [40, 41]. By its nature, the interface stretch exponent gives greater weight to regions of high stretch than does the Liapunov exponent (because regions of locally high stretch preferentially generate new interface precisely in regions of high stretch), and so $\lambda_{\mathrm{is}} > \lambda$ [41]. We define the interface

stretch measure IS to be the final perimeter of the linear blob after two periods of the flow. As in our calculation of ID, we are careful to maintain adequate resolution of the growing interface.

3.3.3. *Lamellar width variance (LW)*

In many applications, mixing is accompanied by chemical reaction between the various species. At present, even on a powerful workstation it remains a significant computational challenge to resolve simultaneous chemical reaction, advection and diffusion in a two-dimensional chaotic flow. However, such simulations may be necessary in order accurately to predict the yield of a multi-stage reaction, since it is known that fine details of the reactant distribution can have a surprisingly great influence on the final yield [2, 31, 42]. While an accurate simulation of reaction–advection–diffusion in the TRM is beyond the scope of this article, we recognise that the statistics of the length scales that develop will influence the progress and the yield of any multi-stage reaction that takes place in the flow, and that the details may be highly specific to individual reaction schemes. Thus we introduce a mixing measure based on the distribution of lamellar widths, defined as follows. We imagine the linear blob to be the interface between two chemical reactants. This interface is tracked for two periods of the mixer, interpolating where necessary as with IS and ID [15]. In the course of the mixing, a large number of highly contorted striations are generated, and we measure their widths by calculating the intersection points of the interface with the y-axis. The distance between successive intersections gives an estimate of the lamellar widths generated by the flow. Our numerical code uses cubic interpolation [43] to generate accurate values for the intersection points between the interface and the y-axis. Once the lamellar widths have been found, we compute the variance of their distribution, scaled by the square of the mean lamellar width, to give the mixing measure LW. (Whether one considers good mixing to correspond to high or low values of LW may depend on the details of the chemical reactions imagined to take place in the flow.)

Note that we use a straight line segment as the initial interface, whereas a closed curve should be used to compute genuine lamellar widths. However, we find that this consideration makes little difference to the LW measure, scaled as indicated here; furthermore, any influence of the nature of the initial interface is mitigated by our later presentation of the results as *rankings* of the various measures.

3.4. Effects of diffusion

Thus far we have used the term ‘mixing’ to describe what is essentially a process of pure stirring (*i.e.*, mechanical stretching of an interface in a laminar flow), but it is worthwhile, from the point of view of practical application, to consider the combined effects of stirring and particle diffusion. Accordingly, our final two mixing measures, described below, concern diffusive tracer particles.

3.4.1. *Residence time (RT)*

In a chemically reactive system, the reaction may take place not only in the bulk of the flow, but also through boundary effects. One source of these, appropriate in some applications, is diffusion through the boundary (*e.g.*, in models for kidney dialysis or heat exchangers). A second source is a chemical reaction, often catalysed, at the boundary that removes chemicals from the system. In either case, the action of mixing is crucial in continually introducing new fluid to the neighbourhood of the boundary, and thereby significantly enhancing these removal effects. We therefore introduce a mixing measure related to the residence time; the

more effective the mixing, the shorter will be the average residence time of a tracer particle. Of course, this measure should not be confused with the more common usage of the term 'residence time' to indicate the length of a particle's sojourn in a continuous-throughput mixing device [44].

For this mixing measure, we suppose that the tracer particles, in addition to being advected by the flow, also diffuse, with diffusivity D, and consider diffusion through either wall to be the sole mechanism of removal (*i.e.*, there is no chemical reaction there). Then the advection equations (5) are replaced by the stochastic differential equations

$$\mathrm{d}x = u(x, y, t)\,\mathrm{d}t + \mathrm{d}\mathcal{N}_x, \quad \mathrm{d}y = v(x, y, t)\,\mathrm{d}t + \mathrm{d}\mathcal{N}_y, \tag{13}$$

where $\mathrm{d}\mathcal{N}_x$ and $\mathrm{d}\mathcal{N}_y$ are independent, identically distributed, gaussian random variables with mean zero and variance $2D\,\mathrm{d}t$.

To model diffusion of tracer particles through the TRM walls, we solve (13), at each time step removing from the numerical simulations any particle that happens to have diffused through a wall. We then adopt a straightforward definition of the residence time mixing measure RT by recording the fraction of particles remaining in the TRM at the end of a given time (here two periods of the flow). We use the fine grid of points as the initial condition and, in order to obtain results that significantly distinguish between the various protocols, we somewhat arbitrarily take the diffusivity to be $D = 10^{-3}a_{\mathrm{out}}^2 T^{-1}$.

Of course, this is a rather crude measure, and other, more sophisticated, versions of the residence time measure can easily be conceived. For instance, in practice, the residence time of a diffusive particle initially placed at some given spatial location will depend on the precise details of the random component of its path, so it may be of interest to consider the distribution of residence times for a set of particles initially placed at the same point in the flow domain. Similarly, the mean residence time of diffusive particles will have some spatial distribution, which could readily be computed. It may also be of interest to examine whether there are regions of the boundary from which particles are preferentially removed from the flow domain.

3.4.2. *Diffusive deviation (DD)*

Our final mixing measure serves to quantify the extent to which the mechanical action of stirring enhances the intermingling of species brought about by diffusion. To motivate the measure we begin by noting that in the absence of diffusion the motion of tracer particles is time-reversible: if the mixer is first run forwards and then the boundary motion is reversed, any tracer particle will return precisely to its initial location. However, when $D \neq 0$ this return is not exact, and the particle will ultimately lie some distance from its initial location. Furthermore, in a chaotic flow, this deviation can be significantly greater than that predicted by simple diffusion alone, because the sensitivity of the chaotic flow to initial conditions enhances the diffusion [45, 46].

To compute the diffusive deviation mixing measure, we allow each particle in the fine grid of points to evolve according to (13), with $D = 10^{-5}a_{\mathrm{out}}^2 T^{-1}$, for two periods of the flow forwards in time, followed by two periods backwards in time. The particles diffuse during both forwards and backwards motions of the boundaries. We then calculate the mean distance of each particle from its initial position, giving the mean diffusive deviation mixing measure DD. (The value of D used here is smaller than for the RT measure purely for computational convenience – we find that for this reduced value of the diffusivity we need to take no special precautions to prevent particles erroneously escaping from the mixer.)

A related measure was developed by Aref and Jones [23], who demonstrated for flow in the eccentric annular mixer the interaction between diffusion and chaotic advection. They took an initially circular blob of particles, evolved it for a given number of periods of their flow, then reversed the motion. In the absence of diffusion, all particles returned to their initial locations inside the blob. However, with non-zero diffusion, the return percentage was diminished, and significantly so in chaotic flow, even for small values of the diffusion coefficient. The tremendous enhancement to the diffusion was due to the exponential growth of the blob boundary during the flow, and consequently greatly enhanced diffusive particle escape from the blob. Although we have taken only a single value for the diffusion coefficient D, and a systematic study of the effects of varying this coefficient is clearly warranted, one might reasonably infer likely trends as D is varied from the work of Aref and Jones [23].

4. Results

In the previous section we defined ten mixing measures that can be evaluated numerically without any user intervention (see Table 1 for an aide-memoire of the various measures). We have computed each of these mixing measures for the 252 protocols described in Section 2. Clearly, there is a degree of arbitrariness in some of the measures, largely due to a particular choice of parameters (such as the number of periods of the flow over which to calculate the measure, initial conditions, or the particle diffusivity, where this is non-zero). Furthermore, for certain of the measures we have chosen a crude, but simple, version to calculate, despite the availability of more sophisticated versions. While the details of our results may change if the measures are defined slightly differently, we expect our broad conclusions to be qualitatively unaffected.

Before describing our results, we make some preliminary comments concerning correlations between the mixing measures, and about energy usage. We then move on to describe general trends in our results, before highlighting observations regarding specific protocols and measures.

4.1. Correlations between mixing measures

In an attempt to find 'best' protocols, it is useful to understand the extent to which the various mixing measures are pairwise correlated. Since the measures have very different distributions (some, like RT, vary only slightly across the range of protocols, while others, like IS, vary by many orders of magnitude, for instance), we evaluate their degree of correlation using a distribution-free test based on the respective rankings of the 252 protocols. We use the straightforward Spearman rank correlation coefficient [47], as follows. To calculate the rank correlation coefficient between two mixing measures, we first independently rank the set of 252 protocols according to each of the two measures. Thus the i-th protocol is assigned integers m_i and n_i that are its respective rankings according to the two measures (the worst protocol having rank 1, the best protocol rank 252, in this case). Exceptionally, protocols may have precisely the same value of the mixing measure, in which case they are ranked equally. Using the two sets of rankings, we calculate the Spearman rank correlation coefficient

$$C = \frac{\sum (m_i - \bar{m})(n_i - \bar{n})}{\sqrt{\sum (m_i - \bar{m})^2 \sum (n_i - \bar{n})^2}}, \tag{14}$$

where $\bar{m} = \sum m_i/252 = 253/2$ and $\bar{n} = \sum n_i/252 = 253/2$. By construction, $-1 \leq C \leq 1$. Large values of $|C|$ indicate that the two measures in question are well correlated, while small values of $|C|$ indicate that the measures concerned do not agree on their assessment of 'good' and 'bad' protocols.

The occurrence of a strong correlation between a pair of measures is of particular interest when the measures concerned have different computational requirements (*e.g.*, correlations between IS and EU, which are expensive and cheap to compute, respectively). If a good correlation exists, we might attempt to use the computationally cheaper measure as a substitute for the more expensive measure, thereby saving computational effort when evaluating a given set of protocols.

Equally, it is instructive to be aware of the existence of poor correlations, since these imply that protocols may independently have desirable characteristics according to different mixing measures. A consequence concerns the idea of using a 'sieve' [36] systematically to eliminate poor protocols, through the application of a succession of mixing measures, each more sophisticated and computationally intensive than the last. If two mixing measures used at different stages in the sieving process are poorly correlated, this may lead to the premature elimination of a protocol that would in fact score well with the later measures, simply because it did not also score well with some prior, less sophisticated measure, uncorrelated with the measure of true interest. A further shortcoming of the sieving approach is that it may eliminate protocols that perform well overall, if they do so only in some average sense.

We note that, although the use of a *rank* correlation coefficient allows us easily to compare measures with very different distributions, it necessarily leads to results that depend on the set of protocols chosen for study. Our hope (and expectation) is that the set of protocols here is sufficiently large, and the protocols sufficiently varied, that our general conclusions will hold notwithstanding this feature of the rank ordering.

4.2. Energy usage

In order to appreciate the role of energy in our considerations, we note that for a given TRM protocol, operating in the Stokes-flow regime, and mixing non-diffusive particles, the outcome depends only upon the path taken by the inner cylinder, and not upon the speed of operation of the device. However, the power input to the TRM scales with the square of the speed of the inner cylinder, and hence the *energy* requirement for execution of a given protocol scales in proportion to this speed (since the period of operation varies in inverse proportion to the speed). In other words, the overall energy requirement varies in inverse proportion to the period of time over which the mixing takes place. A further consequence of operation in the Stokes-flow regime is that if we fix an identical period T for all stirring protocols then those with greatest inner-cylinder path length Λ will tend to require the greatest energy input $\mathcal{E}$ per period, simply because they have the greatest average speed of the inner cylinder. In fact, with T fixed in this way, the energy required per period is roughly proportional to Λ^2. However, as Figure 5 demonstrates, there is some spread about this broad trend. Figure 5 also shows that in general the energy requirements of the protocols $\mathcal{P}(3, *, *)$ exceed those of the protocols $\mathcal{P}(2, *, *)$, which in turn exceed those of $\mathcal{P}(1, *, *)$. The reason for this ordering is straightforward: the protocols fix the path of the inner cylinder axis, and the two cylinders approach one another more closely when the inner cylinder radius is greater, thereby incurring greater penalties in terms of viscous drag.

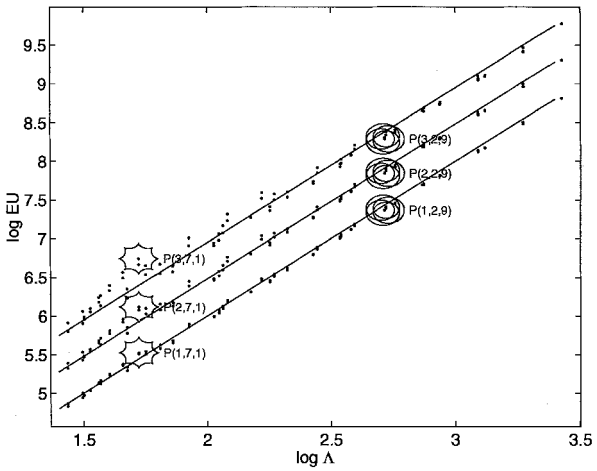

Figure 5. Logarithmic plot of mixing measure EU, from (12), against inner-cylinder path length Λ for each of the 252 stirring protocols. Trend lines with slope 2 are also shown.

4.3. GENERAL TRENDS

Our results, described in detail below, lead us to the conclusion that it is not possible to choose a single 'best' protocol from the 252 protocols examined, since the various measures do not agree in their relative rankings of these protocols. Although there is a strong correlation between some pairs of measures, other pairs are only weakly correlated. However, it is possible to make some general comments, as follow below. In this subsection we shall discount EU (which does not genuinely measure the quality of mixing, although it is important in any discussions of mixing efficiency) and LW (which is only poorly correlated with all other measures; it is also unclear whether large or small values of this quantity are desirable until an application is specified). We therefore discuss here only the remaining eight mixing measures.

Rather arbitrarily we have chosen to focus on the twenty best protocols for each mixing measure. For DD, ID and RT, at least eighteen of the top twenty protocols have $2 \leq j \leq 4$ and $1 \leq k \leq 3$. A similar, but slightly less marked, preference for protocols in this region of parameter space is exhibited by KS, LE and MV. With IM there is a greater spread among top-twenty protocols, although all have $2 \leq j \leq 6$. These results indicate a strong preference for what might be termed 'flowery' protocols. A simple explanation for this is that the flowery protocols tend to have the longest path length Λ and therefore they do more stirring during one complete cycle of the protocol than do other protocols. However, upon further reflection, we see the true story is rather more subtle and that in fact the strong preference for protocols with $k \leq 4$ is intriguing because (see Figure 2) to the eye these protocols look similar to those with $k \geq 10$, whereas the latter protocols do not perform as well as the former (*e.g.*, $\mathcal{P}(*, 3, 12)$ seem as convoluted as, say, $\mathcal{P}(*, 3, 2)$, but are not ranked as highly). In contrast, the top twenty for the final mixing measure, IS, stand alone in including a few protocols from the top half of Figure 2.

Some resolution to this puzzle can be obtained by considering Figure 6, which shows a qualitative picture in the form of iterated mapping plots for protocols of the form $\mathcal{P}(1, *, *)$ (i.e., those with cylinder radius ratio $a_{\text{in}}/a_{\text{out}} = 0{\cdot}05$). From this figure it is clear that the 'best' protocols, with $2 \leq j \leq 4$ and $1 \leq k \leq 3$, have much smaller periodic islands (and therefore exhibit better mixing) than corresponding protocols with $k \geq 10$. Since all measures (except

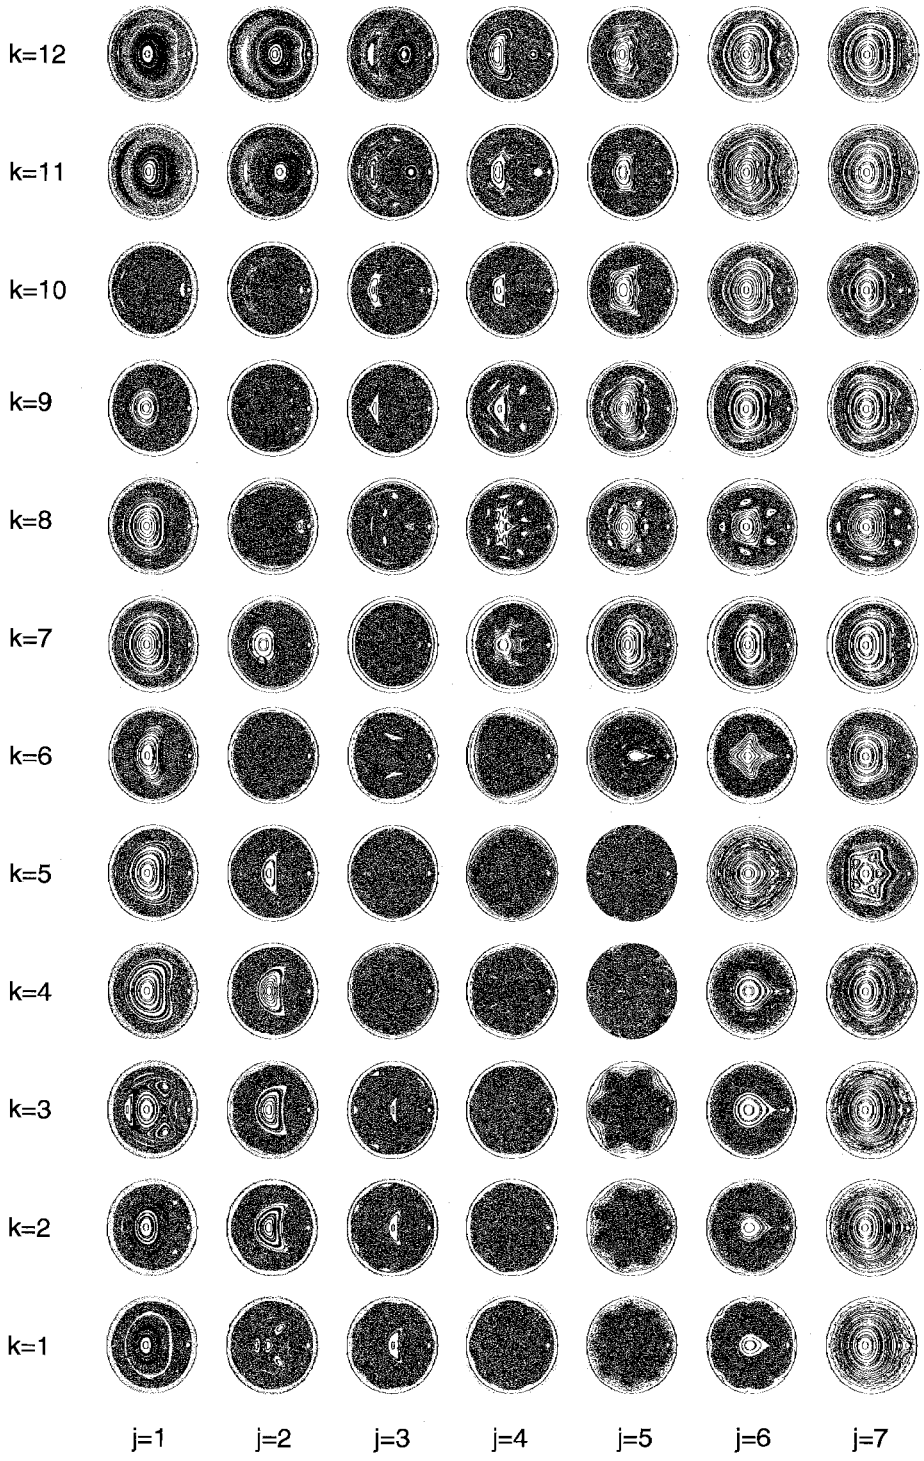

Figure 6. Iterated mapping plots for protocols $\mathcal{P}(1, *, *)$. This figure may be compared directly with Figure 2.

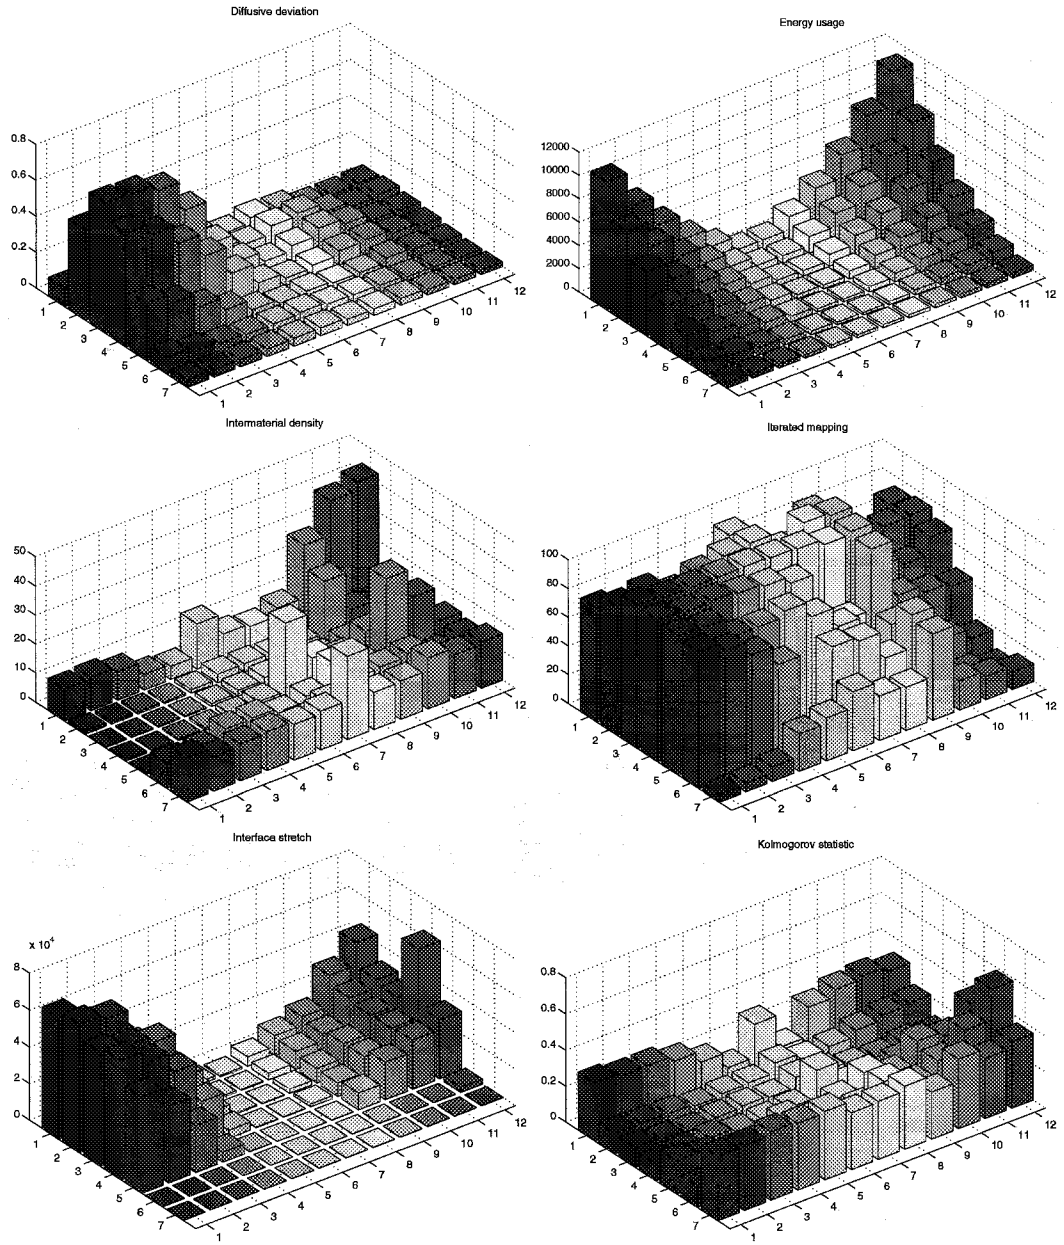

Figure 7. Plots of the 10 mixing measures for each of the protocols $\mathcal{P}(2, *, *)$.

IS) involve some averaging over the fluid domain, the islands exact a penalty on the $k \geq 10$ protocols, which is why they do not appear in the top-twenty lists. The iterated mapping plots also indicate why IS does not discriminate against the $k \geq 10$ protocols: the initial line whose stretch is calculated lies entirely in the chaotic region, and so the existence of significant regular regions does not materially affect this mixing measure.

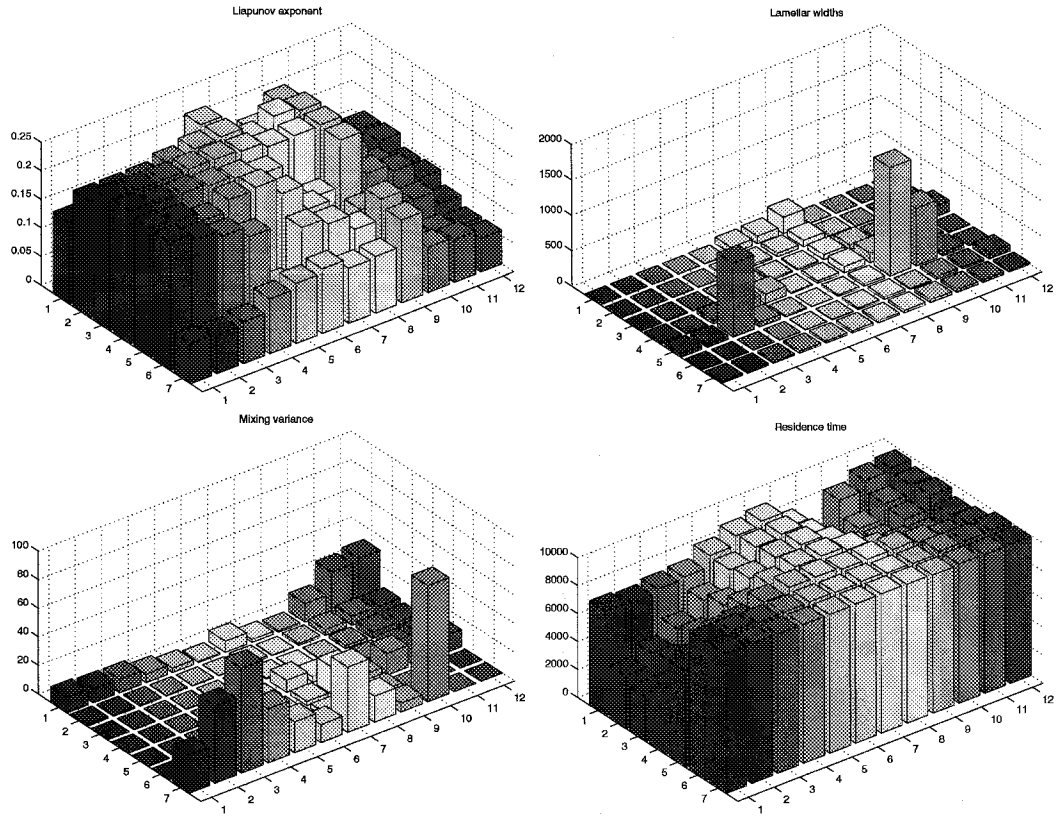

Figure 7. Continued.

4.4. DETAILED RESULTS

Figure 7 shows in detail data for the individual mixing measures, for the protocols $\mathcal{P}(2, *, *)$. As might be expected from Figure 6, many mixing measures show wide variation between the various protocols. For some measures, such as IS, there is a clear clustering of good protocols, while for others, such as MV, the good protocols are more isolated and scattered.

Plots of the pairwise joint distributions of the ranked mixing measures are shown in Figure 8; the corresponding rank correlation coefficients can be found in Table 2. To understand Figure 8, consider as an example the plot in column DD and row EU. This plot contains 252 data points, with coordinates (m_i, n_i), for $i = 1, \ldots, 252$. The m_i and n_i are the rankings of the i-th protocol according to the mixing measures DD and EU, respectively. The remaining plots in the figure are defined similarly. More detailed versions of these plots are shown in Figures 9, 10 and 11, which, in addition, illustrate the stirring protocols that correspond to the various data points in Figure 8.

In the present study, the strongest correlations ($|C| > 0{\cdot}9$), evident in Figure 8 and Table 2, are found between iterated mapping chaoticity, Liapunov exponent and mixing variance mixing measures. Correlation plots illustrating the protocols responsible for some of the data points are shown in Figure 9. MV is significantly simpler to compute than IM or LE, in terms of storage and time, respectively. Thus, for the protocols investigated here, one might reasonably substitute MV for either IM or LE, depending on one's computational constraints. These strong correlations accord with intuition, since each measure heavily penalises those protocols

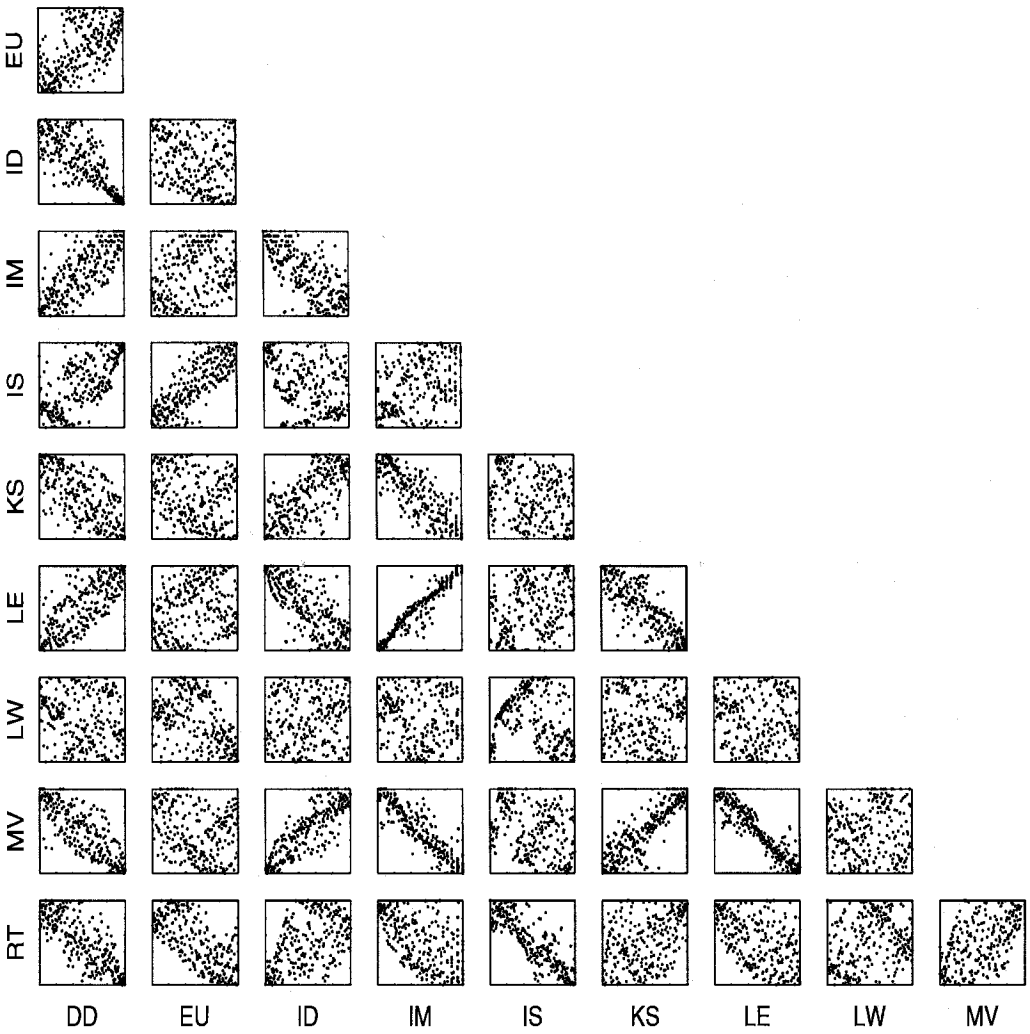

Figure 8. Plots of the joint distributions of the ranked mixing measures.

with large regular regions. We emphasise our earlier point that superficially similar protocols, *e.g.*, $\mathcal{P}(2, 1, 7)$ and $\mathcal{P}(3, 2, 7)$, may lie near opposite ends of the spectrum of mixing quality (see Figure 9).

The next strongest correlations ($0{\cdot}8 < |C| < 0{\cdot}9$) exist between (MV,KS), (LE,KS), (MV,DD), (IS,EU), (RT,DD), (RT,IS), (MV,ID) and (KS,IM). When we consider these pairings (see Figure 10), we note that there is generally a clustering of 'like' protocols (*e.g.*, those we have described above as 'flowery'). However, in some cases there is a sensitive dependence of results on the exact protocol. For example, in the (IS,EU) plot, long flowery protocols, such as $\mathcal{P}(1, 6, 12)$, $\mathcal{P}(1, 4, 3)$ and $\mathcal{P}(1, 6, 11)$ clearly perform well. Clustering of like protocols is also evident in the (RT,IS) plot, where simple protocols like $\mathcal{P}(3, 7, 1)$ perform poorly while broader, flowery protocols such as $\mathcal{P}(2, 2, 3)$ perform better. In the (RT,LW) plot similar clustering is evident. In the (KS,IM) plot the simple protocols exhibit a greater range of mixing quality, and greater sensitivity to protocol details is apparent. For example, the similar protocols $\mathcal{P}(3, 2, 7)$ and $\mathcal{P}(3, 1, 7)$ achieve good and poor mixing, respectively.

Table 2. Correlation coefficient C between pairs of mixing measures.

	DD	EU	ID	IM	IS	KS	LE	LW	MV	RT
DD	1·00	0·64	−0·78	0·79	0·68	−0·67	0·77	−0·18	−0·80	−0·80
EU	0·64	1·00	−0·31	0·37	0·81	−0·25	0·24	−0·44	−0·27	−0·70
ID	−0·78	−0·31	1·00	−0·74	−0·37	0·72	−0·75	0·15	0·84	0·52
IM	0·79	0·37	−0·74	1·00	0·40	−0·83	0·95	−0·12	−0·90	−0·57
IS	0·68	0·81	−0·37	0·40	1·00	−0·26	0·32	−0·23	−0·36	−0·85
KS	−0·67	−0·25	0·72	−0·83	−0·26	1·00	−0·84	0·06	0·87	0·42
LE	0·77	0·24	−0·75	0·95	0·32	−0·84	1·00	−0·05	−0·93	−0·52
LW	−0·18	−0·44	0·15	−0·12	−0·23	0·06	−0·05	1·00	0·04	0·27
MV	−0·80	−0·27	0·84	−0·90	−0·36	0·87	−0·93	0·04	1·00	0·56
RT	−0·80	−0·70	0·52	−0·57	−0·85	0·42	−0·52	0·27	0·56	1·00

In a previous paper [3] we observed from a sample of three protocols the counter-intuitive result that patterns producing highly chaotic iterated mappings tended to generate relatively poor interface stretch (and vice versa). However, the corresponding joint distribution plot of the iterated mapping and interface stretch measures for the larger sample of 252 protocols shown in Figure 11 suggests this is not always the case. It can be seen that, while simpler patterns generally produce a poor interface stretch, the degree of chaoticity of the iterated mapping can depend sensitively upon the pattern shape. A simple thought experiment shows why we should not expect interface stretching necessarily to provide much information about uniformity of mixing. Consider the TRM in the limit as the scaled inner cylinder radius $a_{\text{in}}/a_{\text{out}}$ becomes small (imagine using a wire rather than a stick to stir a tin of paint): although the inner cylinder may still be dragged repeatedly through a material line and thereby stretch it considerably, the thin cylinder transports little of the bulk of the fluid. Thus we should expect a relatively poor correlation between IS and other measures of mixing uniformity such as IM, MV and KS.

As well as showing sensitive dependence upon protocol parameters, some pairs of measures demonstrate a poor correlation ($|C| \leq 0{\cdot}4$) – see Figure 11 for examples. Notably, the energy input, EU, does not correlate well with ID, IM, KS, LE or MV. This indicates that low-energy protocols may exist that perform well according to these mixing measures. For example, both the $\mathcal{P}(1, 2, 6)$ and $\mathcal{P}(1, 5, 5)$ protocols have a desirable low mixing variance and also require relatively little energy input. Desirable low-energy protocols for IM, KS and LE could be $\mathcal{P}(1, 3, 6)$, $\mathcal{P}(1, 4, 6)$ and $\mathcal{P}(1, 5, 5)$, respectively.

The relatively poor correlation between interface stretch and Liapunov exponent ($C = 0{\cdot}32$) may seem counter-intuitive. However, this is an artefact of the crude manner in which we compute the average Liapunov exponent: we do not distinguish between points in the chaotic region (with positive Liapunov exponent) and points in a regular region (for which the finite-time Liapunov exponent is small). This technique tends to depress the computed average Liapunov exponent for flows with a significant area occupied by regular regions, whereas the interface chosen for the stretch calculation lies (by design) largely in the chaotic region, regardless of the particular protocol. In a more sophisticated implementation of our code, it would, of course, be possible to average over only points designated as lying in the

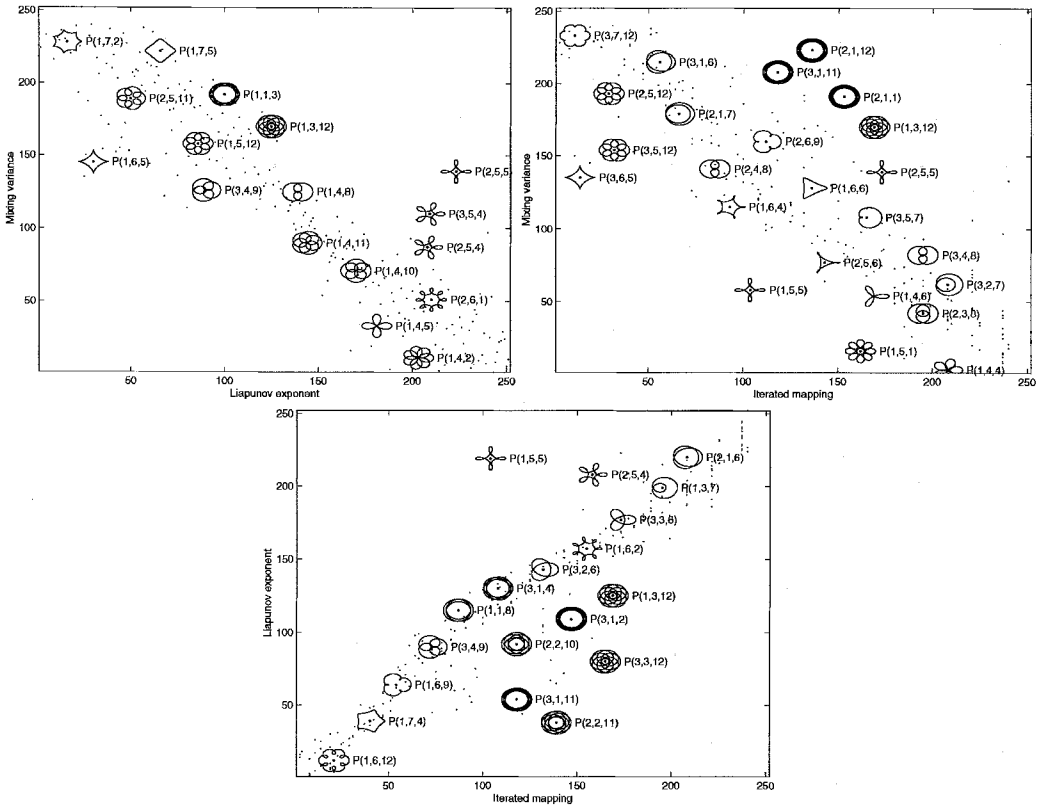

Figure 9. Joint distribution plots for strongly correlated measures, illustrating some of the protocols that correspond to various data points.

chaotic region (using techniques from the iterated mapping chaoticity measure to identify these points).

5. Conclusions

We have described above the numerical evaluation of a range of mixing measures for a large number of stirring protocols in the translating, rotating mixer. Before attempting to extrapolate from these results to the problem of protocol selection, we make some general comments.

First, we note that the preference for protocols with $2 \leq j \leq 4$ and $1 \leq k \leq 3$ suggests a certain robustness to the exact details of the protocol. This robustness is of particular practical importance because in deriving the velocity field for our simulations we have made the Stokes-flow approximation, effectively setting the Reynolds number R of the flow to be zero. Of course, in any real experiment performed at small, but finite, Reynolds number, corrections to this velocity field of $O(R)$ will be required. The robustness of the mixing quality measures to protocol choice suggests (but obviously does not guarantee) a similar insensitivity to finite-Reynolds-number effects. Experiments are underway to investigate whether this proves to be the case (*cf.* [9]).

The measures adopted in this paper were selected because they allow for automated computation (and there is inevitably a good deal of arbitrariness, for example, in our sample of initial tracer locations). However, we should bear in mind that many other measures are

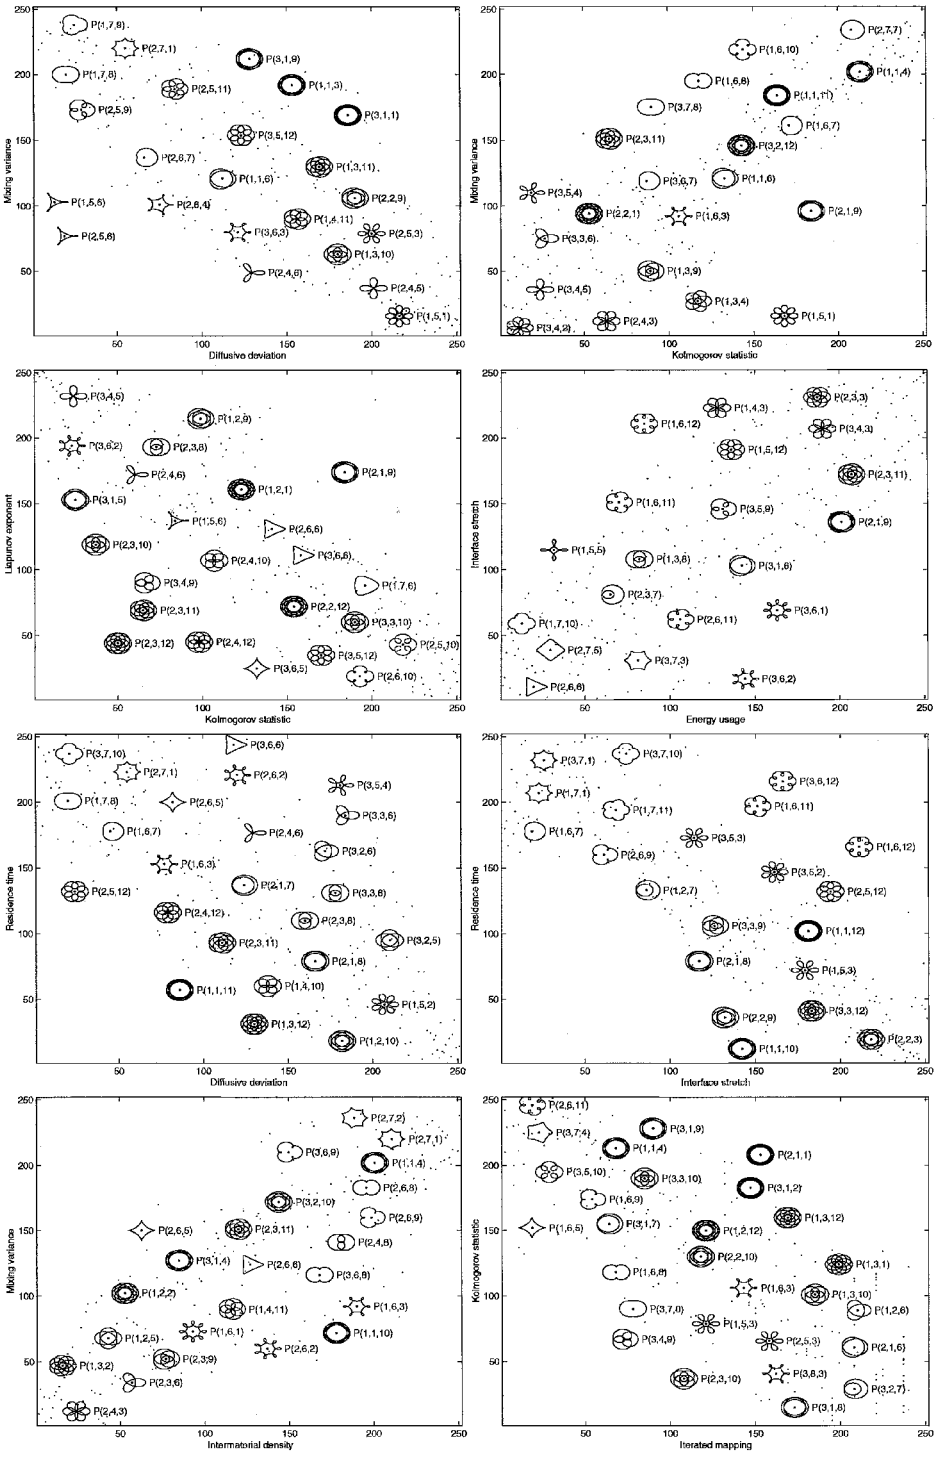

Figure 10. Joint distribution plots for well correlated measures, illustrating some of the protocols that correspond to various data points.

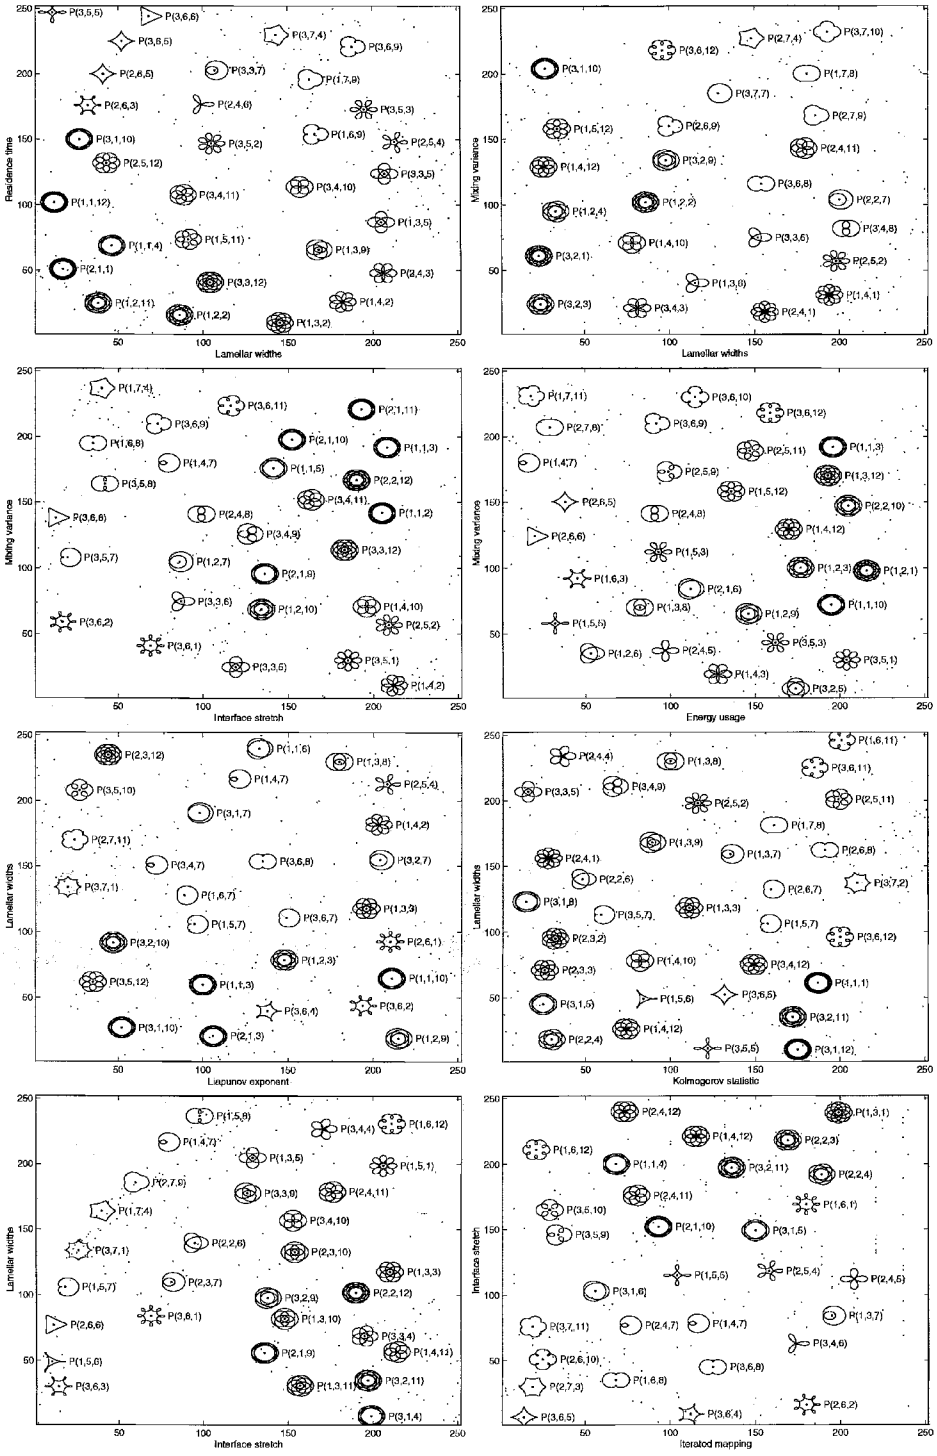

Figure 11. Joint distribution plots for weakly correlated measures, illustrating some of the protocols that correspond to various data points.

inherently analytical, requiring significant human intervention. For example, a major class of measures ignored here concerns the location of periodic points and an understanding of the structure of their associated invariant manifolds [27]. Although periodic points can be found systematically, the requisite numerical algorithms [48] are computationally intensive, and extremely good initial guesses for their locations are sometimes required. Elliptic fixed points, which are found at the centre of periodic islands, are comparatively easy to discern (provided the island is visible to the eye). Hyperbolic points, on the other hand, reside in the chaotic sea and are more difficult to identify. Furthermore, although protocol symmetries can be exploited to find periodic points [12–14, 16, 17], this process often requires significant human intuition, and even then is not guaranteed to locate all such points; it is thus difficult to program robustly. It would, nevertheless, be instructive to investigate the extent to which the automated computational mixing measures contemplated here agree with human analyses and intuitions about real-life stirring devices.

It would also be worthwhile to extend the investigation begun here to three-dimensional mixing flows, either in a batch mixer or in a continuous-throughput device. Some of the mixing measures used here may be applied directly to either case (for example, Liapunov-exponent and energy-usage measures – although in the continuous-flow case the appropriate measure might be the pressure drop); others can be extended by making straightforward changes to reflect the greater geometrical complexity of three space dimensions (for example, intermaterial density and interface stretch measures apply to surfaces rather than to curves). There is a somewhat subtle difference between a time-periodically operated two-dimensional mixer, such as described here, and a continuous-flow device, in which the fluid becomes mixed as it passes down a pipe containing some spatial structure with axial wavelength L (see, for example, [49]). There the analogue of the period-one map $\mathcal{M}$ is the map taking particle positions in the cross-section at $\zeta = 0$ to their corresponding cross-sectional locations at $\zeta = L$, say, where ζ is the axial coordinate.

Finally, we note that, from the point of view of protocol selection, it would also be useful to study the inverse problem to that studied here, *i.e.*, rather than to determine which of a given set of stirring protocols performs best, according to a given mixing measures, to design a protocol to maximise a target mixing measure, subject to constraints on others (such as energy usage). However, such an undertaking is beyond our present scope.

5.1. Implications for protocol selection

We now summarise some of our results as they influence the selection of 'optimal' stirring protocols.

First we emphasise that the best protocols cannot be selected 'by eye': only detailed calculation seems able to distinguish the mixing capabilities of $\mathcal{P}(2, 3, 1)$ (a good protocol) from those of $\mathcal{P}(2, 3, 12)$ (a poor one), or $\mathcal{P}(2, 6, 2)$ (a reasonable protocol) from $\mathcal{P}(2, 6, 12)$ (a bad one).

We have also found that a useful rule of thumb for eliminating some of the poorest protocols is that for good mixing the inner cylinder should move through a large region of the flow domain (*e.g.*, $\mathcal{P}(2, 3, 1)$), rather than on a more limited, nearly circular path (*e.g.*, $\mathcal{P}(2, 7, 7)$). However, this rule is not sufficiently sophisticated to facilitate selection of the best protocols: for example, presumably it would not discriminate between the relative strengths of $\mathcal{P}(2, 4, 1)$ and $\mathcal{P}(2, 4, 12)$. Furthermore, there are exceptions to the rule, as for instance with $\mathcal{P}(2, 6, 1)$ (which mixes well despite the path of the inner cylinder, which explores relatively little of

the flow domain) and $\mathcal{P}(2, 5, 11)$ (which mixes poorly despite the more extensive path of the inner cylinder).

Another design rule of thumb might be simply to choose a protocol with as great a path length Λ as possible, on the basis that the more you stir the better the mixing. To a certain extent this is the case, the rank correlation coefficients of Λ with DD, IS and RT being $0{\cdot}64$, $0{\cdot}89$ and $-0{\cdot}83$, respectively. (The correlation between path length and EU is of course the highest: $C = 0{\cdot}94$.) However, the rank correlation coefficients of Λ with the remaining six mixing measures are all less than $0{\cdot}4$ in magnitude. We therefore conclude that the variations in mixing effectiveness between various protocols are not simply a matter of some protocols having longer paths than others; there is much to be gained from judiciously selecting the manner in which the inner cylinder moves, as well as merely maximising the length of its path.

References

1. J.M. Ottino, Mixing and chemical reactions: a tutorial. *Chem. Eng. Sci.* 49 (1994) 4005–4027.
2. J.M. Zalc and F.J. Muzzio, Parallel-competitive reactions in a two-dimensional chaotic flow. *Chem. Eng. Sci.* 54 (1999) 1053–1069.
3. M.D. Finn and S.M. Cox, Stokes flow in a mixer with changing geometry. *J. Engng. Math.* 41 (2001) 75–99.
4. M.D. Finn, S.M. Cox and H.M. Byrne, Topological chaos in inviscid and viscous mixers. To appear *J. Fluid Mech.* (2003).
5. H. Aref, Stirring by chaotic advection. *J. Fluid Mech.* 143 (1984) 1–21.
6. H. Aref, The development of chaotic advection. *Phys. Fluids* 14 (2002) 1315–1325.
7. D. D'Alessandro, M. Dahleh and I. Mezić, Control of mixing in fluid flow: a maximum entropy approach. *IEEE Trans. Automatic Control* 44 (1999) 1852–1863.
8. H. Aref and S. Balachandar, Chaotic advection in a Stokes flow. *Phys. Fluids* 29 (1986) 3515–3521.
9. D.M. Hobbs and F.J. Muzzio, Reynolds number effects on laminar mixing in the kenics static mixer. *Chem. Engng. J.* 70 (1998) 93–104.
10. T. Hwu, D. Young and Y. Chen, Chaotic advections for Stokes flows in circular cavity. *J. Engng. Mech.* August (1997) 774–782.
11. S.W. Jones, O.M. Thomas and H. Aref, Chaotic advection by laminar flow in a twisted pipe. *J. Fluid Mech.* 209 (1989) 335–357.
12. J.G. Franjione, C. Leong and J.M. Ottino, Symmetries within chaos: a route to effective mixing. *Phys. Fluids* 1 (1989) 1772–1783.
13. J.G. Franjione and J.M. Ottino, Symmetry concepts for geometric analysis of mixing flows. *Phil. Trans. R. Soc. London A* 338 (1992) 301–323.
14. O.S. Galaktionov, P.D. Anderson and G.W.M. Peters, Symmetry of periodic structures in a 3d mixing cavity flow. *Phys. Fluids* 12 (2000) 469–471.
15. T.S. Krasnopolskaya, V.V. Meleshko, G.W.M. Peters and H.E.H. Meijer, Mixing in Stokes flow in an annular wedge cavity. *J. Mech. B/Fluids* 18 (1999) 793–822.
16. F.H. Ling, Chaotic mixing in a spatially periodic continuous mixer. *Phys. Fluids A* 5 (1993) 2147–2160.
17. V.V. Meleshko and G.W.M. Peters, Periodic points for two-dimensional Stokes flow in a rectangular cavity. *Phys. Lett. A.* 216 (1996) 87–96.
18. T. Hwu, Stretches of fluid materials for Stokes flow in circular cavity. *J. Engng. Mech.* 126 (2000) 554–557.
19. F.J. Muzzio and P.D. Swanson, The statistics of stretching and stirring in chaotic flows. *Phys. Fluids A* 3 (1991) 822–834.
20. F.H. Ling, The effect of mixing protocol on mixing in discontinuous cavity flows. *Phys. Lett. A* 177 (1993) 331–337.
21. D. Rothstein, E. Henry and J.P. Gollub, Persistent patterns in transient chaotic fluid mixing. *Nature* 401 (1999) 770–772.
22. D.M. Hobbs and F.J. Muzzio, The kenics static mixer: a three dimensional chaotic flow. *Chem. Engng. J.* 67 (1997) 153–166.

23. H. Aref and S.W. Jones, Enhanced separation of diffusing particles by chaotic advection. *Phys. Fluids A* 1 (1989) 470–474.
24. F.J. Muzzio, P.D. Swanson and J.M. Ottino, The statistics of stretching and stirring in chaotic flows. *Phys. Fluids A* 3 (1991) 822–834.
25. M. Giona and A. Adrover, Global geometry and coarse-grained formulation of the evolution of pointwise intermaterial interface measure in choatic flows. *Chem. Engng. Sci.* 56 (2001) 3387–3399.
26. F.J. Muzzio, M.M. Alvarez, S. Cerbelli, M. Giona and A. Adrover, The intermaterial area density generated by time- and spatially periodic 2D chaotic flows. *Chem. Engng. Sci.* 55 (1999) 1497–1508.
27. J.M. Ottino, *The Kinematics of Mixing: Stretching, Chaos and Transport*. Cambridge: Cambridge University Press (1989) 364 pp.
28. F. Raynal and J. Gence, Energy saving in chaotic laminar mixing. *Int. J. Heat Mass Transfer* 40 (1997) 3267–3273.
29. V.V. Meleshko and H. Aref, A blinking rotlet model for chaotic advection. *Phys. Fluids* 8 (1996) 3215–3217 (Errata in Phys. Fluids 10 (1998) 1543).
30. M.J. Clifford, S.M. Cox and E.P.L. Roberts, Lamellar modelling of reaction diffusion and mixing in a two-dimensional flow. *Chem. Engng. J.* 71 (1998) 49–56.
31. M.J. Clifford, S.M. Cox and E.P.L. Roberts, The influence of a lamellar structure upon the yield of a chemical reaction. *Inst. Chem. Engng.* 78 (2000) 371–377.
32. G. Metcalfe and J.M. Ottino, Autocatalytic processes in mixing flows. *Phys. Rev. Lett.* 72 (1994) 2875–2878.
33. F.J. Muzzio and J.M. Ottino, Dynamics of a lamellar system with diffusion and reaction: Scaling analysis and global kinetics. *Phys. Rev. A* 40 (1989) 7182–7192.
34. F.J. Muzzio and J.M. Ottino, Diffusion and reaction in a lamellar system: Self-similarity with finite rates of reaction. *Phys. Rev. A* 42 (1990) 5873–5884.
35. D.R. Sawyers, M. Sen and H. Chang, Effect of chaotic interfacial stretching on bimolecular chemical reaction in helical-coil reactors. *Chem. Engng. J.* 64 (1996) 129–139.
36. S.C. Jana, G. Metcalfe and J.M. Ottino, Experimental and computational studies of mixing in complex Stokes flows: the vortex mixing flow and multicellular cavity flows. *J. Fluid Mech.* 269 (1994) 199–246.
37. S.C. Jana, M. Tjahjadi and J.M. Ottino, Chaotic mixing of viscous fluids by periodic changes in geometry: baffled cavity flow. *Am. Inst. Chem. Engrs. J.* 40 (1994) 1769–1781.
38. G.H. Wannier, A contribution to the hydrodynamics of lubrication. *Q. Appl. Math.* 8 (1950) 1–32.
39. E.P.L. Roberts and M.R. Mackley, The simulation of stretch rates for the quantitative prediction and mapping of mixing within a channel flow. *Chem. Engng. Sci.* 50 (1995) 3727–3746.
40. A. Adrover, M. Giona, F.J. Muzzio, S. Cerbelli and M.M. Alvarez, Analytic expression for the short-time rate of growth of the intermaterial contact perimeter in two-dimensional chaotic flows and hamiltonian systems. *Phys. Rev. E* 58 (1998) 447–458.
41. M.M. Alvarez, F.J. Muzzio, S. Cerbelli, A. Adrover and M. Giona, Self-similar spatiotemporal structure of intermaterial boundaries in chaotic flows. *Phys. Rev. Lett.* 81 (1998) 3395–3398.
42. F.J. Muzzio and M. Liu, Chemical reactions in chaotic flows. *Chem. Engng. J.* 64 (1996) 117–127.
43. M.J. Clifford, S.M. Cox and E.P.L. Roberts, Measuring striation widths. *Phys. Lett. A* 260 (1999) 209–217.
44. O. Levenspiel, *Chemical Reaction Engineering*. 3rd ed. New York:Wiley (1999) 668 pp.
45. T. Atobe, M. Funakoshi and S. Inoue, Orbital instability and chaos in the Stokes flow between two eccentric cylinders. *Fluid Dyn. Res.* 16 (1995) 115–129.
46. J. Chaiken, R. Chevray, M. Tabor and Q.M. Tan, Experimental study of Lagrangian turbulence in a Stokes flow. *Proc. R. Soc. London A* 408 (1986) 165–174.
47. H.R. Neave and P.L.B. Worthington, *Distribution Free Tests*. London: Unwin Hyman (1988) 430 pp.
48. R.L. Davidchack and Y. Lai, Efficient algorithm for detecting unstable periodic orbits in chaotic systems. *Phys. Rev. E* 60 (1999) 6172–6175.
49. M. D. Finn, S. M. Cox and H. M. Byrne, Chaotic advection in a braided pipe mixer. *Phys. Fluids* 15 (2003) 77–80.